Portland State University

PDXScholar

Dissertations and Theses

Dissertations and Theses

1995

A Taxonomy of Types of Uncertainty

Byrne Elliot Lovell
Portland State University

Follow this and additional works at: https://pdxscholar.library.pdx.edu/open_access_etds

Part of the [Philosophy Commons](#)

Let us know how access to this document benefits you.

Recommended Citation

Lovell, Byrne Elliot, "A Taxonomy of Types of Uncertainty" (1995). *Dissertations and Theses*. Paper 1396.
<https://doi.org/10.15760/etd.1395>

This Dissertation is brought to you for free and open access. It has been accepted for inclusion in Dissertations and Theses by an authorized administrator of PDXScholar. Please contact us if we can make this document more accessible: pdxscholar@pdx.edu.

A TAXONOMY OF TYPES OF UNCERTAINTY

by

Byrne Elliot Lovell

A dissertation submitted in partial fulfillment of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

SYSTEMS SCIENCE

Portland State University

1995

© 1995 Byrne Lovell

DISSERTATION APPROVAL

The abstract and dissertation of Byrne Elliot Lovell for the Doctor of Philosophy in Systems Science were presented May 2, 1995, and accepted by the dissertation committee and the doctoral program.

COMMITTEE APPROVALS:

Harold A. Linstone. Chair

Martin Zwick

Dundar-Kocaoglu

Michael F. McCoy

Ronald C. Cease
Representative of the Office of Graduate
Studies

DOCTORAL PROGRAM

APPROVAL:

Beatrice Oshika, Director
Systems Science Ph.D. Program

.....

ACCEPTED FOR PORTLAND STATE UNIVERSITY BY THE LIBRARY

by _____ on _____

ABSTRACT

An abstract of the dissertation of Byrne Elliot Lovell for the Doctor of Philosophy in Systems Science presented May 2, 1995.

Title: A Taxonomy of Types of Uncertainty

This study considers an expanded meaning of “uncertainty” as it affects decision-makers. The definition adopted is based on a decision-maker who is uncertain, i.e. aware of the insufficiency of her knowledge for the purpose of rationally determining which option to choose. A taxonomy of uncertainties is developed from this definition. The first stage is a Generalized Decision Model, which expands on a standard decision model often assumed in technical works by allowing uncertainty over components of the model that are assumed to be perfectly known in the standard model. These additional potential “subjects” of uncertainty include the feasibility of options, the authority of the decision-maker to effect a choice, membership of and probability distributions over the set of possible future states of the world, and considerations about how the consequences are to be valued.

The taxonomy also describes possible “sources” of uncertainty, dividing them into characteristics of the world (e.g. variability), evidence the decision-maker has (e.g. ambiguity or imprecision), or characteristics of the decision-maker himself. Other important ways in which uncertainties can vary is whether they are hard (irreducible *in principle*) or soft, whether a

decision is unique or repeatable, and the role time has in the decision and in the resolving of the uncertainties.

A finding of this work is that many uncertainties in addition to the uncertainty in the standard decision model over the future state of the world can keep a procedure for implementing rational choice from being decisive, thus requiring another (nonrational) process to complete the selection of an option.

Other insights:

- Deciding is only part of being rational, and in many instances is not the most important part.
- Uncertainty may complicate decision-making, but is by no means always bad for the decision-maker.
- Rationality is inescapably subjective in any implementation.
- True “decision under certainty” does not exist.
- Uncertainties vary sufficiently that no single treatment can be prescribed; it is hoped that this work contributes to a survey of the territory of uncertainty that facilitates Smithson’s (1988) “suburbanization” or subdivision into smaller tracts to be developed individually.

ACKNOWLEDGEMENTS

Life is the art of drawing sufficient conclusions
from insufficient premises.

Samuel Butler

I can claim sole credit only for the insufficiencies herein. Others have given me essential assistance during my process of drawing conclusions, and drawing this project to a conclusion. I want to thank

My wife Peggy, who has borne up under the uncertainty over whether this project would ever end, and has offered me the certainty of her support and love throughout;

My advisor Hal Linstone for his guidance, and his Multiple Perspectives concept — the importance of O and P Perspectives inspired much of this work;

Marg and Richard, my mother and step-father, for giving me sanctuary and sustenance during my times of direst need; and

The Bonneville Power Administration, and Bob Lamb in particular, for logistical support and encouragement, and for a student's summer job that has grown into a challenging career.

TABLE OF CONTENTS

List of Tables.....	xviii
List of Figures.....	xix
I. Introduction: purpose and assumptions	1
I.1. Preamble: once over uncertainty	1
I.1.1. Uncertainty as a problem	4
I.1.2. Increasingly common concepts in technical arenas.....	7
I.2. Statement of the problem.....	8
I.2.1. Rational choice under uncertainty	8
I.2.2. No comprehensive treatments exist	10
I.3. Research objectives and hypotheses	13
I.3.1. Taxonomy: what types of uncertainties are there?	14
I.3.1.1. Objective	14
I.3.1.2. Methodology	14
I.3.2. Implications of the Taxonomy for rational choice	15
I.3.2.1. Objective	15
I.3.2.2. Methodology	16
I.3.2.3. Complications (in Context Chapter)	16
I.3.2.4. Analysis of implications (the Implications Chapter)	17
I.3.3. Guidelines for decision-makers facing uncertainty	18
I.3.3.1. Objective	18
I.3.3.2. Methodology	19
I.3.3.3. Rational Advice Guidelines	20
I.3.4. Hypotheses	21
I.4. Uncertainty defined for this research.....	22
I.4.1. Working definition of uncertainty	27
I.5. Scope: requirements for this research	28
I.5.1. Comprehensive	29
I.5.2. Realistically applicable	30
I.5.3. Errors of hysteresis and hyperrationality.....	32
I.6. The author's perspective	35
I.6.1. Intended audience	36
I.6.2. Use of quotations	37
I.6.3. Use of terms.....	37
I.6.3.1. Use of "I" to emphasize perspectives	37
I.6.3.2. Pronouns and gender	38

I.6.3.4. The term “decision-maker”	38
I.6.3.5. Glossary	39
II. Context: perspectives on uncertainty	40
II.1. Introduction	40
II.2. The Multiple Perspectives Concept	41
II.2.1. The Technical Perspective	45
II.2.2. The Organizational Perspective	47
II.2.3. The Personal Perspective	50
II.2.4. Perspectives on uncertainty found in the literature	53
II.2.5. Science, policy analysis, and the T Perspective	55
II.2.6. Perspectives employed here	59
II.3. The Technical Perspective on uncertainty	60
II.3.1. Introduction to “T”	60
II.3.2. Rationality	61
II.3.2.1. General rationality	61
II.3.2.2. Economic rationality	63
II.3.3. Procedural rationality, bounded rationality and satisficing	66
II.3.3.1. Procedural vs. substantive rationality	66
II.3.3.2. Bounded rationality	68
II.3.3.3. Satisficing	69
II.3.3.4. Prescriptive satisficing	71
II.3.3.5. Strategic vs. parametric rationality	72
II.3.3.6. Metarationality	73
II.3.4. Decision selection processes	74
II.3.4.1. Definition	74
II.3.4.2. Standards for selection processes	80
II.3.4.3. Rational selection processes	81
II.3.5. “Failures” of rationality	82
II.3.5.1. Substantive failures	83
II.3.5.2. Procedural failures	85
II.3.5.3. Metarationality failures, i.e. hyperrationality	87
II.3.6. Non-rational choosing	90
II.3.7. The Standard Model of Decisions and the Expected Utility Hypothesis	92
II.3.7.1. Von Neumann and Morgenstern’s Expected Utility Hypothesis	93
II.3.7.2. Specifics of the Expected Utility Hypothesis	96
II.3.7.3. Statement of the Hypothesis	99
II.3.7.4. Example of uses of tools based on the SDM	100
II.3.8. The relationship between uncertainty and probability	102
II.3.8.1. Historical development of views of probability	102
II.3.8.2. Four major views of probabilities	104
II.3.8.3. Assessing probabilities	108

II.3.8.4. Summary of probability	112
II.3.8.5. Uncertainty that doesn't fit into probabilistic frameworks	116
II.3.9. Possibility Theory	117
II.3.9.1. Belief Theory: A generalization of Possibility Theory	119
II.3.10. Complexity and uncertainty	122
II.3.11. Measurement	125
II.3.12. Tools of the T Perspective	126
II.3.12.1. Decision analysis	127
II.3.12.2. Simulation modeling	133
II.3.12.3. Game theory	135
II.3.12.4. Sensitivity analysis	137
II.3.13. The silver lining for "T"	139
II.4. Organizational Perspectives on uncertainty	140
II.4.1. Introduction to "O"	140
II.4.1.1. Planning and decision-making literatures	143
II.4.2. Organizational rationality	145
II.4.2.1. Economic rationality and the SDM	146
II.4.2.2. Alternatives to economic rationality for organizations	147
II.4.2.3. Satisficing	148
II.4.2.4. Ultra-rationality	149
II.4.2.5. Rationality as news processing	150
II.4.3. Organizational decisions: the myth of the "decision event"	151
II.4.3.1. Overemphasis on selection	154
II.4.4. Decision-making and planning	155
II.4.4.1. A "standard" approach to planning	156
II.4.4.2. Strategic planning	157
II.4.5. Probability in organizational planning	159
II.4.6. Organizational structure as a function of uncertainty	160
II.4.7. The environment of the organization	164
II.4.8. Uncertainty and power	167
II.4.9. Organizational strategies for coping with uncertainty	170
II.4.10. The silver lining for "O"	173
II.5. Personal Perspectives on uncertainty	175
II.5.1. Introduction to "P"	175
II.5.2. Common personal responses to uncertainty	178
II.5.3. Judgment under uncertainty: cognitive strategies and heuristics	182
II.5.3.1. Representativeness	182
II.5.3.2. Availability	183
II.5.3.3. Adjustment and anchoring	184
II.5.3.4. Simulation	185

II.5.3.5. Conservative information processing	186
II.5.3.6. Hindsight	187
II.5.3.7. Confirmation bias	187
II.5.3.8. Overconfidence	188
II.5.3.9. Debiasing	189
II.5.3.10. Neurological bases for heuristics	190
II.5.4. Decision-making under uncertainty: experimental evidence	192
II.5.4.1. The certainty effect	192
II.5.4.2. The reflection effect	194
II.5.4.3. Very low probabilities	194
II.5.4.4. The isolation effect	194
II.5.4.5. The “framing” of decisions	195
II.5.4.6. Utility evaluation effect	196
II.5.5. Alternatives to MEU	199
II.5.5.1. Prospect Theory	199
II.5.5.2. Regret Theory	203
II.5.5.3. Other alternative theories	204
II.5.5.4. Comments on Regret Theory and Prospect Theory	205
II.5.5.5. Janis and Mann’s Conflict Model of Decision	207
II.5.6. The silver lining for “P”	209
II.6. A Sociological Perspective on uncertainty	211
II.6.1. Taxonomy of ignorance	212
II.6.2. Historical progression of attitudes toward uncertainty	213
II.7. Dror’s Technical/Organizational Perspective	217
II.7.1. Concept package for uncertainty	217
II.7.2. Relevance of concept package	221
II.7.3. Fuzzy gambling	221
II.7.4. Planning as fuzzy gambling aid	222
II.7.5. Implications	222
II.7.6. Normative guidelines	224
II.8. Summary of the main three perspectives	226
II.9. Previous taxonomies of uncertainty	226
II.9.1. Knight: uncertainty versus risk	226
II.9.2. Thompson and Tuden	226
II.9.3. Dror	228
II.9.3.1. Type	228
II.9.3.2. Hardness	229
II.9.4. Smithson	229
II.9.5. Klir	230
II.9.6. Friend and Jessop	230
II.9.7. Dubois and Prade	231
II.9.8. Osborn, Hunt and Jauch	232
II.9.10 Morgan and Henrion	233

II.10. Interactions among perspectives	234
II.10.1. Power	235
II.10.2. Risk aversion as an organizational characteristic	236
II.10.3. S.O.P.s	236
II.10.4. Organizational homogeneity	237
II.10.5. Perverse interactions	237
II.10.5.1. Friend and Jessop's study	237
II.10.5.2. Lorraine's study	238
II.10.7. Analysts versus decision-makers	240
III. Taxonomy of uncertainties	242
III.1. Introduction	242
III.1.1. Subjects	243
III.1.2. Characteristics	244
III.1.3. Sources	244
III.1.4. Decision context	244
III.2. Subjects of uncertainty	244
III.2.1. Option uncertainty: which ones are feasible?	247
III.2.2. Authority uncertainty: who decides?	251
III.2.3. State of the world uncertainty	253
III.2.3.1. Known-odds uncertainty	253
III.2.3.2. Likelihood uncertainty	255
III.2.3.3. Possibility uncertainty	263
III.2.3.4. The non-discrete world	265
III.2.3.5. Implications of probability or possibility error	268
III.2.4. Function or model uncertainty	270
III.2.5. Evaluation uncertainty: how desirable is any specific consequence?	277
III.2.6. Multiple-objective uncertainty: what is the relative importance of different dimensions of consequence or value? ..	281
III.2.7. Time distribution uncertainty: how are consequences occurring at different times to be balanced?	286
III.2.8. Stakeholder uncertainty: whose values are relevant?	290
III.2.9. Gamble preference uncertainty: what features of gambles (lotteries) are desired?	293
III.2.9.1. Risk	295
III.2.9.2. Certainty	298
III.2.9.3. Security level and hope level	301
III.2.9.4. Regret	305
III.2.9.5. Gamble evaluation rule	309
III.2.10. A Generalized Decision Model (GDM)	312
III.2.11. Beyond decision models	314
III.3. Characteristics of uncertainties	315
III.3.1. Hardness	315
III.3.2. Locus of uncertainty (data vs. model)	324

III.3.3. Parametric vs. strategic	329
III.4. Sources of uncertainty	333
III.4.1. The world	335
III.4.1.1. Natural world	335
III.4.1.2. Actors	353
III.4.2. Evidence	370
III.4.2.1. Data, measurement	371
III.4.2.2. Linguistic evidence < interaction 2 >	380
III.4.2.3. Evidence from actors	386
III.4.3. The Decision-maker	391
III.4.3.1. Processing capability < interaction 1 >	392
III.4.3.2. Ability to interpret evidence < interaction 2 >	396
III.4.3.3. Mental models	399
III.4.4. Interactions	408
III.4.4.1. Interaction 1: complexity and processing power	408
III.4.4.2. Interaction 2: linguistic uncertainty and linguistic ability	410
III.4.4.3. Interaction 3: incomplete evidence and decisions not yet made	411
III.4.5. Summary of sources	414
III.5. Decision context	417
III.5.1. Unique vs. repeated decisions	417
III.5.2. Equilibrium	420
III.5.3. Surprise-prone environments	423
III.5.4. Low-probability, high-impact events	427
III.5.5. Types of environments	431
III.5.6. The Role of time	433
III.5.6.1. Uncertainty as a function of time	433
III.5.6.2. Speed of our learning and acting vs. rate of change in world	434
III.5.6.3. Un-simplifying the passage of time	436
III.5.7. Once-and-for-all vs. correctable decisions	441
III.5.8. Second-order consequences: anticipating the next decision	442
IV. Implications of uncertainty	445
IV.1. Introduction	445
IV.1. An assessment of rationality in light of uncertainty	446
IV.1.1. Uncertainty as a threat to rationality	446
IV.1.2. "Null Hypothesis": assess significance of uncertainty	449
IV.1.3. Summary of when rationality is likely to be useful	452
IV.1.4. Subjective rationality	453
IV.2. An assessment of rational methods in light of uncertainty	459

IV.2.1. Mathematical approaches in general	459
IV.2.2. Measurement is fundamental	460
IV.2.3. Optimality (and satisficing)	462
IV.2.3.1. "Optimizing behavior" and fuzziness	462
IV.2.3.2. "The best" is nonexistent	464
IV.2.3.3. When is optimizing promising?	465
IV.2.4. Probability	466
IV.2.4.1. Objective, subjective and fuzzy probabilities	467
IV.2.4.2. Distributional nature of probabilities	470
IV.2.4.3. Summary of probabilities	471
IV.2.5. Possibility Theory	471
IV.2.5.1. Probability and possibility compared	471
IV.2.6. Maximization of Expected Utility (MEU)	474
IV.2.6.1. Whole-gamble features	474
IV.2.6.2. Subjective probabilities and probability estimation	475
IV.2.6.3. The distributional aspect of probability	476
IV.2.6.4. Use of mathematical expectation for one-shot choices	477
IV.2.6.5. The meaning of utility	478
IV.2.6.6. Hysteresis and the non-uniqueness of utility functions	478
IV.2.6.7. MEU does not clarify uncertainty	480
IV.2.6.8. MEU does not facilitate learning	480
IV.2.6.9. Information requirements	481
IV.2.6.10. "Small worlds"	482
IV.2.6.11. Completeness of preferences	483
IV.2.6.12. Complexity of MEU assessments	485
IV.2.6.13. MEU versus satisficing	485
IV.2.6.14. Summary	486
IV.3. A Perspective on Normative Decision Theory	491
IV.3.1. Relative emphasis	491
IV.3.1.1. Over-emphasis on selection	491
IV.3.1.2. Under-emphasis on preferences	492
IV.3.1.3. Over-emphasis on state-of-the-world uncertainty	493
IV.3.1.4. The role of decision support	493
IV.3.2. Better tools needed	494
IV.3.2.1. Tools for fuzzy data	494
IV.3.2.2. Displaying uncertainty	496
IV.3.2.3. Possibility Theory and Fuzzy Set Theory	496
IV.3.2.4. Chaos Theory	497
IV.3.2.5. Tools for evaluating consequences	497
IV.3.2.6. Complexity	499
IV.3.3. Smithson's four stages	500
IV.3.3.1. Stage 1: exclusion and banishment	500
IV.3.3.2. Stage 2: inclusion and reductionism	501
IV.3.3.3. Stage 3: second-order strategies	501

IV.3.3.4. Stage 4: suburbanization and relativism	503
IV.3.3.5. Tension: suburbanization vs. comprehensive view ...	505
IV.4. Miscellaneous	506
IV.4.1. The Law of Requisite Variety	506
IV.4.2. Information, attention and the cost of certainty	507
IV.4.3. Total uncertainty and sub-type trade-offs — silver lining?	508
IV.4.4. Probable knowledge	509
IV.4.5. Is social science too hard?	509
IV.5. Future research	511
IV.5.1. Uncertainty = $f(\text{Complexity} - \text{Knowledge})$??	511
IV.5.2. Testing risk aversion vs. security level	513
IV.5.3. “Optimizing behavior”	514
IV.6. Potential applications of this work to theory	515
IV.6.1. T Perspective help for O Perspectives	515
IV.6.2. Alienation	516
IV.6.3. Gamblers and natural selection in the organization	517
V. Guidelines	518
V.1. Introduction	518
V.2. The context of the Rational Advice Guidelines	520
V.2.1. Purpose of decision support	520
V.2.1.1. Rational decision support	521
V.2.1.2. Assessing the quality of decisions made under uncertainty	522
V.2.1.3. Decision-makers’ rational need for information	525
V.2.2. The Rational Advice Guidelines	526
V.2.2.1. Meta-rationality guidelines	527
V.2.2.2. Substantive rationality guidelines	528
V.2.2.3. Procedural rationality guidelines	528
V.2.2.4. Summary of the RAG	533
V.2.3. Diverging from economic rationality	534
V.2.3.1. Adoption of MEU	534
V.2.3.2. Consistency of preferences	534
V.3. The problems of uncertainty	537
V.3.1. What can be recommended?	538
V.4. Advice for problems of fuzziness	539
V.4.1. Reduce fuzziness, obtain more information	541
V.4.1.1. Better measurement	541
V.4.1.2. Experiment	541
V.4.1.3. Function uncertainty	541

V.4.1.4. Values clarification	542
V.4.1.5. Fuzzy time comparisons	543
V.4.1.6. Multiple objective evaluation uncertainty	543
V.4.1.7. More powerful comparison methods	545
V.4.2. Social choice processes	545
V.4.2.1. Agree on process	545
V.4.2.2. Agree on representative or on government	546
V.4.3. Choosing among fuzzy options	547
V.4.3.1. Adopt Null Hypothesis	547
V.4.3.2. Reduce fuzziness	547
V.5. Advice for problems of gambling	547
V.5.1. Reduce gambling aspect of options (de-gamble)	548
V.5.1.1. Improve options	548
V.5.1.2. Mitigate risk	549
V.5.1.3. Improve response capability	550
V.5.1.4. Aggregation of decisions	551
V.5.2. Choosing among risky options — smart gambling	552
V.5.2.1. Elimination of dominated options	552
V.5.2.2. Consider all relevant dimensions of difference across options	553
V.5.2.3. Decision rules and how to decide	553
V.6. Advice for problems of fuzzy gambling	555
V.6.1. Factors nearly always causing fuzzy gambling	556
V.6.2. Improve options	557
V.6.2.1. Feasibility of options	557
V.6.2.2. Authority of decision-maker	557
V.6.3. Mitigate risks	558
V.6.4. Prediction — state-of-the-world and function uncertainty	559
V.6.4.1. Uncertainty due to actors	560
V.6.4.2. Awareness of possible error or possible deception	561
V.6.4.3. Conflicting data	562
V.6.4.4. Conflicting, incomplete or improvable models	563
V.6.4.5. Probabilities	565
V.6.4.6. Complexity	565
V.6.4.7. Fundamental dynamism	566
V.6.5. Surprise-proneness	567
V.6.6. Low-probability, high-impact	568
V.6.7. Modify values	569
V.6.8. Choosing among fuzzy gambles	569
V.6.8.1. Assess gambles as accurately as possible	570
V.6.8.2. In the end, you must gamble	570
V.7. Bounded, meta- and hyper-rationality advice	571
V.7.1. Bounded rationality, limitations of the decision-maker	571

V.7.1.1. Improve processing capability	571
V.7.1.2. Time and resource limits	571
V.7.2. Considerations of time	572
V.7.2.1. Longer term approaches	572
V.7.2.2. Some time is available	572
V.7.2.3. No time available	573
V.7.2.4. One-shot vs. correctable decisions	573
V.7.2.5. Unique vs. repeated decisions	573
V.7.2.6. Delay decision until uncertainty is resolved	573
V.7.2.7. World changes fast compared to us	573
V.7.3. Meta-rationality advice	574
V.7.3.1. Focus on the big picture	574
V.7.3.2. Allocate efforts sensibly	574
V.7.3.3. False certainty	574
V.7.3.4. When uncertainty is high	575
V.7.4. Hardness and softness of uncertainties	575
V.7.4.1. Soft uncertainty	576
V.7.4.2. Hard uncertainty	576
V.7.4.3. Virtually hard uncertainty	576
V.7.4.4. Malleable uncertainty	577
V.7.5. Hyperrationality	577
V.7.5.1. Employ a Null Hypothesis	577
V.7.6. Other considerations	578
V.7.6.1. Utility value of consequence	578
V.7.7. Clarify extent of uncertainty for informed decision- making	578
V.8. Schematic diagram (overview guide)	579
V.8.1. Metarational considerations	579
V.8.2. Treatments of uncertainties from different sources	583
V.9. An irony	583
VI. Conclusions	586
VI.1. Background	586
VI.1.1. Definition of uncertainty	586
VI.1.2. Decision selection processes can fail to be decisive in two ways	587
VI.2. Taxonomy	587
VI.2.1. Subjects of uncertainty	587
VI.2.2. Characteristics of uncertainty	589
VI.2.3. Sources of uncertainty	590
VI.3. Implications	591
VI.3.1. Uncertainty can thwart rationality as a selecting process	591

VI.3.2. Meta-rationality and the Null Hypothesis	593
VI.3.3. Rationality is inescapably subjective.....	594
VI.3.3.1. “Objective rationality” is a myth	594
VI.3.3.2. Rational choice of rational methods is subjective	594
VI.3.4. “Decision under certainty” doesn’t exist.....	594
VI.3.4.1. Virtually all decisions are fuzzy gambles	595
VI.3.4.2. Some choices may be clear nonetheless	595
VI.3.5. Measurement	596
VI.3.5.1. Measurement is central	596
VI.3.5.2. “The best choice” does not exist	597
VI.4. Contributions.....	597
VI.4.1. Fundamental contribution: the Taxonomy, a conceptual scheme.....	597
VI.4.2. This work as a survey of the territory of uncertainty	598
VI.4.3. Useful for tool builders	600
VI.4.4. Useful for decision-makers	600
VI.4.5. Useful to other researchers	601
VI.4.6. Potential applications.....	602
VI.4.6.1. Uncertainty sub-type trade-offs	602
VI.4.6.2. Risk aversion and security level preference	602
VI.4.6.3. Complexity as a function of information and uncertainty	603
VI.4.7. Insights into a limit of rationality	603
VI.5. Limitations of this work	604
VI.5.1. It is a first cut	604
VI.5.2. It is untested	605
VI.5.3. Design limitations	605
VI.5.4. Technical focus.....	606
VI.6. Review of initial hypotheses	606
VI.7. Future work to extend this research.....	608
VI.7.1. Continue refining the survey of the territory of uncertainty	608
VI.7.2. Identify specific needs and attempt to develop better tools	609
VII. References consulted	612
VII.1. Cited works	612
VII.2. Quotations without specific citations	629
VIII. Glossary	632

LIST OF TABLES

Table	Page
1. Certainty vs. Uncertainty	23
2. Features of the T, O and P Perspectives	53
3. Features of science and policy analysis practice	56
4. Features of science and policy analysis function	57
5. Optimizing and satisficing compared.....	72
6. Definition of terms for electric utility DA example	130
7. Consequence values in DA example.....	131
8. Utility values in DA example	132
9. Expected utility values in DA example	133
10. Modified expected utility values in DA example	138
11. T, O and P Perspectives on uncertainty.....	227
12. Thompson and Tuden's decision modes	228
13. Expected utility example	257
14. Modified expected utility example	258
15. Probability and possibility uncertainty.....	266
16. Error in assessing possibility uncertainty	270
17. Effect of discount rate	288
18. Risk vs. security level.....	514
19. Certainty vs. uncertainty	633

LIST OF FIGURES

Figure	Page
1. Unambiguous permission	5
2. Ambiguous permission	6
3. Perspectives on uncertainty.....	55
4. Actions and rationality	62
5. Unambiguous selection	76
6. Ambiguous selection.....	76
7. Meta-, substantive and procedural rationality	84
8. The Standard Decision Model.....	94
9. Utility example decision tree	134
10. A payoff matrix.....	136
11. A typical weight function π	201
12. A typical value function v	202
13. Smithson's Taxonomy of Ignorance.....	212
14. Option uncertainty	247
15. Authority uncertainty	251
16. State-of-the-world uncertainty	253
17. Likelihood uncertainty	255
18. Possibility and probability uncertainty.....	267
19. Possibility uncertainties	269
20. Function or model uncertainty	271
21. Evaluation uncertainty	277

22.	Multiple objective uncertainty	281
23.	Time distribution uncertainty	286
24.	Stakeholder uncertainty	291
25.	Sample utility curve	296
26.	A Generalized Decision Model	313
27.	Hardness of uncertainty	317
28.	Locus — data vs. model	326
29.	Three sources of uncertainty	334
30.	Uncertainty from the world	335
31.	Uncertainty from natural world	336
32.	Uncertainty from actors	354
33.	Uncertainty from un-made decisions	355
34.	Uncertainty from evidence	371
35.	Uncertainty from data	372
36.	Uncertainty from linguistic evidence	380
37.	Uncertainty from evidence from actors	386
38.	Uncertainty from the decision-maker	391
39.	Uncertainty from processing capability	392
40.	Uncertainty from ability to interpret evidence	396
41.	Uncertainty from mental models	399
42.	Uncertainty from all sources	415
43.	Un-simplified passage of time	440
44.	Meta-rational guide to rational choice	580
45.	Treatments of uncertainties from different sources	584
46.	A partial map of the territory of uncertainty	599
47.	An example utility curve	638

I. INTRODUCTION: PURPOSE AND ASSUMPTIONS

In a predestinate world, decision would be *illusory*; in a world of perfect knowledge, *empty*; in a world without natural order, *powerless*. Our intuitive attitude to life implies non-illusory, non-empty, non-powerless decision... Since decision in this sense excludes both perfect foresight and anarchy in nature, it must be defined as choice in face of bounded uncertainty.

G. L. S. Shackle

I.1. Preamble: once over uncertainty

Some call ours the Age of Uncertainty (Laszlo 1988:482). It is becoming more and more common to give explicit consideration to the uncertainty we face when attempting to make good decisions or good plans but have imperfect knowledge of the relevant factors. We often realize that to be sure what the “best” decision is we need to know something — tomorrow’s oil price, next month’s sales, next year’s real estate appreciation — that we do not actually know. There are many reasons for the increasing prominence of uncertainty. One is a dissatisfaction with methods that have, for the sake of simplicity, expediency, or tractability, assumed the availability of knowledge we do not possess; we are demanding greater sophistication from our decision methods. The spectacular increase in automated computing power we have seen in the last two decades has made it possible to employ complex solution methods on some large problems for the first time. The solutions produced by these methods have not always completely solved the original problem. For instance, very large econometric models are now computationally very feasible, but still cannot

predict up-turns and down-turns in the economy as well as we would like. The very size of the programs that are now feasible to execute poses another problem: huge programs inevitably contain errors, and the errors in such highly complex programs are much harder to find. We are not able to write huge, error-free programs. Breaching computational barriers has emphasized the hardness of the problems, and has dashed hopes that increased computing power will allow us to eliminate uncertainty.

Another reason is that many of our decision situations are actually becoming more complex and consequently more uncertain. In some cases, we have solved the easier problems, and that has increased the relative importance of the unsolved, more difficult problems. For instance, we can cure many diseases: if you contract pneumonia, we know how to cure it; we have even eliminated smallpox. We often can tell how you got such a disease. But if you develop cancer, we have treatments that offer only a “probability” of success, often depicted as a 5-year survival rate. And how do people get such a disease? Exposure to certain chemicals carries with it a probabilistic risk, perhaps estimated in per cent risk per unit of exposure per year — but some of the unexposed may also get it. Instead of the simplicity of the contagious diseases where you are safe unless exposed, and even then safe if you are treated, we face only differences in likelihoods. You have a small chance of getting cancer no matter what, and the chances merely increase with exposure to agents that are usually invisible, making even the exposure uncertain.

The increasing power of some of our measurement tools can lead to direct confrontation with uncertainty. The author and his wife were told by

their obstetrician that the level of alpha fetal protein (AFP) in their 15- or 16-week fetus was a little bit low. Low levels of AFP are have been associated statistically with an increased risk of Down's Syndrome. This risk is ordinarily 1 in 333 for a woman of the mother's age. If the fetus were 16 weeks old, the low AFP level would be associated with a 1 in 250 risk, high enough to qualify for free amniocentesis, which could detect Down's Syndrome with near certainty. However, if the fetus were only 15 weeks old, the level of AFP would not be quite so serious, and would indicate a risk of only 1 in 280, not serious enough for free amniocentesis. In addition, the process of obtaining the amniotic fluid carries with it a 1-in-1000 risk of inducing a spontaneous abortion. In such a situation, there is no rock of certainty to hold on to anywhere, and the decision tree that the author dutifully constructed was not useful in this choice, even for a probabilistically sophisticated decisionmaker. One problem is the great difficulty of comparing the utility of consequences such as a child being born with Down's Syndrome or the inducing of a spontaneous abortion — the mathematical machinery is more advanced than the technology for assessing quantitatively the personal meaning of non-financial impacts.

The increasing potency of our technologies is said to be “making the world smaller”, which means that the effects of any actions may carry farther around the globe, faster and faster. For example, *program trading* by large institutional investors can quickly magnify a small change in the stock market into a large one as many different computer trading programs all indicate “sell” at once. This increase in the uncertainty of the market arises from an implicit link in the decision-making of previously

more separate investors via the common elements of trading program design. Whereas trends in human decision-making in the stock market work at human speeds and can be monitored and regulated, trends in computerized decision-making happen at super-human speeds, and are much more difficult to moderate or interrupt. The growing interconnectedness increases the chances for an action to affect other parts of the globe, other groups of people, other parts of our own economic or social systems, etc., in ways we cannot predict.

1.1.1. Uncertainty as a problem

“Uncertainty” is a natural term, understood reasonably well by all speakers of English:

1. the state of being uncertain; doubt; hesitancy.
2. an instance of uncertainty, doubt, etc.
3. unpredictability; indeterminacy or indefiniteness. (Random House 1980)

Uncertainty may refer to any state between complete ignorance and possession of sure knowledge. We might extend this continuum from knowledge toward ignorance even farther, into “negative knowledge”, i.e. error — a feeling of certainty about a proposition that is incorrect. But for the current work, intended to help someone who seeks guidance, the awareness of the uncertainty is important. True error, of which an agent is unaware, will be treated only sparingly. The emphasis here will be on the first of the three connotations of uncertainty described above.

Most of the time we perceive our own uncertainty to be a problem. A common circumstance is that we feel a need to make a decision or take action, and are uncertain about what is best to do. Frequently we can

identify an action that is likely to turn out well should a particular state of the world develop, but that would turn out poorly if a different state of the world were to develop. This makes the decision difficult and unpleasant.

There are times, however, when uncertainty can be perceived to be an asset. President Reagan was disinclined to take steps to reduce American emissions of air pollutants that Canadians claimed caused acid rain in their country. Reagan said that scientific uncertainty about the nature of acid rain made such steps inappropriate at the current time. For Reagan, the scientific uncertainty was advantageous. And of course, the uncertainty your adversary experiences is an advantage (von Neumann and Morgenstern 1953:147) (unless you are linked to him in a doomsday machine and his error could cause you great harm).

There is a sense in which uncertainty and freedom are linked. When one has only one course of action available, there is no uncertainty and no freedom. When one is certain which of several actions is best, one has little freedom to do otherwise.

Within organizations, uncertainty (indefiniteness) about the rules can permit greater freedom of action than the same set of rules without uncertainty would allow. As many observers of organizational politics have noted, "It is better to ask forgiveness than permission." Consider the drawing here of an idealized division of possible

Figure 1.
Unambiguous permission

actions into “Permitted” and “Forbidden”, shown first without any uncertainty about which actions belong in which category (figure 1).

There is often a zone of ambiguity between the acts clearly permitted and those clearly punishable, as shown in figure 2. Asking permission is a request to find that an act is in the permissible category; performing the act and then asking for forgiveness is a request to find that the act is not in the punishable category. The implicit emphasis on conservatism will lead a bureaucrat to define any category as narrowly as possible. In the first case, the category to be defined is *permissible*, which is then minimized, and the zone of ambiguity is counted in the punishable category.

In the second case, it is the punishable category that is defined conservatively, and the zone of ambiguity is implicitly in the permissible category.

There is also a sense in which uncertainty is essential to our ability to find personal meaning in life:

If life were eternal all interest and anticipation would vanish. It is uncertainty which lends it satisfaction. (Kensho Hoshi)

The main but not sole focus here is on the problematical aspects of uncertainty.

Figure 2. Ambiguous permission

1.1.2. Increasingly common concepts in technical arenas

“Uncertainty” and “uncertainties” are terms heard more and more frequently in planning circles as planners and decision-makers struggle to design and choose courses of action when the future in which those actions will play out cannot be adequately known. Market shares and sizes, technological capabilities, regulatory constraints, fuel prices, or production time can have major impacts on the success of an organization’s actions. However, these factors cannot always be known well enough to allow accurate prediction of what the outcome of the actions will be. These factors are increasingly often treated as explicit *uncertainties* within planning processes. A major aim of this research is to demonstrate which analytical techniques are appropriate for decision-makers to use when treating different types of uncertainty.

“Uncertainty” is the subject of a substantial and growing body of literature.¹ A great deal of this work is based upon a mathematical treatment of the subject, with the goal of stating and decisively proving hypotheses. This has often required adoption of fairly narrow initial

¹ e.g., Aiginger 1987, *Production & Decision Theory under Uncertainty*; Allais & Hagen 1979, *Expected Utility and the Allais Paradox—Contemporary Discussion of Decisions Under Uncertainty with Allais’ Rejoinder*; Carter & Ford, (eds.) 1972, *Uncertainty and Expectation in Economics*; Dror 1986, “Planning as Fuzzy Gambling: A Radical Perspective on Coping with Uncertainty”; Ford 1983, *Choice, Expectation and Uncertainty*; Hart 1951, *Anticipations, Uncertainty and Dynamic Planning*; Hey 1979, *Uncertainty in Microeconomics*; Kahneman, Slovic, & Tversky (eds.) 1982, *Judgment under Uncertainty: Heuristics and Biases*; McKenna 1986, *The Economics of Uncertainty*; Munier (ed.), *Risk, Decision and Rationality*; Raiffa 1968, *Decision Analysis: Introductory Lectures on Choice under Uncertainty*; Shackle 1949, “Probability and Uncertainty”; & Southern California Edison Company 1988, “Planning for Uncertainty”.

assumptions or axioms, which can restrict the range of real situations to which the conclusions apply. At the core of this literature is the Expected Utility Hypothesis² (EUH) of von Neumann and Morgenstern (1947), which was built upon foundations laid much earlier in the work of Bernoulli (e.g. 1713) and others who analyzed the popular gambling games of their time. The decision principle of maximization of expected utility (MEU) which EUH supports has been treated as both normative and descriptive by various authors, and a great many criticisms have been made of its suitability in either capacity. The subject of expected utility is examined more closely below in section II.3.7., the Standard Decision Model and the Expected Utility Hypothesis.

I.2. Statement of the problem

I.2.1. Rational choice under uncertainty

Acceptance in our culture of rationality as a fundamental principle for decision-making is so deep that it is seldom explicitly acknowledged or questioned. The essence of this rationality, which we can call means-ends rationality, is “that we seek things which we want, which may be our own or other people’s pleasure, or anything else whatever, and our actions are such as we think most likely to realize those goods.” (Ramsey 1926:35). Rationality is the fundamental decision-making principle in this study, and it is examined in some detail below in the Context chapter. When it is possible to determine which actions (means) are feasible, and to determine which is most likely to realize the desired goods (ends), this is a straight-

² EUH is described in the Technical Perspective section in the Context chapter below.

forward, even simple, guideline. However, this rule can be very difficult to apply when we lack critical information about a factor in this equation — i.e. when we are *uncertain*. At such times, rationality may be indeterminate. This is a theme Jon Elster develops extensively in

Solomonic Judgments:

The central argument is that rationality itself requires us to recognize this limitation of our rational powers [i.e. the possible indeterminacy of rationality], and that the belief in the omnipotence of reason is just another form of irrationality. (Elster 1989a:vii)

Failure to recognize the indeterminacy of rational-choice theory can lead to irrational behaviour. (Elster 1989a:2)

One of the primary aims of this research is to explore closely when this indeterminacy might occur, how to identify it, and what a rational person can do when it does. It will develop that indeterminacy of rationality can be caused by a great many different types of uncertainty.

We could be uncertain about what actions are feasible. We could be uncertain about the consequences of actions. We could be uncertain about the desirability of those consequences in light of our values. We could be uncertain about the relative importance of different values, or how to combine conflicting valuations of a set of consequences when more than one agent participates in the choosing. Most contributors to the literature on choice under uncertainty have focused on uncertainty of consequence, and have often explicitly assumed that the set of choices was fixed in advance, that the valuing of consequences presents no problems (and, in fact, that the function that maps consequences to values, the utility

function, possesses attractive mathematical properties³), and that value is a uni-dimensional attribute. There are exceptions to this generalization: multiple criteria decision-making (MCDM) receives attention⁴, as does the problem of reconciling different preferences among members of a group (social choice theory⁵), and there are theories that propose disaggregations of the attractiveness of a consequence, e.g. into rejoicing and regret⁶. The main problem addressed here, then, is how to decide or plan rationally when the usual interpretation of rationality is complicated by one or more types of uncertainty.

1.2.2. No comprehensive treatments exist

While many authors have written about uncertainty, there are very few sources that deal explicitly with the widely varying nature of uncertainty and the implications of that varying nature for planning and deciding. There is an extensive methodological literature to be found in the area of Decision Analysis, Economics, and Planning. These works typically focus on a method, or perhaps a comparison of a few methods, for application to a particular type of uncertainty. The literature on Scenario Analysis illustrates this: all of the articles in this field listed in the References chapter claim that Scenario Analysis can be helpful in a planning process,

³ E.g. continuity, monotonicity or differentiability.

⁴ Zeleny [1982] provides a 37-page bibliography of MCDM up through 1980.

⁵ E.g., Barry & Hardin [1982], Blair & Pollak [1983].

⁶ See prospect theory [Kahneman & Tversky 1979] and regret theory [Loomes & Sugden 1982].

but few of them indicate the kinds of planning uncertainty for which Scenario Analysis is recommended and the kinds for which it is not, and none of them suggest alternatives for the cases where Scenario Analysis is not the preferred approach.

Similarly, Howard Raiffa, one of the most respected authors in Decision Analysis, devotes the 300 pages of *Decision Analysis: Lectures on Choice Under Uncertainty* (1968) to an exposition of Decision Analysis, and does not even touch on the subject of when to use it and when to use another approach! The book does have a few pages on the “pros and cons of decision analysis”, but even this section is completely non-specific with regard to the type of problem for which the tool may be used. Raiffa’s discussion of different types of uncertainty consists of two paragraphs in which he allows that, indeed, uncertainty about the quality of the next item to roll off a much-studied assembly line is different from uncertainty about the state of the American economy 20 years from now, but that once our judgment about the economy is translated into probability estimates, we can use the same techniques for either situation.

The planning literature contains many well-written discussions of planning that include specific reference to uncertainty. A worthy example is *Strategic Planning and Forecasting* by William Ascher and William Overholt (1983). The index refers the reader to 16 different parts of the book for more detail on uncertainty, and the authors have a full and sensitive appreciation of the importance of uncertainty in planning. Donald Michael (1973) has also written insightfully about the role and purpose of planning in an uncertain environment. However, these authors

do not at any one time talk about the different types of uncertainty that planners may face, nor do they discuss how a planner might select treatments for different types. The primary subject is planning, and uncertainty enters the picture as a complication.

There is much less written on a theoretical plane about planning uncertainty. Yehezkel Dror (1986a, 1986b, 1987) is nearly alone in discussing *types* of uncertainty with more sophistication than the simple certainty-risk-uncertainty trilogy of decision-making modes that dates to Knight's (1921) distinction⁷. Dror has described several distinctions among uncertainties that he thinks must be included in a minimum "concept package" for discussing uncertainty, though he does not propose any concept package himself. My Taxonomy draws heavily on his writing. All of his work, though, and especially his work on uncertainty, is highly abstract, and provides little practical assistance to a real-world planner.

There is, then, no integrated treatment of uncertainty in planning that I have been able to find, and I will attempt to provide one. This dissertation presents a concept package (a Generalized Decision Model, the Taxonomy of Uncertainties and the Rational Advice Guidelines), and uses

⁷ F. H. Knight divided decision situations into Certainty, Risk, and Uncertainty. In the former, the outcome of each decision is known with certainty. *Risk* implies that there are two or more outcomes for at least some decision options, and the probabilities of each are known. *Uncertainty* again postulates a set of possible outcomes for some decision options, but the probability distribution is not known. As Dror points out, there is a fourth category of uncertainty which he has called *Qualitative Uncertainty* (distinguishing it from the traditional Uncertainty, which he calls *Quantitative Uncertainty*). In this fourth type of uncertainty, the possible outcomes themselves are not known. This situation, of course, can present problems even for techniques that are adequate for treating Quantitative Uncertainty.

that concept package as a framework for an analysis of the applicability of available tools to specific types of uncertainty. Uncertainty itself is the primary focus, with decision-making and planning serving as the background context. I provide an analysis of the types of uncertainty and methods for treating them that indicates which tools and approaches can be useful for which types of uncertainty.

1.3. Research objectives and hypotheses

This research is intended to accomplish three major objectives. A very brief summary of each is followed by sections describing them in greater detail.

1. *Taxonomy.* What types of uncertainty are there that can complicate the task of deciding or planning rationally? What distinctions can we draw among the kinds of decision and planning factors about which we can be uncertain that will help to clarify the nature and appropriate handling of those uncertainties?

2. *Analysis of Implications.* What are the implications of the variety of uncertainties described in the Taxonomy for rational choice, for the most important traditional technologies of rational choice, and for the field of decision support?

3. *Guidelines.* What guidelines can be provided for decision-makers facing uncertainty that can be used to clarify the characteristics of the uncertainty and identify some tools that may be useful, given the nature of the uncertainty and the context of the decision or planning task?

1.3.1. Taxonomy: what types of uncertainties are there?

1.3.1.1. Objective

When a decision-maker or planner speaks of facing “uncertainty”, there are a great many aspects of the deciding or planning task to which he may be referring. The traditional analytic literature for the most part addresses only uncertainty about which state of the world will obtain. Many other types of uncertainty can complicate planning and decision-making. My objective is to identify all of the major types of uncertainty that are relevant to planning and deciding tasks.

Describing the various types of uncertainty permits an analysis of their characteristics — the kind of information that is lacking or uncertain, the typical sources of such information, means of eliminating or reducing the uncertainty, etc. This accomplishes two purposes. It provides a framework for a more comprehensive analysis and discussion of uncertainty than can be achieved with many of the typical treatments in the current literature. Secondly, it permits a corresponding analysis of the available tools and approaches in order to match the uncertainty problems with uncertainty treatments. This can help to guide people who must decide or plan under uncertainty to choose more appropriate treatments.

1.3.1.2. Methodology

To lay the groundwork for answering this question, I present an extension of the standard technical model of a *decision* in the first part of Chapter II. This Generalized Decision Model (GDM) extends the usual model of decisions to explicitly encompass uncertainty regarding sets of choices, alternative states of the world, utility functions, the passage of

time, and other complications. This generalized model helps to show the various components of a decision task, which facilitates identifying separate types of uncertainty that can complicate rational choice.

The main part of this objective will be met by creation of a functional Taxonomy of Uncertainties, a conceptual scheme created from the results of investigations into the kinds of distinctions that are useful to draw. Many examples are presented, drawn in large part from the electric utility industry. A conceptual scheme that helps to distinguish among different types of uncertainty will be useful in assessing which treatments (approaches, tools) may be helpful for dealing with each type, and also for considering how much of the entire range of uncertainties the current decision support field addresses.

1.3.2. Implications of the Taxonomy for rational choice

1.3.2.1. Objective

This is the technical version of the normative portion of this work — the Guidelines chapter presents the essence of this material in a form more directly intended for use by decision-makers. The primary question to be treated here is, “what insight into the possible benefits and limitations of rationality for decision-makers under uncertainty does a fuller depiction of types of uncertainty yield?” The heart of this section is analytical, assessing the characteristics of rationality and many rational approaches, including the dominant implementation of rational choice, maximization of expected utility (MEU), and assessing the implications for them of types or extents of uncertainty.

1.3.2.2. Methodology

Meeting this objective requires two preliminary steps. First, the subject of decisions and decision-making needs to be reviewed, with as full and broad a perspective as possible. This is a main purpose of the Context chapter. The “complications” which arise in decision-making that make all instances more complicated than any neat model can suggest need to be surveyed. These will include the natural responses to uncertainty made by individual human beings and organizations and the use of subjective probabilities. Second, the Taxonomy of Uncertainty Types is required; this is chapter III, described above.

1.3.2.3. Complications (in Context Chapter)

When an approach to handling uncertainty that is attractive in theory is implemented, many complications can arise. A major source of complications is the characteristics, especially limitations, of the person or organization making the decision. I review the literature on the natural responses to uncertainty made by both individuals and groups. Kahneman and Tversky and others have identified many general biases in human decision-making that are important to consider. Others have described the typical responses organizations make. I also review the natural responses made by people steeped in a technical perspective.

Existing research on normative and descriptive theories of individual and organizational responses to uncertainty is the major source of information used in the Context chapter. The (natural) technical responses are inferred from the general tenor of treatment suggested in articles in management science, economic, or industry journals such as the journal of

the Electric Power Research Institute, and general management and planning journals, such as *Long Range Planning* and the *Harvard Business Review*. Specific attention is given to the Expected Utility Hypothesis and to different views of the meaning and appropriate uses of subjective probability. This review describes the natural tendencies which my later recommendations will have to take into account in order to ensure that they are feasible. Another task of this section is to examine the conditions under which use of subjective probabilities can be helpful, and the conditions that make attempts to maximize expected utility appear to be reasonable.

1.3.2.4. Analysis of implications (the Implications Chapter)

The implications chapter presents conclusions about the potential usefulness and limitations of rationality and rational methods, based on the context and taxonomy chapters. The theoretical bases and justifications for the methods are assessed, and compared to the characteristics of the types of uncertainty and in consideration of the purposes of planning activities. The result is a set of conclusions regarding the circumstances in which the methods can be helpful for dealing with each type of uncertainty.

Next, the implications of the broadened spectrum of uncertainties for the field of decision support are considered. A four-stage model developed by Smithson describing natural phases in the evolution of mainstream treatments of uncertainty is applied to decision support and used to put the current work into a developing historical context.

Finally, some miscellaneous results of the analysis are presented, along

with some ideas for further research.

1.3.3. Guidelines for decision-makers facing uncertainty

1.3.3.1. Objective

The purpose of the Guidelines is to help a decision-maker who faces uncertainty select some approaches or methods that are likely to be helpful for specific types of uncertainty (and to help the decision-maker avoid those approaches that are unlikely to be helpful). It is based on the results of the Analysis, and has been constructed to assist real-world planners, taking into account the typical setting in which most explicit deciding and planning is performed. Important characteristics of such settings include the resistance to change that is usually displayed by both individuals and organizations, and the kinds of expectations that other parties may have regarding the practices of the organization. Some kinds of change in uncertainty treatment will require only technical change, essentially substituting new methodologies at the level of the analyst within an organization, whereas other kinds of change involve not changes in how the analysts do their work but rather the purposes their work is expected to support, and some of these changes can only be accomplished by senior management.

Finally, since I want to know which methods will be useful, I need to establish criteria for determining whether application of a specific method in a particular situation will be useful, which I am taking to mean that it will generally promote “better decision-making”. I am calling these criteria the “Rational Advice Guidelines” (more on them below).

I.3.3.2. Methodology

The Guidelines in Chapter V present the results of the technical analysis ordered by types of uncertainty a decision-maker can face, with additional analysis of meta-rational concerns.

Chapter V is intended to be useful to a real-world decider or planner who is able to employ it with judgment: the decider is assumed to be more like the “prudent man” of common law than like a computer⁸. The specifics of any choice situation are very important, and taking them all into account can make any decision appear to be very complicated. As Zadeh states in his Principle of Incompatibility (section II.3.10), complexity can make precision and relevance impossible to obtain simultaneously. The implications for this work are that it is impossible to make recommendations that are relevant to all decision-makers and are still precise enough to provide specific guidance.

Therefore, the Guidelines are not intended to identify which particular decision option a decision-maker ought to choose, but to point to decision support actions that are likely to be helpful to the decision-maker as he approaches his choice.

⁸ It is ironic that my work, which is intended to be useful for people, may have more in common with the work of some computer programs than with the traditional literature on Decision under Uncertainty: attempts to create artificial intelligence applications are often employing a kind of knowledge engineering that relies on heuristics instead of hard-and-fast rules. I think the crucial distinction in this area is not between human and computer users of the rules, but between artificial and real worlds as the context in which the rules will need to be interpreted. The real world contains too many “shades of gray” to permit exact definition or specification of all guidelines.

I.3.3.3. Rational Advice Guidelines

The Rational Advice Guidelines (RAG), at the beginning of the Implications chapter, lay out general criteria for assessment of the rationality of decision procedures. My research indicates that there are usually two criteria defined for rationality: 1) means/ends rationality, and 2) completeness and consistency of preference⁹. As my focus is on the kinds of rational *action* that are available to real-world decision-makers instead of on the characteristics of rational *actors*, I have replaced preference consistency, which refers to actors, with a prohibition of arbitrariness of choice, which refers to actions. As I discuss in more detail later, real preferences may be inconsistent (intransitive and/or cyclic). What is lost by giving up the consistency criterion is analytic power: permitting only consistent preferences makes it more likely that the situations we do consider can be resolved by rationality alone, though this means that there are actual situations that we cannot even consider. Omitting it, we can address a wider range of decisions, but are faced with the possibility that some decisions cannot be decided solely by rationality. This is not a surprising result for people with experience in the real-world contexts of human decision-making, but it is not highlighted in the technical literature on rationality or on decisions under uncertainty. It is similar to Zadeh's Principle of Incompatibility, just cited.

The implementation of means/ends rationality remains a major issue. David Gauthier (1982) has claimed that "in our society the received

⁹ Discussed below in section II.3.2.2 on Economic Rationality in the Context chapter.

conception [of rationality], which most persons do accept initially given their socialization, does identify rationality with maximizing activity.” While accepting means/ends rationality, I have not accepted its identification with maximization of expected utility (MEU). I do want to impose another criterion, reflecting the prescriptive nature of this work: a procedure for deciding should not be arbitrary, that is, it should not point to one option instead of another due to the influence of factors that have little or no bearing on the utility of the options. While this may make any given decision procedure less decisive, it will prevent a kind of outcome that I argue is undesirable. For instance, if our data is poor, and has a margin of error of +/- 50%, we would not want to accept a difference of 1% in estimated utility as decisive: we would say that a difference that small in comparison with the much greater and acknowledged uncertainty in the data is not significant. To decide on such a basis would be committing an error of arbitrariness. The study reviews treatments of “rationality” that are especially relevant for assessing behavior (contrasted to abstract ideals of rationality) in order to identify criteria useful for the present purpose.

1.3.4. Hypotheses

1. Distinguishing among the various ways in which uncertainty can complicate planning and decision-making can help to identify approaches that are useful to the planner or decision-maker facing a particular type of uncertainty.
2. Approaches that are helpful in some circumstances can be unhelpful or harmful if used in other circumstances.
3. The definition of uncertainty usually assumed in technical

treatments of uncertainty can be expanded to include more of the types of uncertainties that managers encounter in their work.

4. The technical approach can be shown by rational arguments to be inadequate if used alone.

1.4. Uncertainty defined for this research

In the broadest view, uncertainty is the flip side of knowledge, whose study has engaged the best human minds for millennia. In order to delimit the current topic, the working definition of uncertainty used here will be restricted to those situations in which the agent (one or more people or organizations needing to decide or plan) has some awareness of the need for some kind of information, and an awareness of the currently deficient supply of that information. There are three different levels of uncertainty that we can distinguish here: at the most basic level, the state of the world may be variable or indeterminable; at the next level, our information about the state of the world could be incomplete; and at the last, a psychological level, we could be aware of the incompleteness of our information. It is this third sense in which I use “uncertainty”; clearly, the first two levels are very important, as they are major causes of our feeling uncertain.

Consider the last two levels: one concerns what we actually know; the other, what we think we know. If we think we know something and actually know it, our state of knowledge is *true certainty*, as illustrated in table 1. If we think we know something but are not actually in a position to know, and therefore “ought” to be uncertain about it, we are in *false certainty*. (If we think we know something, are in a position to know it, but are wrong, then we are simply in *error*: we have drawn improper

inferences, perhaps, when proper inferences could have been drawn.) Techniques that disguise the true variability of the world may lead us into false certainty, as can an excess of simplifying assumptions. (Making an assumption may simply be a way of defining an uncertainty out of our awareness without changing our actual information, leaving us blind to risks that could be severe.)

If we believe that we do not know, but actually do, we are in *false uncertainty*. (This case must be distinguished from

Table 1. Certainty vs. Uncertainty

	We have adequate information	We have inadequate information
We think that we know	<i>True certainty</i>	<i>False certainty (Overconfidence)</i>
We think that we don't know	<i>False uncertainty (Underconfidence)</i>	<i>True uncertainty</i>

believing that one doesn't know, guessing anyway, and finding out eventually that the guess was correct. The point is that at the time the guess was made, *it was a guess* — and therefore true uncertainty, not false uncertainty.) It is easy to think of psychological examples using individual persons who are un-self-confident and actually know more than they realize. Examples in the utility business are much harder to come by. (Underconfidence does not appear to be a very common characteristic of modern organizations or the people who run them.) Lastly, if we believe that we do not know, and truly do not, we are in *true uncertainty* — the focus of this work. The distinction between false certainty and true uncertainty lies not in our state of knowledge, nor in any variability in the world, but rather in our belief about our knowledge, and obviously affects

our choice of actions importantly. Will Rogers identified the real danger: “It isn’t what we don’t know that gives us trouble, it’s what we know that ain’t so.” (Weinberg 1975:1) By becoming more aware of uncertainty and its consequences for rational action, we can lessen the danger Rogers identified.

The essence of this distinction is whether we have enough information to allow “reasonable” trust in our conclusions, or whether we believe that our conclusions ought to be treated as guesses that may very well be wrong. This is a distinction that can easily be lost through the ordinary communication processes within organizations, or may be gradually forgotten over the course of time. For instance, a forecaster, asked a question about the future value of some factor, may respond that 30 is his best guess, but there are so many factors that it could easily be as low as 20 or as high as 50—indicating that he is quite uncertain¹⁰. The person receiving the answer may pass it on to someone else, omitting the qualification: “Forecasting says it will be about 30.” The answer now *appears* to be a certain one, though the apparent certainty is clearly “false” in the sense described here. The picture may be even worse, of course, as the forecaster may be overestimating his own knowledge in the first place, perhaps due to reliance on estimates from other persons whose own uncertainty has already been inadvertently lost or forgotten, or to the effects of simplifying assumptions that have been made.

¹⁰ But it also indicates, as is common, that the forecaster is not in a state of *complete ignorance* — he is able to estimate the bounds of variability or error of his forecast.

Even for the uncertainty-sophisticated person the question of how much trust to put in an estimate is very difficult. What inferences can be “reasonably” made from the data available at any time? It isn’t likely that any absolute answers can be given, though considerations about the process used to draw inferences can illuminate the issue, as suggested by the following questions: Is all the relevant information being considered? How do we know? Which is being left out? Are only people within the same industry culture being asked to make judgments about what is relevant, or are several different perspectives (e.g. rate payers, environmentalists, community residents, share holders, futurists, politicians, utility analysts, utility executives, etc.) being heard? Are efforts being made to seek new information? Are early warning signs of changes in current trends being looked for or being seen? Is the level of accuracy of the data measuring processes taken into account?

An illustration of false certainty is provided by the consensus in the Pacific Northwest (PNW) electric utility industry on forecasts of loads in 1970, which was essentially that we knew that loads would continue to increase at about 7% per year. While we can acknowledge that the 30-year historical trend up to then had indeed been very consistent, there were other signs, and in fact warnings from people within the industry (people often thought of as rebels, iconoclasts, or doomsayers) that the likelihood of a serious decline in the rate of load growth was increasing rapidly. The growth rate plummeted from 7% to nearly 0%¹¹, and ratepayers in the

¹¹ Total firm sales of electricity within the Pacific Northwest region grew by only 0.26%
— continued on next page —

region are paying far higher rates today because such overconfidence led to the ill-fated plan to build five nuclear plants at the same time.¹²

Even today it is informative to look closely at the expectations of an industry that “learned its lesson”. Was the lesson, after load growth plunged in the late 1970s, that load growth had become less predictable and planners ought to be uncertain about it and therefore make plans that are less vulnerable to shifts in load growth patterns (i.e. more robust)? No, it appears that many in the Northwest concluded that in the late 1980s we were merely on a “low load-growth path”¹³, and assumed that low growth rates would continue. They examined the painful experience of over-estimating load growth, and concluded that our trouble came not from false certainty but from mere error (believing one proposition when another was true). They adjusted to the notion of lower load growth, and would then be surprised if load growth is high. (There have been notable advances in forecasting: BPA moved to use of a range of forecasts rather than a single forecast. This directly acknowledges that any forecast is uncertain. BPA’s resource planning began to include contingency options

from 1979 to 1987.

¹² The infamous “Whoops” plants were planned under the auspices of the Washington Public Power Supply System, or WPPSS. Of these five plants, only one is producing electricity: two have been canceled, and two have been “mothballed” after substantial construction work was performed (approximately 1/3 of one and 2/3 of the other was completed).

¹³ Mathematical models of utility systems can be programmed to simulate periods of high or low (or in-between) load growth, either at random or as specified by a user, and the trajectory of loads over the course of the simulated years is often called a “path”. Analysts might then look at what happened during the run that used the “low load-growth path”. Application of this term to the real world, though, suggests considerably more inevitability to the continuation of a trend than appears to be justified.

as part of its mainstream approach to ensuring adequate supplies under a range of possible loads. The ranges, however, reflect primarily the variability that is consistent with structural stability, and largely neglect uncertainty effects of possible surprise events or trends.)

Southern California Edison, a large electric utility based in Los Angeles, undertook to change its approach to resource planning in light of the uncertainty of their environment and the recognition that their forecasts were seldom accurate enough to permit exact matching of loads and resources (So. Cal. Edison 1988). They noted that while their resource acquisitions had all been designed to be optimal for the single future Edison saw as most likely at the time the resource acquisition decisions were made, the futures that materialized differed enough from the one envisioned that Edison's resources were generally far from optimal. In a scenario planning project, Edison examined a variety of possible futures, and concluded that the best approach was to have a large variety of smaller resources available. While this set of resources might not be optimal in any specific possible future, it would be flexible enough to be satisfactory in many.

1.4.1. Working definition of uncertainty

These are the two ways in which I will use the term "uncertainty":

1. A felt or consciously acknowledged, relevant lack of satisfactory knowledge.
2. A specific factor, variable or aspect of the decision-making or planning situation that is not known as well as the planner would like to know them, e.g. fuel price, the nature of the feedback between two elements in a system, the rate at which a particular change will occur over the next decade. (While this usage refers to a "factor" in the world itself, it is the uncertainty

on the part of the planner about that factor that qualifies it as an uncertainty.)

The focus is on the psychological condition of uncertainty because decision-makers can use only the knowledge they possess. There are relatively few ways in which it makes sense to say that some factor in the world is itself uncertain: some processes are inherently stochastic and unpredictable (e.g. the time when a single radioactive isotope will decay); some processes are apparently chaotic and in principle not amenable of such precision of measurement to permit prediction (e.g. weather systems). Both for the above situations and for others where the world may be deterministic but the decision-maker cannot predict the future, the relevance for this work is that the decision-maker is uncertain about the future. This perspective is similar to Smithson's conception of ignorance, which holds that it is not

objectively measurable but understandable mainly in subjective terms. The unknown may well be 'out there' in the environment, but its representation as a form of ignorance resides in the mind. (Smithson 1988:292-93)

Uncertainty can act as a mass noun, "We face increasing *uncertainty*", or as a singular noun, "We will deal with the fuel price *uncertainty* by creating a financial buffer, and with the precipitation *uncertainty* by using a Monte Carlo simulation model to calculate the expected value." The context will generally make clear which meaning is intended.

1.5. Scope: requirements for this research

To serve planners and decision-makers, this work needs to meet two major criteria. The first is that the coverage of uncertainties must be comprehensive, and include any type of uncertainty that planners and

decision-makers might encounter. The second is that the approaches and tools recommended must be realistically applicable and potentially helpful in the situations planners and decision-makers may find themselves in. A third goal is to avoid hysteresis in the application of recommended methods, that is, to guard against the danger that our conclusions may be driven by factors that are substantively irrelevant.

1.5.1. Comprehensive

A thorough treatment of uncertainty ought to be able to answer a decision-maker or planner who asks, “I need to make a plan or to decide, but I am uncertain about *X* (and *Y* and *Z* ...); what would it be helpful for me to do?”¹⁴ I will not try to assist decision-makers who are uncertain due to false uncertainty (i.e. underconfidence, discussed later), or who do not have enough information but don’t realize it (i.e. overconfidence).

This dissertation is intended to provide the skeleton of a very broad normative heuristic that can guide decision-makers or planners who face uncertainty. Since that amounts to the proffering of advice, it will be restricted to consideration of situations in which there is a party potentially capable of and interested in making use of advice on applied rationality, typically a person or an organization. An unorganized crowd, a group of people without avenues of communication, or a person who adopts a fundamental choice principle other than rationality (e.g. divine inspiration) may face uncertainty, but this dissertation will not address

¹⁴ Note that the question “... what *should* I do?” tends to elicit a different kind of answer. It both adds a moral obligation and the sense that there is a single correct action.

their needs. As mentioned earlier, this dissertation will not treat (except as a danger to avoid) the very difficult problems of false uncertainty (feelings of uncertainty accompanied by actually adequate information) or error.

Granted the existence of a rational party, though, this dissertation needs to address any uncertainty that might be encountered in the Rational Advice Guidelines, whether one of choice, consequence, or value. The Taxonomy of Uncertainties must be sufficiently comprehensive, as must the Implications, that any such uncertainty can be assessed. The time scale addressed will range from negligible, where the outcome of a decision is known almost immediately, to infinite, where the outcome is virtually never known with certainty¹⁵. The uncertainty confronted may arise from a single, imperfectly known factor, or may be an environment characterized by overwhelming unknowability in nearly all factors. Both one-time decisions and identical, repeated decisions are included.

1.5.2. Realistically applicable

The fundamental standard against which possible guidance will be measured is that it must be helpful, and this requires both substantive and procedural rationality (see the T Perspective section in the Context chapter). Substantive rationality requires that the guidance promote the realization of the decision-maker's values as well as possible. Procedural rationality requires that the guidance be feasible in actual planning and

¹⁵ Robin Pope [1988] suggests that this is the situation facing a researcher deciding among competing approaches to a problem: scientific hypotheses are accorded belief in continuous degrees, rarely at the extreme of total certainty. Thus, in many cases, the uncertainty is never completely resolved.

decision-making contexts; it must not demand such exquisite skill in execution that misapplication can easily be more hazardous than the uncertainty it treats; its adoption must promote awareness and discourage blindness; it must not require knowledge that is unattainable¹⁶; it must not restrict the decision-maker's preferences except in trivial ways (e.g. it must not require her to have a twice-differentiable utility function¹⁷). This guidance need not be comforting: it is better for a planner to help the decision-maker to *do* better throughout the future than to *feel* better in the present moment. And, though it is of course to be hoped that such a conflict would not (or even *could* not) arise, it is more important in this context to demonstrate convincingly that the guidance is genuinely helpful than it is to be able to prove analytically that it is perfectly optimal. This comports with Dror's calls for a "philosophy of policy-reasoning" and his

¹⁶ Economists generally assume that a great deal of information is available as they model the Firm in Perfect Competition, as illustrated by John Hey's comments [Hey 1979]: "uncertainty about the amount the firm is able to sell at the prevailing price, or about the volume of inputs it is able to buy at their given prices ... not only are contrary to the letter of the conventional framework, but also violate the whole spirit of the model" [pp. 126-127]. Even when extensions have been made in modeling the firm under uncertainty, "virtually all the models assume that the probability distributions of any stochastic variables are known to the decision-maker. Clearly, this is a very strong assumption" [p. 166]. Von Neumann and Morgenstern recognized that they "cannot avoid the assumption that all subjects of the economy under consideration are completely informed about the physical characteristics of the situation in which they operate and are able to perform all statistical, mathematical, etc., operations which this knowledge makes possible." [1953:30]

¹⁷ For example, Karl Aiginger [1987] assumes that utility functions $u()$ are known and differentiable, with $u'' < 0$ or $u'' > 0$ everywhere. He notes that he accepts the assumption that firms are von Neumann-Morgenstern utility maximizers "not because of its supposed closeness to reality but because of its operationality and its strict demands on the rationality of the economic units" [p. 6]. For some of his work, Marc Fleurbaey [1988] needs to postulate that $u()$ is increasing, strictly quasi-concave, and cost and demand functions are twice continuously differentiable with first derivatives strictly positive.

suggestions for evaluating it:

As contrasted with the philosophy of science, which accepts criteria of truth and peer acceptance independent from potential utility, philosophy of policy-reasoning should be evaluated in the main according to its potential contributions to improvement of policymaking. (Dror 1988b:143)

1.5.3. Errors of hysteresis and hyperrationality

The approach taken in this work is to identify uncertainty and arbitrariness as clearly as possible. Methods should not be relied upon in situations in which their results are driven by factors that are substantively irrelevant, such as the order in which alternatives are compared, or which of several substantively equivalent technical procedures are employed in estimating numerical quantities.

The point is not that we should discard belief statements that are the outcome of a causal process; all belief statements are causally produced. Rather the point is that we should exclude *belief-irrelevant* causal processes, i.e. causal processes where the causes include not only *reasons*, but also irrelevant causal factors such as mental fatigue ..., a preference for symmetry (why not 50 per cent?) or an aversion to symmetry (surely 50 per cent is too simplistic). ... [W]e should not hold to or act upon a belief if we suspect it to be the result of such belief-irrelevant processes. (Elster 1979:130)

This criterion follows directly from the decision to require both procedural and substantive rationality. Whereas we accept the outcomes of courtroom trials because we agreed beforehand on the procedure to be used, in normative decision theory we must both follow an acceptable procedure and insist that a procedure actually performs the functions it purports to, and additionally avoid hysteresis, occurring when irrelevant procedural details have a large influence on the outcome.

We ought to be rational about how we go about being rational.

Otherwise we are guilty of hyperrationality,

defined as *the failure to recognize the failure* of rational-choice theory to yield unique prescriptions or predictions. As in Kant's critique, the first task of reason is to recognize its own limitations and draw the boundaries within which it can operate. The irrational belief in the omnipotence of reason is what I call hyperrationality. (Elster 1989a:17)

Elaboration of what we can reasonably expect from applied techniques of rational choice is a major theme of this work.

While uncertainty is in one sense the flip side of information, this treatment of uncertainty is not the flip side of Information Theory or Decision Theory. Information Theory starts from abstract principles, and bridges can be made towards the real world to allow applications. Decision Theory, and economic theories such as the Model of the Firm under Uncertainty, begin with a highly restricted model, generally deterministic, and attempt to broaden it by introducing uncertainty into one (or rarely more) factor in the model. The current treatment of uncertainty takes as fundamental the perspective of the human decision-maker in the real world, who may not have any genuine certainty, and attempts to develop principles that will be helpful.

The two approaches are essentially *dual* to each other: Information Theory and other abstract theories start in an analytical context of optimality, and can be extended toward feasibility; this treatment of uncertainty begins with the requirement of feasibility, and attempts to work toward optimality. To some extent, this approach is *top-down* in contrast to the *bottom-up* approach that the economic and mathematical schools take: I want to define and bound the arena of uncertainty, divide it

up according to relevant dimensions, and determine what rational advice can be given as of the mid-1990s to decision-makers facing each of the types of uncertainty; the mathematical approach is to define very carefully a set of initial conditions that allow derivation of results that must hold, with the implicit intention of eventually defining enough sets of initial conditions to span the space of uncertainty types.

I think this approach is in the spirit of John Hey's conclusions about what must be done to move economic modeling from the domain of equilibrium to the broader domain of disequilibrium:

Rather than extend indefinitely optimality analysis to increasingly complex environments, we should be modelling true disequilibrium through the use of *reasonable* 'rules of thumb'. I stress the word 'reasonable', as I believe actual behavior is generally reasonable: people do not go around using stupid 'rules of thumb' — certainly, not for very long! Naturally, I would like to say what 'reasonable' means, and how one models 'reasonable' behaviour; unfortunately, my thoughts along these lines are insufficiently well-developed to commit to print. (Hey 1981:242)

Of course, clarifying "reasonable" is both essential and difficult; I hope to make a contribution to that effort in this work.

The Rational Advice Guidelines are brief and general, and should appear to be reasonable; I do not intend to prove them to be the best or only such guidelines. The level of rigor in the analysis again is intended to establish the reasonableness of the recommended actions; specification of conditions that would be precise enough to permit analytical proof of each recommendation in an analysis of such a broad range of uncertainties would be prohibitively encyclopedic. The level of support should be enough to convince most open-minded, if skeptical, readers. I have opted to emphasize relevance at the expense of precision and certainty of result,

since this is a complex subject, making precision and relevance incompatible (see Zadeh's Principle of Incompatibility in the sub-section on Complexity and Uncertainty in the T Perspective section in the Context chapter, section II.3.10).

I.6. The author's perspective

The choice of perspective employed in this work reflects several considerations. As this work is intended to be useful as a guide to decision-makers, the focus is primarily prescriptive. The type of rationality treated here must not be so demanding that non-idealized decision-makers are disqualified.

Much of the literature on rationality has employed a very generalized or idealized representation of the decision-maker and decision situations. Generalization is a major tool of science: the less-important features of a particular instance are ignored, leaving the core features that are shared by all similar instances. In order for this method to succeed, it must be the case that there exists a critical subset of features that can be abstracted and that represent enough of the situation that conclusions derived from the critical features are still relevant to the particular instance when the less-important features are once again included in the analysis. Whether or not a particular instance of rational analysis is successful in producing a unique prescription depends in large part on the details of the situation. In this way, science and policy analysis or systems analysis work in opposite directions: science induces from the details to the law; systems analysis deduces from the law to the particular situation. This is discussed further in the Context chapter.

I summarize the literatures, and then in the later chapters with my original work, challenge the T Perspective (Technical Perspective: see section II.2 below) by pushing it to its limits in the service of providing advice to uncertain decision-makers. I believe this technical or rational perspective shows itself to be incomplete, much as Jon Elster uses rational arguments to show that rationality has limits, and that expecting too much of it is, itself, a form of irrationality.

This will be done by emphasizing a technical perspective of decisions, and then expanding the scope of what is viewed by that technical perspective to include types of uncertainty that have traditionally been given little attention in the technical literature but are significant in any application of rationality in human decision-making. Such applications are virtually always complicated by the specifics of the situation; these specifics and the inevitable ways in which an abstract and simplified view of people and decisions omits key details can be illuminated by organizational and personal perspectives. The limitations of a purely technical view and of technical rationality will be made more clear by considering the impacts of the diverse uncertainties described below in the Taxonomy.

1.6.1. Intended audience

This work is geared for a technically-minded reader. The Appendix provides definitions of some of the terms that are key to following the arguments here. While the Multiple Perspectives concept will be used to enhance the analysis, there are others with much wider and subtler grasps of Multiple Perspectives than I. My emphasis will be on rational and

technical grounds for broadening one's horizons to include more non-rational and non-technical considerations. In this light, any contributions I can make to the Multiple Perspectives world will be more in the way of helping to strengthen the foundation and show more clearly the need rather than by elaborating the concept itself.

1.6.2. Use of quotations

The following quotation describes self-referentially the attitude I have adopted towards the use of quotations vis-à-vis paraphrasing the references I use:

I find distinct advantages in using many quotations. Sometimes they emphasize a point I have made; sometimes they make the point. I see no reason to paraphrase someone else's observations if they have made them well. Quotations provide a change of stylistic pace; they show that others share the position I am developing (at least to the extent quoted); and they imply there is further understanding to be gleaned from the quoted reference. (Michael 1973:12)

1.6.3. Use of terms

1.6.3.1. Use of "I" to emphasize perspectives

The T Perspective generally assumes that there is a single "objective" view to which all rational people should subscribe upon receipt of sufficient information. Perhaps that is part of the reason that the first person is so seldom used in technical writing: it would suggest that there is something unique about the author's perspective, and this uniqueness could only be a source of difference from the correct, objective viewpoint. Outside of that technical perspective, however, in organizational, personal or other perspectives, it is very difficult to maintain that there is any single perspective with a unique and normative claim on our loyalties. When

many perspectives are possible, and any individual or group must be employing one or more, intellectual fairness requires us — here I feel I must shift from the generic “us” to the particular “me” — to be open about the assumptions that are affecting how I see the world and the conclusions I draw from that view. I will use first person often throughout this work: these words are not “the truth”; indeed, a major implication of these words is that there is no single truth. Rather, they represent the viewpoint of a single person who has immersed himself in this subject and has drawn some conclusions. Though I feel I have not demonstrated that any reasonable person *must* accept these conclusions, or that there are no other conclusions which a reasonable person may derive, I invite my readers to try them out and see if they seem reasonable and useful to you, as they do to me.

1.6.3.2. Pronouns and gender

Using only masculine pronouns for generic references is arbitrary and, more significant, systematically biased. Claims that this practice has deleterious effects on our culture seem plausible to me. Substituting either “s/he” or “he or she” for “he” has a deleterious impact on readability, which in a document like this one is precarious enough. I have chosen to alternate more or less randomly use of feminine and masculine pronouns. This certainly results in arbitrary usage, but with reduced systematic bias.

1.6.3.4. The term “decision-maker”

The frequent references to the “decision-maker” are not meant to imply that there is a single, clearly-identified person with complete and sole authority to make the decision. It is a stylistic convenience to use that

term, which is already bulky enough, instead of more awkward phrases that suggest each time that there may be more than one person, or more than one group or organization, sharing some aspect of the decision-making. Except where specifically noted, as in the discussion of the Standard Model of Decision (section II.3.7), there may actually be more than one party making the decision, any of which may be organizations, and the values to be used may be those of the “decision-maker” or those of other parties (including postulated stakeholders such as nature, endangered salmon, the environment, future generations not yet born, etc.) on whose behalf the “decision-maker” acts.

1.6.3.5. Glossary

A short glossary appears after the References chapter, describing some of the key terms from the literature for convenient reference and some of the terms important for understanding examples from the electric utility industry.

II. CONTEXT: PERSPECTIVES ON UNCERTAINTY

Uncertainty is an uncomfortable state; but
certainty is a ridiculous one.

Voltaire

II.1. Introduction

This chapter describes the general outlines of the literatures that are relevant to uncertainty in planning and decision-making. A great deal of work has been done in several quite different fields, and I will highlight each of them and attempt to provide a framework for considering them simultaneously.

The core of the integrating framework is the Multiple Perspectives concept (e.g. Linstone 1984). The first section of this chapter introduces the concept and defines the major perspectives that will be used in reviewing the literature. These perspectives are then matched against the identifiable fields within the literature. Following the reviews of each perspective's literature on uncertainty, I summarize the strengths and weaknesses for given purposes of the different perspectives and the tools that each provides.

An important theme running through this work is the distinction between prescriptive and descriptive approaches. The main business of the natural sciences is to describe what the world is like. The social sciences attempt to describe what individuals and groups of people are like and how they behave. But more than that, the social sciences are often asked by our society to speak to the question of how individuals and groups *ought* to

behave. Normative theories of rationality and normative theories of decision are attempts to answer that question.

Jon Elster notes that theories can fail because they are indeterminate, and do not give unique predictions, or because they are inadequate, and their predictions are false (Elster 1989a:1). Elster thinks the latter problem is more serious. Rational choice theory is primarily a normative theory, secondarily explanatory, according to Elster: "It tells people how to act and predicts that they will act in the way it tells them to" (Elster 1989a:1). If the theory fails due to indeterminacy, it will not produce a unique prediction. When it does produce a unique prediction, people may fail to behave in the predicted way, and it is important to distinguish between the possible reasons. People may not know how to behave in the predicted way, though they can be taught to do so. In this case, the normative theory may or may not be correct. But another possibility is that the behavior is not within the capability of members of the human species. In this case, the normative theory is wrong for prescribing an impossible behavior.

II.2. The Multiple Perspectives Concept

There is no single 'rational' approach to decision making.

Milan Zeleny

A major analytical tool used in this work is the Multiple Perspectives concept developed by Linstone (e.g. Linstone 1984) after work by Allison (1969, 1971) and others, including Steinbruner (1968, 1974) and Anderson (1977). The essential idea is that any person considering a particular situation brings to that consideration a specific point of view or *perspective*,

and that each perspective accentuates some aspects of the situation but makes others impossible to observe. People who work in a specific field long enough adopt consciously and unconsciously the lenses and filters of the perspective of that field. Economists think of individual human beings as economic agents who make rational choices to maximize their utility; psychiatrists think of those same individuals as battlegrounds between unconscious forces; evolutionary biologists see individuals as physical mechanisms for combinations of genes to express themselves — and these three perspectives are all fairly similar, sharing the general assumptions and guidelines of science. A child thinks of those individual people as “Mom”, “the bully next door”, and “my best friend”, etc.; an organizational insider will think of people as “the power behind the throne”, “that old back-stabber”, “my worst rival”, etc. People steeped in different perspectives will see different realities when they look at the same scene. While some situations are simple enough that a single perspective can be adequate for the purpose at hand, many situations are too complex to be adequately comprehended through a single perspective. Andersen states his conclusion emphatically:

Multiple perspectives are more than an intellectual curiosity of interest only to the academic who wishes to understand after-the-fact how decisions were made. In fact, there are multiple, coexisting worlds of decision making and the analyst who wishes to effectively manage policy modeling projects must understand each of these worlds and be able to move easily and effectively between the different worlds. ... [T]he existence of multiple coexistent policy realities is an important empirical fact that is often neglected or obscured by researchers bent on hammering their observations into a single, rather flat theoretical perspective. (Andersen 1977:29)

Mannheim noted the same limitation of single points of view in the face

of the complexity of historical events in general:

[I]t is possible to think productively from any point of view, although the degree of fruitfulness attainable varies from position to position. Each of these points of view reveals the interrelationships in the total complex of events from a different angle, and thus the suspicion grows that the historical process is something more inclusive than all the existing individual standpoints, and that our basis of thought, in its present state of atomization does not achieve a comprehensive view of events. (Mannheim 1936:226)

There is no unique number or set of perspectives that is always appropriate to use. Both Allison and Steinbruner, like Linstone, used three, though they all selected somewhat different threesomes. Three perspectives can be different enough to provide a good coverage of the most important aspects of most situations. For this work, I have adopted Linstone's set of perspectives — technical (the "T" perspective), organizational ("O"), and personal ("P") — and define them below. Linstone's use of these three is quite well developed in writing, and furnishes a fruitful analogy with the main literatures relevant to uncertainty.

One stimulus for Linstone's development of the Multiple Perspectives concept was the lack of success the T Perspective had in the 1960s and 1970s in tackling a variety of complex societal problems¹. The technical methods that had enabled scientists and engineers to make remarkable progress in the first half of the century were extended to new areas, such as Operations Research, which had great success during World War II. In the 1950s and 1960s, there was great hope that systems analysis

¹ My coverage of Multiple Perspectives here draws heavily on Linstone's work, especially his 1984 book.

(including management science/operations research, technology assessment, policy analysis, etc.), a generalization of the analytical approach, could help solve some of the large-scale problems that the country's economic success had not eliminated. By the close of the 1970s, it was becoming clear that systems analysis was not achieving the kinds of successes that had been hoped for, that it was not making progress comparable to the increasing pace of scientific and engineering progress.

"Why not?" asked Linstone and others. One part of the explanation involves the nature of the problems systems analysts tackled; another part concerns the paradigm within which analytic work *per se* is conducted. The problems were often complex, difficult to define, hard to bound, and often involved human beings in crucial ways, both as important elements of the problem domains and in the treatments and solutions. A detached and objective observer may be able to draw up blueprints for a new car, but there is no objective observer who can implement a new social program: observers and implementers are inescapably part of the social system being studied or modified.

The people in the system cannot be measured as meaningfully as can the mass, hardness and temperature of physical objects. The rational, analytical paradigm underlying the work was developed in a problem arena that was very different. G. L. S. Shackle pointed this out in 1960: "It is difficult even nowadays to realize how deeply economic theory has lain under the spell of classical (physical) mechanics and its appropriate mathematical ways of thinking, and has accepted calculable nature as a model of human affairs" (Shackle 1960:816). While sensitive analysts have

been able to overcome this limitation, the dominant technical perspective was blind to many of the characteristics of individuals and groups that contributed to the difficulty of the problems.

II.2.1. The Technical Perspective

The T Perspective has been developed largely within scientific and technological arenas. With the successes of science, it has come to be a predominant way of seeing the world not just in the realms of science and analysis (e.g. systems analysis or policy analysis), but in our entire culture. Anyone working in a scientific or engineering field is likely to adopt the T Perspective, often to the virtual exclusion of other perspectives. Professions whose workers are especially likely to adopt the T Perspective nearly exclusively are those with technical subjects, including economics, engineering, and the hard sciences. Fields with a T orientation combined with a subject matter that may bring other perspectives along include the soft sciences, business administration or public administration. People teaching or researching in any field are more likely to employ the T Perspective than are people working on applications in the same field. Linstone (1984) suggests that the following eight themes characterize the dominant T Perspective:

1. The Problem-Solution View. It is implicitly assumed that every problem has a “solution”, as is true in textbooks. Linstone counters that it is “more nearly correct to state that we shift problems rather than solve them.” (1984:9)

2. Optimization. Optimization² is a solution criterion that works well for equations, but does not describe the workings of natural systems. A rational analyst is inclined to design a system so as to minimize the probability of failure; a natural system is more likely to evolve so as to reduce the harm that failure would cause.

3. Reductionism. One way to manage the complexity of a problem is to divide it up into several smaller, less complex problems. This has worked very well in much of science where problems have been found to be “nearly decomposable”, in Simon’s term (Simon 1981:210). If the underlying structure of the problem is not decomposable, or if the decomposing is taken too far, reductionism will not work. For instance, if chemistry problems were reduced not to the level of atoms and molecules but to that of sub-atomic particles, chemistry would be extremely difficult to do.

4. Reliance on Models. The use of models has been key to the scientific approach to understanding the world, and is deeply embedded in the T Perspective. The importance of modeling is twofold: 1) If you haven’t modeled it, you haven’t understood it; and 2) If you have modeled it (simulated it), then you have understood it.

5. Quantification. Modeling and applied mathematical analysis requires quantifying objects of study. In western society, understanding and quantification are virtually identified: “By the end of the nineteenth century the official philosophy was that you understood something only when you could measure it” (Hacking 1987:48).

² The search for the “optimal” or best possible feasible answer, which requires knowing that no other answer could be better.

6. Objectivity. Scientific ideas are not supposed to be dependent on the particular thinker — they are “objective”. The emphasis on removing the thinker or observer from the picture tends to imply that there is only one proper perspective: since my identity is unimportant to an assertion, substituting another person for me would make no difference, and my assertion would still hold.

7. Ignoring the Individual. You cannot be sensitive to each of 250 million individuals and still have a “manageable” view of the situation, so the role of individual people is deliberately suppressed. This is partly a carryover from natural science: the identity of an atom of oxygen is completely irrelevant, as far as we know, and aggregation is legitimate. It is helpful and does not cause any loss of fidelity or of important detail in our representation of natural science problems.

8. Linear, Objective Time. Time is assumed to be a physically constant, objective dimension, always passing at the same speed for all parties. The future and the past can be compared to the present, and differences in preferences for consequences near or far in time can be handled by selecting a particular discount rate.

II.2.2. The Organizational Perspective

Since there are many organizations, there are many organizational perspectives, but they share a great many features with each other. A non-scientist who has worked in an organization for a long time is likely to acquire a thorough grounding in the O Perspective of that organization and its culture. How do things get done here? What are the Standard Operating Procedures (SOPs)? The emphasis in science, the home of the T

Perspective, is on truth: what is true about the world? How can we prove or disprove an hypothesis? The emphasis in the O Perspective is on practicality: how can we get this done fast enough and cheap enough? How soon does he really need this? How are we going to get them to agree to this, or what kind of compromise will we be able to reach? What happens to us if we make a mistake? How does one get ahead here? Unlike the T Perspective, which is really outside of any person or persons and outside of time, the O Perspective depends on particular individuals and groups, and discounts future events more steeply. Factors that are too distant in the future to appear relevant to a particular organization will be ignored.

Most people are part of many organizations or groups, and are familiar with several different O Perspectives. Skilled managers, administrators and politicians are likely to have acquired very full and rich versions of at least one O Perspective. The field of law, with its emphasis on precedent and regularity of process instead of on the ideal of truth, promotes the O Perspective. While experience in business generally facilitates adoption of an O Perspective, academic training in business (e.g. MBA programs), is much more T-oriented.

The characteristics of the O Perspective that correspond to the eight T characteristics listed above are briefly described next.

1. *Muddling through.* Organizations frequently encounter problems that cannot be solved, but must be handled and tolerated indefinitely. Sometimes a solution is known but for any of a myriad reasons not available. Trade-offs are common, and “solving” one problem means making or letting another one get worse.

2. *Compromise.* Optimization is not the typical concept used in finding treatments for problems. Optimization often requires more information than is available. Finding the *right* solution, even the *best available*, is often important, but those terms frequently depend on other people's judgments and priorities in a way that is inconsistent with the objective tone of optimization. Disagreements over both values and facts require negotiation and compromise.

3. *Messy interconnections abound.* It is seldom possible to isolate a single problem as much as might be desired, and seldom possible to reduce the number of variables sufficiently to design a good experiment.

4. *Suspicion of data, models.* Different organizations, even different parts of the same organization may have different versions of the same data: "facts" themselves are disputable and relative. One's own models are accepted, but those of others are suspect.

5. *Quantification desirable but often difficult.* Quantification is generally seen as a benefit, though often difficult to achieve. Being able to show that the "numbers say thus and such" is helpful; being the person who has to publicly commit to a particular number is often risky.

6. *Many points of view to consider.* Each person and organizational unit will have a point of view. Value judgments play such an important role that the notion of objectivity is often irrelevant. However, being able to portray one's position as objective is often useful in jockeying and negotiating.

7. *Individuals often play key roles, have individual motivations.* Individuals cannot be ignored; individuals in key places in an organization

can exert great influence on the organization's culture with their own values, or imprint their perspectives on their organization's perspective. Members of organizations are not mere cogs but have their own goals, their own specific abilities and weaknesses.

8. Time scale varies from department to department; intermittent time squeeze. The time scale used varies across organizations and with individual leaders. Research and planning departments are likely to have longer perspectives than financial departments that report to stockholders quarterly and must make the books balance at the end of the day or week. All of these time scales can be greatly modified by specific situations.

II.2.3. The Personal Perspective

Each person has at least one unique personal or P Perspective, forged by the characteristics of her background, values, mental and emotional processes, and current life situation. All human activity is driven by motives that affect individual people. While members of an organization may share motives, what makes each member behave in a particular way is not the fact that the behavior is valued in the organization's culture, but the fact that the individual member subscribes or does not subscribe to that value.

The specific attributes of an individual person can make a real difference, especially in the cases of people in leadership positions. The personal morals, values and perspective of the leader of a company can exert broad influence on the decisions made at all levels in the firm. A technologically correct idea may be impossible for some individuals to

implement.

In addition to the unique characteristics of individual people, there are general characteristics of human beings as individuals that affect how they perceive the world. They tend to have much shorter time horizons than organizations, and to care much more about the people very near them physically. The T Perspective deals with the problem of simplifying the huge number of individual human perspectives by aggregation and averaging; the P Perspective deals with the same problem by selecting a few people from the mass to care about and discarding the remainder.

Counselors and human-interest reporters are among those whose jobs require them to focus on the details of an individual life, emphasizing the P Perspective.

1. *Need to act.* Individuals can often ignore problems for a long time; when pressure comes to act, they respond to that pressure.

2. *Satisficing.* Individuals typically have much more limited resources than organizations, and are rarely able to inventory all available options, carefully assess their probable consequences, and then, after appropriate calculations, select the single best choice. A more common procedure is for a person to accept the first action or solution that appears to be satisfactory.

3. *Inconsistent reductionism forced by cognitive limitations.* Reductionism used by individuals is more often motivated by an inability to consider all factors simultaneously than by the technical desire to isolate variables and arrive at positive understandings.

4. *Data often unavailable; models mistrusted, not affordable.* For

most decisions that people face, there is not enough information to make a perfect choice, and few affordable ways to get more. Mathematical models are not feasible; conceptual models are used, but these are not often made explicit or publicly critiqued.

5. *Quantification often distasteful, de-personalizing.* For many people, quantifying the factors in the decisions they face is both unnatural and repugnant, a further insult of depersonalization in an uncaring society. Many people lack the technical skills to take advantage of elicitation methods that have been developed, and those methods are generally not intended for self-assessment.

6. *One's own point of view most important.* The T Perspective assumes that there is a single correct point of view, the "objective" one. In most personal decisions, there is again a single point of view, but it is that of the individual decision-maker. Objectivity is not only not sought, it is rejected as inappropriate, as it takes too little account of the preferences and values of the single decision-maker.

7. *Ignoring everybody else.* Matters of the individual are treated as important if the individual is the decision-maker, and as less important or unimportant otherwise.

8. *Very high discount rate common; time varies.* Most individuals use a very high discount rate in their ordinary lives. The ability to tolerate delay of gratification is one that parents hope their children learn so that they can stretch their time perspectives and lower their discount rate. Time varies depending on what one is doing, one's mood, and expectations about the future.

Major characteristics of the three perspectives are shown in summary in table 2.

Table 2. Features of the T, O and P Perspectives

	T Perspective	O Perspective	P Perspective
1	Problem/solution view	Muddling through	Need to act
2	Optimization	Compromise	Satisficing
3	Reductionism	Messy interconnections abound	Inconsistent reductionism forced by cognitive limitations
4	Reliance on data and models	Suspicion of data, models	Data often unavailable; models mistrusted, not affordable
5	Quantification	Quantification desirable but often difficult	Quantification often distasteful, de-personalizing
6	Objectivity	Many points of view to consider	One's own point of view most important
7	Ignoring the individual	Individuals often play key roles, have individual motivations	Ignoring everybody else
8	Linear, objective time	Time scale varies from dept. to dept.; Intermittent time squeeze	Very high discount rate common; time varies

II.2.4. Perspectives on uncertainty found in the literature

There are three major categories of literature that relate to the topic of this work. There is a large body of technical writing on rational decision-making, much of it in economic journals. These works discuss abstract, that is, generalized, approaches to the question of what characterizes a good decision under uncertainty. The emphasis is on prescription, although there has been significant research into the question of whether people or organizations behave in the way that the rational decision-making

literature suggests that they ought to (they don't).

There is a mainly separate literature on planning, which treats problems that are typically much less abstract and richer in specific details. This literature has a mixed focus of prescription and description. Descriptive efforts may be either research projects which describe the behavior of organizations that have been surveyed by the authors, or "how-to" descriptions of a particular technique that was employed by one or more organizations. The prescriptive works are largely rules of thumb created by experienced practitioners. These are on the whole far less mathematical than the rational decision-making works. Planning as a formal activity is conducted nearly exclusively by organizations.

A third body representing psychology focuses on the responses of individual human beings to situations in which they make choices under uncertainty. Two major themes in this body have been (1) comparing individual behavior with the prescriptions of the first body of literature, and (2) the identification of systematic biases that individuals exhibit under such circumstances. While many authors point out methods that could be used to compensate to some extent for some of the systematic human biases, most of this work is entirely descriptive. There are many less technical views of the ways in which uncertainty can affect us personally. Uncertainty about what is permitted may make us cautious, but uncertainty about what is forbidden may make us bolder (discussed earlier in section I.1.1). Uncertainty is often experienced as an agony of indecision at times when action is needed, but it also lends excitement to our lives, and inspires us to act with commitment.

The T, O, and P Perspectives match up very well with the three categories of literature outlined above, as illustrated in figure 3. The rational decision-making literature is a pure form of the T Perspective; most planning literature

Figure 3. Perspectives on uncertainty

discusses organizational goals and means for preparing for the future; and the literature on individual responses to uncertainty is clearly linked to the P Perspective. To a significant degree both the planning literature and the individual response literature share the technical nature of academic literature in general, but they both consider many non-technical aspects of their subjects, which the rational decision-making literature tends not to do.

II.2.5. Science, policy analysis, and the T Perspective

Policy analysis, the formal application of analytical techniques to planning and decision-making within organizations, generally shares the same T Perspective as developed in science. There are some important differences between the practice of science and the practice of policy analysis. Morgan and Henrion identify several of those differences in

practice in table 3 (Morgan and Henrion 1990:22).

Table 3. Features of science and policy analysis practice

Science	Policy Analysis
Empirical testing	Testing often impractical
Full documentation	Documentation typically inadequate
Reporting of uncertainty	Uncertainty usually incomplete or missing
Peer review	Review not standard and in some cases arduous
Open debate	Debate hindered by the above problems

Dror ascribes some of the distinction to different fundamental philosophies and goals (Dror 1988b:143§18.2.4). A primary criterion of scientific work is truth, while in policy analysis the corresponding standard is “preferization”, which refers to the degree to which the effectiveness of policy-making can be increased by the analysis.

Some of the differences in practice between the two fields are attributable to the functions for which they are used, as shown in table 4.

Two of the features of science practice could be readily adapted to policy analysis. Testing of significance, one of the major ways in which uncertainty is explicitly reported in science, would be quite feasible. Policy analyses would not always be amenable to direct, quantitative testing of significance, but some estimation should be possible in virtually all cases. Better documentation is almost always possible with the investment of more time and energy, admittedly often in short supply in decision-driven analyses.

Table 4. Features of science and policy analysis function

Science	Policy Analysis
Explicit test of significance to rule out null hypothesis	No regular testing of significance
Main focus on large issues to which each study contributes a small piece: general perspective; details worked out in application by others later	Main focus on the particular study: specific focus, measurement of details inescapable
Tools developed can help others only with the same (large) problem: nothing marketable developed	Tools developed can help others with similar (small) competing problems: marketable products developed
Physical sciences: problems often well-defined, well-bounded	Problems messy, difficult to generalize, hard to bound
Performed in environment where non-routine problems are expected	Performed in environment that abhors non-routine problems, relies on S.O.P.s
Modeling used to test hypotheses	Modeling used to develop hypotheses

Peer review and open debate are more problematic. The results of a policy analysis tend to be harder to generalize to new situations, which reduces the motivation of peers to spend time critiquing others' work. Research scientists have some freedom to choose problems where they think they can produce statistically significant results. The ability to perceive where the boundary between knowledge and ignorance is weakest is important for a scientist in choosing areas in which research is likely to be fruitful (Ravetz 1987; Stinchcombe 1990). Policy analyses are undertaken because the issues are important at a particular time, not because they are perceived to be soluble.

In some respects, policy analysis is more similar to engineering than to science, in that its purpose is to assist in the creation of products and tools

rather than discovering truths. The scale (both time and scope) for policy analysis problems is often greater than that for engineering problems, and it is harder to build prototypes. These difficulties reduce the usefulness of assessing a policy analysis by its results (the results may be decades away), as does the importance of uncertainty in the role of external factors. Examination of the process by which the results are obtained may be the only basis for assessment of quality.

Another field to which we might compare policy analysis is law. The general perspective of the two disciplines is very different, but they differ from science in the same critical feature: law and policy analysis are supposed to reach conclusions, to make decisions, while science can settle for incremental acquisition of knowledge, leaving many uncertainties highlighted and unresolved. Science reasons inductively, from the particular experiment towards the general rule; law reasons deductively, from the general principle towards the specific case (Large and Michie 1981). Reasoning in policy analysis, as in law, is from the general to the particular. Science aims for substantive rationality, independent of the particular experimental procedure. Policy analysis and law cannot — the specifics are fundamental to their purposes. Because of the different directions of reasoning, the location of much uncertainty assessment is different. In science, the uncertainty of experimental results is assessed as the results are examined; the general principles are enough more abstract than the particulars that measurement uncertainty rarely appears in them. Thus, scientific conclusions may not contain the explicit assessments of uncertainty that are common in reporting of scientific experiments —

supported usually by more than one more than one specific case, they are accepted as “true” until overthrown. Policy analysis, on the other hand, accepts general principles, free from uncertainty, and attempts to establish connections to particular instances. Assessment of the specific measurements of each instance is vital to establishing the validity of the connection in each instance. Thus, the conclusions of policy analysis ought to contain explicit assessments of uncertainty in every case.

II.2.6. Perspectives employed here

This dissertation is itself a technical work. It is written in an academic context for an academic purpose. The literatures reviewed — indeed, any “literatures” — are also quite technical. They do differ significantly in that while the O and P Perspective literatures are technical *treatments* of organizational and personal issues and perspectives, the technical literature is a technical view of a technical, i.e. abstract, subject.

The Sociological Perspective of Smithson is similar in some ways to O Perspectives: it is technical in character, while treating uncertainty in a social, though not exclusively organizational, context. Smithson also employs what might be termed a *meta-perspective* as he examines how conceptions of ignorance and uncertainty have been changing in the last twenty years.

I will not try to give equal weight to the three main perspectives employed here. The main purpose of this dissertation is primarily technical. However, the subject matter — planning and decision-making — is organizational and personal in nature. The O and P Perspectives will be used as much as is necessary for a thorough coverage of the topic, and they

will be used more technically here than they might be in other settings. They are, however, absolutely essential *even on purely technical grounds* for an adequate comprehension of uncertainty as it affects planning and decision-making activities of human beings.

II.3. The Technical Perspective on uncertainty

[*Statistics* includes] the entire science of decision making in the face of uncertainty.

J. D. Freund

II.3.1. Introduction to “T”

The T Perspective on uncertainty is abstract, emphasizing general principles at the expense of the details of particular circumstances. The basic conceptual model is the “decision”. In this context, uncertainty is a “problem” that complicates the decision, and it needs to be “handled”.

In order to understand the strengths and limitations of the T Perspective, several key aspects of it will be reviewed. The first is the assumption that rationality ought to guide decision-making. What it means to make decisions rationally, and whether that is always possible, are examined in order to assess the nature and limits of the role that rationality can play in the full range of types of uncertainty described later in the Taxonomy chapter. The second aspect is an implicit model of decision-making as a task to be performed, and the use of decision-making itself as the appropriate concept for dealing with uncertainty (e.g. compare decision-making qua event with planning qua on-going process). Use of this implicit model illuminates some features of uncertainty, but also constrains the way observers are able to conceive of uncertainty. The third essential aspect of the T view of uncertainty is the relationship between

uncertainty and probability. The historical development of probability is reviewed, and the major current conceptions of probability are described.

II.3.2. Rationality

In the case of uncertainty, the definition of
rationality becomes problematic.

J. G. March and H. A. Simon

II.3.2.1. General rationality

In modern Western culture, rationality is generally accepted as the appropriate principle governing the choice from a set of available actions, as indicated by these representative excerpts:

The theory I propose to adopt is that we seek things which we want, which may be our own or other people's pleasure, or anything else whatever, and our actions are such as we think most likely to realize those goods. (Ramsey 1926:35)

A person acts rationally only if the expected outcome of his action affords him a utility³ at least as great as that of the expected outcome of any action possible for him in the situation. (Gauthier 1975:418)

The idealized concept of rationality assumes maximization of a fixed or relatively stable objective, a known set of relevant alternatives and their outcomes, and a skill in computation that allows one to reach the highest attainable point with respect to the objective. (Zeleny 1982:64)

While Ramsey's theory is descriptive and Gauthier's and Zeleny's definitions are neither uniquely prescriptive nor uniquely descriptive, most authors accept a similar notion of rationality for prescriptive purposes as well. It is generally assumed implicitly that one ought to behave

³ Lovell: "Utility" is an abstract concept representing the desirability of an outcome. It is what economists assume decisionmakers maximize when making rational decisions. It is discussed in more detail later in this work.

rationally.⁴ There are different degrees of prescriptive force that this dictum can provide. A weak but nearly universally applicable degree is to prohibit demonstrably irrational actions, i.e. avoid region B in figure 4. A stronger interpretation is to choose only actions that are demonstrably rational, i.e. restrict choice of actions to those in region A. The weak version will often be indeterminate. The strong version is also indeterminate whenever the set of rationally prescribed actions (region A) is empty or has more than one element.

General Rationality Rule No. 1:
People ought to act rationally.

Weak version: People ought to avoid irrational actions, i.e. choose rationally permitted actions.

Strong version: People ought to choose rationally prescribed actions.

A related, stricter principle is to use only rationality in selecting an action. This clearly subsumes the strong version of the previous rule. It

incorporates the belief that rationality is powerful and complete enough in practice to yield unique prescriptions, and that decisions can be made

Figure 4. Actions and rationality

⁴ This assertion is most valid within the T Perspective. Linstone [1984, p. 68] argues that the P Perspective employs as a "guiding ethic" morality *instead of* rationality. Others might describe morality as the basis for selecting ends, and placing constraints on the choice of means, while rationality is then the basis for selecting the means to achieve the ends.

without recourse to any other decision-making principles. If the belief is false, continuing to expect rationality to settle all decision-making tasks is the form of irrationality Elster terms *hyperrationality*, mentioned in the Scope section of the first chapter. “Otto Neurath characterizes the belief that we can always have good reasons for our decisions as *pseudorationalism*” (Elster 1989a:121). Merely avoiding irrationality is clearly not sufficient if you insist on having a good reason for each decision.

The quotations above from Gauthier and Zeleny show that it is generally assumed explicitly that to choose rationally, one ought to maximize one’s utility.

General Rationality Rule No. 2: The rational choice from among a set of alternative actions is the action that will best serve the decision-maker’s ends.

Faith in rationality is based on the belief that we will be better off *in the long run* if we adopt rationality as a principle for making decisions. Only the *ex ante* perspective is appropriate for judging the quality of decisions. “A decision that turns out to be wrong in an *ex post* sense may nevertheless be the best that could have been made at the time of the choice” (Elster 1989a:117). Or as Dror puts it, instead of “the proof of the cooking is in the eating,” we must settle for “the proof of the cooking is in the cooking” (Dror 1988b:157).

11.3.2.2. Economic rationality

The assumptions that actors will behave rationally and that rationality means utility maximization permit the development of some powerful analytical tools in economics. Much of the literature on rational choice has been developed within the context of economic theorizing in order to

facilitate construction of general tools. In addition to the two rules above, which set the tone for rational-choice theorizing, economists typically impose criteria for the preferences a rational decision-maker may hold. A condition that is important for mathematical tractability is completeness, the assurance that the desirability of any two outcomes can be compared. “Outcome” here means an end state that results from the combination of the decision-maker’s choice and the eventual state of the world about which we are currently uncertain. While the state of the world is still unknown (i.e. uncertain), each alternative is associated with more than one possible outcome. After resolution of the uncertainty, the state of the world is known, as (presumably) is the outcome. The decision-maker may be undecided as to which alternative he prefers, but he is assumed to know which of the possible outcomes he would prefer to see. The pattern of preference relationships among the possible outcomes is called a preference structure.

Transitivity of preferences is generally required as a way of imposing a kind of consistency that conforms with most authors’ intuitions about rationality. It seems irrational to prefer A to B, and B to C, and yet to prefer C to A, and cyclic preferences are accordingly ruled out by the second part of this rule.

Economic Rationality Rule No. 1: A rational person’s preference structure is a complete, transitive ordering over outcomes.⁵

⁵ Let $>$ stand for the ordering relation “is preferred to”. Completeness requires that for all consequences, $c_i, c_j \in C$, either $c_i > c_j$ or $c_j > c_i$ or $c_i = c_j$. While von Neumann and Morgenstern [1953:589] thought the concept of a partial (i.e., not complete) order was significant, Savage thought such a weakening “would prove a blind alley losing much in power and advancing little, if at all, in realism.” [1954:21]

It is not consistent with our notions of rationality to prefer one thing to another now, and a moment later reverse the preference. This oversimplifies human reality, of course: it is also not rational never to change one's mind⁶. The possibility of changing one's mind certainly complicates the matter; apparently so much so that it is virtually never dealt with in the economic literature on rational choice.

Economic Rationality Rule No. 2: The preference structure of a rational agent is stable over time.

Rationality as optimizing is fundamental to the workings of most mainstream economic models. Simon claims that the following theory is shared by "virtually all brands of economists": "All people, or at least all economic decision units in society, such as business firms, consumers and members of the labour force, are utility maximizers" (Simon 1988:2). Some economists make a slightly weaker assumption: "economic agents behave *as if* they were maximizing expected utility" (Aiginger 1987:28, emphasis added). It does not matter, this implies, whether people and firms actually attempt to maximize utility; they behave as if they do, and predictions of their behavior based on the assumption of utility maximization will be valid regardless of the details of the internal choice process employed by the decision-maker. This is a testable assertion, and some evidence regarding its validity is reported in the P Perspective section. For those

Partial order: For all consequences, $c_i, c_j \in C$, at most one of the following relationships holds: $c_i > c_j$ or $c_j > c_i$ or $c_i = c_j$. This explicitly permits two consequences to be incomparable, i.e., when none of the three relationships is known to be true. Transitivity requires that for all consequences, $c_i, c_j, c_k \in C$, $c_i > c_j$ and $c_j > c_k$ implies $c_i > c_k$.

⁶ "The problem is how to sort out the good reasons for changing or not changing one's mind from the bad reasons." [Elster, 1984:147]

who construct macroeconomic models and require a mathematical representation of the behavior of economic units, however, the assumption of utility maximization is often made “not because of its supposed closeness to reality but because of its operationality and its strict demands on the rationality of the economic units” (Aiginger 1987:6).

Economic Rationality Rule No. 3:

Prescriptive version: Decision-makers should maximize their utility.

Strong descriptive version: Decision-makers maximize their utility.

Weak descriptive version: Decision-makers behave as if they were maximizing their utility.

This rule differs from General Rationality Rule No. 2 in that it specifies a particular interpretation of “best” as “maximum utility” and clearly suggests maximizing as the process decision-makers employ.

11.3.3. Procedural rationality, bounded rationality and satisficing

11.3.3.1. Procedural vs. substantive rationality

An economist studying systems comprising rational actors may be concerned about an individual actor only for purposes of predicting the actor’s behavior in order to predict the behavior of the whole system. In this context, the natural question is “What is the rational action to take in this specific situation?” The focus is on the relationship between the chosen action and the external conditions the actor faces. A psychologist or other consultant providing guidance to a particular actor (who desires to act rationally) often has a greater need to attend to the problem of how the actor can go about identifying or creating rational alternatives. The focus

is on the internal process by which the actor selects an action. Herbert Simon terms this the difference between *substantive rationality* and *procedural rationality* (Simon 1981:31-33). *Substantive rationality* is related to the actor's external behavior, and how well that behavior is suited for promoting the actor's aims. *Procedural rationality* is concerned with the specific processes the actor uses for generating and selecting the behavior. The T Perspective is concerned mainly with substantive rationality. The O and P Perspectives sweep in much more detail of specific situations, and procedural matters assume a greater importance. Explicit examination of organizational and individual behaviors provided a major source of Simon's evidence for the importance of procedural rationality.

For many descriptive purposes, procedural rationality may be irrelevant. What is desired is the ability to predict the actor's behavior; understanding the process by which the behavior is generated may not matter. This is how the weak descriptive version of Economic Rationality Rule No. 3 is applied. Economic agents are assumed to act *as if* they were optimizing their utility; how they actually make their decisions — whether any concept of utility or any optimization techniques are employed — is viewed as unimportant.

Normative rational choice theories, on the other hand, require both substantive and procedural rationality. On one hand, guidance that leads to actions that are not in the best interests of the decision-maker fails the requirement of substantive rationality. On the other hand, prescriptions for what an actor *ought* to do need to fall within the range of what the

actor *can* do, or they fail the test of procedural rationality. As Simon puts it,

Rules of substantive rationality that are not backed by executable algorithms are a worthless currency. (Simon 1981:43)

II.3.3.2. Bounded rationality

Many psychologists and economists who are interested in the actual behavior of people or firms have found Simon's concept of *bounded rationality* (e.g. Simon 1976) to be very useful. Bounded rationality is "the meaning of rationality in situations where the complexity of the environment is immensely greater than the computational powers of the adaptive system" (Simon 1981:190).

In reality, objectives are dynamic rather than static, information is seldom perfect, and alternatives, along with human cognition, are incomplete and limited. There are obvious time and cost constraints on the amount of effort that can be spent in the search for the optimum. Alternatives characterized by multidimensional consequences do not lend themselves to a well-organized and easily quantified preference ordering. Environmental forces, implementation issues, and one's confidence in the decision present additional limitations.

An ideal objective rationality is therefore unattainable; the capacity of the human mind to formulate and solve complex problems is inadequate with respect to objectively defined rational behavior. Decision makers can operate on the principle of limited or bounded rationality only. (Zeleny 1982:64)

The common identifying of rational decision-making with optimizing is based on an assumption of *unbounded* rationality: such identification is made without regard to the characteristics of the actual process of optimizing, such as computational costs, information acquisition costs, time requirements, or memory capacity. The search for the most extremely good solution – the global optimum – may require extreme resource

expenditures. Resources are virtually always limited in business or individual decision situations. In such cases, it is often necessary to *satisfice*, that is, to set a limited goal of finding a satisfactory solution in order to ensure that any resource-constrained search process can be completed feasibly.

II.3.3.3. Satisficing

Simon's concept of satisficing (e.g. Simon 1981, 1983) has become the major alternative to the idea that the way to make a rational choice is to maximize some quantity. Simon suggests that people do not exhaustively search the space of feasible choices and then select a unique optimal alternative, but that they stop searching as soon as one alternative is discovered that is "good enough". There is no sense of maximizing in this search strategy: alternatives are examined sequentially, and each is compared to a standard of sufficiency, which may change as the search proceeds. When an alternative is found that meets the current standard, the search process terminates.

One of Simon's arguments for the prevalence of satisficing in human decision-making is that the casting of a decision situation into the formal mathematical structure that is required for the tools of optimization "must discard essential facets of the real-world situations and, worse yet, must postulate parameters and variables that cannot be measured" (Simon 1981:36). Another way in which the actual decision situation may not permit optimizing is that the set of possible solutions may not permit an exhaustive search and comparison of the alternatives with feasible expenditures of resources. When this is the case, it may be easy to find

local optima, but impossible to devise a search strategy that is guaranteed to find a global optimum or to determine whether one of the alternatives examined is the global optimum (Simon 1981). It may be much more feasible to devise and implement criteria for determining whether a given alternative is satisfactory.

It is always questionable whether the optimum procedure is to search for *the* optimum value. (Miller and Starr 1967:51)

When the requisite conditions for successfully applying optimization techniques are absent, other decision tools must be used. Heuristic search, for example, is a tool that can often efficiently find satisfactory solutions. Simon's claim is not that satisficing is better:

Now no one in his right mind will satisfice if he can equally well optimize; no one will settle for good or better if he can have best. But that is not the way the problem usually poses itself in actual design situations. (Simon 1981:138)

Satisficing Rule No. 1: Always optimize when you can.

Satisficing Rule No. 2: When you can't optimize, satisfice — i.e. when you can't be sure of finding the very best alternative, make a "reasonable" effort to find a "good enough" alternative.

Simon's concern for procedural rationality (see below) leads him to accept bounded rationality and satisficing as important varieties of rational decision-making. For others who are not willing to accept bounded rationality, recognition of limitations on the feasibility of optimizing leads to restrictions on the situations in which rational decision-making is possible:

Is utility maximization always possible? If rationality is identified with the aim of utility maximization, and if there are situations in which it is not possible to maximize one's utilities, then there are situations in which it is not possible to act rationally. This

is not a conclusion one would willingly accept. Hence a proof of the possibility of utility maximization, in all situations, is required.

But this task is beyond the compass of the present essay.
(Gauthier 1982:90)

By insisting on this particular interpretation of rationality, Gauthier acknowledges that he severely restricts the domain of applicability of rationality as a guide to decision-makers. Mainstream economic literature generally accepts the same interpretation, though it is unusual to see the implications of this interpretation, combined with the difficulties of “field conditions” identified by Simon, traced out so clearly to Gauthier’s conclusion.

Economic Rationality Rule No. 4: If you can’t optimize, you can’t act rationally.

A form of decision-making that could be considered to be a variant of satisficing is *elimination by aspects* (e.g. Tversky 1972). Several criteria are applied sequentially to the feasible alternatives. Each application of a criterion eliminates some of the alternatives until only one remains, and the decision has therefore been made.

11.3.3.4. Prescriptive satisficing

Cast in the form of the preceding section, satisficing is an element of a *descriptive* theory of rationality. Satisficing may be even more important in *prescriptive* theories of rationality. In the maximizing form of rationality, the decision-maker must obtain information about the alternatives. If he is to be rational about the information gathering, he must obtain an optimal amount. But, as Sidney Winter (1964) asks, how is the decision-maker to determine how much information is optimal? This decision, in turn, requires an optimal amount of information about the first information-

gathering decision, and so on. Recognizing this infinite regress forces one to acknowledge that the pure form of optimization is impossible — in principle! Some compromise of the goal of global optimization is absolutely unavoidable. Satisficing is a credible alternative to optimizing as the way that people *ought* to make decisions, given their inescapable limits on information gathering and information processing.

Table 5. Optimizing and satisficing compared

	Optimizing	Satisficing
Criterion	Best possible	Good enough
Major problem for user	Knowing whether there is another alternative that is better (than a given one)	Determining the appropriate standard of satisfactoriness
No. of alternatives that must be examined	All (all but one must be excluded by individual examination or analytical means)	Only as many as required to find one satisfactory alternative
Feasibility problem	Global search process may never terminate	Alternative(s) meeting standard may not exist

II.3.3.5. Strategic vs. parametric rationality

Elster (e.g. 1984) distinguishes between parametric rationality and strategic rationality. In parametric rationality, each actor takes her environment as a relatively inert state without significant intentionality. In strategic rationality, each actor is aware of the role that other actors and their desires and decisions play in shaping her environment. A parametric environment may be complicated by uncertainties, but choices that other agents will make do not figure significantly in those uncertainties.

[Parametric rationality] usually leads to a unique and well-defined prescription for behaviour, given the initial assumptions and expectations. The problem here is that these assumptions and

expectations may themselves be irrational, because founded upon a naive belief about a constant (or parametric) environment. In strategic rationality this belief no longer operates, and instead comes the realistic insight into the strategic nature of the environment. Here, however, another problem arises: in strategic interaction there may be no uniquely defined mode of behaviour which is singled out as *the* rational choice. (Elster 1984:117)

The information required for making strategically rational decisions is much greater and less likely to be available: one needs not only all of the information required for parametric rationality, but the same set of information from all other actors' perspectives, and a set of predictions about what each other actor will do, given that information.

II.3.3.6. Metarationality

When is it rational to act rationally and when to act otherwise? This is a theme Jon Elster has discussed throughout a series of books (Elster 1984, 1989a, 1989b). In particular, Elster has focused on hyperrationality, that is, expecting too much of rationality. This is one of the main issues in what I call metarationality, the subject of when rationality is appropriate and useful, and how to select and interpret procedures for implementing rationality in particular situations. The upshot is that it is not always rational (i.e. not always *metarational*) to be rational.

When Ulysses was about to sail near the sirens, he foresaw that he would be unable to act rationally, and had himself bound in order to prevent his own subsequent irrational actions (Elster 1984:36). Dror (1980:63) and Hamburger (1979:94) describe circumstances in which a state may find it desirable to appear to other states to be acting irrationally, and the game of Chicken is another such circumstance. Dror suggests other situations in which it may be rational to actually behave

non-rationally (Dror 1987:31). He cites as a specific example the success the Israeli government has had without conducting rational planning. He attributes this to the turbulence of Israel's environment, which gives improvisation advantages over the more cumbersome procedures of rational planning (Dror 1987:39). An actor may rationally prefer to employ a random, non-rational decision-making process in order to prevent an opponent from predicting future decisions (Schelling 1960:175; von Neumann and Morgenstern 1953:147; Dror 1980).

It is not difficult to think of situations in which most of us (say, middle-class Americans in the 1990s) would consider rational choice techniques so inappropriate as to hardly merit mention, such as the choice of a spouse. There are many other situations which are not a priori unsuited for rational choice, but the procedures available or employed for performing the selection are not decisive (see next section). Relying upon such a procedure, when it can be known that it is not decisive, is the form of hyperrationality of most concern here.

II.3.4. Decision selection processes

II.3.4.1. Definition

A decision selection process (selection process) will be defined here as a procedure for moving from a situation in which the decision-maker faces a multiplicity of possible actions to a situation in which a single action has been chosen for implementation⁷. A need for action has been recognized,

⁷ Mintzberg et al. [1976] define a broader *decision process*, encompassing three phases: identification (of an opportunity or a problem), development (of alternative solutions), and selection. Their work is reviewed below in the O Perspective section. The definition of selection process used herein corresponds to evaluation-choice in their selection phase.

two or more potential actions have been identified, and a selection from among the available actions must be made before proceeding. The goal of a selection process is simply that of identifying a single alternative as the chosen action. (This definition deliberately excludes implementation in order to focus on selection.) Decision selection processes are needed only when two or more feasible alternatives have been identified. If none has been found, the decision-maker is in a bind. If exactly one has been identified, she knows what to do, and again, no decision process is needed.⁸

We return to the case of more than one feasible alternative, for which a selection process is required. Decision selection processes act somewhat like filters, in that only some of the alternatives will “pass through” the selection process and remain in the reduced set of alternatives. An ideally successful selection procedure eliminates all but one of the alternatives.

The result of applying a selection process may or may not be ambiguous. An unambiguous result of a selection process is one in which the alternatives are sorted unambiguously into two categories, selected and rejected, regions A and B respectively in figure 5. A less happy outcome is to find that at least one and perhaps all of the alternatives remain in a zone of ambiguity (region C), and can be classified neither as selected nor as unselected, as indicated in figure 6.

⁸ This reasoning supposes an optimizing approach to the selection of alternatives: a solitary “alternative” is always the best of the alternatives. If the decision-maker is satisficing, then she must ask if the solitary “alternative” is good enough; if not, more alternatives must be found or generated. Some form of this adequacy testing is nearly always employed by individuals and organizations.

Decision selection processes, even unambiguous ones, can fail in either of two ways. The first type of failure is that of non-existence of output: none of the alternatives is placed in region A, the set of selected alternatives.

The second type of failure is non-uniqueness: more than one of the alternatives passes through the filter of the selection process and are placed in region A. When there are only two alternatives, the results of non-existence and non-uniqueness are functionally the same: it is not possible to discriminate between the two alternatives. With three or more alternatives, the two types of

Figure 5. Unambiguous selection

Figure 6. Ambiguous selection

errors give rise to two distinct situations. In the case of non-existence of selection output, none of the alternatives is shown to be best: if the results are unambiguous, they are all in region B; if the result is ambiguous, some or all may be in either region B or C. The decision-maker cannot make a decision with the security of picking the (or one of the) “best” alternatives. In the case of non-uniqueness, the outcome of the selection process is at

least partly definitive, with the result that two or more alternatives have been determined to be equally good and both unsurpassed and belong to region A. We know that X is as good as Y, and we know that Y is as good as X. The decision-maker is, or should be, truly indifferent to the selection.

In the case of an ambiguous, non-uniqueness result, there may be alternatives that remain in C, the zone of ambiguity, that have not been shown to be inferior. Since A, the set of selected alternatives, is non-empty, the alternatives in the zone of ambiguity cannot be better than those in the set of selections, but other parties may lobby strongly for their inclusion or reconsideration. In the case of an ambiguous, non-existence result, the set of clearly selected alternatives is empty, but B, the set of non-selected alternatives may be non-empty, indicating that the selection process has at least eliminated some of the alternatives.

In the case of either non-existence or non-uniqueness, the selection process has failed, since no single alternative has been identified. Unless the first process was flawed and it is therefore reasonable to think that a repetition of the process with fewer errors might be decisive, then another, different selection process must be employed to make the selection.

For instance, suppose the first selection process was a rational process that captured all of the value dimensions the decision-maker cares about, and that it was executed reasonably well. Then at this point, a non-rational selection process must be used, since the relevant value dimensions a rational process could incorporate have already been employed unsuccessfully (indeterminately) in the first process. Some form of arbitrary decision must be made in order to effect a selection —

arbitrary in the sense that the new factors were not considered important enough to be included in the first selection process. A man indifferent between two shirts in a store may ask the opinion of the sales clerk, not because that opinion matters to him, but because he intends to buy exactly one shirt and needs to make a selection. A random process such as flipping a coin may also suffice. Americans are averse in many situations to the use of random selection processes (Elster 1989a:37), although they can be highly efficient ways to make selections.

Sporting events and series of events can be viewed as selection processes, and often the regular period, be it a game or a season, ends in a tie — a non-uniqueness failure. The secondary process for deciding baseball games is additional rounds (innings) of the primary process, continued until the process is decisive. Basketball games have additional, shorter periods of regular play continued until decisive. The secondary process in professional football games is different from that in the primary periods: the objective changes from scoring the most points to scoring first. If the additional “sudden death” period ends without additional scoring, the outcome is left as a tie. Football seasons, however, unlike football games, are not permitted to end in ties. Teams qualify for the post-season playoffs based on their record during the regular season. If two teams have identical records, secondary indicators in their records of increasingly arbitrary nature, such as the record against teams within the same division, and total points scored, are used to order the teams.

Non-existence failure is more serious, and arises when the selection process is indeterminate. We do not know if A is as good as B, and we don't

know if B is as good as A. The selection process is unable to distinguish between A and B, but we are not yet convinced that the alternatives are equally good – rather, we are uncertain about the comparison between the two of them. In such a situation, a decision-maker may worry that one alternative could truly be better than the others, although she doesn't know which one. In contrast to non-uniqueness failure, the decision-maker is not at all sure that she is indifferent to the selection. This distinction is crucial in decisions that may be reexamined later by others: if the decision-maker is certain that an alternative is unsurpassed although not uniquely so, there is low danger from later second-guessing. However, if the decision-maker is not sure that the alternatives are equally desirable, there is a risk that one of the choices she rejects will turn out to be superior, and that she will be criticized for having made a decision that in hindsight appears faulty. Although one ought to maintain an *ex ante* perspective when critiquing decisions (Elster 1989a:117), it is difficult, and other parties are likely not to do so. "Nobody cares much about *ex ante* rationality. *Ex post* success is what counts" (Elster 1989a:61).

Failures of this type can easily arise from the uncertainty in calculations of a single value dimension or from comparison problems between value dimensions, and in other ways. For example, electric utility planners may believe that two projects are similar enough that cost alone is the appropriate criterion, but that the uncertainty about the cost, e.g. of natural gas versus coal 10 years from now, is so large that it is not possible to say with any certainty which project will be cheaper over the life of the project. It is not that we believe the costs will be identical, but that we

don't know enough about the future to distinguish between the costs of the projects.

Value comparison problems are illustrated by the choice a utility might have to make today between building a coal plant and a plant that uses solar energy. At 1991 cost levels, a coal plant may be cheaper to build and operate than a solar plant. However, the solar plant would use a renewable resource, and would produce far fewer atmospheric emissions, a fact that is especially important now that worry over global warming has highlighted the hazards of higher CO₂ concentrations in the air. It would be convenient to be able to convert all of the environmental impacts into dollar impacts, but as the previous example shows, cost comparisons can be so uncertain as to be inconclusive, and the uncertainty of converting the Greenhouse Effect risk of one coal plant into dollars will almost surely swamp the cost comparison. A serious problem is that the importance of the atmospheric emissions is extremely difficult to compare directly to dollars or to any other scale of value, even if future CO₂ levels could be accurately forecast.

11.3.4.2. Standards for selection processes

Successfully applied selection processes meet two standards. They are decisive (or persuasive), and they are defensible. They are persuasive when they identify a single option and persuade the decision-maker that it ought to be selected. They are defensible when their results hold up to reasonably close inspection and/or second guessing. Science and law each have well-established principles for these two criteria. In science, the typical question to be settled by a decision process is "Does hypothesis X hold?" The issue of decisiveness is settled by showing that the experiment being

reported does in fact illustrate the hypothesis, usually by refuting the opposite of the hypothesis, and that the results were statistically significant. Defensibility is determined by peers reviewing the work for logical flaws and through attempts at replicating the results. In law, the archetypal decision process is a trial, and the two alternative outcomes are either “guilty” and “not guilty” or “find for the plaintiff” and “find for the defendant”. The issue of decisiveness is determined by the outcome of the trial: unless a jury trial ends with a hung jury, the process is decisive. Defensibility is established through the appeals process.

The difference between the scientific and legal selection processes illustrates a significant difference between the T and O Perspectives. In science, the issues of decisiveness and defensibility are, as much as possible, matters of substantive rationality, determined by the reasoning of the observers; both the decisiveness and the defensibility can be reexamined or reinterpreted at any later time. In law, the criteria are matters of procedure, and when the agreed-upon procedures have been performed, the issues are permanently settled except in the rarest of circumstances. In most policy analysis and organizational decision-making the criteria are less clearly defined. Results of electric utility decision processes, no matter how they are made to be decisive, are often defended explicitly in public hearings that allow other parties to challenge the utility’s decision on either substantive or procedural grounds.

11.3.4.3. Rational selection processes

Rational decision-making processes can fail on both the existence and uniqueness counts, being indeterminate in either case. Hyperrationality

ignores this:

Failure to recognize the indeterminacy of rational-choice theory can lead to irrational behaviour. (Elster 1989a:2)

Using the traditional maximizing sense of rationality, a rational selection process must identify the best alternative. If we cannot identify any alternative that we are sure is (one of) the best, the rational choice is non-existent. If we identify more than one alternative that is maximally rational, the process has failed the uniqueness test. In either case, the rational selection process has not clearly identified a single alternative, which is what was desired.

11.3.5. "Failures" of rationality

There are three ways in which efforts to decide rationally may fail. One is that a decision-maker may choose an action that clearly does not advance his interests as well as other choices could, or to choose two actions that promote mutually incompatible goals. Such actions we would usually call "irrational". Most discussions of rationality in the technical decision-making arena concern this *substantive* aspect of rationality.

The second type of failure emerges from applications in particular situations of specific procedures. Rationality can be brought to bear on any specific situation only with the use of specific techniques or procedures, for example, ways to collect or process information. These processes all have shortcomings. Their implementation requires data, information processing capability, and other resources which may not be available in the amount required by the procedures. As noted in the previous section, procedures intended to identify a unique "best" or most rational alternative will not always do so. Failures of *procedural rationality* are described in the second

section below.

The third failure arises from the inherent limitations of rationality as a tool, and concerns inappropriate uses and interpretations of rational methods. For example, rationality does not speak to the issues of what values one ought to hold, or to the issue of how to generate alternatives. Within its domain, it provides considerable power. However, taking into account its limits, it ought not be employed in every situation. This aspect of rationality is here called *metarationality*. It also pertains to the reasonable interpretation of the results of employing rationality. Hyperrationality is Elster's term for errors of this sort, and is described in the third section below. The distinctions among these levels is illustrated in figure 7 below.

11.3.5.1. Substantive failures

A substantive failure of normative rational choice is the prescription of an option that does not promote the decision-maker's goals as well as another option would. The substantive wrongness of a prescription, in a particular situation, depends on the external relationship between the decision-maker's values, the environment in which the choice is to be made, and the expected consequences of the option chosen, not on the method by which the choice was made. However, there is no way to determine wrongness in this sense except by the use of a decision procedure. (This kind of wrongness is an *ex ante* wrongness of choice, not an *ex post* observation that things did not turn out well, which can happen no matter how ably an uncertain choice is made. The same *ex ante* perspective must be used in identifying this type of error.)

Figure 7. Meta-, substantive and procedural rationality

A suggestion that a choice is substantively irrational may be based upon errors in the implementation of the decision selection procedure, in which case the decision-maker may agree with the suggestion. Frequently, however, such suggestions are made on the basis of a decision selection

procedure other than that used by the decision-maker, and there may be no way to determine conclusively which procedure is better. While the principles of rationality are objective, that is, impersonal and general, any implementation involves subjectivity at least in the choice of a decision selection procedure. Any assessment of the substantive rationality of a decision depends critically on the procedure and on situation-specific information, and is in turn subject to procedural failures.

11.3.5.2. Procedural failures

There are three ways in which a rational decision selection process can fail. The first is that it can be indecisive, that is, fail to be persuasive. This can be due to either non-existence or non-uniqueness. In the former case, no alternatives can be shown to be superior; in the latter, two or more are shown to be superior, but they appear equally superior, and a selection still cannot be made yet. Non-existence failures are sometimes overlooked: one alternative may have a higher value on the selection criterion, and the selection process may be taken to be decisive without an assessment of the significance of the differences among criterion values. Unless one score can be judged to be significantly higher than the next highest score, the procedure is not actually decisive. Such tests of significance are, of course, routine in experimental science, but much less common in decision analysis.

(Even when a rational selection process has failed to yield a unique result, it may have been helpful (Elster 1986:19). In the case of non-existence, it may yet be possible to identify poor alternatives, those we are sure *are* surpassed, leaving the others as candidates for a subsequent,

different selection process to work with. In the case of non-uniqueness, the multiple optimal alternatives automatically constitute the set of feasible alternatives for the subsequent selection process.)

The second type of procedural failure is indefensibility. The procedure employed must be appropriate for the purposes of the decision-maker and for the kind and quality of information available. In some contexts, or for some people, it may be sufficient to show that the decision procedure employed was one that might be chosen reasonably; in other contexts, or for other people, it might be expected that the procedure adopted could be proven to be uniquely best. The appropriate factors must have been considered, a full set of options must have been generated, relevant values must have informed the evaluations of options, deliberations must have been reasonable, etc.

The third type of failure is not logically distinct from the first two but merits mention due to the source of the failure: poor implementation. A procedure may fail to be decisive or defensible because it was executed badly, and if redone, may be satisfactory. Insufficient information may have been collected, calculation errors may be made, monetary values may be used inappropriately in place of utility values, etc. In other cases, the implementation may be adequate, and the indecisiveness or indefensibility cannot be corrected by improved execution.

Using a satisficing type of rationality, a rational selection process must identify one alternative that we are sure meets the established criteria of sufficiency. That may not be possible. Non-uniqueness problems arise only when the selection process continues after the identification of a

satisfactory alternative, or when two or more alternatives are identified at the same time, determined to merit evaluation, and ultimately both prove to be satisfactory. One reason for the former possibility is that a satisfactory alternative may be found so quickly that the possibility of being able to meet a higher standard is raised. This challenges the defensibility of the selection process as based on the earlier criterion, and a revised selection process is then employed.

Generally, uncertainty leads to non-existence problems: either an inability to identify any best alternatives (maximizing), or an inability to be sure that any alternative satisfies the goodness criteria (satisficing). Quite good information is required to demonstrate the equivalence of two or more alternatives, which is necessary for non-uniqueness.

11.3.5.3. Metarationality failures, i.e. hyperrationality

Metarationality is the rational use of rationality. The essence of metarationality failures, that is, of hyperrationality, is improper consideration of the potentials and limits of rationality and the procedures used to implement it. Such concerns about the use and interpretation of rational procedures generally arise either before the beginning of or following the conduct of rational procedures. Elster describes the following five categories of hyperrationality (Elster 1989a:26).

1. Ignore Actual Indeterminacy of Beliefs or Preferences. When people's beliefs or preferences are truly uncertain, it will not be possible to determine with certainty what they ought to do.

2. Look To The Second Decimal When Ignorant Of The First. If the value of the most significant digit is in doubt, the process is surely

indeterminate. Looking to the second decimal may be as satisfactory a way of making a decision as any, as long as the arbitrary nature of this selection is recognized.

3. Reframe A Decision to “Eliminate” Indeterminacy. In most cases of difficult or indeterminate decisions, some aspects are more clear than others. It can be tempting to emphasize those solely to make the decision appear to be determinate and choose to ignore those aspects that make the decision ambiguous.

4. Solve Specific Situation When Finding General Rules Could Be More Economical. Developing general rules that could apply to future decisions may be a better use of resources than deciding in a way that applies only to the decision at hand.

5. Ignore Costs Of Decision-making. More elaborate decision-making procedures may produce decisions that are better than a simpler process could provide, but not enough better to justify the additional costs.

An additional type of metarationality error is selecting the wrong procedure for implementing rationality. (Use of the word “wrong” suggests that a specific metarational procedure might be employed to make the metarational selection of the rational procedure for the original decision. In order to avoid such an infinite regress, I suggest less formal standards for assessment of metarationality, perhaps “reasonable” instead of “rational”.) Attempting to make a decision rationally in the absence of a rational decision selection procedure whose costs are justified by the anticipated benefits is a related error. If it were possible to determine that no conceivable procedure could produce a correct, unique prescription, then

it would be irrational to hope that rationality could produce one.

Indeed, the entire framing of a decision can be considered to be part of metarationality: identifying the opportunity or need for a decision, employing some process for identifying and/or generating alternatives, etc. A larger process such as this developed by Mintzberg, et al. (1976), is discussed below in the O Perspective section.

The five types of errors Elster specified generally involve expecting too much of rationality. Can people expect *too little* of rationality? One possible case is where one agent believes rationality could help and another does not value rationality. Choice of rationality as a decision aid cannot be justified by rationality itself; it is a judgment of value. This case, then, is not an error of metarationality.

When an agent has chosen rationality as a choice principle, there may yet be cases where the agent overlooks the possible advantages of employing a rational procedure. Elster's fourth type of hyperrationality is similar to this error: it may be possible to solve an entire class of problems instead of settling for solving the immediate problem. Innovative scientific procedures that are able to settle types of questions according to the accepted standards of scientific certainty may be overlooked by decision-makers who do not know of them, or by courts that have not yet established their admissibility. An important distinction to be made here is that many instances which appear to be expecting too little of rationality stem from a decision-maker's decision not to rely on rationality. The cases of expecting too much of rationality are distinguished by a full acceptance of and reliance upon rationality. It is specifically rational considerations

that are used to identify the limitations of rationality.

In summary, a substantive error of rational choice is the making of a wrong decision (considered *ex ante*); a procedural error is the use of a wrong procedure to make a decision, or a wrong implementation of one; and a metarational error is wrong degree of reliance upon rationality, given the context of a particular decision task.

11.3.6. Non-rational choosing

While rationality, whether the optimizing or satisficing variety, is the favored principle for choosing actions in the Western academic world, other principles are often used. Max Weber considered three categories of non-rational behavior: traditional, value-oriented, and affectual; Gary Becker described random and traditional (Elster 1979:137). Some other categories of choice procedures that have been suggested are instinctual behavior, heuristics (Nelson et al. 1976), tradition, social norms, religious dogma, impulsivity, “black-box” techniques, intuition, politics, and defective rationality. Intuition is not necessarily non-rational: you may attempt to identify the alternative that most effectively promotes your goals by using your intuition in either an optimizing mode or a satisficing mode. Used in this fashion, intuition differs from the rational choice method in degree of quantification and explicitness of the comparisons. Alternatively, intuition may be used to select the choice that “feels right” rather than to identify explicitly the right choice by considering intuitively how each would promote your goals. In this case, intuition is an alternative to rational choice.

As mentioned in the earlier section on metarationality, it is not

necessarily irrational to select a non-rational choice principle. When rational decision processes do not identify a single rational alternative and a decision is needed, rationality (more specifically, metarationality) requires that another choice principle be used in order to obtain a decision. In addition, there can be strategic reasons for wanting your behavior to be unpredictable.

Elster (1984) describes defective rationality as generally equivalent to weakness of will: a person is assumed to have the capability to act rationally and knows how to, but does not, presumably due to insufficient will. Some affectual motivation overrides the ratiocinative faculties. Similarly, the common-sense meaning of irrational applies only to parties believed to be capable of rational choice.

Some of the non-rational modes are individual modes and some pertain only to groups. Politics, e.g. voting or consensus-building, is generally a technique for a group of people to make a choice jointly. Other modes — intuition and affective choice — are only possible for individuals. Divine inspiration could be viewed as a form of either intuition or affectual choice; impulsivity is affectual. Elster (1986:3) suggests that only individuals, not groups, can have beliefs or desires. He infers from this and from the principle of methodological individualism⁹ that groups cannot be rational (or irrational), but social choice theorists in general believe that “collective rationality” is both meaningful and desirable. Several choice principles can

⁹ “Briefly, methodological individualism is the principle that explanations at the social level are not adequate unless the individual actions that are required for the social explanation to be true can themselves be explained adequately.” [Stinchcombe 1990:101]

be employed by either groups or individuals.

Heuristics, tradition, religion, and social norms are different types of rule systems that could be followed by individuals or groups. Instinctual behavior is rule-like, but much less common in human behavior than among animals. Several authors (e.g. Winter 1974, 1984) have proposed theories of the firm in which the use of heuristics is the primary choice principle instead of the mainstream economic assumption of maximization.

“Black-box” techniques are methods whereby the decision-maker(s) gives control of the choice over to a person or method whose internal details are obscured. Random techniques, use of consultants who do not explain their methods, reading of tea-leaves or the stars, and consulting oracles are examples of black boxes from the perspective of the client/decision-maker. The tea-leaf reader or consultant may be following rules for developing his advice, but the client/decision-maker is not employing the rules.

II.3.7. The Standard Model of Decisions and the Expected Utility Hypothesis

Most of the literature on rational decision-making and decisions under uncertainty explicitly or implicitly employs a standard decision model (SDM). The following version, given by Robert Kast, is highly representative of the literature:

The agent chooses an action a in a given set A . When making this decision, he faces uncertainty which is represented by a set Ω of “states of the world”. The agent has no preferences over the set A of actions otherwise there would be no problems but he has some on the consequences of his actions. Let C be a set of consequences and we assume a function $c: A \times \Omega \rightarrow C$ is given which defines the consequence $c(a, w) \in C$ of action a for each state of the world w in Ω .

When there is no uncertainty (Ω reduces to one element) preferences on C which can be represented by a function $u: c \rightarrow \Re$ are sufficient to define an order on A : a is preferred to a' if $u[c(a)] \geq u[c(a')]$.

In a situation of risk the agent knows a distribution on Ω . (Kast 1988)

In a situation of uncertainty, to use the distinction between risk and uncertainty suggested by Knight (1921), the agent does not know a distribution on Ω . The utility function $u()$ represents the values and priorities of the specific decision-maker, not necessarily anyone else. Figure 8 below illustrates the elements of the standard decision model. I have changed the notation: the set of possible states of the world is named “ W ”, and “ Ω ” is used to refer to the possibly known probability distribution over the elements of W .

In the typical treatment of uncertainty, it is only the future state of the world that is considered to be unknown. The other parts of the SDM — the feasible options, the outcome of combining a particular option with a particular state of the world, the relative priority for different types of consequences and values — are all assumed to be known to whatever extent is needed for the analysis.

11.3.7.1. Von Neumann and Morgenstern’s Expected Utility Hypothesis

The problem with analytical methods for selecting the rational alternative from the set of feasible alternatives, even with the major assumption that all of the components of the standard decision model are known, is that for each choice there is an entire set of possible outcomes, one outcome for each possible state of the world. Assessing the desirability of an alternative requires a procedure for aggregating the set of

desirability measures of the possible outcomes into a single measure of the goodness of that alternative.

Figure 8. The Standard Decision Model

There are several decision rules that can perform this aggregation or selection. The *maxi-max* principle calls for choosing the option whose best possible outcome is better than any other option's best outcome. This ensures the possibility of the very best outcome, and is generally conceived

of as being an optimistic guide. The *maxi-min*¹⁰ principle is more conservative, and selects the option whose worst possible outcome is better than any other option's worst outcome. This establishes a security level; the decision-maker is assured of getting at least as good an outcome as the worst outcome of the maxi-min choice.

Principles such as those suggested above are well-known in game theory. An advantage of them in that arena is that they do not require any probability distribution over the set of possible future states of the world. In game theoretic contexts, human choice by other players is the primary determinant of the future state of the world, and probabilities have rarely been used to describe the likely choices of the other players. What is an advantage in game theory, however, is a disadvantage in decision analysis: these principles are unable to make use of probability data when it does exist. What is needed is a procedure for combining the measure of goodness of each outcome with the likelihood of its occurring, given that a particular alternative has been selected.

In 1947, John von Neumann and Oskar Morgenstern provided exactly this procedure. In an appendix to the second edition of *The Theory of Games and Economic Behavior*, they presented and proved the Expected Utility Hypothesis (EUH) (von Neumann and Morgenstern 1947). With EUH, you can essentially calculate what the rational choice is, given the sets of choices and states of the world, the probability distribution over the

¹⁰ This principle is more widely known as the mini-max principle from contexts in which uncertain cost rather than uncertain utility was considered. Total cost, a negative quantity, is to be minimized, while utility, a positive quantity, is to be maximized — the reversal in the terms is purely terminological, and they both represent the same principle.

states of the world, utility assessments of each combination of choice and world-state, and a few other conditions.

The appeal of this is immediate. If you can satisfy the EUH conditions, you can be sure of identifying the most rational choice. The framework is backed up by sound mathematical deductions from a few quite plausible axioms. You can show others your calculations. The essentially calculating nature of the final process and its publicness lend a considerable air of objectivity to the procedure. Objectivity is appealing to many individuals, and is of much greater importance within organizations, where the need to avoid having one's own judgment identified with a potential mistake, i.e. an unfortunate outcome of a decision, is one of the basic motivations.

Of course, use of EUH does not eliminate the need for subjective judgment, but it is often highly convenient that the judgment that is still needed is applied early in the process of coming to a decision, so that at the time the process concludes and a decision is made (or identified, or calculated), there is less opportunity for debate over the subjective estimates. At the same time, the way in which the judgmental aspects are hidden makes them much harder to review: what is convenient for the person making the judgments is often a liability for the integrity of the analysis.

11.3.7.2. Specifics of the Expected Utility Hypothesis

The axioms required for proving the EUH are described next, and other necessary conditions that are implicitly demanded are clarified. This rendering of the EUH Axioms is from Hey (1979:27ff.).

Terms and Definitions

Outcome: A^i is one of the basic outcomes, $A = \{A^i, i = 1, I\}$. A^1 is the most-preferred outcome, A^I is the least-preferred outcome. A (non-degenerate) choice is associated with more than one possible outcome, which is the source of the difficulty.

" $X \geq Y$ " means that either X is preferred to Y or the decision-maker is indifferent.

Choice: $C^j = [(p^1, A^1), (p^2, A^2)]$ indicates a choice that has two possible outcomes, A^1 and A^2 , that will (or are expected to) occur with probabilities p^1 and p^2 respectively. The terms "lottery" or "gamble" are often used instead of "choice" elsewhere in the literature.

The Axioms (Hey's version):

- 1) **Ordering of basic outcomes:** For all i, j , either $A^i \geq A^j$ or $A^j \geq A^i$
- 2) **Transitivity among choices:** $C^i \geq C^j$ and $C^j \geq C^k$ implies $C^i \geq C^k$
- 3) **Continuity:** For all $i = 1, \dots, I$ there is some u^i ($0 \leq u^i \leq 1$) such that $A^i = [(u^i, A^1), (1-u^i, A^I)]$
- 4) **Substitutability:** If $A^i = C^j$ (where C^j is some choice), and $C^1 = [(p^1, A^i), (p^2, A^2), \dots, (p^n, A^n)]$ and $C^2 = [(p^1, C^j), (p^2, A^2), \dots, (p^n, A^n)]$, then $C^1 = C^2$.
- 5) **Application of probability rule:** If $C^j = [(p^{ij}, A^i), i = 1, \dots, I]$, for $j = 1, \dots, J$, and if $C^1 = [(q^j, C^j), j = 1, \dots, J]$, then $C^1 = C^2$ where $C^2 = [(p^i, A^i), i = 1, \dots, I]$, and $p^i = \sum (j = 1, J) p^{ij} q^j$.
- 6) **Monotonicity:** If $C^1 = [(p, A^1), (1-p, A^I)]$ and $C^2 = [(q, A^1), (1-q, A^I)]$, then $C^1 \geq C^2$ if and only if $p \geq q$.

The axioms of expected utility have been debated and reformulated endlessly. A brief summary of Hey's version of each axiom is nonetheless in order, and a few remarks of a more general nature on the relevance of the expected utility model for decision-making follow.

Axiom 1 states that there is a complete semi-order of preferences over basic outcomes. This is a common assumption in economics (see Economic

Rationality Rule 1 earlier), though some question its use as a postulate of rationality. Von Neumann and Morgenstern themselves note that “the common individual, whose behavior one wants to describe, does not measure his utilities exactly, but rather conducts his economic activities in a sphere of considerable haziness” (1953:20). There are many conditions which could cause the decision-maker’s preferences to be incomplete, such as not knowing with certainty what the consequences of a particular choice and world state would be, not knowing the relative importance of different values, or having more than one party involved in determining the value (utility) of the consequence.

Axiom 2 seems very reasonable, though it, too, presupposes exactness in measurement of preference. If all choices can be measured accurately on a one-dimensional scale of utility, Axiom 2 is obvious. There are situations in which preference is not measured unidimensionally, though, and problems of intransitivity can arise there and under other conditions, as in cyclic voting patterns. Von Neumann and Morgenstern remark on the demands for information their hypothesis requires: “[W]e cannot avoid the assumption that all subjects of the economy under consideration are completely informed about the physical characteristics of the situation in which they operate and are able to perform all statistical, mathematical, etc., operations which this knowledge makes possible” (1953:30).

Axiom 3 states that for any outcome, there is a lottery of the best and worst outcomes that is equally preferred to the given outcome. This establishes the comparability of certain and uncertain choices. It can be questioned that the desired u^i always exists: suppose $A^1 = \text{get } £100$, $A^I =$

get bullet in head, and $A^i = \text{get } £1$ (Hey 1979:29). It is peculiarly removed from chronological processes, as discussed below. It also means that even the utility of certain outcomes will be affected by attitudes towards risk, since certain outcomes are evaluated by reference to specific lotteries.

Axiom 4 establishes that outcomes that are common to two lotteries should not affect the decision-maker's preference between the lotteries. This is often formulated as an "Independence Axiom". Much of the controversy over the descriptive accuracy of EUH concerns evidence that people do not behave in accord with this axiom. The evidence is briefly reviewed below in the P Perspective section.

Axiom 5 provides the specific mechanism for calculating the utility of a gamble as the expected utility of each of its uncertain outcomes.

Axiom 6 says that between two lotteries with the best and worst outcomes as the only two possible outcomes, the one in which the best outcome is more likely must be preferred.

11.3.7.3. Statement of the Hypothesis

The Expected Utility Hypothesis states that whenever the axioms are met, it is possible to construct a decision-maker's *utility function* which allows all choices to be compared with each other whether they are certain choices or gambles. Furthermore, the utility of a gamble is simply the mathematical expectation of the utilities of the various outcomes (the average utility across all possible outcomes, weighted by the probability of each outcome). Since the utility of any choice can be calculated in this manner, the choice with the maximum utility can be identified. This is the basis of the identification in economics and other arenas of rational choice

with maximization of expected utility (MEU). MEU as an approach to choice under uncertainty is assessed below in the Implications chapter.

II.3.7.4. Example of uses of tools based on the SDM

Petroleum prospecting involves expensive exploratory drilling with highly uncertain rewards. Hosseini describes a probabilistic analysis of a choice between two potential sites for a drilling a wildcat well (Hosseini 1986). The potential results of the drilling ranged from substantial profits if it produced oil at a high rate, to a considerable net loss if there were no recoverable oil. Though the outcome was uncertain, considerable information was available about the geology of the two sites under consideration. This allowed estimation of the probability of different degrees of success and failure for each site. Because the possible outcomes varied so widely in their desirability, the wildcatter's preference for risk was assessed. Then the expected utility was calculated for each site, given his preferences over the range of possible financial outcomes. The extensive data collected by previous explorations has made probabilistic analysis of drilling decisions a very useful tool for oil companies (e.g. Newendrop 1975; Smith 1974).

Grayson (1960) discusses expected utility analysis of oil prospecting more fully. The probabilities of various degrees of success in drilling can be estimated by examining national statistics, by modifying them according to more local information, when available, or estimated by codifying the driller's subjective judgments about the odds. Oil drillers are well aware that their ventures are risky; they are in a business where the need to make risky choices is a prominent feature. This makes them less

vulnerable to false certainty, which is a chief hazard of using expected utility analysis with subjective probabilities.

Another case in which expected value calculations are very helpful is in the sizing of production runs when orders are not known long in advance but the quantities ordered per month are known for many previous months (Baird 1989:295 ff.). In this problem, the primary uncertainty is over the demand in future time periods, treated as an unknown state of the world, but with the range of possible states, that is demand quantities, well known. When there is no appearance of a trend in the orders, the historical order quantities can be approximated by a mathematical curve to which probability density functions can be fitted. By taking into account the potential costs of either over- or under-producing for the unknown demand, the production quantity that produces the highest expected profit can be determined. The decision-maker's preference for or against risk can be assessed in order to determine the production quantity that would maximize expected utility.

Similar concerns arise in inventory control. The optimal stocking policies depend not only on cost but on both the benefit to the customer, chiefly ensuring an adequate supply for all future orders, and the benefit to suppliers, mainly by attempting to smooth out variations in order quantities (Brown 1978:173ff.). In this case, there is also a chief uncertainty characterized as being a state of the world over which a probability distribution can be inferred. The relevant values are accessible, postulating a single decision-maker is reasonable, the possible actions or policies are knowable, and the consequences arising from combinations of

states of the world and actions chosen are predictable.

II.3.8. The relationship between uncertainty and probability

Where others say that some things are certain, others uncertain, we, differing from them, say that some things are probable, others improbable.

Cicero

Uncertainty and probability have been associated with each other for millennia. Knight (1921) introduced the following distinction between risk and uncertainty: if a probability distribution is known over the set of possible states of the world, we are dealing with *risk*; if such a distribution is not known, the situation is one of *uncertainty*. Others have referred to this distinction as between Type II and Type III decisions (Type I is a decision under certainty). The condition of uncertainty thus implies complete ignorance of the relative likelihoods of the states of the world, and that almost never obtains. Many authors use the term “uncertainty” to refer to both conditions (e.g. Munier 1988:3), and that practice is followed here.

II.3.8.1. Historical development of views of probability

Given for one instant an intelligence which could comprehend all the forces by which nature is animated and the respective positions of the beings which compose it, if moreover this intelligence were vast enough to submit these data to analysis, it would embrace in the same formula both the movements of the largest bodies in the universe and those of the lightest atom; to it nothing would be uncertain, and the future as the past would be present to its eyes.

Laplace

In the 18th century, the dominant conception of the world was of a completely deterministic clockwork. Any uncertainty about something, even a future event, was due to the observer's imperfect knowledge, not to indeterminacy in the world. Knowledge was absolutely certain; probabilities were matters of opinion or belief (Kyburg and Smokler 1980:9). In 1847, Augustus de Morgan wrote of knowledge as a quantity comparable to length and weight, indicating that degrees of belief (of knowledge) could be measured (Kyburg and Smokler 1980:10). In the late 19th century, Gustav Theodor Fechner showed that there were types of indeterminism that were consistent with the laws of causality, helping to establish a legitimate place for laws of chance that apply to the world (Heidelberger 1987:123).

Emile Borel's concern with measurability led him to reject the meaningfulness of true values even for physical quantities. In his view, even present or past physical quantities were indeterminate to a degree, but an additional indeterminism, independent of measurement, applied to future events (Knobloch 1987:229).

Lorenz Krüger suggests that the reluctance to accept statistical laws as legitimate descriptions of nature until long after statistical applications were commonly employed was due to the attractiveness of determinism, a principle perceived to be incompatible with statistical lawfulness (Krüger 1987:60). There were three main factors supporting determinism. The first is that the deterministic Newtonian mechanics was so successful, especially in astronomy. The second is that certainty had been taken to be an essential characteristic of knowledge at least since Aristotle. The third

is that a main goal of acquiring knowledge was to dominate nature which for efficacy again required fully determined results.

By the 1930s, statistical mechanics and the idea of an indeterminate world had achieved widespread acceptance, notwithstanding Einstein's famous complaint ("God does not play dice with the Universe!"). Statistical laws were developed from observations of frequencies, and the frequentist view of probability was preeminent. This view was challenged in the 1920s and 1930s by Ramsey (1926), Keynes (1921), de Finetti (1937) and others, whose subjectivist views are described in the next section. Use of subjective probabilities has become commonly accepted in many fields now, especially in decision analysis.

11.3.8.2. Four major views of probabilities

The four main conceptions of what probabilities mean are generally termed objective or frequency-based, logical, psychological, and subjective. Kyburg and Smokler (1980) point out that while some authors assert that only one view is valid (e.g. de Finetti claimed any but the subjectivist view is meaningless; von Mises and Reichenback claimed only the objectivist view was useful for science), others (e.g. Carnap and Nagel) have admitted more than one type of probability, depending on the source of the probability.

The frequentists (i.e. objectivists or physicalists) hold that probabilities measure the frequency of occurrence of an event over a series of trials, usually the limit frequency over an infinite series. In the case of trials that are not observed more than once, they postulate a frequency that would be observed over the long run, if the trial were to be repeated. In this view,

probabilities are characteristics of the physical world. We use statistical techniques to develop estimates of these true physical probabilities.

In the view initially propounded by J. M. Keynes (1921), probabilities are logical relationships between sets of propositions. Given a set of premises (evidence), there is a single correct degree of belief that a perfectly rational being would have in the conclusion being considered. Probabilities are properties of neither the world nor the observer, but of the formal system of assertions.

A psychological probability is the degree of belief that a person has in a proposition. Psychological probabilities are opinions, and are not required to have any particular relation to the physical world, to evidence, to sets of propositions, or to each other. This conception is important mainly as a background for the subjectivist view.

The subjectivists (or personalists) view probabilities as opinions that can be neither right nor wrong. Sets of opinions held by the same person, however, must be consistent with each other. If a reasonable person can be shown that her set of probabilities is internally inconsistent, she will change at least one probability so that the set is consistent. The subjectivists hold that she may change the probabilities in any way she likes so long as the end result is internal consistency — there is no unique or correct way to do it. For instance, a person may believe that the probability of A is r ; subjectivists maintain that no amount of data makes that belief correct or incorrect. However, the requirement for consistency then compels the person to believe that the probability of *not-A* is $1 - r$. The notion of “true probability” is devoid of meaning. Probabilities are

characteristics of a person's belief system, and may (but need not) change with the acquisition of each piece of new evidence.

This subjective conception is the one typically employed in decision analysis, often closely linked with Bayesian methods of inference. Any new evidence bearing on the likelihood of an event is used to update previous, perhaps implicit, probabilities (the "priors"). The updated probabilities (the "posteriors") may later be updated by additional evidence. The problem of how to obtain or assess prior distributions is serious, as strong priors have a large impact on the posteriors. As Smithson asks,

After all, if personal probabilities need only avoid the Dutch Book¹¹ problem and otherwise satisfy personal probabilities, then why should one idiosyncratic prior distribution be preferred over another? (Smithson 1988:66)

Many influential writers earlier in this century did not hold that all degrees of belief could be measured, or that all degrees of belief could be compared (e.g. J. M. Keynes and Bernard Koopman (Kyburg and Smokler 1980:11), I. J. Good (Good 1962:138)). Nowadays much of the literature that deals with subjective probability concerns expected utility calculations, where the mathematical tractability afforded by assuming that all probabilities indeed can be compared is extremely useful. Coupling that with the pragmatic fact that elicitation procedures are able to produce apparent comparisons¹², the general tendency in current technical

¹¹ Lovell: a "Dutch Book" is a set of bets with the odds arranged so that no matter what events occur, the bettor is a net loser. For example, someone who bets 3 to 2 against a fair coin turning up heads and also bets the same amount 3 to 2 against the coin coming up tails is sure to lose in either event.

¹² The fact that an elicitation procedure produces numerical or comparative results does not by itself establish that those results are meaningful.

literature is to assume that subjective probabilities constitute a complete order.

Patrick Suppes notes that a view of measurement as potentially exact and numerically unique dates back to Lagrange and Laplace, and continues to dominate even in modern decision science and theories of subjective probability (Suppes 1975:153]. To Suppes, this represents “a remnant of past hopes for an orderly, thoroughly deterministic universe” (Suppes 1975:154). In his conception, approximateness is a characteristic of all measurement, specifically including measurement of probabilities and utilities.

Richard Jeffrey argues in “Probable Knowledge” (Jeffrey 1968) that the concept of sure knowledge rarely holds up under close inspection. In a serious way, all of our knowledge is only probable, although sometimes so close to certain as to make the difference insignificant in practice. Making appropriate inferences and developing appropriate degrees of belief from our experience is a skill we learn but never perfect. Even were the world determinate, imperfect evidence for our inferences would lead to different degrees of belief in various assertions. This supports the claims of the subjectivists that the probabilities we work with are opinions and not characteristics of the world.

On the other hand, the success of statistical mechanics shows that there are sources of indeterminacy in the physical world that are regular enough that they can be described numerically with measurable degrees of precision. Andreas Kamlah suggests using different interpretations of probability depending on the individual case (Kamlah 1987:97ff.). The

meaning of $r = p(B/A)$, the probability of B given A , will vary depending on the nature of the arguments A and B . According to Kamlah, if A is a set of events, or descriptions of events, r is an “objective fact”; if A is a set of thoughts, then so is r . Jeffrey’s argument suggests that “objective fact” is an unrealizable ideal of certainty of belief, but does not obviate the concept that the degree of certainty of the arguments matters and affects how we ought to treat the probability r . Probabilities developed from historical data of a great many repeated trials seem to be different from probabilities developed from a few trials or from other opinions instead of observed trials. Sometimes there is a solid statistical or logical basis for calculating the degree of belief that a reasonable person “ought to have”; other times there is none, and a great many situations lie between these extremes.

11.3.8.3. Assessing probabilities

Statistical techniques for estimating probabilities from samples generally include methods for measuring not only the probabilities themselves but also the confidence we ought to have in their measurements. Assessment of appropriate degrees of confidence in objective probabilities (frequencies) is an established technology.

Probabilities that are estimated by reference to a model, i.e. logical probabilities, such as the idealized concept of a “fair die” or “fair coin”, are generally treated as if they were exact. In these situations, the issue of precision lies not in the area of calculating the probabilities, but rather with the exactness or relevance of the logical model to the specific case. For example, everyone except the most extreme subjectivist would agree that the true or correct odds of rolling a “6” with a fair die are 1 in 6; the real

issue is what degree of confidence we ought to have that any particular die is indeed fair. In many cases there is a plausible logical model but no realistic procedure (such as repeated experiments with the same die) for testing the correspondence between the model and the specifics of the case.

Probabilities that reflect pure opinion and not repeated trials are more difficult to assess. Pure subjectivists, such as de Finetti, argue that the concept of a “true probability” is empty and even deceptive (de Finetti 1937). It turns out that the concept of a “true subjective probability” or “true personal probability” is as impossible to operationalize as “true objective probability”. As Savage notes,

[I]t is by no means easy to elicit your own probabilities. Vagueness is a major obstacle, and your first reactions are often greatly modified when you reflect upon their implications. Those who have experimented on themselves and on others generally feel that frequent practice with proper scoring rules and with other probability elicitors helps a person combat vagueness and to arrive more promptly and accurately at his personal probability. This proposition might be difficult to investigate experimentally and even seems difficult to state with precision, but its promise of benefit is great. (Savage 1971:178)

The common usage of “elicit” in the literature implies that people already, in some sense, “have” personal probabilities in their heads. As Morgan and Henrion observe, “More recently it has become apparent that in most cases experts and laypersons do not carry fully formed probability values around in their heads, but must synthesize or construct them when an analyst asks for them” (Morgan and Henrion 1990:108).

Even personal probabilities are available only through some specific assessment technique, such as fixed probability, fixed value, equivalent prior sample, and ranking. Morgan and Henrion conclude that the

evidence on the relative merits of techniques studied to date is mixed, with none appearing to be best (Morgan and Henrion 1990:128). They warn that “whatever elicitation procedure is adopted, subjective judgments are inherently prone to bias and must be treated with care.” This finding parallels the research into methods for assessing utilities: the particular technique makes a big difference, and none can be shown to be “right” (McCord and de Neufville 1983). It is not possible to have direct or accurate access to subjective probabilities or utilities that produces unique numerical results.

After 66 years of analysis and methodological innovation, Ramsey’s assessment still appears to be true: “First, some beliefs can be measured more accurately than others; and, secondly, the measurement of beliefs is almost certainly an ambiguous process leading to a variable answer depending on how exactly the measurement is conducted” (Ramsey 1926:31). Several approaches have been suggested to gauge the accuracy of measures of degrees of belief.

One approach is to seek probability distributions for the probabilities themselves. Some authors, including some subjectivists (e.g. de Finetti 1977), restrict the arguments of probabilities to events that can, at some time, be observed to have obtained or not. They would not accept probabilities about probabilities. In their view, $p(A) = .25$, where A = “the London spot price of West Texas crude oil will be at or above \$60 per barrel on January 31, 2006” is meaningful, while $p(.1 \leq p(A) \leq .4) = .5$ is not, since there is no conceivable experiment that could confirm or refute that the argument of the probability function has occurred or not. Good (1962:144)

suggests calling the first example above a *Type I* probability, and calling probabilities that refer to probability statements *Type II* probabilities. He notes that this typology is potentially infinite, and would rarely suggest using Type III probabilities. Even Type II probabilities are cumbersome to deal with.

A more convenient though less rich approach is to use *upper* and *lower* probabilities to refer to bounds on a range in which we believe the probability lies. This method has been proposed by Kyburg, Smith, Shafer, Dempster and others (Kyburg and Smokler 1980:18). Suppes (1975) presents an axiomatization of approximate beliefs and gambles making use of upper and lower probabilities and utilities. A person may believe that it is at least as likely as X% that event A will occur, but no more likely than Y%. A narrower range indicates greater sureness about the estimated likelihood.

Reliance on point estimates of probabilities poses another problem for Bayesian inference:

This practice does not enable the Subjectivist to distinguish between, say, a uniform prior that rests on considerable accumulated empirical evidence and one generated in blind ignorance by the Principle of Indifference. (Smithson 1988:66)

Whether ranges, Type II probabilities, or other indications of accuracy are adopted, Stigler argues for their importance:

But measurement alone is not enough — it is essential that the measurements be susceptible to comparison, and that, in turn, requires some understanding of the accuracy of the measurements. (Stigler 1987:287)

The use of probabilities in decision analysis centers on the possible future states of the world; “uncertainty” may not be taken to refer to any

other factor in the decision situation. Any specification of a probability distribution over the set of possible futures implicitly defines the probability of an unforeseen future to be 0 (Turner 1978:169). Smithson suggests that this is often overlooked, or taken to be too difficult a problem:

Some of the more sophisticated [probabilistic risk assessment] experts acknowledge this problem but almost invariably place it in the 'too hard' basket. In an article on the estimation of probabilities for risk assessment, Winkler (1982: 352) expresses the conventional tendency to bracket this concern by declaring it irrelevant: "Although the identification of possible outcomes is an important problem, the concern in this paper is with the probabilistic aspects of risk assessment." (Smithson 1988: 87)

11.3.8.4. Summary of probability

Proponents of logical and physical interpretations of probability assert that probabilities refer to something external, either a system of propositions or the physical world as revealed through observed frequencies. Subjectivists argue that people use the concepts of more or less likely (more or less probable) to refer to events about which neither logic nor frequencies speak clearly, and that such uses are meaningful. They suggest further that in the absence of logical or physical bases for probabilities, we often do have *some* information, and subjective probabilities offer the only way of making use of it. Jeffrey's argument, noted above, that all knowledge is merely probable and not certain, adds further legitimacy to the use of subjective probabilities. Ramsey's and Stigler's points, however, emphasize the importance of assessing the degree of precision or certainty associated with each subjective probability. This is often overlooked in presentations of methods (e.g. decision analysis

or simulation) for evaluating uncertain prospects.

A central aspect of the technical view of uncertainty, especially with regard to decision-making, is the intense focus on uncertainty about the future state of the world, to use the terminology introduced earlier in the standard model of decisions (SDM). The probabilistic nature of knowledge that Jeffrey speaks of is generally acknowledged only with regard to this component of the SDM, and it is assumed implicitly that we have certain knowledge about the other components in the model. This assumption is cast into serious doubt by Jeffrey's arguments, and it is explored in more detail in the Taxonomy and Implications chapters below. Those other SDM components are significantly different from the state of the world in some respects. One key difference is the more conspicuous dependence on the role of the human observer/decision-maker: if one wants, one can generate new options or change one's values. There is an obvious and unavoidable subjectivity to these components. The state of the world seems in some sense to be "out there" (although Jeffrey would argue that we must settle for indirect evidence of the state of the world).

Probability is currently used primarily as a measure of the imperfection of our knowledge about the future state of the world. We don't use probabilities generally to express our imprecision of knowledge about the feasibility of options, the form of the function mapping options and world states into consequences, or the translation of consequences into utility. Morgan and Henrion have written a recent, well-regarded book on uncertainty that represents the state of the art in sensitive technical treatment of uncertainty. In it, they explicitly recommend against

probabilistic treatment of defined constants, decision variables (e.g. plant size), value parameters (e.g. discount rate), or model domain parameters (e.g. time horizon or increment) (Morgan and Henrion 1990:51).

Probabilistic treatment is reserved for what they call empirical parameters and chance variables¹³. While they acknowledge that “experienced analysts often argue that uncertainty about structure is usually more important, and more likely to have a substantial effect on the results of the analysis ... than uncertainty about the value of a quantity” (Morgan and Henrion 1990:67), they devote only one and a half pages out of 323 total to this form of uncertainty.

Still another indication of the identification of uncertainty with probability in the T Perspective is the index Morgan and Henrion term the Expected Value of Including Uncertainty (EVIU) (Morgan and Henrion 1990:308-11). The EVIU is defined as the difference between the expected value of the decision minimizing expected loss when treating a variable probabilistically and the expected loss when using only a single measure (e.g. the median or mean) to represent that variable. This “inclusion of uncertainty” clearly applies only to those types of uncertainty for which probabilistic representation is appropriate, that is, to empirical parameters. A more accurate term for the index would be the Expected Value of Probabilistic Analysis, since there is no measure of the value of

¹³ Such parameters or variables are at least in principle susceptible to being measured. The key to the unique suitability of an empirical parameter to probabilistic representation is that it “is both uncertain and can be said to have a *true*, as opposed to an *appropriate* or *good*, value.” [Morgan & Henrion, 1990:52] Empirical parameters will be found in the SDM in the description of the state of the world.

including uncertainty about model form, option availability, the decision-maker's values, etc.

Within decision analysis, probability represents an imperfection in our knowledge of the state of the world, but not imperfections in our knowledge about ourselves (e.g. our values) or potential creations (e.g. alternative actions). Most decision analysts subscribe to the subjectivist view of probabilities in which they represent opinions only, not characteristics of the world. Such probabilities are opinions then, but only opinions that refer to the world. And exact opinions, no less, which surprises Suppes:

It is surely surprising if not literally paradoxical that the view of ideal measurement characteristic of classical determinism has been carried over almost untouched to modern theories of partial belief and decisions. It is, for example, a characteristic requirement of Savage's (1954) theory of decisions that the axioms on preferences between decisions be strong enough to lead to a numerical representation in terms of a *unique* probability distribution on beliefs and a utility function on consequences unique up to a linear transformation. (Suppes 1975:153)

If probabilities are to be reserved for describing degrees of belief that potentially observable events will occur (e.g. De Finetti 1977), how can we quantitatively represent our degree of belief in opinions, such as our degree of belief that the probability of X is r , or our degree of belief in an opinion that the form of a model (i.e. the function mapping alternatives and states of the world into consequences) is thus and such, or our degree of belief that our values are described in a particular pattern?

Problems that remain for justification of probabilities include the dependence of objective probabilities on the "in the long-run" argument, whether conditional or absolute probabilities are more fundamental, and

the current inability of probabilistic frameworks to represent ignorance (Smithson 1988:209). Another concern is that assertions about probabilities are unfalsifiable, except in a probabilistic sense (although Jeffrey argues, with his “probable knowledge” concept (1968), that all assertions must be considered to be merely probable, not absolute).

11.3.8.5. Uncertainty that doesn't fit into probabilistic frameworks

For some authors, probability and associated techniques are adequate for dealing with any type of uncertainty. J. D. Freund (quoted at the beginning of this section on “T”) would probably agree with L. J. Savage: “It may be said to be the purpose of this book, and indeed of statistics generally, to discuss the implications of reasoning for the making of decisions” (Savage 1954:6). Lindley goes even further:

The only satisfactory description of uncertainty is probability. By this I mean that every uncertainty statement must be in the form of a probability; that several uncertainties must be combined using the rules of probability; and that the calculus of probabilities is adequate to handle all situations involving uncertainty... probability is the only sensible description of uncertainty and is adequate for handling all problems involving uncertainty. All other methods are inadequate... Anything that can be done with fuzzy logic, belief functions, upper and lower probabilities, or any other alternative to probability can better be done with probability. (Lindley 1987)

Others, however, point out several types of situation in which probability seems insufficient, inappropriate, or inadequate. Dror (1988:250) mentions the state of extreme uncertainty in which we are essentially *ignorant* of the likelihoods as being inappropriate for probabilistic analysis. This is especially likely and serious, Dror contends, when we are not or cannot be sure what states are even possible. Except in

a few artificial cases, the *principle of insufficient reason* that assigns equal probability to each of the possible states of the world is inappropriate, because the problem of defining discrete outcomes is too subjective (see also Elster 1989a:58). Consider approaching a fork in the road with three branches that could be taken: two go uphill and one goes down. Is this three possibilities, or only two (up versus down)?

Risk analysis is an area where probabilistic considerations alone are often unsatisfactory. Proponents of nuclear power point out that the American nuclear industry has not been proven to be responsible for any deaths, while deaths in the coal industry are relatively common. The expected fatalities from nuclear power, based on both historical probabilities and logically combined probabilities based on rates of component failure, are lower than the expected fatalities from coal power. Yet the public's much greater worry about the hazards of nuclear power cannot easily be dismissed — the probability calculations miss important aspects of risk perception. One interpretation of this is that the public perceives the utility of a single incident causing, say, 10,000 deaths as being much worse than the sum of the utility of 10,000 one-death events. This interpretation argues again against expected *value* analysis and in favor of expected *utility* analysis, and highlights the complexities involved in assessing utility. Many risk studies use expected loss of life, an expected *value*, to represent societal utility, and this may be a misleading criterion.

II.3.9. Possibility Theory

Several authors have suggested that the knowledge or intuitions we have about the uncertain occurrence of events is more similar to judgments

of the *possibility* or feasibility of an event than to judgments of the *probability* of occurrence. G. L. S. Shackle (e.g. Shackle 1949a, 1949b, 1949c, 1961; and Ford 1983) argued for a similar view throughout his career, asserting that what people are actually able to make judgments about is the *surprise value* of a contemplated event. Surprise is essentially the inverse of possibility: the more possible you believe an event to be, the less surprised you would be if it occurred.

One of the first formal expositions of a possibility theory was Lofti Zadeh's "Fuzzy sets"¹⁴ as a basis for a theory of possibility" (Zadeh 1978). The following example illustrates how Zadeh built his concept of possibility on fuzzy sets (Zadeh 1978:6): suppose that the age "28" is thought to have a grade of membership in the fuzzy set "young" equal to 0.7. Zadeh interprets this to mean that age 28 has a degree of *compatibility* of 0.7 with the concept young. Then, upon hearing that "John is young," Zadeh would infer that the *possibility* that John is 28 is 0.7.

More particularly, suppose that A is a fuzzy subset of X , and that $\mu_A: X \rightarrow [0, 1]$ is the membership function for A . Then Zadeh defines a *possibility distribution function* $r_A: X \rightarrow [0, 1]$

$$r_A(x) = \mu_A(x)$$

which he then uses to define a *possibility measure* $\text{Pos}_A: P(X) \rightarrow [0, 1]$

(where $P(X)$ is the power set of X)

¹⁴ Lovell: *fuzzy sets* differ from orthodox sets in how membership in the set is measured. For any orthodox set and any element of the *universal set*, the element is either a member of the set or not. For fuzzy sets, membership is not an all or nothing matter: each element has a *degree of membership* ranging from 0 to 1. Orthodox set theory allows only degrees of 0 and 1, nothing in between.

$$\text{Pos}_A(B) = \max_{x \in B} r_A(x)$$

Later expositions of possibility theory, such as that of Dubois and Prade (1988), also define a necessity measure:

$$\text{Nec}(B) = 1 - \text{Pos}(\neg B)$$

where $\neg B$ is the complement of B .

Dubois and Prade (1988) describe the basic relationships of possibility theory, summarized here. Possibility and necessity are super-additive and sub-additive, respectively:

$$\text{Pos}(B) + \text{Pos}(\neg B) \geq 1$$

$$\text{Nec}(B) + \text{Nec}(\neg B) \leq 1$$

Any event that is even a little bit necessary is completely possible; any event that is less than completely possible is not at all necessary, which implies a third condition that the necessity of an event and the necessity of its complement cannot both be non-zero:

$$\text{Nec}(B) > 0 \Rightarrow \text{Pos}(B) = 1$$

$$\text{Pos}(B) < 1 \Rightarrow \text{Nec}(B) = 0$$

$$\min[\text{Nec}(B), \text{Nec}(\neg B)] = 0$$

The possibility and necessity of conjunctions and disjunctions are calculated using max and min operators:

$$\text{Pos}(A \cup B) = \max[\text{Pos}(A), \text{Pos}(B)]$$

$$\text{Nec}(A \cap B) = \min[\text{Nec}(A), \text{Nec}(B)]$$

II.3.9.1. Belief Theory: A generalization of Possibility Theory

At the same time that Zadeh was developing possibility theory, Shafer (1976) was creating a theory of evidence, based in part on earlier work by

Dempster (1967), now often referred to as Dempster-Shafer theory, evidence theory or belief theory. One of his goals was to separate the notion of degrees of belief from that of chance. The subjective probabilists have taken degrees of belief to be fundamental, and physical chance to be meaningless; the objectivists treat only observable frequencies, and deny the utility of subjective probabilities.

The chances governing an aleatory experiment may or may not coincide with our degrees of belief about the outcome of the experiment. If we know the chances, then we will surely adopt them as our degrees of belief about the outcome of the experiment. But if we do not know the chances, then it will be an extraordinary coincidence for our degrees of belief to be equal to them...Chances, then, must be conceived of as features of the world. They are not necessarily features of our knowledge or belief. And it would be quite untenable to claim that a chance is *merely* a feature of our knowledge or belief. (Shafer 1976: 16-17)

The basic components of belief theory are succinctly summarized by Klir (1989: 351):

Let X denote a universal set under consideration, assumed here to be finite for the sake of simplicity, and let $P(X)$ denote the power set of X . Then, the Dempster-Shafer theory is based upon a function

$$m: P(X) \rightarrow [0, 1]$$

such that

$$m(\emptyset) = 0 \text{ and } \sum_{A \subseteq X} m(A) = 1.$$

This function is called a *basic assignment*; the value $m(A)$ represents the degree of belief (based on relevant evidence) that a specific element of X belongs to set A , but not to any particular subset of A . Every set $A \in P(X)$ for which $m(A) \neq 0$ is called a *focal element*. The pair (F, m) where F denotes the set of all focal elements of m , is called a *body of evidence*.

Associated with each basic argument m is a pair of measures, a *belief measure*, Bel , and a *plausibility measure*, Pl , which are determined for all sets $A \in P(X)$ by the equations

$$\text{Bel}(A) = \sum_{B \subseteq A} m(B),$$

$$\text{Pl}(A) = \sum_{B \cap A \neq \emptyset} m(B).$$

Belief and plausibility measures are connected by the equation

$$\text{Pl}(A) = 1 - \text{Bel}(\neg A)$$

for all $A \in P(X)$, where $\neg A$ denotes the complement of A .
Furthermore,

$$\text{Bel}(A) \leq \text{Pl}(A)$$

for all $A \in P(X)$.

In the special case that every focal element A is a singleton, the belief and plausibility functions collapse into a single function which can be seen to be a probability function. The basic assignment is performed by a *probability distribution function*

$$p: X \rightarrow [0, 1]$$

as follows:

$$m(x) = p(x).$$

A second important special case is when the focal elements are *nested*, i.e. they can be ordered by set inclusion. Here, much as in Zadeh's formulation above, the basic assignment is performed by a possibility distribution function

$$r: X \rightarrow [0, 1]$$

as follows:

$$\text{Pos}(A) = \max_{x \in A} r(x).$$

Belief theory provides a sufficiently general framework that both probability and possibility can be viewed as special, non-overlapping cases. The greater generality of belief theory gives it greater power to describe

different types of uncertainty, but brings with it a higher cost in computational complexity (Klir 1990b:271). Resconi, Klir and St. Clair (1992) have examined the relationship between current mathematical formalisms of uncertainty (probability, possibility, belief, and fuzzy set theories) and modal logic¹⁵. From this analysis, they propose a “hierarchical uncertainty metatheory” which suggests additional formalisms which have yet to be developed.

Possibility theory is not the only recent field to have emerged from the large and diverse world of fuzzy sets that provides additional mathematical tools for treating uncertainty. Notable among those with potential relevance for decision-making is *fuzzy measure theory* (Sugeno 1977).

11.3.10. Complexity and uncertainty

The principle of incompatibility. Stated informally, the essence of this principle is that as the complexity of a system increases, our ability to make precise and yet significant statements about its behavior diminishes until a threshold is reached beyond which precision and significance (or relevance) become almost mutually exclusive characteristics. (Zadeh 1973)

Still more informally, Zadeh’s Principle of Incompatibility states that complexity makes precision and relevance incompatible.

Zadeh’s Principle implies that for descriptions or representations of complex systems to be relevant, they must include uncertainty, and the uncertainty must increase with the systems’ complexity. With fuzzy sets,

¹⁵ Modal logic is an extension of classical logic: in addition to the classical logic operators of implication, disjunction, conjunction, negation, and equivalence, the two *modal operators* of possibility and necessity are defined.

Zadeh hoped to offer a tool that would explicitly permit imprecision in system descriptions and still allow rigorous analysis, thus preserving the chance for relevance.

Not all complexity is problematic for analyses. Weaver (1968) articulated three different complexities. *Organized simplicity* refers to systems with a small number of relevant variables with complex relationships among them. Formal analytical or mathematical methods are often very successful for dealing with organized simplicity. However, increasing the number of variables can pose intractable computational and analytical burdens with this type of system. *Disorganized complexity* is at the far end of a continuum, pertaining to systems with large numbers of relevant variables that are related in simple ways. Statistical methods often work well for such systems. Increasing the number of variables and increasing the degree of randomness of their interaction actually improves the accuracy of statistical tools.

The third type is *organized complexity*, in which the variables are too numerous for analytical methods, and yet they are not amenable to statistical assessment, e.g. due to the importance of individual variables. For example, statistical methods work well for forecasting the electrical load placed on a grid by electric water heater usage: a very large number of customers make individual decisions about when to use hot water, but the aggregate usage is quite predictable. We are left with very large uncertainty about the decisions of any one customer, but the uncertainty does not matter. Uncertainty about an individual decision-maker matters much more when there is a small set of decision-makers whose actions

have a large impact. An example of this sort is the industrial customers of BPA. There are about a dozen such customers, each of which makes decisions about whether to put on or remove one or more shifts, or shut down a plant, based on the world price of aluminum, the current cost of electricity, and other factors, including corporate factors that BPA cannot discover. The impact of a decision by one of these customers has a larger load impact than the total of the hot water usage decisions, but is less predictable.

Problems of the more difficult, third type of complexity abound in medicine, the soft sciences, and interdisciplinary work, according to Klir (1985:136). One way to treat such problems is to simplify the system (Klir defines a system as an abstract representation of a set of real world objects and relationships). Generally, simplifying a system will increase the uncertainty in its description of its components. Klir proffers a *principle of minimum uncertainty* to guide the simplification process: the revised system should contain the minimum additional uncertainty possible (Klir 1992a:71). Klir suggests that allowing additional imprecision into the data representation via the use of fuzzy sets may be an even more important strategy for dealing with organized complexity (1985:138).

When a system is “nearly decomposable”¹⁶ (Simon 1981:210), as many hierarchically structured systems are, a reductionist approach may work well, as it has in many of the hard sciences. This is not quite a

¹⁶ In a “nearly decomposable” system, made up of a number of sub-systems, the relationships among the sub-systems are far less important than the relationships within each sub-system. Much of the behavior of such a system can be understood by studying the sub-systems independently.

simplification strategy, since a set of smaller systems, not a single, simpler system, replaces the original system. Each of the smaller systems is then analyzed separately, and the insights from the subsystem analyses are relevant to the original system in the degree to which that system is decomposable.

The Incompatibility Principle suggests that it would be useful to be able to measure complexity of an object. Klir, after Ashby¹⁷, defines a *system* as an abstraction, a set of variables that have been chosen to represent an object (Klir 1985:132). While there are virtually unlimited variables that could be defined for any object, a particular human observer selects a typically small set from those variables that she can perceive and in which she is interested.

It is not operationally meaningful to view complexity as an intrinsic property of objects. This does not mean, however, that I deny the existence of complexity of objects in the ontological sense. I only recognize that the notion of object complexity is epistemologically and methodologically vacuous, in contrast to the notion of systems complexity. (Klir 1985:132)

Thus, while it can be meaningful to compare the complexity of objects, those comparisons are relative to the representations chosen by an observer and therefore inherently subjective. Objective measurement of complexity of objects is not possible.

II.3.11. Measurement

By the end of the nineteenth century the official philosophy was that you understood something only when you could measure it. (Hacking 1987:48)

¹⁷ Ashby, W. R. 1956. *An Introduction to Cybernetics*, John Wiley, New York;
— 1973. "Some peculiarities of complex systems," *Cybernet. Med.*, **9**, pp. 1-7

To a statistician, the whole of the history of science revolves around measurement: all sciences (and the social sciences are no exception) require at least the possibility of quantitative measurement. (Stigler 1987:287)

Unlike empirical work in science, normative decision theory writing rarely mentions the accuracy of its basic measurements, those of utility and subjective probability. Suppes views this as a result of a “mythology of ideal measurement,” a remnant of the determinism of Simpson, Lagrange, and Laplace, which by the very discussion of “errors” of measurement implies that perfect measurement is possible (Suppes 1975:153). However, as Suppes notes, if perfect measurements are in principle impossible, then the uniqueness of measurements is also impossible, and that is problematic for axiomatic theories of decision such as Savage’s which require unique probability distributions, and utility functions unique up to linear transformations.

Indeed, experimental evidence suggests that different techniques for eliciting utility curve data produce incompatible results, with no normative basis for identifying the unique, “true” curve (see the passage below on the Utility Evaluation Effect in section II.5.4.6).

II.3.1.2. Tools of the T Perspective

Three major normative approaches deriving directly from the T perspective are briefly described. Two of the approaches, decision analysis and simulation, are used mainly in cases where the uncertainty is primarily about the state of the world that will eventuate, i.e., the uncertainty is *parametric*. The third approach, game theory, concerns uncertainty due to the importance of decisions to be made by other actors whose interests are interdependent. This is a *strategic* uncertainty in Jon

Elster's distinction between parametric and strategic rationality (discussed earlier in section II.3.3.5.). Lastly, sensitivity analysis is described.

II.3.12.1. Decision analysis

Decision analysis is an analytically oriented approach to evaluating alternative courses of actions under environmental (i.e., parametric) uncertainty. Its subject is games against nature — situations where the decision-maker is not interdependent with other decision-makers. It is essentially utility theory put into practise, and shares the major assumptions and concepts of that theory. In particular, the standard decision model (SDM) described earlier is a fundamental conceptual model in decision analysis. Once the sets of elements of the SDM have been described, the utility of each outcome is determined. The utilities of the outcomes possible for each alternative are weighted by the estimated probabilities for each of them, and the expected utility for each alternative is calculated. This value, according to expected utility theory, is the appropriate index of the desirability of the alternative, and the alternative(s) with the greatest expected utility is the rational choice.

The probability distribution over the states of the world plays a critical role in this analysis, and has received much attention. Von Neumann and Morgenstern stated explicitly in their original exposition of the expected utility hypothesis that the probabilities they considered appropriate were based on long-run frequency, and rejected subjective probabilities (von Neumann and Morgenstern 1953:19). Savage, however, regarded probabilities as fundamentally personalistic in his treatment of expected utility (Savage 1954). Frequency-based probabilities are often unavailable,

and decision analysts generally accept subjective probabilities as an appropriate measure of the information available to the analyst or decision-maker¹⁸ about the future state of the world (e.g., Morgan and Henrion 1990:49; Roberstshaw, et al. 1978:133). In some cases it is possible to perform experiments to acquire additional information and, using Bayes's Rule, modify the original probability estimates. There are methods for calculating the expected value of perfect information in order to decide whether further experimentation may be worthwhile.

Much of the apparatus of decision analysis has been constructed to facilitate the estimation of utility values and subjective probabilities, and there are many studies of the relative merits of various assessment procedures (some of them will be described below in the P perspective section). The evidence to date indicates that there is no procedure that is "best", and that different procedures may be irreconcilable for estimating probabilities (Morgan and Henrion 1990:128) or for estimating utilities (McCord and de Neufville 1983:301).

Electric utility DA example. An electric utility's load is growing, and the utility will soon have to acquire 1,000 megawatts (MW) of increased generating capacity. There are three actions the utility is considering to meet the increased demand: natural gas turbines, conservation, and cogeneration. The consequences, or measures of goodness, that the utility

¹⁸ Whether the subjective probabilities represent the analyst's beliefs or the decision-maker's is important. The procedure is supposed to maximize the decision-maker's utility, and most methods for assessing or eliciting utilities involve probabilities. All of the probability judgments used in the process should be consistent with each other. This will not automatically be true if the analyst's probabilities are used.

cares about are cost and how well the new generation will fit into the pattern of load the utility expects. The cost of any of the three actions is the sum of the capital costs required to purchase the resource and the operating costs once the resource is on-line. Conservation costs will be fairly constant no matter what the gas price is: conservation costs are all capital costs paid up front. The actual costs of cogeneration depend on the price of fuel, but some cogeneration contracts put all of the fuel price risk on the cogenerator and fix the cost paid for the electricity generated. We will assume the utility will bear only a small portion of the fuel price risk. Contracts for cogeneration energy typically are fixed so that the utility must purchase the amount the cogenerator produces whether it is needed or not. Gas turbine costs, however, cannot be known ahead of time with certainty, since the price of the natural gas fuel they burn is quite variable. Analysts have decided to describe the range of possibilities in terms of high, medium and low gas price, though the possible gas price is continuously variable, this discrete approximation is convenient for the analysis. The sets of available actions, possible future states of the world, and relevant consequences are summarized in table 6.

Cogeneration plants are typically run at the discretion of the owner of the heat process (which is generally not the utility), often at a constant rate around the clock. The utility has no ability to dispatch the power to match its hourly pattern of load, which is generally much higher during the day than during the night. Conservation “produces” power by reducing demand, and is not dispatchable either. However, the pattern of load reduction from conservation generally matches up fairly well with the

typical utility load shape. Gas turbines are very flexible, and their output can be matched to a changing load pattern very well. If cogeneration is selected, matching the utility's total generation to its load will be difficult, and might require purchasing additional flexibility of some sort.

Table 6. Definition of terms for electric utility DA example

Actions	
a_1	Natural gas turbines
a_2	Conservation
a_3	Cogeneration
States of the World	
w_1	High gas price
w_2	Medium gas price
w_3	Low gas price
Consequences	
c_1	Cost: $c_1 = f(a, w)$
c_2	Load following: $c_2 = f(a)$

To complete this example, let us define for each of the consequences a scale ranging from 0 (worst) to 10 (best), and give each of the alternatives a rating. For load following flexibility, the gas turbine is excellent, and gets a 10. Conservation is quite good, though not flexible, and it gets a 7. Cogeneration is quite poor in this regard, and gets a 0. Cogeneration is thought to be the cheapest on average, however, and is rated at 7, 8, or 9 for cost, reflecting the fraction of fuel price the utility pays. Conservation is almost as cheap, getting a 7. Gas turbine costs vary greatly with the price of gas. If the "most likely" forecasts of gas price are borne out, gas turbine costs would be about 5. The natural gas consumed in the U.S. is all produced in North America, but its price is closely tied to the world price

for oil, which could rise sharply and suddenly, as we have seen in the past. The high gas price rating for gas turbine costs gets a 0. On the other hand, the world price of oil in 1991 was quite low, and gas was plentiful; the low gas price could turn out to be very low, so in that possible state of the world, turbine cost gets a 10. The consequences for each option under each state of the world are shown in table 7.

Each of the cells with consequences in it represents a *basic outcome*, and it is assumed by the SDM that the decision-maker is able to make preference comparisons between basic outcomes. Among the entries in the

Consequences (Cost, Flexibility)	$\Leftarrow$ States-of-the-World $\Rightarrow$		
	Low Gas Price	Medium Gas	High Gas
$\Downarrow$ Alternatives $\Downarrow$			
Gas Turbines	(10, 10)	(5, 10)	(0, 10)
Cogeneration	(9, 0)	(8, 0)	(7, 0)
Conservation	(6, 7)	(6, 7)	(6, 7)

column representing Low Gas, the comparison is simple, since (10, 10) dominates the others. The other two columns contain sets of three outcomes each with no dominant outcome, and the decision-maker must have a way to compare the two-dimensional consequences.

Table 7. Consequence values in DA example

The standard technique in decision analysis is to assume the existence of a *utility function* and then to estimate its parameters. Let us assume that for our example utility, cost is three times as important as load following characteristics, and construct a simple linear function of utility

on cost:

$$u(c) = 3 * \text{cost} + 1 * \text{flexibility} = 3 * c_1 + c_2$$

With this formulation of the utility function, the table of consequences can be simplified, as portrayed in table 8.

Table 8. Utility values in DA example

Utility Values $U = 3 * c_1 + c_2$	$\Leftarrow$ States-of-the-World $\Rightarrow$		
$\Downarrow$ Alternatives $\Downarrow$	Low Gas Price	Medium Gas	High Gas
Gas Turbines	40	25	10
Cogeneration	27	24	21
Conservation	25	25	25

Suppose that we consider the medium gas price the most likely — perhaps it has a 60% chance of being realized. With our current information, a high gas price looks quite unlikely — say 15%. The low gas price then has a 25% chance. Using these probabilities of the states of the world, we can calculate the expected utility of each action thus:

$$\bar{u}(a_i) = \sum_k p(w_k) * u(a_i, w_k)$$

The results of the calculations are depicted in table 9, indicating that — given the assumptions made about utility values and probabilities of the states of the world — the Gas Turbine option has the greatest utility, and would be recommended as a result of this analysis.

Table 9. Expected utility values in DA example

Utility Values	⇐ States-of-the-World ⇒			Expectation
	Low Gas	Medium	High Gas	
Probability ⇒	25%	60%	15%	•
⇓ Alternatives ⇓				
Gas Turbines	40	25	10	26.5
Cogeneration	27	24	21	24.3
Conservation	25	25	25	25.0

Decision Trees. An alternative way to view decision analyses is through *decision trees* (see figure 9 below). The analyses and computations are the same, but a graphical portrayal may be preferred for explanatory or other purposes. The paths are called *branches* and the *nodes* are referred to as *decision forks* (the square nodes) or *chance forks* (the circles). The utility of each *terminal node* must be assessed, and then the analysis works backward to determine the most desirable branch to take at the first decision fork.

Related fields with largely similar sets of technical assumptions are operations research and cost/benefit analysis. Cost/benefit analysis does not focus as much on the assessment of uncertainty. Operations research often does, but it is generally a statistical uncertainty arising from repetitive processes where data on the frequency of the events of interest can be gathered, rather than uncertainty due to one-of-a-kind future events.

11.3.12.2. Simulation modeling

Simulation techniques share most of the core assumptions and conceptual models of decision analysis, but rely much more on

computational power and somewhat less on formal analytic procedures.

Figure 9. Utility example decision tree

They are a primary tool in the application of operations research, and are increasingly used in policy analyses. A simulation model is typically constructed by modeling the components of the system individually, often using probability distributions to describe their behavior, describing in computer code the relationships among the components, and using the computer to simulate the interactions of the components and provide an overall picture of the behavior of the system (Hillier and Lieberman 1980:642). When it is not possible to describe the system behavior directly (due to, e.g., complexity of convolving the probability distributions of components, lack of data, system hasn't been constructed yet, etc.), but the elements of the system and their interactions can be described, data about the system's behavior can be obtained by exercising the simulation model

repeatedly.

As the number of probabilistic inputs increases, the complexity of a decision tree analysis generally increases exponentially, while the number of runs of a simulation model required for a given degree of output accuracy increase only linearly (Morgan and Henrion 1990:199). Greater precision in results can be gotten by performing more sampling with the model. This sampling can be treated much like an experiment with the system of interest, and a large technology of statistical tools is available for assessing the sampling (for example, determining confidence intervals, or the number of runs required for a particular degree of uncertainty in outputs). Computational costs have fallen rapidly since computers became widely available, and this has made simulation techniques increasingly attractive. While statistical measures can be derived from the simulation output to any desired precision, this does not guarantee any correspondence to the behavior of the real system, even if the components have all been modeled carefully. The overall behavior of the model should be compared to historical data from the system, if available, or from field tests if possible, or at least to the judgments of knowledgeable people (Hillier and Lieberman 1980:656).

11.3.12.3. Game theory

Incorporating the effects of choices made by other actors moves us from parametric rationality to strategic rationality (Elster 1984), and to game theory. A *game* consists of a set of two or more *players* who are able to make actions or *moves* as permitted by the *rules of the game*; their moves jointly determine, under the rules, the consequences or *payoffs* to each

player.

The payoff structure of a game is often exhibited in a matrix such as the one in figure 10. The two players, A and B, each have two move available. If A chooses 1 and B chooses 4, A receives t , the payoff in the lower left half of the cell at the intersection of row 1 and column 4, while B receives u . When the payoff

		Player B	
		3	4
Player A	1	r s t u	
	2	v w x y	

Figure 10.
A payoff matrix

structure is such that one player's gain is exactly offset by another's loss, the game is called *constant-sum* or more often simply *zero-sum*. In a non-zero-sum game, there is the possibility of an outcome where the total of all players' payoffs is higher than in other outcomes.

A move whose effect is certain may be called a *personal* move, such as moving a chess piece, while one that depends partly on probability may be called a *chance* move, such as deciding to draw another card. A *strategy* is a decision rule that completely specifies what a player is to do under different circumstances. An *optimal* strategy is one that has been shown to be as good as possible. A *pure* strategy specifies with certainty what the player is to do. A *mixed* strategy specifies a combination of pure strategies from which one is chosen probabilistically for each playing of the game. A game for which the optimal strategies have been found for all players is said to be *solved*. A game has an *equilibrium point* if each player has a strategy that is optimal even when the players know each other's strategies, i.e., upon finding out each other strategies, there is no advantage for any player to adopt another strategy. Two-person, zero-sum

games with perfect information with only finitely many possible strategies always have solutions¹⁹, often using mixed strategies.

The relevance of game theory depends on several factors. A game, as a formally defined system, is a model, and may more or less accurately portray any given situation. Formal models of the actual play of chess and rock-scissors-paper can be trusted to be very faithful portrayals; modeling an arms race as a game of chicken requires serious assessment of the usefulness of the analogy. Computational considerations may be important. Chess is a well-defined game, with a large but finite number of strategies possible, but has not been solved. If it had, as von Neumann and Morgenstern observed in 1944, “there would be nothing left to play.” (von Neumann and Morgenstern 1953:125) Lacking a complete solution, we depend on heuristics, and even the advances in computing power in the last half-century are overmatched by the computational impediments of a complete solution. The concept of a complete solution to any game is extremely attractive, and may have colored the research into rational choice in general and decision analysis problems in particular, which are sometimes characterized as games against nature.

11.3.12.4. Sensitivity analysis

Sensitivity analysis is an approach for systematically discovering how sensitive the results of an analysis are to changes in specific parameters, assumptions or inputs used in the analysis. It is usually performed after obtaining a preliminary conclusion, and can be used whether or not the

¹⁹ John von Neumann proved this *Fundamental Theorem of Game Theory* in 1928 in “Zur Theorie Der Gesellschaftsspiele,” *Math. Annalen.*, Vol. 100 (1928), pp. 295-320

preliminary analysis explicitly incorporated uncertainties. Different values of the uncertain factors are substituted into the analysis to find out if the preliminary result would still hold under the new values. Sometimes a range of indifference for each of several parameters is determined by moving the parameter as much as possible in either direction until the result changes (e.g. a different option appears to be best). As Baird (1989:375) puts it, the essence of sensitivity analysis is to ask, "How important is this assumption or that value in a particular decision situation?"

For instance, the variables in the example in section II.3.12.1. might be uncertain. The probabilities of future gas price levels are subjective, and it would be useful to know how depended on those estimates the recommendation for the Gas Turbine is. This options fares very well under low gas prices, and very poorly under high gas prices. If the probabilities for the low and high gas futures were reversed, the results would change to those give in table 10.

Table 10. Modified expected utility values in DA example

Utility Values	⇐ States-of-the-World ⇒			Expectation
	Low Gas	Medium	High Gas	
Probability ⇒	15%	60%	25%	•
⇓ Alternatives ⇓				
Gas Turbines	40	25	10	23.5
Cogeneration	27	24	21	23.7
Conservation	25	25	25	25.0

The change in probability makes the Gas Turbine now the last-place option. It could be useful to go on to explore the effect of varying the cost

factors and the formulation of the utility curve.

Often the effect of a change in one factor depends on the values of other factors. In such cases, it is useful to exercise two or more variables at the same time. This method is only practical, however, with a few variables and a few values for each, for both reasons of both computation and comprehension. Morgan and Henrion observe that varying 10 variables simultaneously with only three values considered for each would yield a total of 59,049 combinations to evaluate (1990:178). They conclude “the results of a parametric analysis with more than three or four dimensions are almost impossible to display in a comprehensible way, and so it is typically of little practical use.” (Morgan and Henrion 1990:178).

While sensitivity analysis is a standard component of decision analysis, it can be usefully applied to game theoretic analyses, the results of computer simulations, qualitative work, or any other assessment in which key factors are not known with certainty.

II.3.13. The silver lining for “T”

In the T Perspective, being uncertain has no advantages. In situations of strategic uncertainty, with other agents competing for the same resources, it can be advantageous to increase the uncertainty of other parties. “The essence of randomization in a two-person zero-sum game is to preclude the adversary’s gaining intelligence about one’s own mode of play” (Schelling 1960:175), — i.e. to maximize his *uncertainty*. Von Neumann and Morgenstern believed the matter crucial in zero-sum games: “We have placed considerations concerning the danger of having one’s strategy found out by the opponent into an absolutely central position”

(Von Neumann and Morgenstern 1953:147).

Interdependence (i.e. the realm of non-zero sum games) is another matter: in an arms race, an opponent's uncertainty may translate into higher anxiety, and may precipitate dangerous, irreversible actions.

II.4. Organizational Perspectives on uncertainty

This book, however, has been written ... with the rationality of organizations being postulated. Note that this is not a postulate that organizations do not make errors. Instead it is a postulate that one can explain variations in the social structure of different parts of organizations on the presumption that the structure of the information and decision system will grow toward those sources of uncertainty whose rational management will matter most to the organization. The evolutionary principle here is that the organs of an organization grow better adapted to their environment over time — in other words, that the rationality of an organization improves. (Stinchcombe 1990:348)

II.4.1. Introduction to "O"

The O Perspective on uncertainty is diffuse and non-abstract. It arises in contexts where the problems are forced upon people, not chosen, and the problems bring their own urgency that may overturn the priorities of the people and organizations they affect. These people and organizations exist in dynamic worlds where time is an important dimension, and therefore planning — designing actions and sequences of actions in anticipation of sequences of consequences — is more central to the O Perspective than is decision-making. The T Perspective on uncertainty focuses on a single decision, abstracted from time; the O Perspective focuses on processes, including planning and decision-making processes, that unfold and develop only through time. In that view, *uncertainty* often refers to a kind of information that is needed continually or recurrently but is not adequately

available. For instance, a T Perspective uncertainty might be the total amount of electric load to be served next month. The corresponding concern in the O Perspective would be loads in general, encompassing the need to stay up to date on how much electricity is sold in each of many time periods, with the accuracy of information about past loads and the time and expense required to obtain it as much a concern as the need to forecast loads for next month, and the knowledge that when the next month arrives, it will be important to forecast the loads for the following month.

There is no single, objective perspective that is adequate for all needs. Different organizations and parts of organizations deal with different kinds of uncertainties and have different goals. While the T Perspective typically assumes for convenience that the decision-maker is a single person, many separate people or groups are often involved in organizational decisions. Frequently the person performing much of the work is an analyst who is not privy to the values and goals of the ostensible decision-maker, a member of management.

Organizational uncertainty varies not only with the inevitable multiplicity of perspectives of different individuals within an organization, but with the diversity of organizational environment, goals, size, structure, etc. It is very difficult to say anything non-trivial that is relevant to all organizational uncertainties (Zadeh's Incompatibility Principle again). Thus, the particular emphases of students of organization are as multifarious as this subject. Among the authors reviewed in detail for this work, Stinchcombe probes organizational structure as a function of

uncertainty, Crozier investigates uncertainty as the foundation of interpersonal power, and Dror emphasizes the irreducibility of strategic uncertainty in organizations' environments and the implications for planning. These three points of view do not even approximate a thorough coverage of the relationship between organizations and uncertainty; they omit the work of behavioralists (e.g. Simon, March, Cyert, Jauch or Kraft) on how decisions and decision rules are derived within organizations, Emery and Trist's treatment of organizational environments, Smithson's sociological examination of the social creation of ignorance, of which uncertainty is one sub-type, and that of many others. I have chosen to review only a subset of the available works in the hope of presenting some views in enough detail to be meaningful.

Reflecting the diversity of organizational issues and perspectives, this section will be more diverse and less unified than the section covering the T Perspective. A contributing factor is that while the T Perspective assumes that there is a single "objective" perspective, recognition of the importance of many perspectives and the relativity of any particular perspective is part of the O Perspective.

The mixed picture that emerges of how organizations use knowledge and ignorance is even more vexing [than how individuals do]. Clearly organizations are not always certainty maximizers, and even when they do maximize certainty it is not in the same way individuals do. Unfortunately, the organizational literature does not provide anything like a comprehensive overview of these matters. Instead, it contains a hodge-podge of studies, suggestions, and conceptual schemas in which ignorance (mainly uncertainty) appears fitfully and plays ambiguous roles. (Smithson 1988:240)

In spite of the lack of coherence in the organizational view, there is a theme that much of this work shares:

The classical view from the organizational literature is strongly normative, as is the classical view in the psychological/decision-making school. It views uncertainty as external to an organization, objective, and harmful (e.g., Cyert and March 1963, Burns and Stalker 1961, Jauch and Kraft 1986). (Smithson 1988:242)

Much of the decision-making that explicitly strives to be “rational” is performed within organizations. Implementation of rational decision principles is typically conducted by specific organizations in order to meet particular goals. The messy details that the T Perspective must discard in order to develop generalized solutions are precisely the details that must be attended to in dealing with specific organizational issues.

11.4.1.1. Planning and decision-making literatures

The literature on planning differs strikingly from the literature on decision-making under uncertainty. The planning literature largely comprises descriptions of procedures and lessons learned in particular applications in specific fields rather than treatments of planning per se (Dror 1987:25). It tends to be contributed to by people who perform planning; the decision-making literature focuses on theoretical treatments, and is dominated by academicians. Perhaps one reason for the major difference in emphasis is that planning is a recognizable field: some people have had experience in that field, most have not. Decision-making, however, is such a pervasive human activity that there is no identifiable subset of humanity which might claim the domain as its own by virtue of being the only practitioners.

Stinchcombe (1990:175) points out a distinction drawn by Mannheim between sociological theories that is relevant to the distinction between

planning and decision-making theories:

The basic argument of [*Ideology and Utopia*, Mannheim 1936] is that theories of activities already embedded in ongoing practices tend to be disintegrated theories about how to solve particular problems, emphasizing the wisdom of institutions and the value of experience, with no overall scheme of how the whole thing works. Theories associated with movements for innovation ("utopias"), in contrast, tend to cover great areas of social life with one simple theory, one master mechanism to drive it all, involving a cataclysmic crisis or revolution that will quickly create a new society that runs on different principles, with no hard work of building that society up piece by piece involved. (Stinchcombe 1990:175)

The analogy is that decision-making cannot be characterized now as a distinct "ongoing practice", since it pervades virtually all human activities, and T Perspective theories of decision-making share the prescriptive, highly general orientation of utopian theories.

The scope of the two activities differs: while planning includes decision-making, it includes many other complications. This makes it much messier and less amenable to abstraction, generalization, and mathematical formalization, which are critical tools in most scientific work.

Another factor differentiating the two subjects is that decision-making, even decision-making under uncertainty, does not necessarily involve predictions about the future, whereas planning inescapably does. The aspects of the future that concern planners are rarely small and neat — they are generally vague, messy, and difficult to quantify. They can only rarely be associated with historical data in order to develop frequency-based probabilities. In contrast, there are many decision-making tasks for which statistical data is available.

Early work on the effect of uncertainty on organizations emphasized

the difficulty it poses for maintaining internal control and its threat to the equilibrium of an organization (e.g. March and Simon 1958). The recent trend in the field has been away from attempting to eliminate uncertainty and towards efforts to understand and manage it (Smithson 1988:250). For example, Miles and Snow (1978) have identified organizations that deliberately seek opportunities provided by uncertain environments, terming them “prospectors”.

Terminological Note. Different authors use a variety of terms to refer to essentially similar fields, and the following terms are used herein more or less synonymously: planning, policymaking, systems analysis, policy research, policy analysis, policy science.

II.4.2. Organizational rationality

Economic theory often treats organizations as if they are each monolithic entities engaged in utility or profit maximization. This view accepts all of the assumptions about rationality that are implied by the SDM. Planning perspectives also generally assume that organizations “intend” to act rationally (Dror 1987:30), though it is a less restrictive notion of “rationally” than economic theory employs. Winter (1971) describes two alternatives to rationality as maximizing that groups of authors have considered, managerialism and behavioral theories, and Dror suggests that a more inclusive concept is needed, which he terms ultra-rationality.

Much of the technical work on rationality, and much of the research in cognitive psychology as well, assumes that the consequences of decisions can be known fairly well, at least to the extent that the possible outcomes

and a probability distribution over them is known. For a member of an organization, this assumption is often untenable, which may help explain why organizational perspectives are less concerned with “perfect” rationality than are technical perspectives.

Rationality necessarily involves an analysis of the future, because the consequences that give purpose to acts are necessarily in the future. Thus all rationality is based on predictions of one kind or another, not on knowledge. Assuming that actors are perfectly rational, of course, implies that they are certain what the future holds, that in all relevant respects our notions about the future constitute knowledge. (Stinchcombe 1990:1)

11.4.2.1. Economic rationality and the SDM

The mainstream economic view of organizations (taken to be synonymous with firms) is based on the SDM: organizations act as unified decision-makers whose objective is maximization of utility, taken to mean maximization of profits, with some “reasonable” distribution of profits throughout time assumed, allowing for decisions with short-term expected losses in anticipation of long-term gains. Uncertainty appears in this view as a complication for the task of rationally selecting optimal actions.

Optimal selection is not seen as the essence of rationality by organizational theorists as it often is by writers in the T world. McCall writes, “*having* managerial talent to run the business is just as significant as making optimal choices.” (McCall 1988:10). This perspective not only points to a different focus of attention — ensuring the supply of managerial talent instead of making optimal selections — but optimization is disqualified for the new focus. McCall cites Kotter’s research which suggests that as much as 10 to 20 years are required to “grow” a general manager (Kotter 1982). Business projections are not accurate over spans

half that long, and it is not possible to grow optimal assortments of managerial talent. Rather, McCall argues, an organization ought to attempt to acquire a wide enough variety of talent to deal with diverse possible challenges (pursuit of *robustness*, not *optimality*).

In many decisions within organizations, there is tension between the staff's sincere desire to remain faithful to principles of substantive rationality and the need by their managers to reach some decision in a specific time frame. Rational procedures are not always decisive; there is often pressure to force them to be in order to facilitate the making of a decision. Rational procedures are often seen within organizations as the best evidence that a decision is defensible via the claim that it is not merely a judgment but a calculation of what is best.

11.4.2.2. Alternatives to economic rationality for organizations

Winter (1971) distinguishes two main streams of criticism of the orthodox theory of the profit-maximizing firm. The first is the managerialists, who say that the largeness and diffuseness of typical ownership of firms give the managers of the firm enough discretion that profit is hardly the only motive. They may focus instead on growth of total revenue, and their personal motives can conflict with maximization of total profit (Winter 1964:26).

Jauch and Kraft (1986:778) advance an argument against the view that uncertainty is a threat to organizational equilibrium²⁰ and therefore is to

²⁰ March, J. G., & H. A. Simon 1958. *Organizations*, Wiley, New York; Burns, T., & G. M. Stalker 1961. *The Management of Innovation*, Tavistock, London; Chandler, A. 1962. *Strategy and Structure*, MIT Press, Cambridge, MA; Cyert, R. M., & J. G. March 1963. A

be reduced. This view is based on the more fundamental assumption that equilibrium is, itself, desirable. Jauch and Kraft point out that specific performance measures may be more useful than equilibrium as a goal. They suggest in particular “*achievement of organizational effectiveness through the accumulation of resources*” as a alternative goal. They note that “accumulation of resources ... create[s] disequilibrium and uncertainty. ... Seeking equilibrium, in contrast, may not serve capitalist values” (1986:782).

11.4.2.3. Satisficing

Winter’s second stream of criticism comprises the behaviorists, who say that the specifics of decision processes within the firm are important enough, and conflict enough with the assumptions of profit maximization, that the orthodox model is fatally flawed. He asserts that most observers would agree that short-term decisions are typically made on the basis of established guidelines, not through explicit optimization procedures. Winter proposes a behavioral theory incorporating that assumption: “Firms *satisfice* with respect to decision rules. That is, if existing rules are functioning well, the firm is unlikely to change them; if not, search for better rules will be stimulated” (Winter 1971:245).

Mintzberg, Raisinghani and Théorêt propose a satisficing scheme based partly on Soelberg: On some criteria, the decision-maker seeks merely satisfactory performance. On others, usually one, never more than three in Soelberg’s view, he seeks to get as much as possible. In screening, the

Behavioral Theory of the Firm, Prentice-Hall, Englewood Cliffs, NJ; Emery, F. E., & E. L. Tryst 1965. “The Causal Texture of Organizational Environments”, *Human Relations*, 18, 21-32

secondary constraints are used to reject alternatives. The alternatives that remain are then rated as acceptable, unacceptable, or marginal in terms of the primary goal's dimensions. The acceptable ones enter into an "active roster" where they are later compared with each other, unless an "outstanding" alternative is found, in which case search is terminated. In making this comparison, the decision-maker prefers a dominant alternative, one that is best along all the primary goal dimensions. If none can be found, he uses crude internal scales such as "significantly better" and "a little better," to compare alternatives (Mintzberg et al. 1976:259).

Dror offers an illustration of a very common occasion for satisficing (1988b:155). If a decision-maker finds that none of the available options is adequate, he is very unlikely to select by rote the best of a bad lot of options no matter how poor the prospects. Instead, a search for at least one "good enough" option will be undertaken, and only if that effort proves unsuccessful will an optimizing (least-bad) decision be made.

II.4.2.4. Ultra-rationality

Dror (1987) argues that planning requires a conceptual framework that goes beyond rationality to encompass *ultra-rationality*. It must include self-binding to prevent foreseeable failures of will; behaviors that are not part of rationality but are necessary for planning, such as creativity; the ability to explain and utilize paradoxes and varieties of irrationality; an understanding of irrationality in human behavior and the implications it has for planning; ways to resolve apparent conflicts between individual and collective rationality, and deliberate examination of value judgments and attempts to improve value judging capabilities. A major element of

ultra-rationality is “debugging”, which Dror defines as “striving for reduction of decision errors” (Dror 1987:32). He notes that while debugging requires some sense of what “nonerroneous” decision-making is, it can make a significant contribution before a mature, fully developed theory of ultra-rationality has been developed. Three specific problems to be tackled are incompetence, perversions of policy-reasoning, and motivated irrationality (Dror 1988b:162). Incompetence refers to the lack of or poor use of skills that are actually available. Perversions of policy-reasoning include known human biases, group-think, and the bias “rulers” face because their status leads to receiving overly positive feedback. Motivated irrationality subsumes errors due to strong emotions and ideologies, such as denial of reality as a way to reduce stress that makes future mistakes more likely, and errors made due to the “hot nature of policymaking.”

Another component of ultra-rationality is making sure that the right problem or decision is being addressed. When an organization faces a situation of a continuously declining curve, for example, making small, incremental decisions in the best possible manner is not ultra-rational: a radical move of some sort is required, and the futility of making even the best of the available small steps must be recognized (Dror 1991:9-11). In the terminology of optimization, this is pursuing a local optimum that is far inferior to the global optimum.

11.4.2.5. Rationality as news processing

Stinchcombe argues that “the core of organizational rationality is efficiency in news processing, not efficiency in maximization” (1990:348). To him, “rationality for an organization consists mainly of reacting

accurately to recent uncorrupted news” (1990:349). He suggests that this view makes sense when dealing with organizations but not with individuals. The complexities of individual lives and the scarcity of time and resources prevent us from shaping our news processes around uncertainties. Organizations are often able to compartmentalize their news problems and develop specialized systems.

Stinchcombe asserts that people are generally held to a higher standard of rationality when working for organizations than in their personal lives, and that organizations come closer to rationality (1990:347).

If one imagines that what people and organizations are doing to be rational is solving a maximization problem with given data, there is no reason to believe that an organization should do this job better than an individual. If one imagines instead that rationality is efficient processing of news about uncertainties, it is obvious that organizations are likely to do it better than individuals, and that individuals are likely to do it better in their organizational roles than in their lives. Thus, the idea that organizations are more rational than individuals is intimately tied up with the conception of rationality as processing and reacting to news about uncertainties. (Stinchcombe 1990:350)

11.4.3. Organizational decisions: the myth of the “decision event”

It can be natural to think of decisions as acts or events, but organizational decisions are often lengthy processes, “‘stretched’ over time rather than being ‘made’ at some definite point” (Smithson 1988:239). For example, Kunreuther et al. (1983) reviewed four instances of decisions about where to site liquified energy gas facilities, and found that these decisions were not events at all, but processes taking between 3 and 7 years. During each of the processes, the sets of facts thought to be relevant changed and the values brought to bear on the decision shifted as the

many interested parties battled to shape the outcome. At almost no time could the relevant components of the SDM be identified unambiguously (sets of completely feasible options, complete sets of possible states of the world, clearly elaborated utility functions, etc.).

Two reasons for the longevity of the myth that decisions are events can be found in the T and P Perspectives. The T Perspective on uncertainty and on decision-making employs for its central metaphor a decision abstracted from time (see the section on assessment of MEU in the Implications chapter). Not only is the decision itself condensed to an instant, the totality of the consequences are generally considered explicitly to occur in a second, later instant (von Neumann and Morgenstern 1953:19).

Additional support comes from our personal lives. Most people have had one or both of the following experiences. (1) In a specific act of will, apparently occurring instantaneously, we “decide” to do something, and are consciously aware of the moment in which we make the commitment. (2) We realize suddenly that what had been a puzzling decision has been resolved, and we now know what to do. An instant earlier we did not, and it may appear to us that an instantaneous decision had been made. In both of these cases, however, we are likely to be misled by our emphasis on conscious mental processes. It is very likely that unconscious processing took place over an extended period of time, and what was instantaneous was in the first example a conscious ratification of the unconscious decision, and in the second, our becoming conscious of the results (including ratification) of the unconscious processing.

Mintzberg, Raisinghani and Théorêt examined 25 cases of strategic decision processes. Their general assessment agreed with a characterization provided earlier by Witte²¹: it makes sense to define separate phases of decision-making processes, but not to expect them to occur in a stable sequential order (Mintzberg et al. 1976:252). Mintzberg et al. differentiated three main phases: (1) identification; (2) development; and (3) selection.

The identification phase comprises two routines ...: *decision recognition*, in which opportunities, problems, and crises are recognized and evoke decisional activity, and *diagnosis*, in which management seeks to comprehend the evoking stimuli and determine cause–effect relationships for the decision situation. ... Development may be described in terms of two basic routines, search and design. Search is evoked to find ready-made solutions; design is used to develop custom-made solutions or to modify ready-made ones. This distinction is fundamental—the difference between what psychologists call convergent and divergent thinking. ... The normative literature²² describes the selection phase in terms of three sequential routines: determination of criteria for choice, evaluation of the consequences of alternatives in terms of the criteria, and the making of a choice. In reality, selection seldom allows a neat delineation of these three routines, and our study suggests that it is more appropriate to describe it in terms of screen, evaluation-choice, and authorization. (Mintzberg et al. 1976:252-57)

The evaluation-choice routine of Mintzberg et al. corresponds closely to the T Perspective view of the entire decision. Mintzberg et al. distinguish three bases for making the choice — judgment, bargaining, and analysis (Mintzberg et al. 1976:258). Of these, judgment is fastest and easiest, and

²¹ Witte, Eberhard 1972. “Field research on complex decision-making processes—the phase theorem”, *International Studies of Management and Organization*, 156-182

²² Lovell: the “normative literature” generally employs the T Perspective; Mintzberg et al. are describing the O Perspective they developed.

most preferred, especially in strategic decisions. The importance of groups in addition to the decision-making group, either outside or inside the organization, made bargaining more likely. Analytical processes, which receive virtually exclusive attention within technical discussions, were observed surprisingly rarely.

A more pragmatic rendering ... sees the analyst presenting his factual analyses of the consequences of various alternatives to the manager who determines the value trade-offs in his head and thereby makes a choice. (Mintzberg et al. 1976:258)

Identification and development, which do not figure in the technical model, so dominated the decision process examined in this study that the entirety of the selection routines "often appear to be just a trimming on the overall decision process" (1976:274).

The hypothesis with the strongest support in our study is that the organization designs only one fully-developed custom-made solution.... only one decision tree was followed to its ultimate conclusion. (Mintzberg et al. 1976:256)

Whereas in the SDM of the T Perspective, the alternatives are assumed to be given before the analysis is begun, an organization approaches a decision problem or opportunity with no alternatives given. The options must be created, and since their elaboration consumes resources, it is performed selectively. While it would be possible to craft a recursive SDM with miniature optimization problems for deciding how far to develop each alternative, given the information available at each point in time, the conceptual clarity of the SDM, one of its greatest strengths, would disappear, and the problem of creation of mini-alternatives would remain.

11.4.3.1. Overemphasis on selection

The organizational literature is not as preoccupied with the selection of

alternatives as the technical literature:

We live in a culture that focuses on decisions and decision makers. Our status ladder's top rung is for the executive; our highest rewards are for those who make our choices. In emphasizing the opting, we too frequently ignore the options. We relegate to a subsidiary role the generating of the alternatives among which the choice must lie and the testing of whether the candidate options adequately exhaust the possibilities open to us or do justice to our creative powers. In emphasizing the optors, we also too frequently ignore the criteria by which the choices are made. (Wilson and Wilson 1970:203)

Another factor that decreases the salience of the selection phase in organizations is the distribution of information over time:

In general, the further a commitment to a line of action extends into the future, the larger the proportion of all the information entering into the final decision will be that is news coming in over the course of implementation of the decision. (Stinchcombe 1990:4)

More attention should be given to option generation, an "extra-rational" process, because "no screening can arrive at a better option than the best of those considered" (Dror 1987:40).

II.4.4. Decision-making and planning

Decision-making is characterized in the technical literature as a one-time selection of the available action whose consequences are most desirable. Aside from the moment of selection, only a single instant, in which all the consequences are deemed to occur, is considered (von Neumann and Morgenstern 1953:19). There is no reconsideration or course-correction. As Hey puts it,

[An] essential feature of von Neumann-Morgenstern utility theory [is] that it conceives of choice as being a once-and-for-all affair. (Hey 1979:30)

Considering only one dimension of difference, that of time, planning could

be described as making a decision, or series of decisions, so as to be in the best position for future decision-making. This enlarges the time scale considerably, bringing in more than the two instants implicit in the T model of decision-making (selection and the subsequent occurrence of consequences). This view of planning explicitly includes cognizance of further opportunities for action that will only arise in the future, after the choice is made. In Dror's view, the requirement in planning for not only the synchronic coherence of a single decision but also for diachronic coherence²³ "constitutes its single most important unique characteristic" (Dror 1987:33).

11.4.4.1. A "standard" approach to planning

Ascher and Overholt (1983) describe a "standard" approach to strategic planning: identify groups of *interests*, a fixed *environment*, and then develop a *strategy* for achieving objectives. But in long-range planning, uncertainty about these components intrudes. Their recommendation for stating interests is to do so at a level of generality sufficient to ensure that they will remain applicable throughout the planning period. This can be seen to be an application of Zadeh's Incompatibility Principle: the relevance of the interest statement is obtained by trading off details. This is possible, they explain, because we know enough about our interests to be able to generalize usefully. Typically, however, generality will not work for uncertainty in the environment, since we know too little about the details of the future environment (Ascher and Overholt 1983:24). While trend

²³ A concern for the reasonableness of the relationships of opportunities, costs and benefits as they occur across time.

forecasting may be of some use, “most important planning issues concern an environment that contains a very large irreducible minimum of uncertainty” (1983:26).

11.4.4.2. Strategic planning

Planning has a dual nature, preparing on one hand to adapt to uncertain futures and attempting on the other hand to shape those futures (Dror 1987:38). By contrast, most decision-making perspectives, especially technical ones, epitomized by the SDM, assume that the future environment is fixed, not to be affected by the decision-maker. Dror elevates this to a main goal of what he calls “grand policy analysis”:

To put the matter in its most ambitious form: grand policy analysis deals with policies aiming at significant impact on the future. (Dror 1991:10)

It is possible to extrapolate from decision-making through planning to something “beyond”, where the purpose of planning is conceived to be to get into the best position for future planning. Perhaps this is a definition for *strategic* planning: it is not merely a longer time scale than other planning (tactical, or operational, etc.) but a change in level of abstraction, so that the objective is defined not in terms of future consequences but in terms of future planning opportunities and resources. One way to do that is to alter the environment; another is to alter the organization.

At Royal Dutch/Shell, for instance, the key to an organization’s ability to respond effectively to its environment is thought to be how well its senior managers can apprehend events in its environment and devise appropriate actions (De Geus 1988:70). Therefore, De Geus asserts, the goal of planning is learning:

The real purpose of effective planning is not to make plans but to change the microcosm, the mental models that these decision makers carry in their heads. (De Geus 1988:71)

All of us carry mental models around in our heads, and it is these images of reality, not reality itself, to which we react — that is why improving the mental models is so important. A natural inference to draw from De Geus's perspective is that when an organization is more uncertain about its environment (due to complexity or volatility, for instance), the ability of the organization as an institution to learn will be more important. When there is less need for news, information gathering and learning abilities will be less important. De Geus reports that when Royal Dutch/Shell could control the major aspects of an industry, for instance, when it is highly integrated and the environment is stable, they are able to focus on optimizing instead of learning (1988:71).

Ascher and Overholt add that “the central role of the forecaster is to clarify the policy maker's thought about the future environment” (1983:27). Dror believes uncertainty increases the importance of this goal:

Accelerated and improved learning becomes a main response to hard uncertainties²⁴, with more emphasis on policy pilot testing, constant monitoring of ongoing decision implementation and improvement of underlying models. (Dror 1988:260)

For planning to be effective in promoting learning, planners must take the “radical” step of openly acknowledging and sharing their own uncertainties (Michael 1973:124).

²⁴ Lovell: “hard” = not reducible; see Taxonomy later in this chapter or in Chapter III: Taxonomy.

11.4.5. Probability in organizational planning

With the longer time frame that is typical of planning and other organizational pursuits, probability is seen in a different light. The T Perspective focuses on a single decision, and this makes it appear that the important aspects of the future state of the world can be captured in a set of discrete possible futures. The O Perspective's longer view makes it more difficult to be satisfied that all the possibilities are, or can be, known. Dror calls this uncertainty over the very membership of the set of possible futures "qualitative uncertainty", distinguishing it from the "quantitative uncertainty" represented by a particular probability estimate for one possible future (Dror 1986b). A non-distributional uncertainty is needed for situations of qualitative uncertainty; probability is fundamentally distributional²⁵. Possibility, described in the earlier T Perspective section of this chapter, holds some promise, not much realized yet, in this area. Decision analysts often support the use of subjective probabilities as the only cogent way to make use of information (such as intuitions, judgments, etc.) about what kinds of futures are more likely. Planning theorists, however, concerned about possible futures that have not even been thought of yet, often worry that subjective probabilities will be more deceptive than helpful:

[A] pervasive means for reducing uncertainty is to translate a situation of uncertainty into one of risk; that is, to act as if one knew what was *really* involved well enough to assign probabilities to outcomes. This gratuitous assignment of probabilities to situations

²⁵ The probabilities of a set of exhaustive and mutually exclusive events must sum to 1.0; i.e., there is fixed "quantity" of probability equaling 1.0 that must be distributed over the set. Possibility values for the same events, on the other hand, can sum to more than 1.0.

that do not justify them is, in the United States, an endemic means for self-delusion and deluding others. (Michael 1973:119)

But subjective probabilities in all their forms and uses must be rejected for most decisions, and consequently the modes of using all other methods and their relative importance change and additional methods must be developed. (Dror 1988a:259)

Some environments appear to be “surprise-prone”, which makes the hazards of using subjective probability even greater (Dror 1991:4).

But organizations strive to develop standard operating procedures (Lorraine 1991:6) and this striving opposes the views just recounted. The mechanisms of MEU are a very attractive set of candidates for standardizing on, seeming as they do to help to “ungamble” decisions. Theoreticians in the planning field face a serious challenge in arguing against over-reliance on standardized procedures such as institutionalized use of MEU methods.

Ascher and Overholt (1983:267) offer three reasons why forecasting groups may find it preferable to use verbal rather than quantitative representations of probability. Many trends are qualitative and difficult to quantify. Studies have shown consistently that people are overconfident in estimating probabilities, with more results falling outside the range than estimated. The third consideration is that a small difference in the estimate may matter greatly to some subgroups, leading to intense internal struggles, and these could perhaps be avoided by use of non-quantitative indicators.

11.4.6. Organizational structure as a function of uncertainty

Arthur Stinchcombe (1990) argues that a primary influence of

uncertainty within organizations lies outside the metaphor of decision. In his view, the kinds of uncertainty that are important for an organization to deal with are determined by the sources of important news, and these shape the structure of the organization. His conception relies on rationality in a significantly different way than does the T Perspective:

We do not connect the postulate of rationality to decisions (a mass of ethnography of decisions shows that they are not very rational), but to social structures that gather information and, more or less routinely, make decisions. The postulate of rationality, then, is much more a *guiding principle of growth of the structure* than it is a notion that a single utility is everywhere maximized in each decision. (Stinchcombe 1990:242)

In a similar vein, Crozier claims that the particular expertise of those promoted to top management positions (e.g. finance, production, law) often identifies the sorts of uncertainty that are currently the most problematic for the organization (1963:164).

Three major structural ways organizations respond to difficulties caused by uncertainty are to decentralize or divisionalize, to replace markets with hierarchies, and to increase the skill level of its employees in order to deal better with non-routine problems (Stinchcombe 1990:29).

Stinchcombe suggests that the basic need for divisionalization²⁶ arises when different markets have uncertainty that ramifies back into manufacturing and engineering (1990:117). When the uncertainty in the different markets is sufficiently different, it requires developing separate information systems. This makes a divisional structure more efficient than

²⁶ Divisionalization is dividing up some of the functions of the central office and creating smaller clones for each division. The central office functions more like a holding company than a command center.

one with information systems centralized at headquarters.

The ramification of uncertainty back into engineering and manufacturing is especially acute when a firm introduces innovations (Stinchcombe 1990:155). Major market uncertainty affects engineering and manufacturing, and they need systems to deliver information as quickly as possible. Consumer problems need to be worked out early in the production runs. The uncertainties about engineering, manufacturing and marketing are more interdependent than usual, and information systems organized around those three functions will not respond fast enough. A mini-division structure facilitates moving information as rapidly as is needed.

Although economists consider markets to be generally more efficient than hierarchies, Stinchcombe suggests that uncertainty can give the advantage to hierarchies (e.g. direct employment or ownership instead of a contracted relationship). He notes that when key features of performance that an organization needs are uncertain (e.g. costs, final performance criteria, quantities, timing), or “when many adjustments will have to be made during the course of contract performance, the transaction costs of negotiating and enforcing contract[s] rise” (1990:194). Changes made within a hierarchical structure are simpler and incur lower transaction costs. Stinchcombe notes that another accommodation to this condition is the cost-plus contract. Frequent changes can be expected when it is hard to define performance requirements ahead of time, when costs and quantities are unknown or likely to vary, or when the performance of several parties is interdependent and their performance is difficult to measure.

Contracts with individual employees are particularly prone to be

uncertain, since it is difficult to predict the quality of an employee's contribution and also difficult to predict for long periods of time the task the organization will want to assign the worker. Organizations can increase their flexibility by hiring employees hierarchically and obtaining greater freedom to shift each worker to the most important task (Stinchcombe 1990:200). Under conditions of change and uncertainty, hierarchies (including vertical integration) can be more efficient than markets. On the other hand, if it is not the type of work required but the total volume of work that is uncertain, contracting may provide an organization with more freedom to adapt to changing conditions.

The third structural response to uncertainty is elicited by non-routine problems. A primary way for organizations to reduce uncertainty is to rely on creating standard operating procedures, but these don't work well on non-routine problems (Lorraine 1991:6).

A nonroutine problem exists when there is uncontrolled or unidentified variation in the tools and materials or in the processes used in the solution. Problem solvers cannot describe in detail the sequential steps required in the solution, and they are unable to determine the consequences of action. (Lorraine 1991:3)

Non-routine problems cannot be handled as easily with fixed structures and low-skilled employees. In such situations, organizations often hire skilled employees and professionals, who are more able to tackle new problems and develop solutions instead of relying on previously designed solutions (Stinchcombe 1990:66). Michael sees a risk of this strategy in that each profession has its own lenses for viewing the world, and these will obscure as well as highlight various features of the environment. (1973:120). A decision-maker who relies on experts may become less aware

of the relevant uncertainties.

Smithson observes that the value of diversification, long seen as a way of reducing the impact of uncertainty, is being reassessed (1988:247-8). Diversification of investments has been thought to be an effective hedge against risk, and Lawrence and Lorsch (1967) reported a positive correlation between the differentiation or diversification of organizations and the uncertainty of their environments. However, Smithson reports that recent studies (e.g. Bowman 1980; Fiegenbaum and Thomas 1988) have not verified the earlier results.

11.4.7. The environment of the organization

A seminal work on the influence and nature of environments in which organizations exist is Emery and Trist's "The Causal Texture of Organizational Environments" (Emery and Trist 1965). In their research on organizational change they encountered situations in which the environment appeared to be undergoing change. Much of the existing literature on organizations assumes a stable environment (for example, many economic theories assume that the environment is at equilibrium). They developed the following typology of environments comprising "four 'ideal types', approximations to which exist simultaneously in the 'real world' of most organizations, though the weighting varies enormously" (Emery and Trist 1965:30).

- a. [Type 1] In the simplest type, goals and noxiants are relatively unchanging in themselves and randomly distributed. This may be called the placid, randomized environment. A critical property is that there is no difference between tactics and strategy, and organizations can exist adaptively as single, and indeed quite small, units.

b. [Type 2] The next type is also static, but goals and noxians are not randomly distributed; they hang together in certain ways. This may be called the placid, clustered environment. Now the need arises for strategy as distinct from tactics. Under these conditions organizations grow in size, becoming multiple and tending towards centralized control and coordination.

c. [Type 3] The third type is dynamic rather than static. We call it the disturbed, reactive environment. It consists of a clustered environment in which there is more than one system of the same kind, i.e. the objects of one organization are the same as, or relevant to, others like it. Such competitors seek to improve their own chances by hindering each other, each knowing the others are playing the same game. Between strategy and tactics there emerges an intermediate type of organizational response — what military theorists refer to as operations. On the other hand, stability may require a certain coming-to-terms between competitors.

d. [Type 4] The fourth type is dynamic in a second respect, the dynamic properties arising not simply from the interaction of identifiable component systems but from the field itself (the 'ground'). We call these environments turbulent fields. The turbulence results from the complexity and multiple character of the causal interconnexions. Individual organizations, however large, cannot adapt successfully simply through their direct interactions. An examination is made of the enhanced importance of values, regarded as a basic response to persisting areas of relevant uncertainty, as providing a control mechanism, when commonly held by all members of a field. This raises the question of organizational forms based on the characteristics of a matrix. (Emery and Trist 1965:30-31)

Emery and Trist believe that the environment of modern organizations possesses "a greater degree of system-connectedness" due to in large part to three trends. The first is that some organizations have grown large enough as they meet the challenges of Type 3 environments that they can cause new forces or processes within the environment to materialize. The second is an increasing interdependence between economic and other aspects of society. The third is that research and development efforts as a

whole are now inducing a continual change gradient in society. These conditions lead to increasing prominence of turbulent fields, a type of environment Emery and Trist had not found previously described in the literature.

For organizations, these trends mean a gross increase in their area of *relevant uncertainty*. The consequences which flow from their actions lead off in ways that become increasingly unpredictable: they do not necessarily fall off with distance, but may at any point be amplified beyond all expectation; similarly, lines of action that are strongly pursued may find themselves attenuated by emergent field forces. (Emery and Trist 1965:26)

In the terms of the T Perspective discussion of rationality, the Type 1 environment, with random features and little regularity, is one in which rational prediction is of limited use, but characterization of the random features through probabilities may be in determining cost-effective responses. Type 2, placid but clustered, is appropriate for parametric rationality. Type 3, with the prominence of similar, competing organizations in a context of regularity, demands strategic rationality. Type 4 environments need more sophisticated processes that include strategic rationality, such as the “ultrarationality” Dror calls for. These environments are characterized by the fact that at any given moment, there is regularity that is useful to know about, but the aspects of that regularity change within relevant time spans. In game theoretic terms, Type 3 generally corresponds to zero-sum games of pure competition and to positive-sum games of mixed cooperation and competition (i.e. interdependence), while Type 4 may involve positive-sum games, but the possibility that the rules of the game may change diminishes the power of game theory for this type of environment.

Emery and Trist suggest that in Type 4 environments, organizations are unable to ensure their stability through their own actions alone. The emergence of “*values that have overriding significance for all members of the field*” is a powerful means for dealing with “persisting areas of relevant uncertainty” (1965:28). When it is impossible to “trace out the consequences of their actions as these are amplified and resonated through their extended social fields,” organizations can develop shared values, which have the effect of simplifying the turbulent environment and reducing the intensity of the dynamics.

In contrast to Stinchcombe, who suggests that hierarchies have advantages over contractual relationships under high uncertainty, Emery and Trist suggest that in turbulent conditions organizations can enhance their chance of survival by increasing the redundancy of their response capabilities by shifting toward non-hierarchical, multiple-participation structures.

11.4.8. Uncertainty and power

Crozier uses *power* as the main lens in his examination of French bureaucratic organizational phenomena (Crozier 1964). He adopts Dahl’s definition of power:

The power of a person A over a person B is the ability of A to obtain that B do something he would not have done otherwise.²⁷

Freedom of action that person A has constitutes uncertainty on the part of person B about A’s actions. When A’s actions are important to B, B is motivated to attempt to influence A to produce more desirable behavior:

²⁷ Dahl, Robert 1957. “The Concept of Power”, *Behavioral Science*, II, July, pp. 201-15

this gives A power over B.

Most work in the bureaucracies Crozier studied was highly rationalized (routinized); two exceptions illustrate how uncertainty creates opportunities for power. On the highly routinized work floor, work stoppages are a major source of uncertainty. They require the specialized, non-routine skill of the maintenance workers, are unpredictable in duration, and clearly affect production volume. The only people who can solve them, the maintenance workers, therefore have considerable power, and enjoy a higher status than the other workers.

The second illustration is that of the chief engineer, who is also working in an area of low certainty, and less rationalization of their work. The successful chief engineers do not strive to create regularity and routine in their work, because they could then be replaced by a less capable person. They have no incentive to devise rules to govern their work, as that would diminish the source of their organizational power. Instead, their innovation is restricted to relatively unimportant matters (Crozier 1964:131).

Crozier points out that if the organization were perfectly rationalized, each person would have but one best way to do things, and everyone would know this. Each person would be predictable, and able to predict others. There would be no opportunity to exercise power, no bargaining (Crozier 1964:156). But this ideal is unobtainable:

Strategical analysis makes it possible to discover that power can neither be suppressed nor ignored. It stems from the impossibility of eliminating uncertainty in the context of bounded rationality which is ours. (Crozier 1964:158)

Management would like to rationalize the work, so that workers'

activities are governed by more precise rules and are therefore more predictable. This reduces the skill needed by the workers, and reduces their bargaining power. Supervisors, however, would like to retain some discretion in the interpretation and enforcement of the rules so as to give them more power to obtain increases in productivity from their subordinates, for example, by overlooking minor transgressions. The subordinates, therefore, oppose rules governing their behavior, but are in favor of rules that govern their supervisor's behavior. They would prefer, for example, clear-cut policies on hiring and promotion practices, such as the use of seniority or standard tests (Crozier 1964:161).

Those who hold the greatest formal power in an organization, the top managers, often reflect in their backgrounds as financial specialists, production experts, or budget analysts the kind of uncertainty, and thus the kind of expertise, that has been most important for the organization in the recent past (Crozier 1964:164). Experts in general hold a fragile kind of power: any success they have in understanding and rationalizing their field reduces their own power. Crozier disputes the common belief that "complexities of our technical age" will give rise to a new ruling class comprising technocrats. He asserts that the accelerated pace of change makes it impossible for the expert to fight the rationalization of his own field and the resulting diminution of his power (Crozier 1964:64-65). (Here he may be confusing individuals and groups: the rationalizing process reduces the power of all non-experts, thereby increasing the power of the whole class of experts. It is just that each expert finds that the transformation of her expertise into rules makes it difficult for her to

remain a member of that class. The power of the class of experts may increase while membership in the class becomes less stable.)

II.4.9. Organizational strategies for coping with uncertainty

Strategies previously mentioned

- Reduce internal uncertainty by rationalizing; create more rules and standard procedures (Crozier 1963:156; Lorraine 1991:6).
- Decentralize or divisionalize (Stinchcombe 1990:65).
- Diversify (Smithson 1988:247).
- Develop better “news” systems (Stinchcombe 1990:224).
- Use cost-plus contracts (Stinchcombe 1990:194).
- Increase the skill level of employees when routinization does not work (Stinchcombe 1990:66).

Additional strategies:

- Avoid uncertainty — choose a stable environment (Stinchcombe 1990:65).
- Buffer some parts of the organization from environmental uncertainty so that they can be routinized. (Stinchcombe 1990:65)
- Convert uncertainty into risk by assigning probabilities:

Often the need to reduce uncertainty is so great that the person justifies the assignment of probabilities by that most tenuous and attractive of rationalizations—“It’s the best we can do,” or “It’s better than nothing.” ... One of the most serious consequences of this behavior is that the resulting emotional relief encourages all involved to believe they really do know what is important, what is going on. (Michael 1973:119)

- Convert risk into certainty — gambling casinos and insurance companies are able to pool large numbers of situations, each of which individually is risky, into aggregate situations of statistical certainty.

Uncertainty is not present when the consequences of the action are actuarially certain. (Shackle 1957:107)

- Anticipation and resilience — Wildavsky (1985) suggests that anticipation and resilience are fundamentally different approaches to coping with an uncertain future. These correspond to the cybernetic concepts of feedforward and feedback, respectively (e.g. Robertshaw et al. 1978:24). He maintains that anticipation can work well when the environment is stable, but that when it is less stable, and especially when it is turbulent, resilience provides much greater security. He avers that people are more likely to choose anticipatory approaches when they perceive that risks are high or when they fear losing something they value. Anticipatory strategies involve attempts to reduce or eliminate uncertainty; planning based on resilience requires thorough consideration of all uncertainties, to the extent of ensuring common language for comparing different people's subjective perceptions (Smithson 1988:292).

- "Tightening up"

Another means for attempting to reduce uncertainty is by various managerial and administrative devices aimed at "tightening-up" the org. The ends sought in this way include increasing internal incisive action, controlling what goes in and out of the organization that could rebound as more turbulence, and reemphasizing the role and ability of management to reduce uncertainty by resort to the traditional tightening-up ritual. (Michael 1973:120)

- Impose controls to reduce managerial anxiety.

Parameters [may be] controlled not because they are relevant but because they are measurable. Their function is to create an illusion of certainty as a means of coping with intolerable uncertainty. (Michael 1973:120)

Quantification can also serve this function of reducing *apparent*

uncertainty.

- Openly acknowledge uncertainty. Michael terms this the most radical and yet most necessary approach (1973:120). This offers the possibility of sharing the discomfort of uncertainty, encouraging people to face it, and from that acknowledgment, to attempt to learn more about their situation and ultimately reduce their uncertainty. It can reduce overconfidence and over-commitment to single perspectives, and foster trust, openness, creativity and intuition.

In our culture, however, such acknowledgment within organizations is considered foolish. It runs counter to the need to assign blame. It highlights the limitations of rational techniques, which is threatening to both consumers and producers of rationalistic products. It demands personal skills that are not common. Linnerooth suggests three short-term organizational needs that militate against open acknowledgment of uncertainty: 1) the need for control to maintain authority; 2) the need to defend policy decisions using analyses that appear certain; and 3) the need for narrow, not broad analyses (Linnerooth 1984).

- Accept contradictions:

One of the main specifications for handling of deep uncertainty involves perceiving, accepting and processing contradictions. (Dror 1991:33)

The spatial and temporal heterogeneity of organizations permit contradictions to persist indefinitely. Contradictions may involve conflicting values, beliefs, known “facts”, instructions given by different managers, simultaneous but incompatible goals, goals at different times, or future goals and current realities. Many of the same factors that make

uncertainty difficult to acknowledge also work against explicitly recognizing contradictions.

II.4.10. The silver lining for “O”

As the T Perspective recognizes, it can be helpful to increase competitors' uncertainty. Organizational research and development plans and results are often closely guarded. Pharmaceutical firms sometimes patent unproductive research to increase their competitors' uncertainty about their actual progress and plans (Jauch and Kraft 1986:777).

Uncertainty can help to hide contradictions and help opposing parties reach agreements (Thompson and Warburton 1985). Politicians, for example, may obscure some of their own views so as to keep loyalty of groups with conflicting positions. Organizations may have very different preferences regarding uncertainty within themselves and uncertainty among heterogeneous organizations:

Organisations maintain themselves by maintaining cultural homogeneity; they generate *shared understandings* of the problems they face....

But, in inter-organisational decision making, such cultural homogeneity is, by definition, unattainable. If one cultural orientation were to succeed in establishing its hegemony over the rest then that would signal the creation of a single all-enveloping organisation ... Until that happens each organisation has its own definition of the problem: one that contradicts all the others and, as the decision process gets under way, is increasingly thrown into contention with them. These are the *contradictory certainties* and, like the organisations they support, they do not go away. (Thompson and Warburton 1985:4)

Thus, while organizations may attempt to reduce uncertainty on goals and definitions within themselves, they may need to actively create a systemic uncertainty that is large enough to mask the “contradictory certainties” in

their relationships with other organizations in order for the heterogeneous system to be able to persist. Resolution of the systemic uncertainty would require reducing the contradictions, which could only be done by sacrificing some of the perceptions, goals and perhaps even legitimacy of some of the institutions.

Jauch and Kraft (1986) point out that organizations may want to increase their own uncertainty along with that of others if they believe they have an information-gathering advantage over their competitors: actions that uniformly increase uncertainty may work to their advantage by emphasizing the importance of their information skills. An organization may want to increase its internal uncertainty in order to challenge its strategic assumptions and reduce the likelihood of its own *overconfidence*, or generate internal conflict in order to deal with problems such as incompatibilities between goals (Jauch and Kraft 1986:783). In order to increase its information-handling capacities and thereby reduce the impact of external uncertainty, an organization may adopt a more organic, less rule-governed internal information structure. By reducing reliance on rules, this may have the side effect of increasing its internal uncertainty (Jauch and Kraft 1986:783). Mintzberg et al. suggest that managers and organizations may seek situations of uncertainty when they offer the risky opportunities and when the organization is not already beset by too many crises (1976:254).

While uncertainty itself may pose problems for organizations, attempts to reduce it may beget worse problems. As Smithson (1988:248) points out, “not all adaptive strategies are risk-averse...organizations that reduce

feedback, induce conformity, and reduce internal competition begin to behave like closed systems and thereby make themselves even more vulnerable.” March and Olsen (1976) assert that actions designed to avoid uncertainty harmfully restrict the learning ability of organizations. Efforts to reduce internal uncertainty can lead to excessive conformity and the blindness of groupthink (Janis 1972).

II.5. Personal Perspectives on uncertainty

33. All important decisions must be made on the basis of insufficient data.

34. Yet we are responsible for everything we do.

42. What do you know...for sure...anyway?

Sheldon Kopp

II.5.1. Introduction to “P”

As soon as questions of will or decision or reason or choice of action arise, human science is at a loss.

Noam Chomsky

The great decisions of human life have as a rule far more to do with the instincts and other mysterious unconscious factors than with conscious will and well-meaning reasonableness.

Carl Gustave Jung

The intensity and immediacy of our personal perceptions of the specific demands of particular situations often overwhelm the cool appeal of general principles such as rationality. The experience of uncertainty often includes strong emotions such as excitement and fear. The ordinary anxiety accompanying everyday uncertainty can increase and turn into panic when the need to act is immediate and we don’t know what to do.

In the arena of personal perception and action it is much harder to

ignore details in order to focus on consistent rules than in the realm of organizations. The P Perspective differs yet more greatly from the abstract realm of the T Perspective, where generalization and the forgetting of details is not just common but imperative. Thus it is natural that the P Perspective would, by attending to a wider range of influences, leave a relatively smaller role for rationality and optimization (or alternatively, allow more diverse interpretations of “rational”).

Proponents of normative rational approaches would claim that Jung’s and Chomsky’s comments are descriptive, and only go to show how much we need powerful prescriptive methods to improve the deplorable condition of human decision-making. However perfectible human reason may be, my own reasoning, compelled to provide an answer by 4:00 some Friday afternoon, remains limited in ways I am only too aware of.

Nonetheless, many are willing, as is Jon Elster, to begin with the presumption that people generally behave rationally:

Not all human behaviour is rational or intentional, but there is a well-grounded *presumption* that this will typically be the case.
(Elster 1984:ix)

The question of just what behavior is rational is difficult indeed when examined on a personal basis. To describe another’s behavior as irrational is tantamount to admitting that you don’t understand it. Freud was able to understand behavior that baffled others, and saw it as at least making sense if not actually being rational. Phobias, for example, could probably be translated quite readily into the framework of rational choice theory, with suitable formulation of a set of values and a utility function. If an encounter with a snake is the most awful event a person could experience,

is it irrational to walk miles around a forest that others safely penetrate? Is the intense dislike of snakes irrational? Rational choice theory is mainly mute on the subject of which values to hold.

While American economists and psychologists have generally adopted a conception of rationality as a maximizing activity, French psychologists such as Allais tend to take “the view that rationality consists of choosing appropriate means to pursue reasonable ends,” according to Lopes (1988:413). Lopes views the American reoccupation as a result of the extraordinary impact of von Neumann and Morgenstern’s presentation of utility theory.

In some ways this is a frustrating perspective to work on. This work is largely technical, and it attempts to generalize its material so as to be more widely applicable, as much scientific effort does. Yet it is difficult to capture much of the essence of any, let alone many, of the Personal Perspectives within a technical, general framework. Little can be found in technical literatures that is relevant to personal perspectives. This problem stems in part from the necessarily general approach any technical study takes, which leaves out details precisely in order to be general. “God is in the details,” it is said, and so are personal perspectives. Personal perspectives can only really be made real and vivid in case studies which illuminate particular personal perspectives, and then the specificity limits their generality. What can be done here is to describe some of the ways that personal perspectives are important.

Two fields are especially relevant to P Perspectives: psychology, the study of the characteristics of individual human beings, and art, the

expression of often intensely personal perspectives. The relationship of uncertainty to art is at least two-fold: literature is an effective way to communicate attitudes and experiences of uncertainty to others, and uncertainty as a theme often enlivens writing and gives it suspense. Yet this connection is difficult to generalize about, and will appear here mainly in the silver lining. Most of this section can be called psychological.

The first two subsections will describe common human responses to and attitudes toward uncertainty, and experimental findings from psychologists and economists on the ways uncertainty affects human thinking and decision-making. Following next is a summary of research into maximization of expected utility (MEU) as a descriptive theory of rational decision-making, and another subsection briefly describing some alternative theories that attempt to improve on MEU. Lastly there is a subsection that discusses uncertainty's silver lining in the P Perspective, drawing in large part on literature.

II.5.2. Common personal responses to uncertainty

For the most part, people in our culture associate the term uncertainty with stress and complication, and they would prefer to avoid it.

The first point to be noted is that viewing their decisions as largely fuzzy gambles is anathema to most decisionmakers. (Dror 1988:262)

Examination of business decision making and governmental policy making suggests that, whenever possible, decision makers avoid uncertainty and the necessity of weighting and combining information or trading-off conflicting values. (Slovic 1971)

People want to feel certain when they take an action that it is better than the alternatives, that it will lead to the most favorable

consequences for motive satisfaction. (Mills 1965)

Our philosophical tradition, from the Greeks, through Descartes and Laplace, and on through the economics of the twentieth century, treats uncertainty as an impediment to rational ordering of human affairs. American culture in particular emphasizes the possibility and value of individual achievement and doing, and uncertainty about what to do can block the drive to act. We admire people who take risks successfully, recognizing that grappling with uncertain prospects is often frightening and requires bravery.

I shall argue that we have a strong reluctance to admit uncertainty and indeterminacy in human affairs. Rather than accept the limits of reason, we prefer the rituals of reason... Generally speaking, we tend to see uncertainty as an unmitigated ill. Uncertainty prevents us from planning for the future. Even more important, it prevents us from making choices that we can justify to ourselves and others as grounded in reason, whence the various tactics for uncertainty avoidance and uncertainty reduction. (Elster 1989a:37-38)

As Michael puts it, "the most direct way to avoid uncertainty is to avoid situations rich in uncertainty" (Michael 1973:118). A related psychological tactic is to deny the perception or significance of uncertainty. The medical malpractice situation indicates how unwilling we are to accept the fact that medical treatment, performed by fallible human beings, is not certain, and therefore that treatment failures are not necessarily the result of culpable error. Perhaps part of the reason that cancer seems to be more feared than cardiovascular problems, though it causes fewer deaths, is that it is more uncertain. It is less clear how to prevent it (with a few exceptions), the treatments are less certain, and finding out the success or failure of the treatment takes longer. It is often described in terms of

probability, which for many people symbolizes not just uncertainty but cold-blooded uncertainty.

Our discomfort in uncertainty keeps us from using random procedures for making decisions, such as child custody cases, even when those involved would all agree that the “right” decision is not knowable:

The use of lotteries is associated with uncertainty, indifference, indeterminacy and incommensurability. (Elster 1989a:38)

I want to suggest that visibly arbitrary chance is often repulsive. Even when we have no reason to decide one way or another, we would like the outcome to be determined by reasons. (Elster 1989a:55)

Yet attitudes towards uncertainty vary hugely. In some cultures, the drive to dominate nature is much weaker, and uncertainty about the consequences of one’s acts is less disturbing. Fatalistic cultures may emphasize accepting rather than affecting the future and therefore be unconcerned about uncertainty. But perhaps one of the motivations for maintaining a strong religious faith is to subdue one’s uncertainty about specifically what will happen with a belief that progress in the big picture is a certainty, or that heaven for the believers is a certainty. Individuals may attempt to decrease their uncertainty-based anxiety by adopting an ideological perspective that makes the situation appear more certain (Michael 1973:118).

The variation between individuals is equally great. Some find risk essential to their lives, others insist on it only in recreation. Keats thought attitude towards uncertainty very important:

At once it struck me what quality went to form a man of achievement, especially in literature, and which Shakespeare possessed so enormously — I mean *negative capability*, that is, when a man is capable of being in uncertainties, mysteries, doubts,

without any irritable reaching after fact and reason. (John Keats)

Responses to different types of uncertainty vary as much as different individuals' responses do. Monat, Averill and Lazarus (1972) investigated two different conditions of uncertainty about a potentially painful stimulus. They found that when the uncertainty was about when the threatened stimulus would occur, subjects' arousal decreased over time; when subjects were certain about the timing (whether or not the stimulus was sure to occur), their arousal increased. Janis and Mann (1977:228ff.) found that uncertainty about known consequences can be especially aversive, in some cases outweighing carefully thought out and apparently decisive lists of pros and cons. They assert that people often would, as Hamlet noted, "rather bear those ills we have than fly to others that we know not of." This type of uncertainty is not treated by MEU or the research inspired by it, or work on biases and heuristics: all of this work focuses on specific consequences that are described clearly and merely not certain to occur; the "unknown rub" Janis and Mann studied may loom larger because of its lack of description and the potentially enormous consequences that cannot be ruled out.

Uncertainty may be an unbearable burden:

There is no more miserable human being than one in whom nothing is habitual but indecision. (William James)

And people often want help bearing the burden of uncertainty:

[Planning,] intelligence collection and assessment are expected to reduce uncertainty and "ungamble" decision making, but this, too, is a chimera. Good intelligence in the face of hard uncertainty should explicate and map the gambling elements of decisions and not hide them, but this product is very unwelcome among decision makers. (Dror 1986b:30)

11.5.3. Judgment under uncertainty: cognitive strategies and heuristics

Tversky and Kahneman's famous 1974 article in *Science*, "Judgment Under Uncertainty", has stimulated a continuing interest in examining how human subjects go about making decisions under uncertainty. They found that people showed systematic biases in their responses that cast serious doubt on the ability of human beings to make rational decisions when information is limited — or at least on their ability to act in the way that economists had defined as being rational. Subsequent studies have often shown that experts in particular fields, and even experts in statistics and other disciplines focusing on reasoning and decision-making, are vulnerable to these biases. This makes it much harder to attribute the findings to the irrationality of naive subjects who could be trained to perform better.

Common cognitive strategies or heuristics and the biases that the heuristics cause are described next. The first three, representativeness, availability, and adjustment and anchoring, are from Tversky and Kahneman's 1974 article.

11.5.3.1. Representativeness

The representativeness heuristic is a cognitive rule that causes subjects' estimates of the likelihood that A comes from B to be influenced by how similar A is to B. This heuristic misconceives chance: when asked to judge how likely a series of 6 coin tosses is, people associate likely with random-appearing. HTHTHT looks more "random" than HHHTTT, and is judged to be more likely, though in fact all sequences of 6 coin tosses are equally likely. Similarly, HHHHHT appears to be un-random and is

judged to be less likely.

Representativeness also appears to override sample size. Subjects were asked how likely a man is to be an engineer or a lawyer when given a brief description. The man was one of a group of 100 comprising 30 lawyers and 70 engineers. When the description was neutral and was equally representative of stereotypical lawyers and engineers, subjects ignored the prior odds and said it was 50% likely to have been an engineer. When no such irrelevant data was presented, the prior odds were taken into account.

The representativeness of an event and its actual likelihood are especially likely to diverge under two conditions (Tversky and Kahneman 1982a). The first is when the judgment is based on questionable evidence. They cite as example the confidence job interviewers have in their ability to identify good prospects through the interview process. There are so many factors that cannot be assessed in the interview that may affect the interviewee's performance that interviews are demonstrably unreliable for this purpose. The second condition is when the event is highly specified. Making an event more specific does not reduce its representativeness, though it does reduce its likelihood.

11.5.3.2. Availability

Events are judged to be more likely to occur when subjects find it easier to think of an example of the event. Characteristics that can make an event easier to think of include (1) retrievability, (2) effectiveness of search set, and (3) imaginability. Subjects were read a list of male names and one of female names, equal in length. Afterwards, the list with more famous

names (thus easier to retrieve) was judged to be longer. The recency of an experience, and its vividness (e.g. a verbal report of a fire compared to a photo of it) also affect retrievability. When asked if there are more words that begin with the letter “r” or that have “r” as the third letter, subjects find it so much easier to search for words by the first letter that they report erroneously that words with “r” in the first position are more numerous.

When subjects were asked how many 2-person or 8-person committees could be formed from a pool of ten people, unsophisticated subjects’ answers averaged 70 and 20, respectively. The correct answer is that the number of combinations of 10 items taken eight at a time, or $10C_8$, is the same as the number of combinations of 10 items taken two at a time, or $10C_2$, which is $10!/(8!*2!) = 45$. Committees of 8 are difficult to imagine, and underestimated.

11.5.3.3. Adjustment and anchoring

Subjects’ estimates of quantities are often anchored to the first value that is suggested to them. They then make adjustments to that value, but typically the adjustments are insufficient. When subjects are asked to estimate the 10th and 90th percentile probability values²⁸, they typically provide much too small a range. Some studies have indicated that the range between subjects’ estimates of the 1st and 99th percentiles, which should capture 98% of the true values, misses the true value as often as

²⁸ For example, in estimating the length of a river, the 10th percentile value is one indicating that the subject believes there is only a 10% chance the river is shorter than that value; the 90th percentile is a value that the subject thinks has only a 10% chance of being exceeded.

30% of the time. This has been demonstrated for sophisticated as well as naive subjects, even when they are given incentives for accuracy. Subjects apparently anchor on the first point estimate they produce, whether it is a mid-point or an end-point, and then make inadequate adjustments to it in estimating other values.

Subjects anchor on the probability of a simple event, and then do not adjust adequately for compound events. The probability of a conjunctive series (logical *and*) is overestimated, and that of a disjunctive series (logical *or*) is underestimated. For instance, they would rather bet on drawing 7 white balls in a row (replacing the ball each time) from an urn containing 90 white balls and 10 red balls than bet on drawing a single white ball from an urn filled equally with red and white balls. The true odds of drawing 7 whites in a row, a conjunctive event, are $.9^7 = 47.8\%$, less than even odds. Similarly, they would rather bet on drawing a single white from the fair urn than bet on drawing at least one red in seven tries from the 90% – 10% urn. The true odds of the last event, a conjunctive one, are $1.0 - .9^7 = 52.2\%$. Since it is very common that a plan's success will require the success of each of a series of steps, this bias is dangerous for planning, as Tversky and Kahneman remark below.

11.5.3.4. Simulation

Tversky and Kahneman later elaborated on the availability heuristic (Kahneman and Tversky 1982a). They proposed that mental images become available through two means. One is recall from memory, in which retrievability and search strategy would matter. The other is construction, or *simulation*, as they name this heuristic, which typically involves

creation of a scenario. Scenarios that depend on a few dramatic and distinct events would be judged more likely than scenarios that depend on gradually changing circumstances. The former would be more distinct and vivid, and therefore more available.

A second bias of simulation is that scenarios that are based on a single, plausible happening are judged more probable than scenarios that could occur through many different but less likely routes.

Such a bias could have especially pernicious consequences in a planning context, because it produces overly optimistic estimates of the probability that the plan will succeed. By its very nature, a plan consists of a chain of plausible links. At any point in the chain, it is sensible to expect that events will unfold as planned. However, the cumulative probability of at least one fatal failure could be overwhelmingly high even when the probability of each individual cause of failure is negligible. Plans fail because of surprises, occasions on which the unexpected uphill change occurs. The simulation heuristic, which is biased in favor of downhill changes, is therefore associated with a risk of large and systematic errors. (Kahneman and Tversky 1982a)

11.5.3.5. Conservative information processing

Edwards (1982) characterizes human information processing as “conservative” in that people do not revise their original estimates of event probabilities as much as they ought to when they receive new information, according to “an abundance of research”. Bayes’s theorem provides the mathematical model for the normatively proper updating of prior probabilities.

A convenient first approximation to the data would say that it takes anywhere from two to five observations to do one observation’s worth of work in inducing a subject to change his opinions. (Edwards 1982)

Clearly this conservatism of incorporating new information can

reinforce the anchoring effect discussed earlier.

11.5.3.6. Hindsight

The hindsight heuristic is a consistent tendency to overstate what could have been known or predicted in a past situation.

[People] not only tend to view what has happened as having been inevitable, but also to view it as having appeared “relatively inevitable”. By encouraging us to exaggerate the extent of our knowledge, this bias can make us overconfident in our predictive ability. Perception of a surprise-free past may portend a surprising future. (Fischhoff 1982)

The hindsight bias is diametrically opposed to the *ex ante* perspective on judging the result of decisions that Elster and Dror argue for (cited in the earlier T Perspective section of this chapter), which preserves all of the uncertainty surrounding a decision at the time it was made. The hindsight bias discounts important particulars of a decision, with effects similar to those of the *fundamental error of attribution*.²⁹

11.5.3.7. Confirmation bias

Preliminary beliefs we have about a subject affect our process of searching for additional evidence.

When making predictions or forecasts, we often lead toward one perspective, and the natural tendency is to seek support for our initial view rather than to look for disconfirming evidence. Unfortunately, the more complex and uncertain a decision is, the easier it is to find one-sided support. (Russo and Schoemaker 1992:12)

This particular bias can be reduced by choosing to seek disconfirming

²⁹ Fundamental error of attribution: we attribute successes and failures of other people to their enduring characteristics, but attribute our own (especially failures) to specific circumstances (e.g., luck).

evidence. Even use of this counter-bias strategy leaves major weaknesses in search strategies, since it depends on the particular “initial views”. Evidence for or against a view that is not being considered may not be sought.

11.5.3.8. Overconfidence

When people are asked to assess the quality of the incomplete information they have on a matter, they have a strong tendency to overrate their knowledge (Fischhoff 1982).

A typical study might show probabilities of .75 to be associated with a ‘hit rate’ of only 60% and expressions of certainty ($p = 1.00$) being correct only 85% of the time. ... [I]t is not uncommon to find true answers falling outside of their 98% confidence intervals 20% to 40% of the time. (Fischhoff 1982)

Russo and Schoemaker (1992) refer to knowledge of a particular subject as *primary knowledge* to distinguish it from *metaknowledge* — knowledge about *the quality of the primary knowledge*. Some respondents to a survey of theirs claimed that since the questions dealt with material outside of the requirements of their job, the overconfidence the respondents displayed in their knowledge did not matter.

We strongly disagree. Whether you know a lot or a little about a subject, you are still responsible for knowing how much you don’t know. (Russo and Schoemaker 1992:9)

Many studies have reported similar results of widespread overconfidence, including those using professionals from disciplines where statistical or psychological sophistication characterized the subjects. It is interesting to note that *primary knowledge* of statistics or the psychological literature on heuristics and biases does not seem to automatically lead to better *metaknowledge*.

Judgments made by groups are vulnerable to *groupthink*, in which the perspectives of the members of the group “use their collective cognitive resources to develop rationalizations supporting shared illusions” (Janis and Mann 1977:129). Groupthink seriously exacerbates the problem of overconfidence. Russo and Schoemaker compared the overconfidence of a group to that of its members, and found the group to be significantly less overconfident (1992:15). Many of the strategies suggested below for counteracting overconfidence appear to be well-suited for groups, especially those that are able to maintain an open atmosphere that encourages assumption challenging.

11.5.3.9. Debiasing

Little research has been reported to date showing how well training can compensate for natural bias. Fischhoff (1982) reported his attempts at “debiasing” subjects to overcome the effects of hindsight and overconfidence. His results were that both biases were very robust and resistant to the efforts of the experimenters to reduce the strengths of the biases. Russo and Schoemaker (1992) report promising results using several strategies to reduce overconfidence. One technique asks subjects to think of one argument against their view before they rate their confidence in it. Another asks them to contribute new possibilities to a list of ways in which a plan might fail. Imagining alternative future scenarios has helped to heighten planners’ awareness of the variability of the political or economic environment and reduce overconfidence (Wack 1985a; 1985b). Weather forecasters as a group are extraordinarily well-calibrated (i.e. they are neither overconfident nor underconfident); they receive rapid

feedback and are compensated partly on the basis of their accuracy, and oil geologists working for Dutch Shell are also well-calibrated (Russo and Schoemaker 1992:10).

Future clarification of the potential for debiasing is important to normative decision theory. If observed biases are correctable, existing normative models (especially Bayes's theorem and the Expected Utility Hypothesis) can be retained; the discrepancies between prescribed and observed behavior will be reduced by improving the behavior. If the biases are not correctable, then the models must be modified: prescribing impossible behavior is no service.

II.5.3.10. Neurological bases for heuristics

Research in neurophysiology and brain/mind studies shed some light on how basic characteristics of human beings would be likely to affect decision-making. Ornstein identifies four major principles that appear to guide the human neurological system, from perceptual and sensory organs to the most complex centers in the brain (Ornstein 1986:27):

1. Extreme sensitivity to recent information; attention is focused on recent problems and then drawn away to the next problem;
2. Focus on changes in situations — “news” — rather than on stable or unchanging conditions; sharp changes are salient while gradual changes may be undetected;
3. Relativity of judgments, with important dependence on the reference to which something is compared;
4. Massive simplification with most incoming information discarded.

All of these have direct implications for how people will make decisions. The sensitivity to recent information calls into question the assumption

that people have a stable preference structure. Emotions are an especially volatile or transitive form of reaction, and often constitute important parts of the information used by people in making decisions. The bias towards recent information would emphasize such intense but short-lived “information”, probably contributing to the problem of *akrasia* (weakness of will).

Focusing attention on news could help account for part of the simulation heuristic that makes scenarios with vivid events appear more probably than ones that depend on gradual changes. The principle is elegantly phrased by Gregory Bateson: “Information is news of difference” (Bateson 1979:32).

The preponderant relativity of human judgment is not represented in the mechanics of MEU, which in most treatments evaluates only final outcomes, and postulates an absolute utility scale for all outcomes (for a given decision-maker). It is accommodated in the two alternative theories discussed below in section II.5.5. In prospect theory, outcomes are evaluated in terms of gains or losses relative to the current asset position. In regret theory, outcomes are evaluated through a comparison with other outcomes that could have occurred with the same state of the world but under a different decision. The importance of the framing of a decision (Tversky and Kahneman 1981) again illustrates the relativity of judgment.

The fourth principle, simplification of inputs, was proposed to be one of the editing processes in prospect theory. In general, cognitive limitations and the problems of handling non-simple information are often suggested as the explanation for observed “defects of rationality”.

These human characteristics must be accounted for in an adequate descriptive theory of decision-making. Their relevance for normative theories is less direct: prescriptions for human decision-making must be such that when beings constructed in such a fashion attempt to follow the prescriptions, the decision-making improves. Procedural characteristics, such as the utility evaluation effect, clearly bear directly on the normative suitability of approaches that depend on the procedures.

11.5.4. Decision-making under uncertainty: experimental evidence

In addition to the research reported above, much of which focuses on judgments of probability, there is by now a large volume of research on the decisions people make under uncertainty. For the most part, the experiments in this field have presented subjects with choices among lotteries or gambles with specified outcomes and probabilities of the outcomes. Most of the experimenters have accepted MEU as a normatively correct model of how risky choices ought to be made, and they have sought to identify regular violations of MEU in naive and sophisticated subjects. Several such systematic violations have been found and replicated.

11.5.4.1. The certainty effect

This refers to a thoroughly documented tendency of subjects to prefer sure gains more strongly than MEU predicts that they would. The classic experiment on this was performed by Maurice Allais in what has become known as the Allais Paradox. His initial experiment involved professional colleagues who must be termed sophisticated subjects. When the violation of MEU implicit in the patterns of choice they had just made were pointed out to them, some of his subjects changed one of their preferences in order

to comply with MEU, but others did not.

Here is a simpler experiment that demonstrates the effect, conducted by Kahneman and Tversky (1979:266), based on Allais' work. In each problem, subjects were asked to select the alternative they preferred. The bracketed numbers are the percentages of respondents choosing the alternative.³⁰

“Problem 3:

A: (4,000, .80)	or	B: (3,000)	
[20]		[80]*	N = 95

Problem 4:

C: (4,000, .20)	or	D: (3,000, .25)	
[65]*		[35]	N = 95”

Problem 3 shows that subjects prefer a sure 3,000 to an 80% chance of 4,000. Using $u(\cdot)$ to indicate the utility or value of an outcome, this indicates that

$$u(3,000) > .8 * u(4,000).$$

However, Problem 4 indicates that subjects strongly prefer alternative C over D, which indicates that

$$\begin{aligned} .25 * u(3,000) &< .2 * u(4,000), \\ \text{or } u(3,000) &< .8 * u(4,000), \end{aligned}$$

contradicting the results of Problem 3. The effect certainty has on the valuing of a prospect has been demonstrated with non-monetary outcomes as well.

³⁰ Notation: a choice (x, p) is defined by x , the possible gain, and p , the probability of its occurrence. When p is omitted, it is 1.0, meaning a sure gain of x .

11.5.4.2. The reflection effect

When subjects are presented with gambles involving losses instead of gains, however, the impact of certainty is reversed. Gains that are certain are more attractive than MEU predicts, but losses that are certain are less unattractive (people prefer sure gains to gambles with equivalent expected gains, but they do not prefer sure losses to gambles with equivalent expected losses). The result is apparent risk aversion for gains and apparent risk seeking in the domain of losses.

11.5.4.3. Very low probabilities

Kahneman and Tversky presented subjects with the following pair of choices (1979:267):

“Problem 7:

A: (6,000, .45)	or	B: (3,000, .90)	
[14]		[86]*	N = 66

Problem 8:

C: (6,000, .001)	or	D: (3,000, .002)	
[73]*		[27]	N = 66.”

They interpret the results as indicating that subjects evaluate probabilities as small as those of Problem 8 as indicating only the *possibility* of an outcome. Since both are mere possibilities, the more favorable outcome is preferable.

11.5.4.4. The isolation effect

Complex gambles can be simplified in more than one way, with varying impacts on patterns of preference, as shown in the following experiment (Kahneman and Tversky 1979:271):

“Problem 10:

A: (end of game, .75; go to Stage 2A, .25)	B: (end of game, .75; go to Stage 2B, .25)
Stage 2A: (4,000, .80)	Stage 2B: (3,000)
[22]	[78] N = 141.”

The mathematical summary of Problem 10 is that alternative A gives a chance of .80 for a gain of 0, and a chance of .20 for a gain of 4,000, which is alternative C of Problem 4 above. Alternative B gives a chance of .75 for a gain of 0, and a chance of .25 for a gain of 3,000, which is alternative D of Problem 4. The preferences have been reversed. People respond to Problem 10 as if they note that the first stages of the two alternatives are identical, and then ignore them. They focus only on the comparison between the second stages, isolating that aspect of the decision, and choose as in Problem 3.

11.5.4.5. The “framing” of decisions

The isolation effect is one way in which the particular presentation of a decision can affect patterns of preference. Another example: when people are asked to assume that they have just been given 1,000, and then are asked to choose between a sure gain of 500 more and an even chance at 1,000 more, they choose the former — effectively a sure 1,500. When this problem is “reframed” by asking them to assume they have just been given 2,000 and must choose between a sure loss of 500 and an even chance at a loss of 1,000, they reject the first option, which is the same sure 1,500 they

prefer in the first version (Kahneman and Tversky 1979:273).

Framing effects are quite varied. Whether outcomes are framed as gains or losses, whether alternatives are viewed as one-stage or multi-stage, whether outcomes are framed in absolute terms or relative to a particular reference point will affect the choices people make. While consistency (or coherence) of preference is the common and sole standard of “technical” rationality, Tversky and Kahneman believe that the framing of decisions affects the quality of decisions enough that it deserves an explicit role in normative rationality (Tversky and Kahneman 1981:458). In particular, they endorse framing outcomes in terms of final positions rather than changes to a very recent reference point, and attention to how satisfied the decision-maker would be in the future with a particular outcome instead of concentrating only on immediately present attitudes towards the outcome.

11.5.4.6. Utility evaluation effect

Evidence from many studies, including work by Allais in 1952 and by Karmarkar in 1973, indicates that when utility curves are assessed by the use of a series of choices among gambles, the particular probability values contained in the gambles has a clear impact on the shape of the resulting utility curve (Munier 1988:16). Since MEU is predicated upon a utility curve for each decision-maker that is unique up to a linear transformation, the inability to produce a unique assessment of such a curve would be a serious failing in normative use of MEU.

To test this, McCord and de Neufville (1983) attempted a painstakingly careful examination of this effect. The experimenters were experienced,

skilled practitioners of decision analysis. All of the subjects were sophisticated in dealing with probabilities. The procedures were conducted with great care to rule out irrelevant causal factors. Three different assessment methods were used, two using different probability values, and a technique for assessing the deterministic value of outcomes (value function approach). The two probability-based utility curve methods are unambiguously required by MEU to produce identical results.

They did not. The curves were very different. No two methods were even close. The utility and value functions were not axiomatically required to be the same, but they were no more different than the different versions of the utility functions. A major note is that the differences were systematic, and subjects differed by whether they were risk seekers or risk averters on the direction of the differences.

The empirical result is, in summary, that very different utility functions are obtained depending on the probability distribution used in the assessment. This phenomenon is quite incompatible with the axioms, and places the practice of expected utility decision analysis in a quandary. (McCord and de Neufville 1983:300)

One way out of the quandary would be to say that the axioms are valid but our assessment techniques are flawed. McCord and de Neufville used state-of-the-art methods with great care, so this exit is unpromising. Of course, it is always possible that new assessment techniques will be developed that can circumvent this problem.

The other way out is to conclude that the axioms are wrong — maybe the substitution axiom, which permits excluding the probability distribution from the assessment of utility. This explanation leaves MEU analysis out in the cold.

The conclusion is that the justification of the practical use of expected utility decision analysis as it is known today is weak. Specifically, the whole computational side of the method, which seeks to define a best choice by means of calculus based on a description of a person's utility, does not appear valid: either we do not know how to encode this von Neumann-Morgenstern utility or it does not exist. (McCord and de Neufville 1983:301)

McCord and de Neufville note that many of the other violations of expected utility have been interpreted in two different ways. Critics of MEU have found the violations to be very damaging to MEU. Defenders have seen the violations as proof of the need for MEU, and claim what McCord and de Neufville term a "normative immunity" from the evidence (1983:280). The critics are rejecting MEU for its poor predictive power; the defenders are claiming that the evidence is actually damning human decision-making and pointing out the need for better normative guidance. The utility evaluation effect, however, cannot plausibly be ascribed to faulty human judgment, since McCord and de Neufville's subjects were very sophisticated and the differences among evaluation methods were so consistent. This effect therefore constitutes serious evidence against MEU as normative tool.

McCord and de Neufville believe that decision analysis can still be helpful in clarifying values, preferences, possible outcomes, etc. What this suggests is that expected utility decision analysis is not satisfactory when used by itself to make decisions, but that it can be helpful when used as a method of decision support (when the making of decision is performed in another manner and decision analysis merely provides some of the information used by that unspecified other method).

11.5.5. Alternatives to MEU

In attempts to retain the general framework provided by MEU and conform more closely to the growing body of experimental evidence such as that cited above, new theories have been developed to describe how people choose between risky prospects or gambles. Two of the better known alternatives, *prospect theory* and *regret theory*, are described next, followed by a summary of a very different theory proposed by Janis and Mann.

11.5.5.1. Prospect Theory

Unlike MEU or regret theory, Kahneman and Tversky's prospect theory (Kahneman and Tversky 1979) has two parts. They hypothesize that people begin their choosing with an editing phase in which several types of processing are performed, and then move into the evaluation phase. One of the editing processes is *combination*: different branches of a gamble that have the same outcome will be combined into a single branch. The branches of the gamble (100, .25; 100, .5) would be combined into the gamble (100, .75).³¹

When branches share a feature, the process of *segregation* is used to factor out the common feature, leaving a consequence that is certain and a modified gamble. Kahneman and Tversky suggest a gamble with a gain of 200 common to two branches, such as (300, .80; 200, .20), would be segregated into a certain gain of 200 and a residual gamble of (100, .80),

³¹ Notation: Kahneman and Tversky represent a gamble in the form $(x_1, p_1; x_2, p_2; \dots x_n, p_n)$, which indicates that outcome x_1 will occur with probability p_1 , and so on, with $p_1 + p_2 + \dots + p_n = 1.0$. A branch x_k, p_k where $x_k = 0$ can be ignored, so (100, .5; 0, .5) will be represented as (100, .5). A gamble consisting of a single, sure outcome x is written (x) .

and $(-400, .40; -100, .60)$, with a common loss, would be segregated into a certain loss of 100 and remaining gamble $(-300, .40)$.

In addition to combination and segregation, which are applied to individual gambles, people use processes that operate on more than one gamble. The effect of *cancellation* is to remove features of two or more gambles that are common to each, such as a branch with the same outcome and the same probability, or an identical first stage in different multi-stage gambles.

For example, Kahneman and Tversky propose that cancellation would transform a choice between $(200, .20; 100, .50; -50, .30)$ and $(200, .20; 150, .50; -100, .30)$ to a simpler choice between $(100, .50; -50, .30)$ and $(150, .50; -100, .30)$. *Simplification* is a process of rounding off outcome or probability values. They hypothesize that $(101, .49)$ would be simplified and treated as if it were $(100, .50)$. One of the most important effects of simplification is to eliminate very low probability outcomes altogether. *Detection of dominance* is the final editing process identified, eliminating any gamble dominated by another. For example, $(101, .30; 51, .50; 10, .20)$ dominates $(100, .30; 50, .50; 10, .20)$, and the latter gamble would be rejected. Editing processes themselves account for many “anomalies of preference”. The order in which the various editing processes are used can affect the result, and is affected by the context and presentation of the decision.

The evaluation phase determines the value V of a gamble, based on two scales. A weight for each branch of a gamble is calculated from the probability of the branch: $\pi(p)$. The decision weights are not additive: while $p_1 + p_2 + \dots + p_n = 1.0$, $\pi(p_1) + \pi(p_2) + \dots + \pi(p_n) \leq 1.0$, and typically < 1.0 .

A modified value of the branch is calculated from the outcome: $v(x)$.

Outcomes are evaluated in comparison to a reference position (the current asset position of the decision-maker) and gains and losses may be treated differently from each other.

Restricting consideration to “simple prospects” of the form $(x, p; y, q)$, the value of a prospect or gamble is evaluated in one of two ways. If the gamble is strictly positive (negative), i.e. $x, y > 0$ ($x, y < 0$), and $p + q = 1.0$, then editing (segregation) is performed first. All other, “regular” gambles are evaluated as follows:

$$V(x, p; y, q) = \pi(p)v(x) + \pi(q)v(y).$$

When $p + q = 1.0$ and $x > y > 0$ or $x < y < 0$ (non-regular gambles), then

$$V(x, p; y, q) = v(y) + \pi(p)[v(x) - v(y)].$$

In their theory, $v(0) = 0$, $\pi(0) = 0$, and $\pi(1) = 1$. v measures outcomes, V measures gambles, and $V(x, 1.0) = V(x) = v(x)$. A typical weighting function illustrated in figure 11 shows that $\pi(p) < p$ for most values of p , but $\pi(p) > p$ for p near 0. The

importance of the weighting function is that it can be used to

explain several anomalies. This shape, for example, matches the typically observed over-weighting of low probabilities, and may be interpreted to indicate that extremely low probabilities are treated as if they were equal

Figure 11.

A typical weight function π

to 0.

Some anomalies of decision are explained by the typical value function proposed by Kahneman and Tversky. Their suggestion is that value functions are usually concave in the region of gains, convex for losses, and steeper for losses than for gains, as shown in figure 12. The convexity explains why the certainty-equivalent of a gamble is usually less than the

Figure 12.
A typical value function v

expected value of the gamble. The greater slope of the value function in the loss region corresponds to findings that losses generally appear to have greater impact on decisions than gains of the same magnitude.

Kahneman and Tversky clearly propose prospect theory as a description of how people behave rather than a norm for them to follow:

The equations of prospect theory retain the general bilinear form that underlies expected utility theory. However, in order to accommodate the effects described in the first part of the paper, we are compelled to assume that values are attached to changes rather than to final states, and that decision weights do not coincide with stated probabilities. These departures from expected utility theory must lead to normatively unacceptable consequences, such as inconsistencies, intransitivities, and violations of dominance. Such anomalies of preference are normally corrected by the decision maker when he realizes that his preferences are inconsistent, intransitive, or inadmissible. In many situations, however, the decision maker does not have the opportunity to discover that his preferences could violate decision rules that he wishes to obey. In these circumstances the anomalies implied by prospect theory are

expected to occur. (Kahneman and Tversky 1979:277)

11.5.5.2. Regret Theory

Loomes and Sugden (1982) begin their exposition of regret theory with a description of “choiceless utility function”, $C()$, which identifies the value a consequence would have for the decision-maker “if he experienced it *without having chosen it*” (Loomes and Sugden 1982:807). If the decision-maker chooses action A_1 instead of A_2 , and the j th state of the world eventuates, then he will experience consequence x_{1j} instead of consequence x_{2j} . The major point of regret theory is that the impact to the decision-maker depends not only on the utility of x_{1j} , the consequence that he experienced, but also on the knowledge that consequence x_{2j} was *not* experienced. If x_{1j} is preferred to x_{2j} , he will experience *rejoicing*; if x_{2j} is preferred to x_{1j} , he will experience *regret*. The *modified utility* of choosing A_1 instead of A_2 when the j th state of the world occurs is then

$$m_{1j}^2 = M(c_{1j}, c_{2j}) \text{ where } c_{1j} = C(x_{1j})$$

Their theory is that people choose not to maximize expected utility, but according to maximize “*expected modified utility* E_i^k of action A_i , evaluated with respect to action A_k , by:

$$E_i^k = \sum_{j=1}^n p_j m_{ij}^k$$

(Loomes and Sugden 1982:809)

Regret theory does not specify how the rejoicing or regret is to be incorporated into the modified utility: some people may be relatively unaffected by regret, or they may be very sensitive to it in some circumstances or for some types of decisions and not for others. Loomes and Sugden appraise a limited version of $M(.)$ in which “the degree of

regret or rejoicing depends only on the difference between the choiceless utility of 'what is' and the choiceless utility of 'what might have been'. This allows us to define a *regret-rejoice function* $R(\cdot)$ which assigns a real-valued index to every possible increment or decrement of choiceless utility, and then to write:

$$m_{ij}^k = c_{ij} + R(c_{ij} - c_{kj}) \quad "$$

(Loomes and Sugden 1982:809)

Loomes and Sugden postulate that $R(0) = 0$ for all decision-makers, and for some it may be that $R(m) = 0$ for all m , in which case regret theory would coincide with MEU. Revisiting the choice considered earlier, they assert that a person would prefer A_1 to A_2 when

$$\sum_{j=1}^n p_i [c_{1j} + R(c_{1j} - c_{2j})] > \sum_{j=1}^n p_i [c_{2j} + R(c_{2j} - c_{1j})].$$

11.5.5.3. Other alternative theories

Munier briefly reviews additional alternatives to MEU (Munier 1988:20). In the same year that Loomes and Sugden published their article on regret theory, Bell (1982) published a very similar theory, based also on psychological principles, and Fishburn (1982) published an axiomatic theory that is now known as Skew-symmetric bilinear utility theory (SSB theory). Machina (1982) developed a generalized expected utility model that relies on multiple utility curves that depend on the distribution of probabilities in a particular gamble. Yaari (1987) and Allais (1988) introduced fairly similar models that incorporate distributions of probability differently. All of these theories are mathematically more complex and less easily grasped than prospect theory and regret theory. They each have specific explanatory strengths. None yet seems to have

been acclaimed as champion. The increasing complexity of new theories is testimony to the complexity of the subject, notwithstanding the seductive parsimony of MEU.

11.5.5.4. Comments on Regret Theory and Prospect Theory

Both of these theories attempt to modify MEU by adjusting how gains and losses are treated, and in particular by emphasizing the importance of relative instead of absolute evaluation. Regret theory begins with an absolute “choiceless utility”, but the modified utility at the core of the objective function compares with each other the outcomes of two choices under a single state of the world. The major contribution is to incorporate the *ex post* experience of regret (rejoicing) that another choice that had actually been available would have been better (worse). The relativity of prospect theory comes from evaluating possible outcomes in comparison with the decision-maker’s current position. Each theory is able to predict several of the effects that appear to be deviations from MEU.

Kahneman and Tversky are endorsing the normative model of MEU. All of the deviations from the model they report could be corrected within MEU. They do highlight the importance of the framing of a decision, but given a framing, they accept the principle of expectation as performed with the decision weight π instead of raw probability p . Valuing outcomes in terms of gains or losses does not inherently conflict with MEU, though it emphasizes the role that the framing of a decision can have.

In contrast to prospect theory, regret theory incorporates some aspects of decisions that have no place in MEU. The utility of von Neumann and Morgenstern depends only on an outcome and its probability. Perhaps it

could be separated from its probability, but the methods traditionally used to assess utility involve choices among gambles, and so depend deeply on the probabilities of outcomes. Kahneman and Tversky propose value as a function of outcome and decision weight as a function of probability, but these are still functions only of one branch of the gamble. The riskiness of an alternative is a function of all of the branches of the alternative, not of only a single branch. Neither prospect theory nor regret theory explicitly incorporates risk. Loomes and Sugden, however, do explicitly allow a function of the entire decision, that is, a function of all alternatives and all of their branches, when they describe regret.

Unlike Kahneman and Tversky, and indeed, unlike most of the experimenters in this field, both economists and psychologists, Loomes and Sugden do not accept MEU as the normative standard for rational choice:

[W]e shall challenge the idea that the conventional axioms constitute the only acceptable basis for rational choice under uncertainty. We shall argue that it is no less rational to act in accordance with regret theory, and that conventional expected utility theory therefore represents an unnecessarily restrictive notion of rationality.

Regret theory rests on two fundamental assumptions: first, that many people experience the sensations we call regret and rejoicing; and second, that in making decisions under uncertainty, they try to anticipate and take account of those sensations.

In relation to the first assumption, it seems to us that psychological experiences of regret and rejoicing cannot properly be described in terms of the concept of rationality: a choice may be rational or irrational, but an experience is just an experience. As far as the second assumption is concerned, if an individual does experience such feelings, we cannot see how he can be deemed irrational for consistently taking those feelings into account. (Loomes and Sugden 1982:820)

They do not claim that a person behaving according to MEU is

irrational: it just indicates that for that person $R(m) = 0$ for all m . The implication of their argument, though, is that MEU is incorrect as a normative approach since it forbids behavior that Loomes and Sugden assert is rational.

In fact, even a decision-maker whose own regret-rejoice function is identically 0 may care about the reactions of others who do experience regret and rejoicing, and MEU prohibits the decision-maker from assigning utility to those reactions. This is discussed below under Interactions Among Perspectives.

11.5.5.5. Janis and Mann's Conflict Model of Decision

Janis and Mann's theory is a behavioral or psychological description of the *process* an individual or group goes through while making a decision (Janis and Mann 1977). Unlike MEU, regret theory and prospect theory, it does not focus on the evaluation of alternatives. After receiving some warning of potential danger, the decision-maker is hypothesized to ask a series of four questions and eventually selects one of five possible responses.

- Q1. Are the risks serious if I don't change?
If No, *unconflicted compliance* is the result.
If Maybe or Yes, continue.
- Q2. Are the risks serious if I do change?
If No, *unconflicted change* is the result.
If Maybe or Yes, continue.
- Q3. Is it realistic to hope to find a better solution?
If No, *defensive avoidance* is the result.
If Maybe or Yes, continue.
- Q4. Is there sufficient time to search and deliberate?

If No, *hypervigilance* is the result.

If Maybe or Yes, *vigilance* is the result.

Unconflicted compliance is continuation of the status quo. *Unconflicted change* is the adoption of the most promising of the actions. (Janis and Mann allocate fairly little attention to the evaluation of the alternatives, which receives the bulk of attention in most other treatments.) *Defensive avoidance* is a deliberate ignoring of signs of danger in order that no action would appear to be needed, since there is apparently no satisfactory alternative. *Hypervigilance* is a state of very high arousal in which people are prone to making snap decisions (panic is a type of hypervigilance).

These four responses generally lead to defective decisions, though in any particular circumstance, any might happen to succeed. If a threat never materializes, unconflicted compliance will turn out to have saved effort. Unconflicted change will be adaptive when the danger is well enough known that the available solutions are good. Defensive avoidance may help to bolster spirits during prolonged periods when no solutions are available. Hypervigilance may allow a decision-maker to take actions that would ordinarily not be considered. Careful search for new alternatives and information and careful appraisal of all information, whether confirming or disconfirming current beliefs and hypotheses, characterize vigilance. Vigilance is the response that most fully satisfies Janis and Mann's criteria for good decision-making and has the lowest "probability that the decision maker will make serious miscalculations that jeopardize his immediate objectives and long-term values." (Janis and Mann 1977:13).

11.5.6. The silver lining for "P"

When we are not sure, we are alive.

Graham Greene

The centuries-old philosophical tension between free will and determinism is echoed within individuals. Complete determinism is repugnant to the Western desire for self-efficacy and especially to the American emphasis on individual achievement. A completely certain world would be a completely deterministic one. There can be no freedom without uncertainty.

Two other extremes are equally unappealing. In a random universe, without consistent laws of causation, we might have the freedom to act, but our actions would be meaningless. In a world where we have complete freedom to create our lives within a framework of strong external causality, many of us would find the burden of such total responsibility impossible to bear.

As in many other aspects of our individual roles in the world, the right amount of uncertainty is somewhere in the middle of the range of possibilities. (The concept of the best being a mid-range value conflicts directly with the essence of maximization, though guidelines may be drawn from other cultural settings, such as the golden mean of the ancient Greeks, or the prominence of balance, for instance between Yin and Yang, in the Orient.) With no uncertainty, we would find life uninteresting, probably unpalatable. We would have no choices; at best, we would have mere calculations to make, at worst, we would be automatons. With too much uncertainty, our ability to make sense of our environment and to act within it is overtaxed. A folk wisdom says that variety (read uncertainty)

is the spice of life: we all prefer food when augmented with our favorite spices, but any spice can be overdone.

Without the element of uncertainty the bringing off of even the greatest business triumph would be a dull routine, an eminently unsatisfying affair. (J. Paul Getty)

Maintaining uncertainty or ignorance may make life bearable in the face of burdens that would otherwise overwhelm us, such as terminal illness or the threat of nuclear war (Smithson 1988:211). Smithson suggests that this may be an example of the “utility of being ‘nonrational’”.

The person who frames a situation may be able to use uncertainty to her advantage. This is what makes it better to “ask forgiveness” than to “ask permission” in an organization: asking forgiveness, after the fact, allows the actor to get the benefit of any uncertainty about whether his act deserves punishment (see also section I.1.1 earlier). This is one source of power in being the one to frame issues.

Uncertainty is a key feature of many recreational pursuits. Much of the “utility” of gambling arises from the uncertainty the gambler feels before the result of a gamble becomes known. The suspense that many appreciate in spy novels and movies depends on uncertainty about what will happen next. Most sporting events would be far less entertaining if the outcome were known in advance. It may be that some of these experiences offer us a riskless uncertainty: we are able to act as if we care about the outcome, but we are for the most part not seriously affected by them, and therefore do not experience genuine risk. Teams or characters with whom we identify do, and the vicarious uncertainty and risk is often perceived as a safe and diverting stimulation.

One of the computational resources required by human decision-makers is attention. Unlike the potential supply of information, the supply of attention is limited by human characteristics:

In a world where information is relatively scarce, and where problems for decision are few and simple, information is almost always a positive good. In a world where attention is a major scarce resource, information may be an expensive luxury... we cannot afford to attend to information merely because it is there. (Simon 1978:13)

Simon's point implies that tolerating uncertainty is a strategy for conserving attention. For example, there have been suggestions that one aspect of schizophrenia is an inability to disattend to irrelevant stimuli.

Another way in which uncertainty can be desirable is that an uncertain harm is often preferable to a certain one. There are times when the tension of an unresolved uncertainty may be harder to tolerate than any of the possible resolutions. But there are also many situations where the uncertainty — the possibility — of a bad outcome implies the possibility of a good outcome as well. For example, being aware of the possibility that their plane could crash is very upsetting to many airplane passengers, but few would trade that uncertainty for a certainty that the plane would crash.

II.6. A Sociological Perspective on uncertainty

In his 1988 book, *Ignorance and Uncertainty: Emerging Paradigms*, Smithson treats uncertainty as one of many types of *ignorance*, albeit one of the most-studied. Ignorance, like knowledge, is socially created, having no meaning without reference to a group or individual. It is not absolute but relative in Smithson's definition:

“A is *ignorant* from B’s viewpoint if A fails to agree with or show awareness of ideas which B defines as actually or potentially valid.” This definition ... incorporates anything B thinks A could or should know (but doesn’t) and anything which B thinks A must not know (and doesn’t). (Smithson 1988:6)

II.6.1. Taxonomy of ignorance

The first distinction in Smithson’s taxonomy (1988:9), shown below in figure 13, is between the deliberate act of ignoring and a merely passive ignorance, which he terms *irrelevance* and *error* respectively.

Figure 13. Smithson’s Taxonomy of Ignorance

Irrelevance may be due to a *untopicality*, a perceived lack of relevance, *taboo*, a social proscription on certain types of knowledge, or *undecidability*, insolubility or unverifiability. Error may be due to wrong or to insufficient “information”, termed *distortion* and *incompleteness*. Distortion in degree is *inaccuracy*, while distortion in kind is *confusion*. Incompleteness, which has received more attention than other forms of

ignorance, can again be divided into incompleteness in degree, or *uncertainty*, and incompleteness in kind, labeled *absence*. Smithson goes on to differentiate three types of uncertainty: *ambiguity*, *probability*, and *vagueness*, and to partition the latter into *fuzziness* and *nonspecificity*. (Smithson does not define the last five terms more precisely.)

II.6.2. Historical progression of attitudes toward uncertainty

Smithson maintains that since 1970 or so there has been a

...widespread but unobvious trend in Western intellectual culture [involving] recent changes in how we think about and respond to ignorance. We are currently witnessing the greatest creative work on ignorance since the decade of 1660 (when probability theory emerged). (Smithson 1988:vii)

New approaches represent fundamental re-examination of our ways of dealing with uncertainty. They rely less on probability, and focus on managing the consequences of uncertainty rather than attempting to obtain certainty (Smithson 1988:93).

Perhaps the most strikingly novel characteristic of recent formal approaches is that they seem to accept varieties of uncertainty and even ignorance that heretofore were either ignored (both in the sense of being declared irrelevant and constituting part of our meta-ignorance) or eliminated. More than passively admitting those varieties into the picture, the new approaches prescribe methods for managing and understanding matters such as vagueness, ambiguity, nonspecificity, and possibility. (Smithson 1988:146)

Increasing technological complexity and better understanding of the ways our actions harm the environment are frequently put forth as reasons why uncertainty is receiving more attention today (Smithson 1988:264). Smithson cites a study by Douglas and Wildavsky (1982) in which the authors argue that life today is neither less safe nor more

unhealthy than in previous eras. They do claim that Western society is becoming more risk-averse, and that increasing distrust of large institutions is one reason³².

Smithson (1988:28-40) describes four historical stages in coming to grips with ignorance and uncertainty, with illustrations from mathematics:

1. Exclusion (or banishment of uncertainty);
2. Inclusion (control, often through reductionism);
3. Second-order strategies;
4. Suburbanization (subdivision, relativization).

1. *Exclusion.* The goal during the exclusionary stage is to banish uncertainty, often by defining problematic issues as irrelevant. When non-Euclidean geometries were introduced in the middle 1800s, the certainty and absolutism of important mathematical concepts were threatened. Declaring new geometries to be non-mathematical would rescue that certainty, but the elimination of useful concepts and tools was too great a cost for pragmatists.

2. *Inclusion.* The subsequent inclusionary response is to tolerate the new topics and their uncertainties, but search for ways to at least tame the problems they bring with them. The most common tack is reductionism, seeking more fundamental bases, such as new systems of axioms, that could control the complications. In mathematics, the logicism of Frege, Whitehead, Russell and others was an inclusionary effort to build a

³² Distrust is a second-order uncertainty, a matter not of whether our information is complete but whether we can trust it, however complete it appears to be.

foundation of logic underneath mathematics.

I wanted certainty in the kind of way in which people want religious faith. I thought that certainty is more likely to be found in mathematics than elsewhere...if certainty were indeed discoverable in mathematics, it would be in a new field of mathematics, with more solid foundations than those that had hitherto been thought secure. (Bertrand Russell)

Even with the assistance of banishment efforts (e.g. the *Theory of Types* of Russell and Whitehead), inclusion did not rescue certainty.

3. Second-order strategies. Second-order strategies involve analyzing ignorance and uncertainty explicitly. In mathematics, these efforts focused on completeness and consistency, and proofs of their attainability, illustrated by Hilbert and formalism. Smithson claims this approach was significantly different from inclusion, in that it focused on a second-order certainty, a “meta-mathematical methodology” (1988:35).

4. Suburbanization. In 1931, Gödel’s *Incompleteness Theorem* demolished hope of restoring certainty to mathematics. What has followed is a subdivision of the field into nearly independent areas (“suburbanization”). Approaches to uncertainty vary from one area to the next, typically selected by aesthetics or utility since even a hope of absolute truth has had to be abandoned. This stage is characterized by an acceptance of uncertainty and a relativization of treatments and approaches to uncertainty.

Smithson notes that in some ways the new efforts are traditional, such as the use of traditional mathematical machinery:

Fuzzy logicians do not prove their theorems in degree less than 1; in fact they frequently resort to proofs by contradiction and

thereby utilize the strong version of the Law of the Excluded Middle³³...the old uncertainty reduction impulse being established at the second-order level: If we cannot have first-order certainty, let us attain certainty about uncertainty. Far from being the province of reductionists like Hisdal or Eytan, this impulse is spelt out by some of the most ardent revolutionaries. (Smithson 1988:147)

Smithson places himself in the camp of suburbanites. Noting that “engineers almost always quantify their probabilities while lawyers and judges generally resist quantification,” he argues for an open discussion of relative normative standards — guidelines that vary across situations (Smithson 1988:302). His sociological perspective highlights the human context in which ignorance and uncertainty are created. This implies a need to consider more individual, social, and situational factors than the dominant perspective from economics, decision analysis and cognitive psychology has permitted, and a framework that is less general and more case-specific.

Smithson identifies several key points from his work (1988:300):

Wildavsky’s Maxim. “System resilience requires toleration of ignorance;” attempts to control and eliminate ignorance and uncertainty weaken adaptability (Wildavsky 1985).

Rationalist Quandary. “The more specialized and normatively restricted the language for representing ignorance, the more people are excluded from discussion or debate, and the less effective the language

³³ Lovell’s note: the Law of the Excluded Middle states that either **S** or not-**S** is true for every proposition **S**. It is rejected in both fuzzy set theory (an element may have a degree of membership in a set that is neither 1 nor 0) and in possibility theory (the sum of the possibilities of an event and its negation may be greater than 1.0; the sum of the corresponding necessities may be less than 1.0).

becomes for communicating about ignorance to nonspecialists.”

Zeleny's Tradeoff. “The more context-free (hence unambiguous) the representation of ignorance, the lesser are its representational capabilities. The more context sensitive, and therefore descriptively rich, the more difficult it is to attain clarity of expression. As more than one wag has put it, we have a choice between being vague and wrong” (Zeleny 1984).

II.7. Dror's Technical/Organizational Perspective

One of the few authors to write at length specifically about a more inclusive interpretation of uncertainty is Yehezkel Dror. His perspective blends technical and organizational aspects: his work is very technical and often abstract, but it is motivated and informed by his involvement with planners, the subject of planning, and organizations conducting explicit planning efforts. A good illustration of many of his points is his 1986 article “Planning as Fuzzy Gambling: A Radical Perspective on Coping with Uncertainty.” The five major sections of his paper, which has been a major inspiration for this dissertation, are summarized next.

II.7.1. Concept package for uncertainty

Dror contends that an adequate theoretical approach to uncertainties requires what he terms a *concept package* for distinguishing uncertainties from each other. He notes that while in some contexts even the certainty/risk/uncertainty scale (described earlier in the section on the Standard Model of Decision) is not used, more comprehensive frameworks are needed for a sophisticated treatment of uncertainty.

The first additional distinction is between what Dror calls “quantitative” and “qualitative” uncertainty. In the former, the possible

future states of the world are known, though their odds are not; in the latter, even the possible states of the world are unknown. The SDM recognizes only quantitative uncertainty by assuming the set of states of the world is known.

The second distinction is between situations in which alternative trends or tendencies can be identified to some degree and others in which even trends and tendencies appear to be unknowable. Examples of the latter type of situation are those that would be described as “explosive”, “surprise-prone”, or where there is a “high probability of low probability events”. As Dror adds in a later article, low probability events with high impact present special problems (Dror 1988a:257). He recommends developing lists of potential high-impact surprises and having possible responses ready, since time may be short if and when they occur. An extreme of one sort is turbulence, when Dror says that policy analysis is most needed but when fine tuning of policies is often a delusion: robust, not optimal, policies are needed (Dror 1991:12). This echoes Wildavsky’s recommending resilience over anticipation under turbulence, mentioned earlier in section 11.4.9.

A third and central characteristic of an uncertainty is whether it is “hard” or “soft”. A hard uncertainty is one that is not reducible, even hypothetically. Chaotic systems, whose behavior we have grounds for believing to be unpredictable, present the difficulties of hard uncertainty. Dror comments that

this is a very important distinction, because hard uncertainty imposes an absolute limit on predictability, and on the outputs of

even ideal prediction methods, which cannot supply more than a hypothetical 'perfect witness'³⁴ could know. (Dror 1986b:26)

The fourth of the distinctions in Dror's preliminary concept package is whether we even have enough information to apply, with any degree of reliability, the first three "mapping categories" or not.

In decision situations, Dror believes it is essential to draw a distinction among different "dimensions of output-goals". By this he means that assessment of substantive goal consequences is one major dimension, usually incorporated into analyses of decisions, but in addition preferences regarding the lottery characteristics of the choice and preferences over the distribution through time of the consequences will always be factors that can affect the decision (Dror 1988b:139). While MEU denies the meaning of preferences over lottery characteristics³⁵, the empirical evidence that many people have such preferences for or against risk is so strong that the terms "risk seeker" and "risk averse" are a standard part of the terminology of the field. People may have very strong and specific preferences regarding the degree of risk or uncertainty their lottery choices will yield them, as illustrated by one writer's personal comments:

Sometimes, however, we welcome an element of uncertainty, and

³⁴ "The concept of 'perfect witness,' as referring to an entity that (a) knows all the present facts; and (b) knows all the laws of change dynamics, is helpful here. In the face of chaotic, random, and similar processes, a perfect witness will not say 'I do not know,' but will state 'I know that one cannot know.' These two statements have quite diverse implications for planning ..." [Dror, 1986b:endnote 13]

³⁵ The common identification in economics of the utility of a lottery with the weighted average of the utilities of the possible outcomes is said to eliminate any preference for or against risk. The economic explanation for risk aversion is the concavity of the utility function, and that there are no genuine preferences for or against risk *per se*. People who believe gambling to be immoral, or to be exciting, disagree.

even create it if necessary. It is true that uncertainty makes it difficult to plan for the future, but without uncertainty we might not even want to plan for the future at all. ... For myself, however, I am quite sure that I would prefer a shorter life expectancy with a larger spread to a longer life with no spread at all. If I had the choice, on this issue, between an unfair lottery and a sure thing, I would take the lottery. (Footnote 7: Formally, this is a form of preference for risk. The underlying reasons, however, are quite different. If I prefer the smaller average with the larger dispersion, it is not because I gamble on a long life but, as stated, because the certainty would be intolerable.) (Elster 1989a:38)

Discussions of time-pattern preferences are more rare, though the issue cannot be said to be any simpler or less important. Von Neumann and Morgenstern explicitly avoided it:

However, it would be an unnecessary complication, as far as our present objectives are concerned, to get entangled with the problem of the preferences between events in different periods of the future. (von Neumann and Morgenstern 1953:19)

Elster concurs that this is a complex and significant dimension:

Against the economists, I believe that it is irrational if one knowingly does what will make one's life as a whole turn out worse than it could have been. I have no argument against [Derek] Parfit's view that a person who expected his future states to be weakly connected would not be irrational in discounting the future. I believe, however, that a person who takes his future states as given, rather than as something to be created, is fundamentally irrational. Moreover, I believe a person will be better off by striving for connectedness, since only then will he be able to form the long-term plans that are a condition for living a meaningful life, *even in the present*. (Elster 1986:11)

Most discussions of time-pattern preferences are restricted to debate over which positive *discount rate* to use after assuming the use of an exponential decreasing function. Some have suggested, however, that a negative discount rate ought to be used, meaning that the consequences of our actions that our children will bear are to be counted more heavily, not

less heavily, than the immediate consequences we will bear.

II.7.2. Relevance of concept package

Dror assesses the uncertainty of major variables on which large-scale planning work depends, including impacts of technology, global demographic trends, the possible impacts of surprise events, political ideologies in developed countries, and not only ideologies but specific regimes in less developed countries.

It seems justified to conclude that most planning must cope with a lot of quantitative and qualitative uncertainty, much of which is 'hard,' that is, built into the dynamics of relevant phenomena and, therefore, irreducible by any prediction methods. (Dror 1986b:29)

II.7.3. Fuzzy gambling

Any decision whose outcome is less than fully certain constitutes a gamble from the decision-maker's point of view, whether the uncertainty is due to randomness, the unpredictability of other agents, psychological or other factors. In Dror's view, however, most "real-life decision making" is much more complicated. The rules of the game, the payoffs, the roles of the other players, and so on — these are generally known imperfectly, and all of them can change *during the time frame of the decision and its results*. This is what "fuzzy gambling" comprises: the decision-maker cannot count on knowing the basic rules of the game. Dror's position is that most decision-making and most planning should be viewed as fuzzy gambling. Still, this is better than what he refers to as

chaotic policy gambling, when nothing whatsoever about relevant futures can be said on a level above purely imaginary speculations. (Dror 1988b:150)

Even leaving aside the fuzzy or chaotic aspects, the notion of

policymaking as a gambling activity is a radical departure from the mainstream of thinking in planning, economics, and other technical fields (Dror 1991:35). No matter how carefully and optimally an uncertain decision is made, the outcome is still uncertain, and therefore constitutes a gamble. There is no choice that can guarantee success.

II.7.4. Planning as fuzzy gambling aid

Dror's conception of decision-making as fuzzy gambling leads him to conceive of the purpose of planning as either fuzzy gambling itself or as an aid to fuzzy gambling activities. While top managers who lead organizations are often aware of the fuzzy gambling nature of decision-making and planning, many other users of planning and many professional planners think that planning will enable them to "ungamble" decisions, and they are likely to find fuzzy gambling to be a radical view. They are likely to want planning and analysis units to reduce uncertainty, not emphasize it: this poses a moral challenge to the planning profession (Dror 1987:37).

II.7.5. Implications

Dror's radical reconceptualization of planning as fuzzy gambling prompts him to ask if planning is even useful. (A similar question will be asked below about rational decision-making in the Implications chapter's section on rational methods.) His answer is that planning is in many cases an unavoidable activity when long-term projects are contemplated, and that in other cases it does appear that planning has been helpful. In yet other situations, planning provides a means for not merely coping with but actually shaping future events.

One implication of this reconceptualization is that a major purpose of planning may be to emphasize the prevalence and nature of fuzzy gambling so that people are less likely to be misled into false certainty. In fuzzy or chaotic environments, continuing current policies is still gambling, even if those policies have been successful, since there is no assurance that the current conditions will hold and the policies will continue to meet with success (Dror 1988b:150). Another function Dror suggests is to “debug” decision-making processes to help compensate for known errors and biases in human thinking (such as those described earlier in the P Perspective section of this chapter).

In fuzzy gambles, unexpected events will occur — inevitably. “Very often good planning may fail and sometimes very bad planning may succeed” (Dror 1987:37). This means that the quality of planning and decision-making cannot be assessed by first-order results. (Dror is agreeing with Elster’s rejection of the *ex post* perspective for such assessments (Elster 1989a:117).) Clients may not accept this and insist on holding planning accountable for results. The practical reality of this “intellectually unjustified” position is important for planners and analysts (Dror 1987:37).

At the same time that he believes this view of planning is essential, Dror also understands it entails risks:

Even the very understanding of policies as fuzzy gambles can have negative results, for instance by encouraging recklessness. The Talmudic rule on some types of knowledge which should not be diffused may well apply here. The problem needs consideration within [a] philosophy of policy-reasoning. (Dror 1988b:164)

II.7.6. Normative guidelines

Dror has expanded the implications of his fuzzy gambling perspective in his 1988 “Uncertainty: Coping with It and Political Feasibility”. In general, he sees the inadequacy of numerical methods in the face of hard uncertainty as requiring greater reliance on qualitative tools (Dror 1988a:259). Among the approaches he advises are sensitivity testing without probabilities, scenario analysis, explicit consideration of surprise possibilities, and shifting emphasis from risk reduction to resilience and robustness of planning. Akin to emphasizing resilience is recognizing the distinction between striving to achieve a positive situation and striving to avoid a negative situation, and giving additional attention to the latter (Dror 1988b:137). Two reasons for noting the distinction between “good achievement” and “bad avoidance” are that it may be easier to reach a consensus on which bads to avoid than which goods to pursue, and even brief reviews of history find many examples of bads the avoidance of which would have been of great benefit (Dror 1991:27). Utility and preference functions used to determine “best” fall short when complex value judgments must be made under qualitative uncertainty: it becomes necessary to “rely more on ‘acts of will’ by decisionmakers” than to rely on calculations of optimality (Dror 1988a:259).

He notes that research in individual psychology, group dynamics, organizations, and cultural studies all reveal systematic biases and errors in responding to uncertainty (Dror 1988a:261-62). The cultural biases may pose special problems for systems analysts, who are embedded within the culture, to overcome or “jump out of” (Dror references Hofstadter 1979:ch.

XV). The systems analysis units must constrain the decision-makers against their own biases either by having institutionalized freedom to criticize them or by helping decision-makers binding themselves as Ulysses did before the Sirens (Dror references Elster 1979; Schelling 1984:ch. 5)

Better tools for communicating about uncertainty are needed:

There is little hope of conveying the decision-gambling idea and applying it correctly, and in a form comprehensible and useful to nonprofessionals, without developing novel modes of presenting complex analyses to busy clients ... in a way that clearly explicates major uncertainties and the judgments required from them. ... I am not sure whether decisionmakers will welcome such presentations, which clearly put on them the onus for hard choices. (Dror 1988a:263)

The style of appropriate planning depends on the organization's environment, especially when it approximates Emery and Trist's turbulence:

Planners must be aware of the limits of planning when uncertainties become extreme. Under such conditions, crisis management and improvisation become preferable modes of policy-reasoning with far-reaching revisions of the roles of planning. ... Under conditions of far-reaching uncertainty, as in Israel, good improvisation is far superior to traditional planning, which assumes much predictability. (Dror 1987:39)

Extreme uncertainty may be engendered by environmental complexity. "Very complex approaches are essential for very deep complexity" (Dror 1991:33). This recalls Ashby's Law of Requisite Variety, which states that in order to be able to keep a system's primary variables constant and protected from disturbances in the environment, a regulator, or controlling mechanism, must have at least as much variety in its responses as the amount of variety in the environment (Ashby 1954). The complexity of

approach Dror counsels includes use of multiple perspectives, inclusion of both numerical and non-numerical representations, and extension to literary and artistic imagery and mind sets.

II.8. Summary of the main three perspectives

Table 11 below displays and contrasts salient characteristics of the three main perspectives described above.

II.9. Previous taxonomies of uncertainty

II.9.1. Knight: *uncertainty versus risk*

The distinction between risk and uncertainty is traced to Knight (1921). The uncertainty involved, as in the SDM, concerns only the possible future state of the world.

- Decision under ***certainty***, or Type I: the consequences of each of the options is known with certainty.
- Decision under ***risk***, or Type II: the set of possible futures, in which the consequences will occur, has more than one member; a probability distribution over these possible futures is known, and therefore a probability distribution over the (known) set of possible consequences for each option is also known.
- Decision under ***uncertainty***, or Type III: as in risk, the set of possible futures is known and has more than one member, but the probabilities are not known.

II.9.2. Thompson and Tuden

Michael (1973) cites work by Thompson and Tuden (1959) elaborating a four-part schema of organizational decisions. They define two dimensions: “causation”, referring to the ability to predict the consequences of a choice, and “preferences”, referring to which consequences are preferred and to what degree. On each dimension there may be agreement or not, that is, the decision-making group may be certain or uncertain, giving $2 * 2 = 4$

Table 11. T, O and P Perspectives on uncertainty

	T Perspective	O Perspective	P Perspective
Uncertainties considered	Focus is on uncertain factors or variables, esp. future states of the world	On-going sources of uncertainty, information flow	Environment, options, values, existence of choice all uncertain, perceptions uncertain
Human context	Archetype is one decision, one event, in one instant	Planning and decision-making processes, S.O.P.s	All aspects of life permeated by uncertainty, doubt, insecurity
Disadvantages	An impediment to rationality	A complication, reason to be more cautious	Source of anxiety, excitement, lack of confidence
Advantages	None — one's own uncertainty is purely an ill	Provides freedom for maneuvering, for compromise	Spice of life, excitement, can add meaning to life
Assumptions about future	Future assumed to be external, not affected by decisions	Modifying future is often an explicit goal of planning	Hope of changing future varies widely; (correlated with hope?)
Goals of main approaches	Determine the "right" answer	"Muddle through", develop consensus to permit action	Attempt to reduce anxiety, restore appearance of certainty
Common techniques	Generalization, simplification, reductionism; calculation of the "right" answers	Emphasis on option generation for compromise, elaboration of single plan, delegation	Simplify problems, choose from reduced set of solutions, ignore uncertainty
Meaning of "best" course of action	"Best" identified by formal optimization	"Best" identified by satisficing or rote exercise of optimization procedures; political judgment in satisficing	"Best" identified intuitively, often by reference to "gut feeling" or other signs of lowered anxiety
Role of quantification	Required; non-quantitative factors discounted or ignored	Tolerated but suspect; may be used defensively ("the computer said...")	Resisted, seen as dehumanizing, depersonalizing
Role of values	Values are external to "T", assumed to be "given" somehow; static references	Emergence of shared values can reduce problem of evaluating actions; can be instrumental	"Inherited" personal values, possibly modified by experience, are the predominant mode

possibilities. Thompson and Tuden assert that there is a characteristic

mode of making decisions for each of these situations, shown in table 12.

When there is agreement on both causation and preferences, i.e. when the consequences of each option and the desirability of the

Table 12. Thompson and Tuden's decision modes

Preferred decision-making mode	Causation agreement (certainty)	Causation disagreement (uncertainty)
Value agreement (certainty)	Computation	Judgment
Value disagreement (uncertainty)	Compromise	Inspiration

consequences are known, the decision-making process is one of computation. When there is agreement about causation but disagreement about preferred outcomes, the process is compromise. When there is disagreement about causation but agreement about preferences, decisions are based on judgment. But when there is disagreement about both causes and outcomes, decisions are likely to be made by 'inspiration'.

According to Michael, "inspiration would seem to be the most likely process for many future-oriented goal-setting decisions" (1973:149). On the other hand, computation is the method prescribed by the traditional "rational" model, embodied in the SDM, for both causation certainty and causation uncertainty. Since the SDM assumes that values are exogenous to the decision process, it implicitly assumes value certainty in all cases.

II.9.3. Dror

II.9.3.1. Type

Dror (1986b) asserts that a common situation facing planners is that the set of possible states of the world is not known. He terms this "qualitative uncertainty", or Type IV, to contrast it with the traditional

Type III case (see section II.9.1), which he labels “quantitative uncertainty”, in which the possible states are known but their (quantitative) likelihoods are not:

- Type I: certainty; future state of the world is known.
- Type II: risk; possible future states of the world are known and their probabilities are known.
- Type III: quantitative uncertainty; possible future states are known but their probabilities are not.
- Type IV: qualitative uncertainty; possible future states of the world are not known (thus no probability distribution can be known or estimated either: probability is distributional, so a total of 1.0 of probability must distributed over the set of possible states, but even the number of these states is unknown).

II.9.3.2. Hardness

Some uncertainties can be reduced, at least in principle, by some action; Dror terms these “soft uncertainties”. Others cannot; even a witness with the best possible knowledge “would declare in such a situation ‘[not] only do I not know, but I know for sure that it is impossible [to] know.’” (Dror 1988b:149); these are “hard uncertainties”.

- Soft uncertainties: potentially reducible.
- Hard uncertainties: known to be irreducible.

II.9.4. *Smithson*

As described in more detail in the section above on A Sociological Perspective, Smithson divides uncertainty into three types, with one subdivided further:

- Ambiguity: several possible and distinct interpretations are possible.
- Probability: the likelihood of a possible event.

- Vagueness
 - Fuzziness: typically applied to linguistic predicates.
 - Nonspecificity: lack of precision

II.9.5. Klir

Klir (1992a:82) distinguishes four types of uncertainty:

- fuzziness or vagueness
- nonspecificity or lack of informativeness
- dissonance or conflict, and
- confusion or pure and potential conflict.

Fuzziness and nonspecificity are roughly equivalent to imprecision of degree and imprecision of kind. “The food was hot” is non-specific (often called ambiguous), since “hot” could refer to either temperature or spiciness. “The stove is quite hot” is fuzzy, since “quite hot” does not measure the temperature very precisely. Dissonance is the kind of uncertainty that probability can represent. Klir suggests that this is a useful way to characterize bodies of evidence when some evidence supports one belief and other evidence supports another belief(s) — when there is conflicting evidence. When some of the evidence supports none of the beliefs, there is an added element of “potential conflict” which is better captured by measures of confusion.

Klir’s taxonomy is highly mathematical, focusing on the construction of mathematical measures of each of his four types of uncertainty rather than on the interpretation of the types or responses to them.

II.9.6. Friend and Jessop

As reported by Michael (1973), Friend and Jessop devise a three-part

classification of uncertainty in an examination of a rebuilding project in Coventry, and assess the typical treatments and their disadvantages (Friend and Jessop 1969).

- Class UE includes environmental factors, largely technical and technological;
- Class UR pertains to related fields, matters of choice, the decision-making system, and how these are to be viewed;
- Class UV comprises value issues.

II.9.7. Dubois and Prade

Dubois and Prade (1988:3-5) refer to imprecision and uncertainty as complementary aspects of imperfection of knowledge. An “item of knowledge” has four components: attribute, object, value, and confidence (illustrated below). Imprecision, they say, refers to the value, uncertainty to the confidence. Degrees of uncertainty can be measured with words such as probable, possible, necessary, plausible, or credible. Probable, much studied for 200 years, has two different meanings: one is physical, based on statistics and frequencies of observed events, the other epistemic, pertaining to subjective judgment. Plausible and credible are epistemic only: credible means deducible from a corpus of knowledge; plausible means does not contradict a corpus of knowledge.

Dubois and Prade use the terms vague, fuzzy, and ambiguous somewhat synonymously, referring to the “absence of a clear boundary to the set of values attached to the objects it refers to” (1988:3). Precision is a matter of whether the value component can be subdivided further. They illustrate fuzziness thusly:

An imprecise, nonfuzzy proposition:

$(x = y \text{ to within } \epsilon) \equiv (\text{equality}, (x, y), \text{to within } \epsilon, 1)$

An imprecise fuzzy proposition:

$(x \text{ approximately } = y) \equiv (\text{equality}, (x, y), \text{approximately}, 1).$
(Dubois and Prade 1988:3)

Precision and uncertainty are opposed in the sense that making a proposition more precise tends to make it more uncertain, and the uncertainty of information tends to make less precise any conclusions to be drawn from it (Dubois and Prade 1988:4).

Their treatment of uncertainty is very technical, couched in the argot of possibility theory. On the other hand, it is not exclusively mathematical, and they consider verbal measures of certainty and precision to be important.

- Vagueness, imprecision — lack of precision in a measurement or assessment of a quantity.
- Fuzziness — imprecision due to the inexactness of language.
- Uncertainty — the degree of confidence associated with a measurement or assessment.

II.9.8. Osborn, Hunt and Jauch

Conceiving of uncertainty largely as *change* in their book on organization theory, Osborn, Hunt and Jauch (1980) describe four dimensions:

- Disparity (heterogeneity);
- Volatility (rate of change or velocity);
- Degree of change (force); and
- Predictability of change (directional deviation).

II.9.10 Morgan and Henrion

Morgan and Henrion offer a taxonomy of sources of uncertainty in *empirical variables*, which must be at least at some time potentially measurable (1990:50 ff.). These are distinguished from *decision variables*, the values of which the decision-maker may choose, *value parameters*, representing the values the decision-maker will use for comparing the utility of consequences, *index variables* and *model domain parameters*, which have to do with the mechanics of operating a model, and *outcome criteria*, or evaluation criteria. While they acknowledge that uncertainty about the value parameters to be used, or about the form of a model to be used, only empirical variables are considered in their taxonomy, presented next.

- Random error and statistical variation — resulting from direct measurement of a quantity. This is the source of uncertainty that is most studied and most easily treated with standard tools.
- Systematic error and subjective judgment — biases that have not been noted and corrected in equipment and procedures, and the effects of the many applications of subjective judgment that cannot be avoided. Both of these effects are typically underestimated and underappreciated as sources of uncertainty. “Careful examination of what is involved in measuring physical constants and assessing their uncertainty, reveals that even here, there is an irreducible component of subjective judgment.” (p. 60) Care should be taken not to overlook this possible source of uncertainty in the conclusions of an analysis.
- Linguistic imprecision — the vagueness inherent in use of words, especially when interpreted by different people. Uncertainty from this source cannot be eliminated completely, but can usually be reduced to a tolerable level fairly easily.
- Variability — samples drawn from a population or frequency distribution (*not* a probability distribution). “A common mistake is failure to distinguish between variability due to sampling from

a frequency distribution and empirical uncertainty that arises from incomplete scientific or technical knowledge.” (p. 62)

- Randomness and unpredictability — either inherent randomness or unpredictability due to limitations, such as on computational power. Here Morgan and Henrion want to “distinguish uncertainty that might be reducible by further research or more detailed modeling from uncertainty that is unlikely to be reducible.” (p. 64)
- Disagreement — disagreements among experts. When the uncertainty between experts is small compared to the uncertainty of each individual’s estimate, mathematical combinations of opinions are reasonable. When interpersonal differences are large, alternative conclusions may be the best way to accommodate disagreement without ignoring its significance.
- Approximations — the imprecision introduced by the simplifications and approximations of any model. The effects of sample size (e.g. Monte Carlo samples) can be estimated easily, but effects of grid size or time scale may be very difficult to estimate without enormous effort.

II.10. Interactions among perspectives

Dogmatism and skepticism are both, in a sense, absolute philosophies; one is certain of knowing, the other of not knowing. What philosophy should dissipate is *certainly*, whether of knowledge or ignorance.

Bertrand Russell

Perspectives are ways of viewing, and they affect not only how we see particular objects, but even what objects or characteristics of them we are able to see. This sub-chapter examines some of the ways that the three major perspectives considered in this work — T, O, and P — can interact, and some effects of this interaction.

It is not sufficient to recognize that there are different aspects of the object that may correspond to Technical, Organizational, Personal issues,

but the investigators ought to use different ways of viewing the different aspects. A T Perspective on the impact of uncertainty on an individual person is quite different from a P Perspective on the same impact. The scientific research of Tversky and Kahneman (e.g. Tversky and Kahneman 1974) represents the technical view of personal responses, assessing common heuristics and the biases associated with each. An O Perspective on personal responses might focus more on the demoralizing effect of widespread uncertainty in the work force, and how managers can affect that demoralization by how they respond to new rumors. A P Perspective would include more detail on how unsettling uncertainty can be (perhaps in a short story or a poem), or what it feels like to be highly uncertain, or the common response of denying the extent of the uncertainty in order to reduce discomfort.

II.10.1. Power

Crozier (1964) asserts that the natural tendency of what we are calling the T Perspective within organizations is to develop rules to control the performance of work, that is, to “rationalize” work. The rules act to reduce the flexibility of all organizational actors, which reduces the uncertainty others have about the actors’ behavior, which reduces their power. However, one of the main focuses of the O Perspective and of skillful organizational insiders is the understanding and exercise of the levers of power. There is, then, a direct conflict between O and T, in that implementation of T principles reduces the value of O expertise, and some experienced O players, perceiving this, will resist the T initiatives.

II.10.2. Risk aversion as an organizational characteristic

In many organizations, failures are more salient than success.

But [executives] were all conservative about risking prospects of loss for their corporations perhaps because the penalty in organizational life for being responsible for one bad outcome is so much greater than the rewards for selecting alternatives that do well. (Bass 1983: 74)

Thus, risk-aversion could be a natural result of the organization visibility of failures. Although uncertainty increases both risks and opportunities, a risk-averse culture is more sensitive to the amplified risks, and the dominant effect is likely to be increased caution. This increased caution can also be observed in many public decision-making bodies, where second-guessing is likely. Regret Theory (summarized previously in section II.5.5) is especially well-suited to represent this effect. Two factors whose influence on the visibility of failures could be investigated are whether failures tend to be identified sooner and more easily than successes, and whether successes depend on a larger number of events and people than failures, thus making credit more diffuse than blame.

II.10.3. S.O.P.s

The emphasis on standard operating procedures within organizations (O Perspective) may be partly due to individual responses to uncertainty (P Perspective):

The anxieties generated by the experience that one's knowledge is not always adequate for one's tasks can be calmed by making a ritual of conformity with those procedures with which one has become familiar. (Blau 1963:91)

II.10.4. Organizational homogeneity

Stinchcombe provides an illustration of an impact of a P phenomenon on organizational make-up:

Investment groups have to trust each other with large quantities of money or other resources. There is a strong tendency for people to trust only those others whose behavior they understand well – people of the same culture, often the same people they see socially... One consequence is that people of different cultural backgrounds (for example, blacks, women, or foreigners) may have a harder time breaking into an investment groups than into a technical theory group. (Stinchcombe 1990:178)

Meeting a technical job need, that of handling large sums of money, requires a personal relationship of trust, built on a certainty about how others will act. Dealing with co-workers with different background increases the uncertainty about their behavior, making trust more difficult to develop. This leads to a greater homogeneity among investment groups, which could pose managerial problems for an organization.

II.10.5. Perverse interactions

II.10.5.1. Friend and Jessop's study

Friend and Jessop's study (1969), referred to earlier in section II.9.6, classified uncertainties into three groups: (1) UE: environmental, relating most directly to the T Perspective used here; (2) UR: "related fields", largely a matter of the context of the project and its decision-makers, comparable to the O Perspective; and (3) UV: values, a concern of the P and O Perspectives. Friend and Jessop describe the typical responses to each type of uncertainty, and it is clear that each response creates additional problems for other perspectives.

The typical response to UE uncertainty was increased research, which

prolonged the processes and overemphasized measurable factors. Any protraction of a decision process is likely to cause organizational problems. Emphasizing measurable factors often downgrades the importance attached to concerns of individuals, many of which are subjective and difficult to measure. Quantification itself often alienates individuals.

Typical responses to UR uncertainty included greater efforts at planning and coordination, leading to more linkages among decisions, and thus more difficult decision-making. More complex decisions pose problems for both the analytical tools of the T Perspective and the cognitive capacities of individuals who must complete the processes.

The common response to UV uncertainty was to call for more policy guidance, and the resultant importance of political considerations was often resented by individual workers. This would be especially true for workers who focus on the technical and “objective” aspects of a decision, or who believe that they possess expertise that is neglected in favor of politics.

II.10.5.2. Lorraine's study

Lorraine (1991) reports interactions among the three perspectives that occurred in a medfly eradication effort. The two primary perspectives she examines are the T Perspective, generally employed by scientists on the project, and the O Perspective, employed by the project's managers. She also identifies the effects of several Cognitive (C) Perspectives, substantially the same as the P Perspective utilized here. Each of the perspectives has common ways of dealing with uncertainty, and each systematically ignores specific types of uncertainty. A pernicious effect

that Lorraine observes is that the responses of one perspective to the uncertainty it focuses often creates uncertainty that is especially difficult for another perspective to handle (Lorraine 1991:7).

Common ways for the T Perspective to deal with uncertainty are to focus on sub-systems, ignoring more complex larger-system interactions (i.e. to use reductionism), to ignore information that did not fit the accepted theories, and to ignore many types of idiosyncratic data (anecdotal data, qualitative data, the importance of unique events, the specifics of field conditions at a given time, and the particular actions of individual people or of different organizational units). It assumes that organizations act as coordinated whole entities with consistent sets of goals, and that all prescribed procedures are implemented exactly.

A main technique employed by the O Perspective is to resolve uncertainty through negotiation, sometimes over the details of implementations of technical solutions.

[The O Perspective] allowed [managers] to make choices and take action in the face of conflict and incomplete information by creating the very sources of uncertainty that the rational-technical perspective implicitly assumes will be strictly controlled. (Lorraine 1991:8)

The O Perspective, conceiving of truth as negotiable, underestimates the importance of precise implementation, careful information gathering, and monitoring. It emphasizes standard operating procedures as a way of reducing uncertainty, which makes it less able to take advantage of new technical information that hasn't yet been standardized.

Individual actors have their own P Perspectives and sets of motivations. Individuals reduce the uncertainty they have about project

goals or implementation techniques in individualistic ways which are typically invisible to both the O and T Perspectives. The T Perspective, for example, assumes that all individuals act only rationally, and that they know exactly what to do.

II.10.7. Analysts versus decision-makers

Ascher and Overholt (1983:44 ff.) point out that an organization is far from being a single, integrated entity with a single set of goals and values. In addition to motives that may depend on the characteristics of their individual lives, there are some regular differences of motivation between policy-makers and analysts in the same organization. The policy-maker has several “rational” (T Perspective) needs that must be met by the information provided to him by the analyst, including the amount of uncertainty associated with it, the borderline between information and recommendation, the context and embedded assumptions of the information, and knowledge of which factors in the policy-maker’s world are relevant to the information.

The policy-maker has political or organizational (O Perspective) needs, which Ascher and Overholt characterize as follows:

1. To be a convincing advocate of preferred policies—hence to have access to professional-seeming information.
2. Whenever possible, to be correct (that is, to choose policies that, if implemented, produce positive results).
3. When wrong, now to be disastrously so—thus, to make conservative decisions that avoid major mistakes even at the expense of foregoing certain opportunities.
4. When wrong, to minimize the blame focused on the policy maker.
5. To maintain his policy making discretion. (Ascher and Overholt 1983:45)

The policy-maker also has specific psychological (P Perspective) needs. The first that Ascher and Overholt describe is for certainty, since apprehension of uncertainty is generally stressful. The second is for simplicity, in order to facilitate comprehension and integration with other knowledge. The third psychological need is for information that confirms the policy-maker's other views of the world.

Analysts, they continue, have a straightforward need to conduct competent analyses. This includes accurately portraying the uncertainty in an analysis, which can conflict directly with the policy-makers' need for certainty. They have personal needs for security, influence, and respect, and therefore want their information to be useful, and to downplay the importance of accuracy. Their need for professional recognition and respect may drive them to use sophisticated tools (e.g. large computer models) or innovative techniques when simpler methods would actually work better (and serve the needs of the policy makers better). The rational need of the policy-maker for a clear representation of uncertainty may run counter to the personal need of the analyst to avoid appearing incompetent.

III. TAXONOMY OF UNCERTAINTIES

Nothing is more certain than uncertainties [sic];
Fortune is full of fresh variety: constant in
nothing but inconstancy.

Richard Barnfield

III.1. Introduction

The uncertainties that decision-makers face vary in many ways that affect decision processes. There are different factors about which someone can be uncertain (*subjects*), the uncertainties possess attributes affecting how they can be treated (*characteristics*), and they arise from different underlying causes (*sources*). These peculiarities can affect the possibility or best hope of reducing an uncertainty, procedures for arriving at a decision that would be less severely affected by a particular uncertainty, or avoiding a similar uncertainty in the future.

Every representation or formulation of a decision task contains many assumptions and estimations. The dimensions of uncertainty described in this chapter characterize the varieties of assumptions that must be made, implicitly or explicitly, in any rational decision process. Ignoring the importance of these assumptions leads at best to an overly confident and narrow view of the future. Ascher and Overholt concluded from a review of forecasting efforts that the validity of the core assumptions of the efforts dominated methodological issues in affecting the accuracy of forecasts (1983:105 ff.).

It is certainly not always easy to determine a unique subject or source

of an uncertainty; very often these are multiple and mutually compounding. Still, identifying some of the sources or characteristics may yield valuable insight into ways of responding that might be helpful — or ways of responding that can be seen to be unhelpful.

The topics of this chapter are:

III.1.1. Subjects

The SDM presented in Section II.3.7 is a T Perspective view of the factors involved in rational decision-making. It is a formal model of decisions, an abstraction general enough to be applied to nearly any decision situation. However, each specific decision is made in some combination of organizational and personal contexts, and involves a great many factors that do not appear in the T view of decisions. The details of these other variables are often essential to understanding the personal and organizational significance of specific decisions, and the existence and general characteristics of the other variables are essential to understanding much about the nature of organizational and personal decision-making.

To broaden the view of decisions to include the variables and factors that are important in the O and P Perspectives, this chapter introduces a Generalized Decision Model (GDM). Many of the factors are added because decisions made in either a personal or an organizational context are often complicated by features neglected by the standard technical model. A statistician's viewpoint would require assessment of the uncertainty about all of these factors even if the only problems they presented were those of measurement and the determination of the precision (or error) of the

measurements for significance testing.

III.1.2. Characteristics

To describe more fully the nature of different uncertainties, I define three major dimensions of characteristics of uncertainties — *hardness*, *locus*, and *parametric versus strategic*. These dimensions are orthogonal to each other, not mutually exclusive. Any uncertainty can be characterized by a degree of hardness, a locus, etc. Not all of these dimensions will be relevant in any given decision situation to all of the components of the GDM (options, authority, state of the world, likelihood, consequence function, evaluation, multiple objectives, time distribution, stakeholder, gamble preference). In general, the fullest subset of these dimensions is likely to be observed in the case of uncertainty about the state of the world.

III.1.3. Sources

The next section in this chapter comprises a taxonomy of the possible sources or fundamental causes of a decision-maker's uncertainty. The source can be important in determining if or how an uncertainty can be reduced, or how to describe it.

III.1.4. Decision context

The final section describes several aspects of the context within which a decision is to be made that bear on how uncertainties can be treated.

III.2. Subjects of uncertainty

The Standard Decision Model (SDM) described in Section II.3.7 treats only the uncertainty about consequences ($c_i \in \mathbf{C}$), which stems entirely from the uncertainty about which state of the world ($w_k \in \mathbf{W}$) will obtain. Many other components of the SDM are likely to be known only

imperfectly. These possible *subjects of uncertainty* — the things a decision-maker can be uncertain about — will be described next. Then a Generalized Decision Model (GDM) will be presented which illustrates the much larger extent of the domain of uncertainty with which decision-makers must cope.

The basic question of rational decision-making, “What should I do?” (i.e. “What’s the most rational thing for me to do?”) can be divided into three parts:

1. What can I do?
2. What will happen if I make a particular choice?
3. How desirable are those results, and thence, the choices?

The first question asks what our options are. It is largely a matter of constraints and creativity. Our beliefs about the world enter into this issue explicitly as constraints on the feasibility of potential options or implicitly as constraints on our ability to think of new options, including any requirements to coordinate with co-decision-makers or acknowledge their power to veto our decisions. Our goals and values guide the development of options.

The second question focuses on our beliefs about what the world is like in order to predict the outcomes that would follow a particular choice. This requires predicting a state or set of states of the world, perhaps along with a probability distribution over the possible states of the world. Then it requires predicting what would happen if any of the available choices were made and a particular state of the world eventuated.

The third question is one of preferences and values, translating

predicted consequences into assessments of desirability, which in the SDM is represented by a *utility function* $u: C \rightarrow \Re$ ¹. The nature of the utility function is widely discussed, but except for the role that uncertainty about the state of the world plays, the decision-maker's potential uncertainty about the utility function is not. Even evaluating the desirability of a predictable, one-dimensional consequence is not always simple. Balancing different dimensions of value or consequence is more complicated.

Comparing sets of consequences that occur at different times is yet more difficult. Further, different people or organizations may have different sets of values about each of the preceding issues. Finally, there are several aspects of a decision or gamble that are functions of the whole gamble, and a decision-maker may have important preferences over some or all of these gamble features.

For our purposes here, it may be useful to divide these aspects of preferences and values into two groups. The first group has to do with predicting and valuing the consequences of one specific pairing of option and state of the world. These are aspects of a decision about which a decision-maker could be uncertain even in a “decision under certainty”. The second explicitly moves beyond a single state of the world (decision under certainty) to the whole set of states of the world (decision under uncertainty) that are believed to be relevant to assessing the desirability of

¹ This notation indicates that the function u translates consequences, the elements of set C , into real numbers. $\Re$ indicates the “real number line” that runs from minus infinity to plus infinity, although it is common to assign the worst possible consequence a utility of 0.0 so that all other utilities are positive.

an option.

All of these potential sources of uncertainty are discussed below in the elaboration of the GDM. For each, a description is provided first, usually accompanied by a diagram, followed by an illustration, and then a brief statement of how it is represented in the SDM.

III.2.1. Option uncertainty: which ones are feasible?

The set **A** of available actions is assumed in the Standard Decision Model to be known, and nearly always assumed to be finite. In many real situations, it is far from clear which actions are possible, i.e. the membership of **A** is imperfectly known, illustrated in figure 14 by the gray options. Sometimes the set of actions that is used in the decision-making process is restricted primarily by the creativity of the decision-makers. There may be additional options that would clearly be feasible if only someone would think of them. In other situations, it may be possible to list the options exhaustively but not to know which are feasible.

In some cases, it may be useful to use fuzzy set theory to describe this set: each action may have a “degree of membership” in set **A**, meaning that it has a degree of availability. For example, the legality of many acts is determined only after an attempt has been made to implement the act and another party has brought suit against the party attempting it. The capability of performing some acts may depend on technology that has not yet been proven, and that may or may not work. Many other

Figure 14.
Option uncertainty

uncertainties can arise in organizations or other groups of people (e.g. electorates). An action that is technically possible may be one for which approval by the group is impossible. The inability of Congress to withstand (perceived) public pressure and allow the ceiling on Federal salaries to rise in January 1989, is an example. While Congress had the legal authority to approve the raise, each member of Congress faced political pressure that made it undesirable to go on record as approving it. How this pressure balanced against the desire to receive the raise was not knowable until the vote was taken. Implementing some decisions can require extensive cooperation from other branches within an organization; a decision-maker may not have information about whether there will be a de facto veto from someone somewhere else in the organization.

Political feasibility can be very difficult to assess, and many analysts may wish to avoid this issue.

But, in principle, an option that is not feasible is no good, and there exists no reason to regard political feasibility constraints as less real or less important than, say, economic and technological feasibility constraints. (Dror 1988a:266)

Dror's advice is that systems analysts must be familiar with political constraints, but political assessment, work on political action, and efforts to change political feasibility should be done by others within an organization.

Illustration. Suppose that there are three options available to a utility for additional generating capability — combustion turbines, cogeneration, and conservation. Suppose now that popular sentiment against electric power companies is currently strong, and a state initiative was passed with the intention of prohibiting construction of additional generating

facilities such as the combustion turbines under consideration. The utility's legal staff, however, believes that the Public Utilities Commission (P.U.C.) still retains ultimate authority over new generating facilities in spite of the initiative. Only a court test would settle the issue one way or the other. It is also unclear whether the P.U.C. would be willing to overrule the initiative, even if it believed it had legal authority to do so.

The feasibility of the cogeneration resource is questionable. The initiative was worded vaguely, and it is not clear whether new contracts to purchase electricity from facilities that already exist by adding the cogenerating units would also be precluded.

Another complication comes from the utility's research and development staff, which reports that solar cells are so near to being cost-effective that perhaps only a few major orders of equipment would be enough to significantly increase production volume and thereby lower the per-unit costs. A drawback is that the production facility has never gotten its output rate up to the level that would be required to supply the utility. This option had not yet been considered in the utility's decision process, but some in the R&D department are still attempting to get it onto the list of candidate resource. Thus, this is an option not currently included in the set of feasible options that may actually be feasible.

Thus, determining the membership of set **A** is non-trivial: there is only one option that is firmly available, with three others that are questionable: it is not clear either that the option could be chosen, or that, if chosen, would perform as needed,

SDM Position. In the SDM, the set of options is assumed to be fixed

and known, and it is rarely given any attention at all.

Implications. Uncertainty about the feasibility of alternative actions poses a major problem for the SDM, since it assumes the set of actions is given. Simple treatments of this uncertainty are dangerous. For instance, one might think of using subjective estimates of the feasibility as weights. If an option is thought to have an 80% chance of being feasible, its estimated utility could be multiplied by .8. This misses much of the impact that non-feasibility would have, since relying on an option that turned out at the last minute to be infeasible might leave the decision-maker without any feasible plan. A closer analysis might proceed by making the simplified assumption that the question of feasibility would be settled at some particular time. Choosing this option would constitute a gamble with an 80% chance of getting the utility of the fully feasible option and a 20% chance of having to do something else.

One of the most promising strategies is to attempt to eliminate the option uncertainty. One way to do this is to acquire enough information to determine the feasibility of the options by researching the technology involved, more careful interpretation of relevant laws and regulations, polling other parties who could influence the political success of the option. Another way is to attempt to increase the feasibility of an option.

Another important strategy is to generate new options, perhaps combining features of desirable but infeasible options with those of less appealing but more feasible options. These might include options that are known to be feasible or easier to assess, or contingency options to be pressed into service if one of the original options turns out to be unfeasible.

III.2.2. Authority uncertainty: who decides?

While the decision-maker is always a unitary entity in the technical model, there is often a great deal of uncertainty concerning who will actually decide — who will propose decisions or options, who will select alternatives, who will approve or disapprove the selection, who may object, who may exercise a veto power, who will signal an end to the process, etc. The picture of the decision-maker from the SDM is shown in figure 15 as indistinct or possible multiple.

Illustration. It is often convenient to think of a utility as a single entity, but there are often different factions within any organization that compete for the authority to make decisions such as what kind of resource to obtain. It is very common for one sub-

Figure 15.
Authority uncertainty

organization to have formal authority to make the decision but for another, perhaps an executive committee, to have review authority. Other parties outside the utility may also have formal or informal review authority. Public utility commissions often have clear responsibility to approve or disapprove utility plant construction plans. Regional bodies made up of several utilities may exercise a non-binding review that may nonetheless cause the utility to reconsider its plans. Groups of concerned citizens or environmental organizations may bring lawsuits to halt a utility's plans.

SDM Position. In the SDM, the decision-maker is a single entity, typically a “rational actor” entity, with complete authority to decide.

Implications. When it is not clear whether you have authority to

make a particular decision, the problem is not to *select* properly but to *design* a course of action that will achieve desired goals. If you characterize this uncertainty as *hard* (i.e. irreducible), you can then either delay making the decision, or proceed to act as if you have the authority, and make contingency plans for the possible ways in which your authority may be insufficient. (Within organizations, folk wisdom says it is easier to seek forgiveness than permission - see section I.1.1).

If you characterize it as a *soft uncertainty*, you can make efforts to reduce the uncertainty and explicitly clarify your authority. (Seek permission.) Sometimes the most efficient way to clarify authority is to assume it and see what happens.

There may be steps you can take to increase your authority instead of merely clarifying it — cashing in favors, making more compelling statements of your reasoning, or acquiring more resources.

Designing better options does not necessarily change or clarify your decision authority, but may make it more possible to enact an option that advances your ends by reducing possible obstacles, such as others' opposition. "Win-win bargaining" is one way of seeking better options when two or more parties share authority and have non-identical interests.

Decision processes that require substantial time encounter an additional problem in that the parties involved may change over the course of the process. Some people or organizations may more or less significantly involved, and their level of influence may change. Individuals within relevant organizations may change roles.

III.2.3. State of the world uncertainty

The nature of the state of the world² in which the consequences of a decision will unfold is central to rational choice. Our information about the future state of the world can vary enormously. Figure 16 shows a set of states of the world where the membership is not completely known.

The most satisfying degree of knowledge is to know with certainty what the relevant future state of the world will be. Decisions made with this much information are said to be made “under certainty” (Knight 1921), or as “Type I” decisions.

Figure 16.
State-of-the-world
uncertainty

Implications of certainty. When the future state of the world can be predicted, a variety of analytical techniques are available for assisting decision-makers, ranging from simple lists of pros and cons to cost/benefit analysis and determination of pay-back periods or internal rates of return.

III.2.3.1. Known-odds uncertainty

The next most satisfying degree of knowledge is to know the set of possible states of the world, **W** in figure 16, and also to know a probability distribution, Ω in figure 17, over the set of states. Decisions made with this much information are called “Type II” decisions, or “decisions under risk” (in Knight’s terminology). “Risk” is generally used in a broader sense, in both the popular literature and the field of risk analysis, referring also, for

² “World” here refers to the entire environment or universe outside of the decision-making entity, not merely to our planet.

example, to situations where the possible consequences are severe but unlikely, whether probabilities are known or not. To avoid this ambiguity, this type of uncertainty about the future state of the world is called *known-odds uncertainty* here.

Illustration. Electric utility operators need to ensure that the generating units available at each moment can produce enough power to meet the utility's load, but these generating units are subject to unexpected failure. However, the industry has sufficient experience with most of the generating technologies to be able to estimate the likelihood that an individual generating unit will fail to perform when called upon. With statistical information for each unit in the system, and statistical information about the size of the load to be met, the probability that the total amount of power produced by a given set of generating units will be sufficient for the load at any moment can be calculated. This can then be compared to a desired level of reliability to determine whether the set of units needs to be augmented to increase the reliability, or decreased for better economy.

SDM Position. Known-odds uncertainty is the type of uncertainty that the SDM focuses on. Many tools, especially mathematical tools, have been developed for treating this category of situations. So attractive are some of the methods for dealing with known-odds uncertainty that a great temptation exists to treat situations of probability uncertainty or even possibility uncertainty as if they were cases of known-odds uncertainty.

Implications of Known-odds Uncertainty, or "Risk". This is the type of uncertainty that most of the literature on decision-making under

uncertainty has addressed. Knowledge of a probability distribution over the states of the world allows you to use the apparatus of MEU. Still, if you cannot know the state of the world in which the consequences of your actions will unfold, uncertainty about the consequences cannot be avoided. You must gamble. What MEU tells us is not how to achieve the outcome we desire, but rather what MEU asserts is the *best gamble*. And if you have sufficient opportunities to take the same gamble, your outcome will approach closely the MEU estimate. Probably.

So here the principles of decision analysis can be employed usefully. The problems a decision-maker may suffer from here are largely ones resulting from failing to adequately consider factors or assuming too much certainty about factors. Errors of the first type include omitting the value assessments of relevant parties, or deciding on the basis of expected value instead of the expectation of the utility of the values³.

III.2.3.2. Likelihood uncertainty

A still less satisfying degree of knowledge is to know the set of possible future states of the world, W , but not to know the probabilities of their occurrence. Figure 17 illustrates this

Figure 17.
Likelihood uncertainty

³ “Values” in the sense of levels of variables, not in the sense of valuing or ranking outcomes: expected monetary outcome is expected value; the expected utility of a monetary outcome varies with the magnitude of the money outcome for most people. A choice with a positive expected dollar value may include the possibility of bankruptcy, which might be so undesirable that a different choice, with lower expected dollar value, but a less-severe worst-case outcome, could be preferred.

uncertainty over the distribution, Ω . Knight (1921) referred to decisions made with this level of information as “under uncertainty”, but this term is quite ambiguous (especially in the context of the current work, in which virtually no decisions are seen to be free of uncertainty). In much of the technical literature on decisions under uncertainty, probabilities are assumed to be known, and therefore the decisions are “under risk” in Knight’s terminology. These decisions are sometimes called “Type III” decisions, or in Dror’s terminology (1986b), “quantitative uncertainty”. As the possible states are known but their likelihoods are not, this category will be termed *likelihood uncertainty* here.

In the standard model the relative likelihood of the possible states of the world is represented by a probability distribution, Ω , over $\mathbf{W}$. However, decision-makers rarely have a complete knowledge of the distribution, and are also rarely in complete ignorance of it. This uncertainty *about the distribution* is not captured in the distribution itself. For example, consider the opportunities to play two lotteries. Tickets for either lottery cost \$5, and each prize is \$100,000. A decision-maker knows that the odds of winning the first are 1 in 20,000, but knows only that the odds of winning the second are “slim”, which she interprets to mean that the odds are about 1 in 20,000, but could be anywhere between 1 in 1,000 and 1 in 1,000,000. She thus has enough information to know that the first lottery is a “fair game”, but cannot be sure the second lottery is fair (and would not even if she knew that the median, or the mean, i.e. the most likely, estimate of this parameter of the second lottery were also 1 in 20,000). For those potential players who care about this kind of fairness, the two

decisions are quite different.

Illustration. Suppose a utility firm plans to build additional resources, and has narrowed its choices down to four alternative resource plans. The firm is not sure what the future economic and regulatory climate will be, but has developed four descriptions of possible future environments that are believed to capture fairly well the range of possibilities. The four alternative resource plans would fare quite differently in the divergent futures. Analysts have considered costs, operational flexibility, environmental impacts, and other important attributes in assessing the over-all utility of each plan should each of the futures occur, and they have quantified their estimates of utility as shown in table 13.

Analysts have also estimated the probability of each of the future environments: there is a 15% chance that environment A will eventuate, a 30%

Table 13. Expected utility example

Original Probability Estimates				
Future	A	B	C	D
Prob.	0.15	0.30	0.35	0.20

Utilities of Future/Alternative Pairs						
	A	B	C	D	Exp.Val.	Rank
Alt. 1	0	400	100	560	267	4
Alt. 2	240	275	275	320	279	3
Alt. 3	160	210	525	75	286	2
Alt. 4	600	180	425	0	293	1

chance of environment B, a 35% chance of C, and a 20% chance of D. Using these probabilities as weights, analysts have calculated the mean or *expected* utility of each of the plans, and Alternative 4 appears to be the best.

A vice-president reviews the analysis, and notes that the analysts spent only 90 minutes estimating the probabilities, while the remainder of the

quantification effort consumed three months. She asks the analysts to refine their probability estimates by talking with outside forecasters and examining their other evidence more closely. A week later, the analysis team has finished, and its members all feel quite pleased that their original probabilities were modified by only three percentage points each: their first estimates had been amazingly close to the modified probabilities produced with a vastly greater effort.

So, the new probabilities are plugged into the expected utility calculus as shown here, and the result is a shocker: with only minor changes in the probability estimates, the ranking of the four alternatives has been turned completely upside down. Alternative 4, the apparent optimum with the original probabilities, now appears to be the worst, and Alternative 1 has moved from last place into first place, as shown in table 14.

Technical staff throughout the utility are all jarred by the implications of this. If a change of only three percentage points can upset a major analysis, how much faith can be placed in other expected

Table 14. Modified expected utility example

Modified Probability Estimates						
Future	A	B	C	D		
Prob.	0.12	0.33	0.32	0.23		
Change	-.03	+.03	-.03	+.03		
Utilities of Future/Alternative Pairs						
	A	B	C	D	Exp.Val.	Rank
Alt. 1	0	400	100	560	293	1
Alt. 2	240	275	275	320	281	2
Alt. 3	160	210	525	75	274	3
Alt. 4	600	180	425	0	267	4

utility work? How many of their other probability estimates are accurate to no more than three percentage points?

This fabricated example is clearly highly sensitive to the probability

estimates. Sensitivity analysis may be useful, though it is difficult to incorporate more than a few dimensions (Morgan and Henrion 1990:178, 181), and many electric utility decisionmaking contexts involve many distinct factors whose probabilities are estimated at least partly subjectively.

SDM Position. The standard literature distinguishes between *risk* and *uncertainty* by whether or not a probability distribution over the (known) states of the world is known. The case where no such distribution is known is often deliberately excluded from consideration. The case where partial information is available is often treated as if a distribution were known. Von Neumann and Morgenstern (1953:19) asserted that frequency-based probabilities were the only type acceptable, but Savage (1954) constructs his entire framework around “personalistic”, i.e. subjective, probabilities. In this case, a person’s estimate of the probability, based on whatever information and/or intuitions she may have, is used for the probability. There is seldom a consideration of measuring the errors of the estimate. Usually, SDM treatments assume the distribution is perfectly known by the decision-maker, or that an experimenter knows it and tells it to an experimental subject, who then also knows it, or that an experimental subject knows the distribution at least implicitly and the estimates of that distribution produced by the experimenter’s elicitation procedures capture it accurately. In virtually all of these cases the distribution is thereafter treated as if it were known exactly. Sensitivity analysis in general, and over probability estimates in particular, is more often recommended by application-oriented authors (e.g. Morgan and Henrion 1990:172ff.) than

by those with an experimental or more purely normative focus.

Implications of Likelihood Uncertainty, or “Quantitative Uncertainty”. This type of uncertainty is at the heart of a major controversy over the legitimacy of the use of subjective probabilities in decision analyses or system analyses⁴. The following two views are extreme characterizations of two poles of attitudes. Many authors believe that we cannot really know anything about future occurrences, especially about possible one-of-a-kind events. They assert that in the absence of frequency data we should consider ourselves to be completely uncertain about the likelihoods of the possible futures. The use of subjective probabilities is considered to be deceptive and harmful. In the other camp, those who hold the subjectivist view of probabilities accept subjective probabilities as capturing what subjective information and knowledge we have. In most of the literature of this camp, the subjective probabilities, however elicited, are treated as if they are exact quantities.

Neither of these positions seems reasonable. In virtually every imaginable case, we ought to be uncertain to some degree about our assessment of the likelihoods of the possible states of the world. The important questions are how uncertain are we about Ω (the estimates of probabilities), and how do we indicate that meta-uncertainty? This question calls for a measurement of the *likelihood uncertainty*. Statistical techniques often incorporate measurement of uncertainty, but techniques for eliciting subjective probabilities do so much more rarely. Using

⁴ See section II.3.8 for a discussion of the legitimacy of subjective probabilities.

estimates from experts may result in multiple estimates of Ω , and it is not clear how best to integrate them and yet reveal the inexactness of the resulting figures.

The point emphasized here is that while there are many bases for developing probabilities⁵, the uncertainty attending any development of probability estimates needs to be assessed. In the case of logical probabilities, such as the probability of rolling a six with a fair die, the uncertainty lies not in the logical reasoning, which is mathematically exact, but with the presumption that any particular die is indeed fair, an empirical matter than can be settled only within a margin of error, not with complete certainty. In the case of frequency-based probabilities, the uncertainty stems from the use of samples instead of whole populations, or the possibility that the system, accurately measured yesterday, has changed today, etc. In the case of subjective probabilities without the support of logical or statistical bases, we would expect greater uncertainty about the probabilities — but this, too, should be assessed carefully in each case.

An approach endorsed by the subjectivists when the likelihoods of possible states of the world are not known is to use whatever information is available, however informal, and produce subjective estimates of the likelihoods of occurrence, and proceed with techniques developed for situations of *risk*. Delphi methods can be used to synthesize probabilities from the judgments of a group of people who do not need to be gathered

⁵ See section II.3.8 for discussion of logical, frequency, and subjective probabilities.

together physically. It would be very prudent to include some estimate of the uncertainty surrounding those subjective estimates in an assessment of the sureness of the conclusions drawn from them.

However, the results mentioned in section II.3.8.3 of McCord and De Neufville (1983) indicate that subjective estimates of probabilities are quite dependent on the method used to elicit the probabilities. As the example in section III.2.3 indicates, expected value calculations can be very sensitive to differences in probabilities. Together, these findings reveal that there is a strong possibility that results based on subjective probabilities will be unduly influenced by the choice of elicitation method, which is a form of arbitrariness that should be avoided.

The *principle of insufficient reason*⁶ could be employed to develop probabilities, but the difficulties of properly partitioning the space of possible futures into sub-spaces makes this so arbitrary the results are likely to be more misleading than helpful.

If time permits, there may be forms of research that could narrow the range of probabilities of each alternative state of the world. Research may also uncover relationships between known precedent conditions and the future state of the world allowing the state to be predicted. For example, a probability distribution over the annual precipitation in the Columbia River basin can now be estimated from an historical record of over 50 years. There are indications, however, that it is correlated with El

⁶ The *principle of insufficient reason* suggests that when there are n possible outcomes whose probabilities are unknown, they can be treated as if each has a 1-in- n chance of occurring.

Niño/Southern Oscillation (ENSO) conditions many months earlier. If this relationship is verified and refined, it may make it possible to narrow the range of possible precipitation for the coming year once the ENSO condition for the prior summer has been learned.

Game theory offers several decision rules that require knowing only what the possible states of the world are, not their likelihoods. Maximin recommends the alternative with the most attractive worst-case outcome (the highest *security level*); maximax selects the alternative with the most attractive best-case outcome (the highest *hope level*); minimax regret exposes the decision-maker to the smallest risk of being second-guessed from the *ex post* perspective.

III.2.3.3. Possibility uncertainty

When our information is truly meager, we may not even know which states of the world are possible (e.g. Dror 1986b:26; Smithson 1988:87). Dror suggests terming decisions of this sort “Type IV”, or “qualitative uncertainty” to distinguish it from likelihood uncertainty, which he would call “quantitative uncertainty”. Since it is the possibility of various states about which we are uncertain, *possibility uncertainty* will be used here.

Possibility theory and fuzzy set theory suggest a richer way of characterizing what we know about the set of possible states of the world. When **W**, the set of possible states of the world, is conceived of as an ordinary (“crisp”) set, imaginable states of the world must be considered to be either possible or impossible. If **W** is treated as a fuzzy set, imaginable states of the world are considered to have varying degrees of membership in the set of possible states. This corresponds to treating imaginable states

of the world as having different degrees of possibility, ranging from 0.0 to 1.0 inclusively.

Illustration. Possibility or qualitative uncertainty becomes a more prominent issue as time horizons are stretched forward in time. Suppose a utility believes that an important factor in its investment planning is the nature of the dominant power generation technology 20 years hence. Surveying the current technologies, an analyst might propose that the candidates are nuclear power, standard coal plants, exotic coal plants (fluidized bed, etc.), natural gas combustion turbines or cogeneration. Probabilities for each of these could be estimated.

But 20 years is a long time — long enough for new technologies that are barely in prototype now to be in production; long enough for global warming to have inspired political commitments to cleaner plants; long enough for many currently operating nuclear plants to face retirement and replacement. How big a breakthrough in solar cell technology would be required to cut the cost in half and make it competitive with thermal generation? How soon will commercial fusion power be possible? We don't know whether the current technologies will still be possible, and we don't know which new technologies may have become possible.

SDM Position. In the SDM, the only uncertainty facing the decision-maker is about which of the *given* states of the world will occur. The SDM does not treat possibility uncertainty. Mathematical machinery in general, such as that developed and employed in economics and decision analysis, requires *data* for its operation. Possibility uncertainty reflects a more fundamental *lack of data* than do the previous types of state of the world

uncertainty, and therefore offers less purchase for the levers of rational analysis.

Implications of Possibility Uncertainty, or “Qualitative Uncertainty”. In the case of uncertainty about which future states of the world are even possible, the use of probabilities is very hazardous. Probability is *distributional* in that a total of 100% of probability is to be allocated across the possible alternatives. If it is believed that a surprise alternative may develop, or more than one, how much of the 100% is to be reserved for the surprises? If it *is* a surprise, how can its likelihood be estimated? Possibility theory offers some promise here, as it is non-distributional (e.g. two mutually exclusive states can each be 100% possible).

In this situation, the greatest risk may not be that of choosing the wrong gamble from a set whose risks are known, but that a surprise may occur that is much worse. Therefore, contingency planning can provide extra safeguards. Brainstorming activities to develop more expansive sets of possible future states of the world can also reduce the risk of surprises.

One possibility that is often considered to be outside the realm of decision analysis is developing better options.

Ignoring possibility uncertainty keeps many possible surprises from showing up in our pictures of the future. We are likely to be more certain of our actions than we ought to be, and less well prepared to deal with unexpected developments.

III.2.3.4. The non-discrete world

Many shades of uncertainty are hidden by this partition of

uncertainties into neat categories. "Decisions under certainty" may be highly uncertain: we may know for sure that the state of the world's temperature dimension tomorrow will not be cold, cool, or warm but actually hot; but does that mean 95° or 105°? Even if the state is defined and known precisely, there are many other ways in which a decision-maker can be uncertain.

Likelihood uncertainty is clearly not simply present or absent; we typically know more than nothing about the relative likelihoods of various states of the world, but rarely know the likelihoods exactly. Similarly, our knowledge of the possibility of various states is a matter of degree.

Table 15. Probability and possibility uncertainty

		Possible future states of the world		
Likelihood of occurrence of future states of the world		Unique	Multiple, known	Multiple, not known
	Known	Type I Decision under "certainty"	Type II Decision under "risk"	Meaningless
	Partially known	N / A	Mixed; not treated in the literature	Mixed; not treated in the literature
	Unknown	N / A	Type III Decision under "uncertainty"; known- odds uncertainty	Type IV Decision under "qualitative uncertainty"; possibility uncertainty

Information we have about the future state of the world varies in descriptive power. The most powerful is to know which state of the world will eventuate. This would be a subset of the lower left quadrant in figure 18, where the probabilities are not only known but have the characteristic that they are

Figure 18.
Possibility and probability uncertainty

all zero except for one state with probability of unity. The next most powerful is to know the probabilities of the possible future states of the world. A less powerful but more attainable description is to know the possibilities of the set of possible future states of the world. If it turns out that there is only one possible future state of the world, then we know the possibilities and the probabilities, which describes our state of knowledge, and further, the world about which we have this knowledge is unusually and conveniently simple.

It does not make sense to conceive of a situation in which we know probabilities but do not know possibilities — knowing probabilities requires more information. Any event that is probable at all is certainly possible:

For any event x , $\text{poss}(x) \geq \text{prob}(x)$.

Thus, one quadrant in figure 18 is vacuous.

III.2.3.5. Implications of probability or possibility error

Although this work focuses on uncertainty and generally neglects false certainty and error, the distinction between likelihood uncertainty and possibility uncertainty can be illuminated further by contrasting the effects of errors made in assessing these two uncertainties.

Implications of probability error or false likelihood certainty.

What if we are falsely certain or in error about knowing the probabilities of future states of the world, i.e. if we think we know them and we either cannot know them or should know that they are different from our belief?

In this case, we are presumed to know which states of the world are possible, so we can determine the set of possible consequences attending each decision option. Any of these consequences may arise no matter what the odds are, so knowing the odds correctly does not prevent or guarantee any particular outcome. Error means we may select a gamble that with better probability information we would not select. This doesn't mean we would be worse off, since the outcome of any gamble is uncertain.

Observing the consequences of a one-time gamble does not tell much at all about the accuracy of our probability estimates. Neither does it tell us whether we made a good choice — only whether we made a lucky one. The consequences of erroneous likelihood assessment cannot easily be distinguished from the consequences of being lucky or unlucky in one-shot decisions — a state of the world eventuates, and it may or may not have been the most likely one or the one anticipated. In the case of repeated decisions, erroneous assessments of likelihood will cause the total outcome

of the series of decisions to be probably worse than it could have been, with a probability that increases with the number of the decisions.

Implications of possibility error or false possibility certainty.

What if we are falsely certain or in error about knowing the possibilities of future states of the world? There are two different ways this could come about. We might consider a state of the world, and erroneously judge it to be impossible (an error of judgment, region # 1 in figure 19), or we might fail consider it at all (an error of imagination, region #2).

In either case, there may be consequences of some decision options of which we are totally unaware. In the first case, depending on when the state was assessed to be impossible, we may have considered the state and even consequences ensuing

Figure 19. Possibility uncertainties

from it, while in the second case, not having even conceived of the state, we could very well be caught totally off guard: "To be blind-sided" means to suffer a blow (originally, a blow of a fist) coming from one's blind side, so that there was no warning at all. This type of event typically carries with it not only the physical or logical consequences of the event but also a psychological shock which can make responding to the other consequences more difficult.

Table 16. Error in assessing possibility uncertainty

Re- gion	Pos- sible?	Be- lieved Poss?	Imag- ined?	Description	Consequences
1	√	-	√	Falsely rejected	Potential surprise
2	√	-	-	Never considered	Potential big surprise
3	√	√	√	Properly considered	OK
n/a	√	√	-	meaningless	—
4	-	-	√	Properly rejected	OK
5	-	-	-	Purely abstract	OK
6	-	√	√	False hope False fear	Disappointment Wasted preparation
n/a	-	√	-	Meaningless	—

The two regions labeled “n/a” have no meaning, since they refer to states that are believed possible, but have not been imagined. Regions 3 and 4 contain states that have been imagined, and are either both possible and believed possible, or neither. In these two cases, the possibility assessment is correct, and no harm results. Region 5 has not been imagined or believed possible, and is not possible. States in this purely abstract region have no consequence. Region 6 is significant in that hopes or fears about outcomes involving states in this region are likely, but the states are not actually possible. The possible harm from states in this region amount to disappointment and the loss of any resources used in preparing for states that are not possible. It is clear from this assessment that the greatest risk lies in regions 1 and 2, where either insufficient imagination or erroneous determination of the impossibility of states can lead to being unprepared for and surprised by the future state of the world.

III.2.4. Function or model uncertainty

We may be uncertain about the consequences ensuing from a given

action combined with a given state of the world (i.e. we may be uncertain about the function $c: \mathbf{A} \times \mathbf{W} \rightarrow \mathbf{C}$ and therefore be uncertain how to calculate or obtain $c_{jk} = c(a_j, w_k)$). This uncertainty is illustrated in figure 20, adapted from the SDM. We may not have adequate observations yet to permit deduction of the function c in general, or perhaps either some action a_j under consideration or some possible state of the world w_k is different enough from those we have already observed that we cannot be sure how c behaves in the vicinity of a_j or w_k . Our knowledge of the relationship between options, states of the world and the ensuing consequences constitutes a *model*, whether the model is as explicit as a highly specified computer model, or as implicit as the hazy understandings about how things work that a child carries around in his head, whether the model is as abstract as the equations of general relativity theory or as tangible as a relief map.

In many cases, what is needed is a sufficient understanding of the functional aspects of the system so that the behavior of the system in the future can be predicted. In the case of a system whose behavior can be observed repeatedly, it may be possible to do without this detailed structural information and to model it as a black box. The mapping of options onto consequences is then made by observing past behavior. This is valid only when it is safe to presume that the system will not have changed between the time of the previous, observed behavior and the time

Figure 20.
Function or model
uncertainty

in the future when a prediction of the behavior is needed.

Illustration. BPA is publicly committed to enhancing the viability of populations of fish in rivers that have hydroelectric dams on them. The fish are affected by both the policies that regulate operation of the reservoirs (actions) and the amount and distribution of rain and melting snow flowing in the river (the state of the world). Thus, fish viability (fv) is a function of policies (p), stream flows (sf), other factors ($o...$), or $fv = c(p, sf, o...)$. The policies are endogenous or decision variables: BPA (with other parties) has control over them. The stream flows are exogenous variables: they are outside the control of BPA. The “other factors” include water quality as affected by run-off from agricultural or lumbering activities, etc., fishing activities in international waters, and other such factors outside the control of BPA. Even for a specific policy and specific pattern of river flow, estimating the impact on fish stocks is very difficult. We just do not know enough about what affects fish survival to be able to estimate the impacts of policy choices very well. For policy analysis, it would be very convenient to be able to create formal models of fish survival as a function of policy options, but we are currently very uncertain about the correct form such a model would take.

Another instance of model uncertainty involves the amount of precipitation received in a year. This determines the amount of “fuel” available to drive the hydroelectric generators, which in turn is a major determinant of BPA’s power sales revenue and power purchase costs for the year. The precipitation pattern has been observed carefully for 50 years, which provides a fairly good base of data for estimating how likely it

is that next year will be wet or dry. This prediction is not based on a fundamental understanding of the weather system, i.e. of what factors influence annual precipitation — we are not very close to understanding the weather patterns that well. It is a black-box treatment that describes past behavior and projects it probabilistically into the future, based on historical data. It assumes that the weather system will be the same in the future as in the past. That assumption is probably useful but is certainly open to question.

For one thing, climate studies have shown that there are periods of varying length during which the average temperature (or rainfall) is higher (or lower) than the years before or after, and we do not know enough to predict the onset or length of these periods. More recently there has been great concern that the amount of CO₂ that human activities are releasing into the atmosphere will warm the Earth and thereby change the weather system significantly. But we don't know how it will change it, especially with regard to a specific region such as the drainage basin of the Columbia River, and we do not know how fast any changes could possibly occur. Different climatologists make predictions that in some cases vary enormously from each others', and reconciling them is highly problematic.

Position of the SDM. The SDM assumes that the function c mapping choices and state of the world into consequences is perfectly known, implying the adequacy of mental or other models of the world.

Implications. Function uncertainty poses a serious challenge to rational decision-making, since its presence means that we cannot adequately know what the consequences of choosing an option might be. In

extreme cases, it makes any knowledge of the likelihood of future states of the world irrelevant, and even complicates decisions under certainty, as even with only one state of the world possible, we still cannot predict the consequences of our actions. It renders many mathematical techniques as powerless as does possibility uncertainty about the future state of the world.

Most analytical methods for supporting decisions rely on a single model of how the world works that is assumed to be reasonably accurate.

Techniques based on this assumption are inappropriate when the consequences ensuing from a world state and an option cannot be inferred.

Thompson and Tuden's taxonomy (1959), described in section II.9.2, suggests that when there is agreement in an organization on both the consequences and the desirability of a choice, organizations typically use computation to make decisions. Being able to predict the consequences requires both understanding the function that maps choices and states of the world onto consequences, and being able to predict the state of the world. When the consequences are not predictable, Thompson and Tuden suggest judgmental decision-making is commonly preferred. The SDM deals with the two types of prediction problems quite differently. When the consequences of a pairing of choice and state of the world are known, the SDM prescribes that computation is the method that ought to be used. When it is the function predicting the consequences (of a known choice and a known state of the world) that is unknown, the SDM is silent — it assumes this function is known. Thompson and Tuden's view is descriptive while the SDM is normative; it is possible that computation would replace

judgment as the preferred approach when sufficiently capable computational methods become available.

Implications of Soft Function Uncertainty. Reducing function uncertainty may be possible through research. Data gathering can reduce this type of uncertainty when other characteristics of the function are known, for example, if the relationships between states of the world, options, and consequences are believed to be linear but the coefficients are not well-enough known. When less is known about the function, theory construction efforts may help to build more accurate and useful models; acquiring more data can assist here also.

Structural modeling tools (e.g. ISM, KSIM) may be helpful in making people's intuitive understanding of the important dynamics more explicit. They allow approximate models to be created quickly so that the overall responses of the models can be compared to what is known of the responses of the world. The value of these tools is not in providing more information, but in helping to clarify and possibly formalize information people have, and to allow proposing and examining structural and functional hypotheses.

Implications of malleable function uncertainty. Function uncertainty is most likely to be malleable when the system whose dynamics are unclear depends on human behavior, and order can be imposed by regularizing the behavior. Rules within organizations, families, and schools, laws and norms within society, and voluntary agreements among individuals or groups serve to reduce function uncertainty by regularizing behavioral responses.

Implications of Hard or Virtually Hard Function Uncertainty.

When function uncertainty is not reducible within the time needed for a decision, the possibility of unexpectedly bad results is important to guard against. Alternatives that allow for frequent course corrections, that facilitate learning more about the relevant dynamics, and that preserve options for the future will be especially valuable.

Function uncertainty is difficult, perhaps unappealing to address, but important:

Uncertainty about the form of a model is generally harder to think about than uncertainty about the value of a quantity. But experienced analysts often argue that uncertainty about structure is usually more important, and more likely to have a substantial effect on the results of the analysis. (Morgan and Henrion 1990:67)

Detailed, data-driven models embodying many assumptions about how the world works are inappropriate when there is significant uncertainty about which assumptions are valid — models are too likely to be fundamentally wrong. “Assumption drag”, the danger that core assumptions will be out of date, and the possibility that core assumptions are simply wrong are the two most significant causes of forecasting errors (Ascher and Overholt 1983:106). Most uses of probability are contraindicated under function uncertainty, although probabilities used as subjective weights for different *models*, not states of the world, may be helpful. Human judgment applied with as wide a spectrum of information as can reasonably be obtained may be the best method for choosing in many cases.

III.2.5. Evaluation uncertainty: how desirable is any specific consequence?

Predicting consequences requires an understanding of the world; determining the value or meaning of a specified consequence requires an understanding of the values of one or more beings as illustrated in figure 21. Normative approaches to rationality generally leave decision-makers great latitude for choosing their values. Rational decision-making requires comparing the desirability of different consequences, so the specification or assessment of values is an important part of the process. Even in the simplest case in which only one dimension of consequence, e.g. money, is deemed relevant, this is not necessarily trivial, as the marginal utility of money and other goods typically varies with one's total supply.

For instance, while it is commonly assumed that the utility of an additional dollar declines with the total dollars a decision-maker possesses, the precise rate at which that marginal utility declines can only be assessed through a particular procedure, and no available procedure can be safely assumed to be infallible.

Figure 21.
Evaluation uncertainty

The distinction between assessing consequences and assessing the utility of those consequences demands careful attention. Economic studies often implicitly assume that utility is linear with money and then omit assessing the function mapping money to utility, but the following example

illustrates briefly why that is insufficient.

Illustration. A very important issue for BPA that recurs regularly is whether to propose raising the rates charged to its customers for BPA's electricity. Beginning a rate case means committing massive resources (e.g. personnel and computer time) to the project, and the process is highly contentious, as there are many other parties with strong interests in BPA's rates. BPA is required to set rates that balance its costs and revenues, so a surplus or shortfall of revenue will necessitate a rate change if it falls outside of some range of flexibility (which itself is not known precisely). Within the range of financial flexibility, assuming BPA's utility to be linear with money is probably a satisfactory approximation.

There is a major discontinuity, though, in the vicinity of the bounds of that financial flexibility range. It may not matter much whether BPA ends the year with a surplus of \$x, \$2x, or a shortfall of \$x or \$2x compared to its target, as long as no rate case is required ("x" is somewhere on the order of 20,000,000 to 40,000,000; more when BPA's reserves are high at the beginning of the year). But if a shortfall of -\$3x would trigger a rate case, then the difference between -\$2x and -\$3x is extremely important. Increasing the amount of shortfall further would eventually lead to the likelihood of suspending repayment of loans owed to the U.S. Treasury; that suspension would have highly negative political consequences for BPA, so there is another discontinuity in the utility function. An economic assessment of a proposed plan that aggregates predicted money values instead of utility values for different years and under different states of the world is blind to these discontinuities. Two alternatives with the same

expected money result may actually expose BPA to quite different risks of surpluses or shortfalls that would be large enough to trigger a rate case or undesirable political impacts.

SDM Position. It is assumed in the SDM that a decision-maker is able to make judgments about the relative desirability of any two possible consequences (i.e. the decision-maker has a *complete preference structure*). Without this assumption, the framework of expected utility analysis would be seriously weakened since its comparison of alternatives depends on the ability to compare any two consequences. Yet even a brief appraisal of the matter reveals that comparing consequences can be very difficult, and a person or organization may well be uncertain about such comparisons. How desirable is an additional \$1,000 compared to an additional \$5,000, or \$1 million compared to \$5 million? How desirable is an electrical system reliability of 1 outage in 10 years compared to a reliability of 1 in 20 years?

Although their formal exposition of the expected utility hypothesis largely ignored it, Von Neumann and Morgenstern recognized this difficulty:

It is very dubious, whether the idealization of reality which treats this postulate [the completeness of the preference structure] as a valid one, is appropriate or even convenient. (von Neumann and Morgenstern 1953:630)

Implications of Evaluation Uncertainty. Uncertainty about the desirability of consequences is important to reduce for rational choice: as the old saw has it, if you don't care where you're going, any road will do. If basic aspects of what is important to the decision-maker are unknown, values clarification techniques may help. Exercises to visualize in detail what various consequences would be like may make the decision-maker's

preferences clearer. For organizations, examination of what its basic values or basic mission are may reveal that a consequence is either more or less important than it seemed at first.

Often what needs to be evaluated to make a decision is not an absolute measure of the desirability of any consequence but the difference in the desirabilities of two or more consequences. Sometimes changing the level of detail of the descriptions of consequences helps to understand how significant the difference is. For example, when stated to 7 digits, the difference to a middle class citizen between winning \$1,000,000 in a lottery and winning \$1,101,235 is \$101,235, a substantial sum. But characterized verbally, the two prizes are a million dollars and a little more than a million dollars, and the difference in desirability of the two consequences does not appear as great. This is especially useful when the figures themselves are imprecise or known to be subject to errors.

The desirability of many goods is non-linear with respect to the good, and usually the desirability of an additional unit of the good declines with the quantity already obtained. For example, we would expect that a gift of \$1000 would mean much more to a pauper than to a millionaire. For this reason, it is important to evaluate differences in absolute consequence level after a decision, not only differences in magnitude of change. The starting position before a decision can make a big difference in the importance of a change of a given magnitude along one dimension of consequence.

The axioms of the Expected Utility Hypothesis require that the decision-maker be able to compare all possible consequences but von

Neumann and Morgenstern recognized that individuals may experience evaluation uncertainty:

The common individual, whose behavior one wants to describe, does not measure his utilities exactly, but rather conducts his economic activities in a sphere of considerable haziness. (1953:20)

This haziness can be enough of an obstacle to comparing the desirability of options as to make MEU ineffective or inappropriate.

III.2.6. Multiple-objective uncertainty: what is the relative importance of different dimensions of consequence or value?

Another uncertainty becomes important when more than one type of consequence is important, and the utilities of the different types of consequence have to be aggregated into a single estimate of the utility of that outcome. Most non-trivial decisions result in two or more dimensions of consequence, such as cost, production volume, environmental impact (multidimensional in itself), popularity, option-enhancing or option-reducing, etc. There may be dimensions with unknown degrees of importance, as illustrated by the gray shapes in figure 22.

The import of each dimension may depend on different values. Two or more basic values may imply different choices, and it is frequently difficult to compare these value dimensions. This is the multiple criteria decision-making problem, which virtually precludes treatment by optimization methods, since optimization requires defining a single criterion for maximization (or minimization).

Figure 22.
Multiple objective uncertainty

Illustration. Analysts of the electrical power system of the Pacific Northwest are facing increasingly difficult problems of multiple objective comparisons. Even for strictly power-related attributes of the system, such as reliability and cost, the comparison is difficult. The recent declaration of several species of salmon as threatened or endangered brings in other objectives to be compared: survival of these fish stocks is now a very important objective, and activities that support that goal conflict in many ways with actions for power production and for other uses of the river.

SDM Position. The SDM assumes that decision-makers are able to make any necessary multiple objective comparisons accurately, as discussed above under evaluation uncertainty.

Implications of Multiple Objective Uncertainty. It is rare to know with certainty the relative importance of different dimensions of consequence, and an active field has developed for investigating ways to tackle this problem. This can be a major cause of uncertainty even in decisions made “under certainty” (i.e. the future state of the world is known). There are at least two major approaches to increasing the decision-maker’s understanding of the relative values. The approach typically taken by the decision analysis literature is to posit that the decision-maker has a utility function which translates combinations of different consequences to a one-dimensional utility scale, and then to attempt to reveal the utility function. The decision-maker is presented with a series of forced choices between combinations of consequences. Each answer provides more information about the relative preferability of varying quantities of different dimensions of consequence.

Advantages of the utility function elicitation approach:

- If the decision-maker's preferences are stable, the utility function thus revealed can be used in future decisions;
- The resulting utility function is explicit and can be examined by others, perhaps adjusted with additional information from other parties;
- Will eventually indicate a numerical superiority for one alternative if pushed hard enough;

Disadvantages of the utility function elicitation approach:

- If the decision-maker's preferences are likely to change with time, the utility function must be updated; there is a danger that it will not be, and will be used misleadingly;
- The nature of the utility function is dependent on the particular choice of elicitation method, so utility functions elicited differently are not comparable, and there is doubt as to the accuracy of any elicitations;
- Apparent numerical superiority may be an artifact of the process used; uncertainty surrounding rankings of alternatives may be difficult to include and display;
- The value comparisons and trade-offs are often made early in the process, perhaps by people quite removed from the decision-maker, and may be difficult to review, revisit or revise.

The second approach does not aim to describe the decision-maker's utility function completely, but only to make comparisons among members of a specific set of alternatives. It does not implicitly assume that the decision-maker's utility function already exists in complete detail, needing only to be "revealed". One implementation of this approach is the trade-off methodology developed by the Electric Power Research Institute (EPRI Project #2537 and later projects) (Burke et al. 1988). After the alternatives have been scored on all the dimensions of consequence by whatever means

are available (not necessarily by computer models), the decision-maker is asked to describe two quantities for each dimension. The first is *NSW* (not significantly worse): if two alternatives differ by no more than *NSW* on a dimension of consequence, then neither is significantly worse than the other in that respect. The second quantity is *MB* (much better): if one alternative scores at least *MB* higher than a second on one dimension, then it is regarded as being truly superior on that dimension. One alternative is then considered to be better than another if it is not significantly worse than the second on any dimension, and if it is much better on at least one. All of the alternatives that are not dominated are collected in the decision set. These can be compared to determine how much of one benefit would have to be traded off to gain a particular improvement in another dimension of consequence (as indicated by the set of alternatives that have been evaluated).

This method will not indicate which alternative is better when one is much better on one dimension of consequence but a second alternative is much better on another dimension. However, it highlights that as an issue, and makes it obvious that what the decision-maker would really like to have is an alternative that combines the best of each.

Advantages of the trade-off approach:

- Complete elaboration of the decision-maker's utility function is not necessary;
- Highlights the trade-offs between dimensions of consequence implicit in the alternatives that have been evaluated;
- Determinations of the value of each dimension of consequence can be made at the end of the decision process, are explicit, and can be adjusted fairly easily (for instance, if the decision-maker's

values change);

Disadvantages of the trade-off approach:

- Does not contribute much to resolving comparisons among alternatives that are each clearly superior along different dimensions of consequence;
- Information about the decision-maker's utility function is obtained only for those magnitudes of consequences that are represented in set of alternatives that are evaluated; little is learned that might be useful in future decisions about magnitudes of consequence greater or smaller than any those associated with the evaluated alternatives.

The problem of comparing different dimensions of consequence is not easily amenable to rational solutions. It is inherently subjective.

Rationality can be powerful in choosing *means* to achieve given *ends*; determination of the ends to be pursued is outside its purview. This muteness of rationality applies also to balancing one end against another.

One role that rationality can play usefully is to help guard against employing procedures that appear to be arbitrary. A *support end* can be defined: the purpose of decision support in this context is to assist the decision-maker in making value trade-offs, and to ensure that decisions are made on the basis of actual values rather than on results that are artifacts of processes or implementations. In this way, reason can help to prevent *irrationality* and *hyperrationality* in the choice process even when rationality itself is mute regarding the specific outcome of the choice.

As Zeleny reminds us, "the importance of generating imaginative alternatives cannot be overemphasized." (1982:101) The best way for someone shopping for a car who experiences a conflict between buying a car that accelerates well but guzzles gas and buying a economical but

boring car to escape the dilemma is to find a car that is *both* fast *and* economical. Fuel injection offers precisely this possibility — it presents a different — and more attractive — trade-off curve than carburetor technology, and allows improvements in both power and economy at the same time. (The decision-maker will still be faced with the power/economy trade-off, but can achieve more desirable results on both scales.)

III.2.7. Time distribution uncertainty: how are consequences occurring at different times to be balanced?

The next complication involves the distribution of consequences through time: valuing consequences occurring today relative to consequences occurring tomorrow, a year from now, or a generation from now, or farther in the future, as suggested in figure

There are serious moral issues surrounding the actions we take today that will have consequences for future generations who cannot express any preferences now. Some argue for a heavy weighting on future consequences by saying that we have not inherited the

Figure 23.
Time distribution uncertainty

earth from our parents, but rather are borrowing it from our children. The moral position of an actor making choices that will affect only herself is simpler, but if the consequences will last 30 years, weighing the consequences is still difficult. Her value system is likely to change over the

course of those 30 years in ways she cannot predict as she acquires or loses resources and goes through experiences that shape or change her values.

In quantitatively oriented organizations, time comparison of consequences is often handled by using a discount rate⁷ to weight the importance given to future consequences, especially dollar consequences. Some who are concerned about the morality of imposing an irreversible negative impact on future generations without their consent suggest increasing the prominence of those impacts within today's analyses by using a negative discount rate. Even when most consequences can be measured in dollars, the desirability of the consequences must still be assessed carefully, since an option that has a high net present value in dollars may have dollar consequences in a single future time period that translate into value or utility consequences that are intolerably severe. If the utility of each time period's consequences is not assessed, this intolerable consequence could be melded in with the consequences from other time periods and be overlooked.

A major pathology in time comparison is weakness of will or *akrasia*, an important hazard for decision-makers who would be rational (Elster 1984). Individuals may sometimes choose an action that they know is not optimal for them because they over-value some very short-term consequence, such as sensory gratification or avoidance of anxiety. Organizations may over-

⁷ Let r represent the discount rate, often taken to be the cost of capital. Then the discounted value of a dollar received or paid out n years in the future is $(1-r)^n$. A positive discount rate reduces the importance of future consequences; a large positive rate reduces it more.

value immediate political impact even while having some awareness that it is at the expense of the organization's long-term interests.

Illustration. Suppose two different power plants could supply a utility's incremental power needs over the next 20 years. Plant A is expensive to build, but cheap to operate, while Plant B is cheap to build but expensive to operate (this corresponds to a historical contrast between coal plants and combustion turbines). Suppose the construction costs are \$2 billion and \$1 billion, respectively, and the annual operating costs are \$120 million and \$200 million. Federal agencies often use 3% as the discount rate for calculating the net present value of future costs, but private firms may use discount rates corresponding to their cost of capital. The lower discount rate accords a greater importance to future costs. The net present value over 20 years of these two plants, under both a 3% and 10% discount rate are as shown in table 17.

Table 17. Effect of discount rate

	Capital Cost	Annual Operating Cost	N.P.V. over 20 years, using 3 %	N.P.V. over 20 years, using 10%
Plant A	\$2,000,000	\$120,000	\$3,800,000	\$3,000,000
Plant B	\$1,000,000	\$200,000	\$4,000,000	\$2,700,000

These figures illustrate that when a low discount rate is used, future costs are important, and it is worth while to spend money up front so as to pay lower future costs. When a high discount rate is used, Plant B, with the smaller initial outlay, is preferable.

SDM Position. Von Neumann and Morgenstern avoided this issue in their exposition of the Expected Utility Hypothesis by assuming all events occur at the same time (1953:19). When the SDM is broadened to include

consequences in different time periods, the consequences are usually folded back to the present by using a discount rate to calculate a net present value.

Implications of Time Distribution Uncertainty. Although many arguments have been put forth, the question of what discount rate ought to be used is unlikely to be settled permanently, as it hinges on the relative weight one wishes to attach to present and future consequences. This is an inherently subjective decision, again outside the domain within which rationality can be of primary assistance. It can still be helpful in facilitating a political or personal decision that resolves the time distribution uncertainty.

The most common decision pathology, especially for individuals, is over-weighting of near-term consequences (as judged by similar but less closely involved decision-makers, or as judged by the same decision-maker at a different time). Immediate consequences are the easiest to imagine, lending them a greater psychological salience (the *availability heuristic*, section II.5.3.2.). Deliberate attempts to visualize or imagine more distant consequences and how desirable they would be at various times in the future may help to counter this tendency. American businesses are also accused of this type of short-sightedness as they face pressure to produce quarterly profits for stakeholders and are influenced by B-school grads with techniques of calculation that appear to reveal the short-term consequences of alternatives while saying little about the longer term.

Political arrangements, e.g. an agreement to adopt a particular discount rate for investment decisions, can help to tame this issue, though

they do not settle the matter of what the “right” discount rate is. Moral positions can serve the same purpose, such as a principle that people should leave their children a better world than they themselves inherited. A belief that the world is so uncertain and hazardous that survival beyond one year is unlikely, and therefore consequences farther in the future than a year do not matter, would settle the question with equal dispatch, though some might characterize this as an *immoral* position. Interestingly, Horace’s dictum, *Carpe diem, quam minimum credula postero*⁸, is often cited not in support of ignoring future consequences but to encourage action today, suggesting that the future carries great risks even for those who do not act.

The issue is compounded by the general pattern that more distant consequences are much harder to estimate than nearer ones. This makes the comparison of near and distant consequences even more difficult because the uncertainty surrounding the estimates of the consequences differs enough that it is itself hard to compare.

III.2.8. Stakeholder uncertainty: whose values are relevant?

Whether a decision is to be made by a single person (or entity) or by many, the values used to guide that decision may be those of a single person (or entity) or those of many parties, as suggested by the modification of the SDM shown in figure 24. It is often difficult to know whose values ought to be used in making a given decision. Even if that

⁸ “Seize the day, put no trust in the morrow!” [Beck 1980]

question is settled, the question of precisely how to combine the values of the different parties can be very difficult. If others will bear the consequences of the decision, the decision-maker may be under additional obligation to act conservatively.

Even when each member of a group has a structurally identical value system, each will usually give more weight to consequences affecting his own situation. More commonly, even the value systems will differ. There are large literatures on the subjects of *social choice* and *collective*

Figure 24.
Stakeholder uncertainty

rationality. The technical results obtained so far, such as Arrow's Impossibility Theorem⁹, are discouraging to those who hope for an analytically neat solution. It does not appear that rational methods can provide us with a generalized answer for how to combine values of multiple stakeholders. Utility functions by themselves, even if perfectly known, do not solve this problem, since the meaningfulness of interpersonal comparisons of utility is debatable (e.g. Barry and Hardin 1982:265).

Illustration. There are many stakeholders that BPA may consider when making a decision. Although it is required to be self-financing, it is still a part of the Federal government, and an argument can be made that

⁹ Arrow described five quite plausible conditions we could require for any procedure for obtaining a "social preference" from the set of individual preferences, and then showed that any procedure (e.g., majority voting) must violate at least one of the conditions. See, for example, the essays in the second half of Barry and Hardin's *Rational Man and Irrational Society* [Barry and Hardin 1982].

it ought to use the values of the entire population of the United States in making decisions. It is in many respects, however, a regional organization, and another argument can be made that it is the values of the residents of the Pacific Northwest that should be used. But BPA does not serve all residents: it sells electricity wholesale to retail utilities, largely public power entities, and therefore it may be that the values of the people served by BPA's utility customers should drive BPA decisions. There is growing acknowledgment that major west coast utilities in British Columbia, the Pacific Northwest and California may be able to increase the coordination of their operations and create win-win arrangements: perhaps the values of residents or consumers of electricity up and down the west coast ought to govern BPA's decisions. To some extent, the values of BPA's top management will steer BPA's decisions, but even that term oversimplifies the matter, since "top management" includes different people who may have different value structure.

SDM Position. The SDM assumes that a unique, stable value structure exists and is known, or that the relevant value structure exists in a single person.

Implications of Stakeholder Uncertainty. Once again we come upon a type of uncertainty about which rationality can say relatively little. The question of how to weigh the interests of different stakeholders against each other is an essentially political one. Various principles can be put forward, some of which can lay stronger claims to the moral high ground than others. Aside from pointing out logical internal conflicts within such principles or other forms of irrationality, reason must yield

center stage to political processes, while possibly contributing by helping to improve the structure of the processes, identifying and correcting flaws. Voting, consensus seeking, lobbying, negotiating, law suits and military force are some of the ways that the interests of multiple stakeholders can be aggregated. In the U.S., pure force is generally considered to be an unacceptable method. Since the founding of the country, “no taxation without representation” has stood for the principle that those who are affected by a decision ought to have a voice in the decision, but how to arrange that is a difficult question. Ironically, one of the preeminent products of rational thinking on this subject, Arrow’s Impossibility Theorem, proves that not only must reason yield center stage sometimes, but in the case of combining values of multiple stakeholders it must yield to social choice processes that can be proven to be at best less rational than we would desire.

Stakeholder uncertainty is amplified when a decision process occurs over an extended period of time, since the parties whose values are relevant may change, and the values of those parties may also undergo shifts.

III.2.9. Gamble preference uncertainty: what features of gambles (lotteries) are desired?

There are several features of gambles about which a person or organization may have preferences which bear on which option “ought” to be chosen. (For my purposes here, a gamble is a decision or lottery with two or more options, at least one of which has uncertain consequences. To simplify the discussion, I will assume that a probability distribution is known over the possible outcomes associated with each option.) It is my

contention that a gamble is more than the sum of its parts, and that many of the ways this is so cannot be represented in the SDM or other models based on the Expected Utility Hypothesis. The preferences discussed next cover features of the gamble itself, not features of the particular outcomes comprising the gamble.

Implications of Gamble Preference Uncertainty. These features are global properties of the choice with its entire array of possible states of the world and associated consequences, or of the set of options available in a decision. MEU is based on evaluations of consequences of each outcome (of a single option under a single state of the world), which are then combined solely via probabilities into a weighted average; it cannot consider characteristics of a whole gamble. Preferences over outcomes do not predict preferences over risk, certainty, security level, hope level or gamble evaluation methods.

A major reason for considering the following types of gamble evaluation uncertainty is that in many organizational and individual decisions the outcome is not known for quite a while, and this pre-outcome period may itself present consequences to the decision-maker that are important. Since its exposition by von Neumann and Morgenstern, MEU has treated decisions as if only two moments in time mattered: a first moment in which the decision is made, and a second, later moment in which all consequences occur (Von Neumann and Morgenstern 1953:19). This prevents consideration of the consequence of the level of anxiety, hope, anticipation, etc., during that period in MEU calculations.

(The importance of pre-outcome consequences for gamblers could be

tested by comparing comparable gaming situations with varying lags between the gambler's choice and learning of the outcome. Slot machines, for example, have levers that are pulled; similar machines could be constructed that responded immediately to the insertion of a coin, or that have longer delays between pulling the lever and the payoff or lack of payoff.)

III.2.9.1. Risk

People have widely varying preferences for or against risk. The *American Heritage Dictionary* (1982, Houghton Mifflin, Boston) provides a simple definition of risk, adequate for this brief treatment:

risk *n.* 1. The possibility of suffering harm or loss; danger.

If a gamble were changed so that the possibility of harm increased, it would become riskier¹⁰; if the possibilities remained the same, but the magnitude of the possible loss increased, it would again become riskier. If the possible outcomes diverge more widely from each other in desirability, the prospect is riskier.

A person is said to be *risk averse* if, for example, she would rather have \$30 for sure than to have a 50-50 chance between \$100 and nothing. The expected dollar value of the gamble is \$50, but having a chance to take that gamble is worth only \$30 to her. Her utility curve, then, would be concave (concave downward), as shown in figure 25. Other people are willing to pay more than the expected value of a gamble for the

¹⁰ Although as the likelihood of a bad outcome increases and it becomes not merely possible but highly probable, we would say that the choice becomes not merely risky but bad.

opportunity to play it; casinos and state-run lotteries depend on such *risk seekers*. People may be risk averse for some types of outcomes and risk prone for others, and it is rare to be perfectly certain about assessments of such risk preferences.

Figure 25. Sample utility curve

The risk of an option is sometimes measured by the variance of the utility. The expected utility of an option is the *first moment* of the distribution, and the variance is the *second moment*.

Expected utility or mean utility of option a_i :

$$\text{Exp}(u(a_i)) = \bar{u}(a_i) = \sum_k p(w_k) * u(a_i, w_k)$$

Variance of the utility of option a_i :

$$\text{var}(u(a_i)) = \sigma^2(u(a_i)) = \sum_k p(w_k) * (u(a_i, w_k) - \bar{u}(a_i))^2 = \text{Exp}(u^2) - (\text{Exp}(u))^2$$

Illustration. Many of the customers of BPA state that they want BPA's rates to be stable and predictable, since they must pass on to their customers any increase in BPA rates, and their customers hate large rate increases. Actions that BPA takes that have more predictable but mediocre financial consequences (e.g. frequent but small rate increases) might then be preferred to actions that are riskier in that they might have either very good (no rate increases) or very bad financial consequences (sudden, large increases).

SDM Position. The SDM employs the concept of the utility curve to explain risk aversion or risk seeking. The utility curve of someone who considers \$30 and the 50-50 lottery for \$100 or \$0 equally attractive could be calculated thus: assign a utility of 0 to \$0; choose an arbitrary value for the utility of \$100, say 1.0. Then the utility of \$30 equals the utility of the lottery, which by the expected utility hypothesis equals the expectation of the utilities of the outcomes, or $.50 * 0.0 + .50 * 1.0 = .50$. Therefore, the utility, for this decision-maker, of \$30 is .50, and her utility curve would look like the illustration above.

Implications of Risk Preference Uncertainty. *Riskiness* is a global characteristic of an option in a decision. How to decide what level of risk is desirable? When similar decisions have been made often enough in the past, it may be possible to examine them and their results to determine if the level of risk tolerated in those decisions appears to be appropriate in hindsight.¹¹ In many business fields such examinations have produced general guidelines, in some cases being as formal as specific rules developed by regulatory agencies that oversee the decisions of organizations. Individuals may be able to identify events from their own lives that are sufficiently numerous, or events from the lives of others that are sufficiently similar to the decisions they face, to be able to develop a sense of how much risk to tolerate.

¹¹ This is different from critiquing an individual decision in hindsight, which has been argued against here, in that it concerns evaluation of a decision *policy* — the level of risk accepted — applied to a series of decisions believed to be similar enough to comprise a sample of observations about the success of the policy.

The Expected Utility Hypothesis (in specific, the 5th axiom in Hey's presentation of the axioms, described in the Context chapter) precludes any specific utility or disutility for gambling (Hey 1979:30). This "moral neutrality" of risk, as Vickers puts it, is questionable:

Some men take more chances than others. This is in some a general trait and in others restricted to actions of a certain sort. The most important distinction to make here is that between those risks which affect mainly the agent (mountaineering, gambling for money, fighting when one is dishonored or threatened) and those which significantly affect others (the decision of a physician about how to treat a patient, investing others' money, fighting in a revolution). Men who take too many risks of the first sort are foolhardy, but not in general evil. There seems to be no name for men who take too few risks; we may say that they are too conservative. A man who takes too much or too little risk of the second sort is evil, unwise, or weak. (Vickers 1975:302)

III.2.9.2. Certainty

Risk as described above is a function of both the desirability of consequences and their likelihood or possibility. A person may have preferences for or against certainty that do not depend on the desirability of the outcomes. Jon Elster's preference against knowing exactly when he would die (quoted in Section II.7) is not based on his estimation of the desirability of lives of 51 years, 52 years, etc., but rather on his distaste for this particular certainty.

Some people like surprises, some don't, altogether aside from any risk. If an option has two possible outcomes, and they are equally desirable, then there is no *risk* in the classical sense, since the utility of the two outcomes are the same, and therefore the utility of the outcome is known. That is not quite the same as the utility of the option, since there may be a pre-outcome period during which the decision-maker experiences

consequences. For instance, it may be that a process cannot begin until the outcome of a risky choice is known; this would be a disadvantage of the uncertainty itself, not belonging to any of the possible outcomes.

Illustration. In the time before a presidential election, a utility (i.e. its managers or board of directors) may believe that while the candidates have different positions on matters that affect the utility, no candidate is preferred. The utility could adjust equally well to any of the possible presidential positions. In the technical sense, there is no risk, since the desirabilities of all the possible presidential policies are equal, and there is no reasonable possibility of a very good or very bad outcome. However, the utility may have a very strong preference for knowing which candidate will win the election¹², since any of the possibilities will require adjustments in the utility's strategy, and these adjustments take time.

In 1992, petitioners in Oregon secured enough signatures to place on the ballot a measure that would, if successful, close the state's only nuclear power plant. BPA did not declare a preference for or against the measure, and indeed it is unclear which outcome would have been preferable for BPA. In the time before the election, however, the uncertainty surrounding the fate of a plant that generates a significant fraction of the electricity used in the state posed a problem for the operators of BPA's system that was distinct from risk.

Individuals are more likely than organizations to prefer to be

¹² This example simply illustrates a preference for certainty, not the possibility that the utility may be able to know ahead of time or to determine the outcome of the election.

uncertain. Suppose that two people are about to order a meal at a restaurant where the chicken and salmon are both excellent and equally enjoyed by our hypothetical diners. One person may tell the waiter to surprise him, and not let the diner know ahead of time which he will bring. Such a preference for surprise (i.e. for uncertainty) appears no more irrational than the preference for certainty displayed by his companion who insists on declaring his own choice, even though he values both possibilities equally. In this case, the surprise or the certainty are valued for their own sake, and are not properties of the outcomes (chicken versus salmon).

SDM Position. There is no room in the SDM for preferences for or against certainty. The utility or disutility of uncertainty generally occurs after a choice has been made and before the uncertainty has been resolved. A diner may enjoy looking forward to a chicken dinner, remembering previous chicken dinners, and generally enjoying the anticipation. Von Neumann and Morgenstern explicitly ruled out utility or disutility of this sort:

To avoid misunderstandings let us state that the “events” which were used above as the substratum of preferences are conceived as future events so as to make all logically possible alternatives equally admissible. However, it would be an unnecessary complication, as far as our present objectives are concerned, to get entangled with the problem of the preferences between events in different periods of the future. It seems, however, that such difficulties can be obviated by locating all “events” in which we are interested at one and the same, standardized, moment, preferably in the immediate future. (von Neumann and Morgenstern 1953:19)

Thus, they have ruled out deriving utility during the pre-outcome period when anticipation of the outcome could be experienced. Unless we

are prepared to assert that such preferences are inherently irrational, a comprehensive model must allow for them.

Implications of Certainty Preference Uncertainty. Certainty, or its lack, is another global property of a decision option. The utility of any pre-outcome costs of uncertainty should be assessed. If a decision-maker is uncertain about this preference, discovering his preference by experiencing certainty in some choice situations and uncertainty in others may be helpful.

It would be surprising to find an establishment organization with a preference for its own uncertainty. The benefit of surprise is typically an emotional one, and emotional outcomes are not “supposed” to influence organizational or business decisions. However, there may be strategic reasons to choose an option with a more uncertain outcome so that a rival or adversary could not discover the outcome as easily. Anti-establishment organizations may have a preference for surprise, such as a campus radical group or a terrorist “crazy state” (Dror 1971).

A preference for or against certainty will make a difference only when one of the options considered has a certain or virtually certain consequence and at least one of the others does not. A common result of inadequate testing of assumptions is that some options may be assumed to have consequences that are more certain than they really are, and thus those options are overvalued (undervalued) due to the decision-maker’s preference for (against) certainty.

III.2.9.3. Security level and hope level

Let us call the worst possible outcome for an option the option’s *security*

level, and the best possible outcome its *hope level*. A person may value having a high security level, or having a high hope level, in addition to any preferences for or against risk and uncertainty. A person concerned with worst-case outcomes, who wants to ensure that no matter what the outcome is, it will be tolerable, may place a high value on the security level. On the other hand, someone may want to have at least a chance of coming out exceptionally well. This provides grounds for feeling hopeful about the future. The hope level and security level of a gamble depend only on the desirability of best and worst outcomes; they are affected by neither the likelihoods of the outcomes nor by the desirability of intermediate outcomes.

Illustration. The state of Oregon runs a lottery named the “Daily 4” in order to raise money. If it is a money-maker for Oregon, it must be a money-loser for the players taken as a group, and therefore the mathematically expected return for each player is less than the price of the ticket. In fact, the expected payoff for a \$1 bet is only \$0.4605. For most people, the utility of money declines with increases in money, so that the utility of a \$1500 payoff would be less than 1500 times the utility of a \$1 payoff. Thus, arbitrarily choosing the utility of \$1 to be 1, the expected utility of the lottery ticket would be even smaller than .4605. And yet many people play it. Are they all risk prone? In order for the expected utility of the payoff to exceed 1, the player’s utility function would have to be an (at least partly) increasing function of the quantity of money, which is atypical. Indeed, the existence of a specific preference for gambling is usually considered to be ruled out by the axioms of expected utility theory

(e.g. von Neumann and Morgenstern 1953:632; Hey 1979:30).

An alternative explanation is that the players want to have some hope that their position could improve enormously. As one regular purchaser of lottery tickets explained it to me, "You can't win if you're not in the game."

Similarly, many people spend money on insurance, though this is a money-winner for the insurance industry and therefore a money-loser on average for purchasers of insurance. The usual explanation is that purchasers are risk averse (at least on the loss side of the balance sheet) — the disutility of a large uninsured loss is so great that the expected utility of a certain but small loss (the insurance premium) is preferred to the expected utility of foregoing coverage. However, wanting to have a high security level can explain the behavior equally well. A major difference in the two explanations is that the expected utility argument examines only the utility of consequences after the veil of uncertainty has lifted, while the security level argument incorporates the preferences over different states or possibilities before the veil is lifted.

The safety risks of nuclear power are interpreted in different ways by the people at large, who are quite scared by it, and by the engineers, who calculate the "expected loss of life" as being quite low relative to other power sources. One explanation of the divergent evaluations is that people do not perform expectation calculations in their heads, but simply ask, "What's the worst that could happen?", i.e. "What is the security level?" In the case of nuclear power, the worst that could happen is very bad indeed, even though it appears to be extremely unlikely. (Another explanation is that number of deaths is not an accurate indication of the disutility of an

accident, and that the public views 1,000 deaths in a single event as being more than 100 times as bad as 10 deaths in a single event. If this is the case, then the engineers were falsely certain about the evaluation of consequences.

SDM Position. These preferences, like a preference for or against certainty, are ruled out by the SDM and the Expected Utility Hypothesis, as they are properties of decision options, not of consequences.

Implications of Security Level, Hope Level Preference

Uncertainty. When there is a preference for a high hope level, it is usually desired for emotional reasons, and therefore it is uncommon for organizations to place a high value on hope level. For individuals, however, hope can be an essential component of what is considered to be a good life. While some may play lotteries to enjoy the feeling of risk, knowing that their fate is currently unknown, others play because owning even one ticket provides them with reason to hope they could win. It's not uncommon for people to put off opening a letter that would inform them of their results in a contest, a school application, or job hunt, because until they read the letter, they can hope. After reading the results, they may have good news or bad, but the certainty of their knowledge deprives them of hope.

Security can be valued for both emotional reasons and for prudence. A preference for a high security level can be part of a satisficing approach to screening options. If there is a certain level of utility of outcome that is thought to be the minimum level acceptable, all options whose worst outcomes are worse than such a level can be discarded. Beyond that, if an

option's worst possible outcome is at least as good as the level of outcome that would be considered "good enough", choosing such an option would assure the decision-maker of an outcome that would be at least "good enough" and may be even better, and that is often a very comforting assurance. The computational burdens of such an approach may be far lower than those of an expected utility maximization method.

Differences in preference for security level may explain why the lay public disagrees with the results of some expert risk analyses. For example, in the nuclear power example just cited, engineers calculate that the expected loss of life from nuclear power is quite low compared to, say, coal-fired power. However, the worst that could happen with coal power is something on the order of scores of fatalities, while the worst that could happen with nuclear power is thousands of fatalities, many times more injuries, and permanent poisoning of large parts of the earth, perhaps of whole communities. Therefore, the coal power option has a much higher security levels. Measures such as the security level or hope level that do not depend on probabilities are especially attractive when a person believes that probabilities over the possible states of the world are poorly known or unknown, a condition that renders expected utility results much less meaningful.

III.2.9.4. Regret

Loomes and Sugden (1982) developed Regret Theory (described in Section II.5.5) to accommodate the common human experience of feeling either "rejoicing" or "regret" about a previous risky choice after the consequences of the choice are observed and the actual choice looks either

good or bad in retrospect. Though their attention was focused on an individual, psychological experience, their theory can be applied to any anticipation of a post-outcome re-evaluation of a choice and the effect that such anticipation would have on the making of a choice. Even a decision-maker who would not experience any psychological regret or rejoicing may find that anticipating the re-evaluations of others is rational when their regret-based actions could be significant, since the decision-maker would experience what we could call “strategic regret”. The amount of regret, psychological or strategic, varies greatly across decision-makers and situations. The fact that it is anticipation of a future regret, and therefore impossible to predict precisely, further complicates the weighting of regret in any decision.

Illustration. Some forms of regret are hard to imagine for an electric utility. In particular, the *emotion* of regret does not seem relevant to a corporate context. However, there are considerations of the potential *ex post* responses of other parties that enter a corporate decision-maker’s deliberations in ways that are functionally equivalent to an individual’s anticipation of possible regret. No matter how rationally an organization behaves, and how thoroughly it adopts an *ex ante* perspective in both making and evaluating its decisions, it may be unrealistic to expect that all of the parties whose criticism may be relevant will do the same. The organization’s critics may say, “Look how badly choosing X worked out: you should obviously have chosen Y.”

BPA’s customers are very sensitive to any increases in the rates the agency charges them. BPA’s revenues are very sensitive to the total

annual precipitation in the Columbia River drainage basin. In wet years, the hydro system can produce up to 50% more electricity than in dry years. While this “nonfirm” energy that cannot be counted upon in dry years does not fetch as high a price as the “firm” energy, it can still bring in hundreds of millions of dollars in a wet year, a very substantial amount compared to BPA’s annual budget of about 3 billion dollars.

If BPA were to increase its rates in the expectation that its revenues would not otherwise cover its costs, and then, after that decision, find that the revenues from nonfirm sales were high enough that it ended the year with a large financial surplus, customers could be very upset. BPA’s leaders would be able to claim that at the time of the rate decision, the revenue forecasts were so low that foregoing a rate increase would have been irresponsible. Everyone willing to look at that decision from the *ex ante* perspective would be inclined to agree. But nothing requires BPA’s customers to adopt that perspective, or in fact, to behave according to any particular precept of rationality. And BPA’s managers know this at the time they make rate decisions. They must ask, “What would happen if we raise the rates and then unexpected (unexpected!) good fortune arrives, making the rate hike appear unnecessary?” Thus, potential regret enters into the evaluation of rates no matter how strictly rational BPA’s decision processes are.

SDM Position. As with the other features of gamble evaluation, regret has no place in the Standard Decision Model. Utilities of consequences are to be calculated on an individual basis, not by comparison of one consequence to others. After the utilities are calculated, the utilities of the

consequences are compared. Some variations do permit calculating utilities by comparing each possible outcome to the decision-maker's current asset position, but not with the other possible outcomes.

Implications of Regret Preference Uncertainty. Potential regret is a function of both a particular option and the set of alternative options. While Jon Elster may reject the use of the *ex post* perspective in judging the quality of a decision (Elster 1989a:117), even a decision-maker who is rationally convinced that he made a good choice may still experience the emotion of regret if the “rational” choice works out badly. A decision-maker who can predict that he would feel regretful may be considered non-rational not to take this into account when making a choice. I believe this is a personal choice about how to feel after the results of a risky choice are discovered, and I don't think it can be said that those who factor in the anticipation of possible regret are less rational than those who don't — they are making different choices about how to feel should particular circumstances develop.

How much weight should be given to potential regret depends on many factors that can generally be estimated only via subjective judgment, such as the kinds of actions others may take if they are unhappy with the eventual consequences of the decision-maker's choice, and how important the decision-maker deems these actions to be. Uncertainty over how much weight to give to possible future regret is generally a higher impediment to making a decision than even a very large but known preference against potential regret.

Experience from previous decisions may allow determination of

whether the decision-maker or others might experience regret which would matter to the decision-maker. When there are procedures established for subsequent review of decisions, it may be possible to infer from the procedural rules whether the review will be made from an *ex ante* perspective, thereby precluding regret, or not. One technique that apparently works for some to eliminate uncertainty about regret preference is to resolve to have no regret — the “Never look back” philosophy.

When regret is experienced, people may attribute its cause to their particular choice or their poor luck (or both), since the outcome they find unpleasant resulted from a combination of a particular choice and the eventuation of a particular state of the world. Attributing the outcome to luck probably leads to lower inner turmoil, and is consistent with Elster’s advice about evaluating decisions from an *ex ante* perspective¹³.

III.2.9.5. Gamble evaluation rule

It is assumed in many discussions of rationality that maximizing expected utility is the goal of rational decision-makers (see Economic Rationality Rule # 3 from the Context chapter, section II.3.2.2), especially in the technical literature of economics and decision analysis. The issue of what decision rule to use in evaluating gambles then never arises. However, other decision rules, such as maximin or maximax are legitimate options, and satisficing, which is not a decision rule at all, may be

¹³ This endorsement is made on the assumption that an *ex ante* assessment reveals no reason available at the time of the decision to have selected a different option.

preferable. Not everyone automatically assumes that MEU ought always to be used.

Illustration. Suppose BPA needed to decide on a strategy for operating the Federal dams on the Columbia River, and was interested in several dimensions of consequence: power production, total cost, total revenues, recreational opportunities, survival of several fish species and also numbers of surviving fish (individuals). In order to use maximum expected utility as the decision rule, each option would be evaluated to determine the possible outcomes and the likelihoods of those outcomes. Then each possible outcome would be evaluated along each of the dimensions of consequence, and finally, the utility of each specific type of consequence would be assessed. This is likely to be a highly controversial procedure.

An alternative would be to use the maximin rule, and to determine for each option what its worst possible outcome would be. Then these worst outcomes would be compared, and the option with the least-bad worst outcome would be the chosen alternative. This might involve far fewer comparisons, and thus far fewer specific opportunities for controversy. This is not necessarily a better way to make the decision, but it is a quite legitimate way. Especially in the presence of large or varied uncertainties, decision rules other than MEU may be very useful; the importance in MEU of being able to make arithmetic comparisons means that uncertainty in the quantities being compared can render it indecisive.

SDM Position. The SDM is especially clear on this point: maximization of expected utility is the only rational rule for evaluating

gambles.

Implications of Gamble Evaluation Rule Uncertainty. Of the many rules for selecting one of the options in a risky choice, maximization of expected utility has received the overwhelming endorsement of the analytical community since von Neumann and Morgenstern's original exposition in 1944. Considering that von Neumann and Morgenstern themselves pointed out ways in which their axioms misrepresented how they believed people actually made decisions (1953:19, 20, 30, 630), and that many others have also pointed out aspects of risky decisions about which decision-makers may have preferences that cannot be represented in the MEU model, the dominance of the MEU perspective is quite a testament to the appeal of mathematical formalizations.

Some objections are fairly general:

One of [the Bayesian approach's] main implications is to eliminate the maximin criterion in decisions under uncertainty (i.e. total ignorance). In the first place, we are told, there is no such thing as total ignorance; and in the second place the only use for the maximin criterion is in the solution to zero-sum games where the environment (i.e. the other player) can be predicted with total certainty. Instead of the maximin criterion it is then recommended that we use the expected-utility criterion, or maximize the expected utility on the basis of some probability distribution defined over the set of alternatives. Of course no one has ever denied that this is a good criterion in many cases. The novelty of the Bayesian approach is to argue that it *should* always be used because it *is* always possible to define a set of subjective probabilities ... I would like to question this step from 'is' to 'ought'. (Elster 1984:128)

The mechanics of MEU are very attractive, in that they provide a mathematical approach to identifying *the best* decision option among those that are available. It is tempting to conclude that maximization of expected utility is the decision rule that should be used to make all risky

decisions. However, the prescriptive clarity of the principle of MEU is ineluctably muddled by data gathering procedures that can never be perfect, and by choice situations that differ from the simplifying assumptions von Neumann and Morgenstern made. MEU is considered in more detail in later chapters.

III.2.10. A Generalized Decision Model (GDM)

Below figure 26 shows the full generalized decision model which illustrates a wider variety of factors about which a decision-maker may be uncertain. These subjects of uncertainties are summarized as follows:

- Option uncertainty: How well is the feasibility of each option known? How well is it known that the set of options is as complete as it should be?
- Authority uncertainty: How well is it known who is involved in making or vetoing the decision?
- State-of-the-World uncertainty: What are the possible future states of the world? How well known are the possibilities of imagined states of the world? What are the relative likelihoods of the states of the world? How well known are these likelihoods?
- Function or model uncertainty: How well known are the dynamics of the system? How well is it known what will happen if a particular option is chosen and a particular state of the world occurs?
- Evaluation uncertainty: How well known is the desirability of a particular kind of consequence?
- Multiple objective uncertainty: How well known are the relative desirabilities of different kinds of consequence?
- Time distribution uncertainty: How precisely known are the relative importance of consequences occurring at different times?
- Stakeholder uncertainty: How well is it known whose values are relevant to the decision and are to guide the decision?

- Gamble preference uncertainty: How clear are the relative values of risk, certainty, security level, hope level, regret, and the

Figure 26. A Generalized Decision Model

method of evaluating gambles?

- Decision context: discussed later in this chapter.

III.2.11. Beyond decision models

Even the GDM includes major assumptions about the decision situation that may not hold. It assumes that a decision situation has been identified; often there may be uncertainty about whether a decision opportunity exists at all, or at the other extreme, whether a decision challenge must be faced or if it could be avoided. It assumes a particular framing of the decision, and there may be many perspectives on the matter, each with its own version of the options, the major task, the roster of relevant players, etc. What one party may conceive of as a single decision may appear to others as a sequence of decisions or a continual process.

While the GDM explicitly acknowledges uncertainty about the existence or feasibility of options, it does not address in any way how decisions are made about whether to generate more options or how to generate them. It says nothing about implementation of a decision, which may affect the outcomes more than the particular decision does. In general, decisions about whether to generate more options must be satisficing rather than optimizing decisions. Optimization requires that the options are in a sense “out there” already and can be surveyed more or less simultaneously. When each option must be *created* instead of *discovered*, the question of whether the process has generated enough options yet is one of satisficing, not optimizing.

The state of the world and the set of options appear to be separate in the GDM, but there may be many linkages. The making of a decision, in fact, the mere public intention by one party to make a decision concerning a specific subject, may change the state of the world. The state of the world

can never be described completely, and can only rarely be described accurately even along the dimensions that are known to really matter, and of course, there may be dimensions that are not now or even cannot now be known to be as important as they turn out to be.

Time plays several important roles in making of decisions, and these are difficult to represent in the GDM. The context within which the decision will be made can also affect how a decision ought to be made. These issues are discussed later.

III.3. Characteristics of uncertainties

This section describes *characteristics of uncertainties* that affect the options decision-makers have about dealing with uncertainties as they make decisions. Uncertainties about different decision factors (“subjects”; see section III.2) or arising from different sources (see section III.4) may have the same characteristics, and vice versa.

III.3.1. Hardness

The first distinction in this Taxonomy, between *hard* and *soft* uncertainty, hinges on the possibility of gathering more information and thereby reducing the uncertainty (Dror 1986b). Soft uncertainty can be reduced by our actions, e.g. sampling, polling, measuring, or some other research effort that yields more information about the world. Hard uncertainty is *in principle* irreducible — either there is variability or indeterminacy in the actual situation, not only in our knowledge of it, that we cannot reduce, or else it is known to be impossible to improve our information. While efforts to reduce a soft uncertainty may or may not pay off, efforts to reduce a hard uncertainty are wasted. It isn’t just that we

don't know, but that even a "perfect witness", knowing all facts and dynamics, would say "I know that I cannot know" (Dror 1988a:250).

To Dror's two types a third and fourth may be added, which I will call *malleable* and *virtually hard*. Malleable uncertainty can be reduced by our changing the world, rather than by learning more about it. Virtually hard uncertainty is a type that is potentially reducible, but not by any means that are presently available. There may be hope of developing an effective measurement or prediction technology, there may exist one that is too expensive to be worthwhile, or that would take too long when a decision must be made quickly — it is not irreducible in principle, but for now it is must be treated as if it were. These four categories are illustrated in figure 27.

The current quality of the output of a production line and the opinions of a political public are soft uncertainties for which statistical sampling techniques are well developed. These techniques provide not only a reduction in our uncertainty, but mathematical measures of the remaining uncertainty, and the residual uncertainty can usually be reduced to whatever level is thought to be worth paying for. A soft uncertainty whose reduction via research is considered to be too expensive becomes a virtually hard uncertainty. Some uncertainties will resolve themselves over time: if we can afford to wait for that natural resolution, we can treat the uncertainty as soft; if not, and if we have no means of accelerating the resolution we must treat them as if they were hard.

The distinction between soft and virtually hard depends on one's current information-gathering technology and resources, while the

distinction between virtually hard and hard depends on the nature of the world. The Laplacian universe as conceived of in the Age of Reason was one that contained no hard uncertainty — the entire universe was potentially completely knowable. The assumption was that we would proceed steadily towards total knowledge, converting virtually hard uncertainties into soft uncertainties and then into certainties. Early in the twentieth century the Heisenberg Uncertainty Principle shocked holders of that view with the claim that a limit existed to our ability to acquire information simultaneously about the location and velocity of sub-atomic particles, and that the limit was absolute and independent of the capability of our instruments. More recently, Chaos Theory has again challenged that assumption, suggesting that *sensitive dependence on initial conditions* renders the behavior of some systems inherently unpredictable (Gleick 1987). If the global weather system is chaotic in this sense, then no amount of information gathered today would suffice to predict the details of the weather seven days from now — next week's weather would be a hard uncertainty.

Figure 27. Hardness of uncertainty

There may be many other reasons for our inability to reduce an

uncertainty. A health board, for example, may be prohibited by law from requiring tests for the presence of AIDS antibodies in a certain population, and be restricted to far less effective means of inferring the rate of infection, which thus makes the rate a virtually hard uncertainty. The origin and nature of a particular uncertainty may be deemed to be beyond the bounds of the area in which an organization wishes to invest its research resources, or such research may be considered to be too expensive to pursue.

The term *malleable* is used here to refer to an uncertainty that can be reduced by action that changes the situation itself, not merely our knowledge of it. An industry group might be able to reduce the uncertainty it faces about the severity of possible Federal safety regulations by adopting self-imposed moderate regulations and thus forestalling government action. Preemptive military action is an attempt to convert a situation of dangerous but malleable uncertainty to a situation that is both more advantageous and more certain by changing the world.

Not all useful attempts to modify the real world and thereby advantageously modify a malleable uncertainty will *reduce* the uncertainty. Consider a situation that is likely to be resolved in either of two ways, one more preferable to the decision-maker than the other. The decision-maker believes that the poorer alternative is slightly more likely, and acts to change the world. Her action may change the situation to the effect that the better alternative is then slightly more likely. While the decision-maker's prospects would clearly be enhanced by such an effort, the amount of uncertainty has not been reduced — there are still two main

alternatives, one slightly more likely than the other.

The uncertainty in a great many situations is mixed — partly hard, partly soft or partly malleable. Thus, it may be possible to reduce the uncertainty to a significant but incomplete extent, after which the residual uncertainty must be treated as hard. As Dror points out, the significance of determining that an uncertainty is truly hard is that efforts to improve prediction or other means of acquiring information are pointless, and use of subjective probabilities illusory or delusory (Dror 1986b:26). Such uncertainty cannot be reduced, and we must live with it. A virtually hard uncertainty, however, may be amenable to reduction at some time in the future. Unfortunately, determining when or if this may be possible may itself be difficult or impossible.

Implications of malleable uncertainty. Reductions in malleable uncertainty are made not by gathering information about the world but by actions intended to change the world, and efforts toward that end are more *design* than *research*.

For example, if the variability of a key aspect of manufactured objects is too large, a revised manufacturing process could be designed to produce more uniform output and thereby reduce the output uncertainty. If the future behavior of other actors introduces problematic uncertainty, it might be possible to devise agreements with the other actors that would provide additional certainty.

Implications of soft uncertainty. As with malleable uncertainties, an important, soft uncertainty should be reduced, if the cost of reducing it is smaller than the benefit. (This comparison typically must be made

judgmentally.) If the costs are too great, then the uncertainty must be treated as virtually hard. The nature of the endeavors likely to help reduce soft uncertainty depends on the uncertainty. Generally they are types of research, such as gathering larger volumes of information in the same way smaller quantities have been gathered, developing better measuring techniques or instruments, querying or surveying people, experimentation, demonstration projects, reviewing literature, etc.

When attempting to determine if an uncertainty is soft or virtually hard, it can be important to assess how much reduction in uncertainty would be required to be significant within each particular context. Expensive attempts to reduce uncertainty that leave unacceptably high levels of residual, hard uncertainty are not worth while. For example, a utility may find that its traditional net present value methodology is unable to identify the better of two alternatives for acquiring more generating capability because some of the factors are so uncertain that the estimates for the two options are not significantly different. If the uncertainty for several of the key variables could be cut in half with additional research, but the uncertainty in the bottom line would still prevent identifying which alternative is better, the research would be a poor investment.

Implications of virtually hard uncertainty. Virtually hard uncertainties differ from hard uncertainties in that the latter are known to be irreducible, and the former are merely being treated as hard because ways of reducing them are not known or are too expensive or time-consuming. Circumstances can change, especially with technological

innovations, and many uncertainties that were once virtually hard are now soft. Their reduction may become possible due to increased computing power, expanded technological capabilities (e.g. more uniform manufacturing methods, or a jump to a higher performance curve), more sensitive instruments (e.g. better optics, or ultra-sound), people becoming more sure of their values, agents becoming more willing to be interviewed, longer observation yielding more data, better theories, political agreements, new regulations, economic changes (e.g. cost reductions, or the collapse of monopolistic pricing), memory training or aids, and on and on... Virtually hard uncertainties should not be assumed to be irreducible forever, and those that are repeatedly or continually important should be checked from time to time to see if any new developments have made their reduction practical.

Implications of hard uncertainty. Ah, yes: what to do about hard uncertainties! Well, you can't reduce them. You can only clarify them, and design and make choices as intelligently as possible, working around the uncertainties if possible, preparing for uncertain results if not. A genuine hard uncertainty means that you simply cannot know the outcome of your options; in many situations it will be impossible to *prove* which option is the best.

Assuming that we are facing a choice under uncertainty, does rational-choice theory tell us anything about what we ought to do? The answer is: very little. What it tells us is that we cannot rationally take account of any consequences of an option except the

best and the worst¹⁴. This principle may eliminate some options, but usually the hard choices are left unresolved. (Elster 1986:6)

Rational choice theory starts with the assumption that the decision-maker is facing a decision, and considers what the decision-maker ought to do within that context. Rationality in general is not restricted to as narrow a domain, and many of the actions that would be rationally recommended when facing hard uncertainties have to do with the context of the decision. For example, one of the best strategies here is to design new options that are less sensitive to any irreducible uncertainties (but this is outside of what Elster refers to as rational choice theory).

One of the most fundamental and powerful responses to hard uncertainties is to move from optimality towards robustness. This is a major change in perspective, requiring abandonment of the hope of finding the best option in a world assumed to be known and attempting instead to design options that are at least good in all possible worlds.

Another slant on this strategy is to shift from “good achievement” to “bad avoidance” (Dror 1991:27). As Dror goes on to point out, there is an additional advantage to this approach when multiple stakeholders are able to agree more readily on outcomes to be avoided than on those to be pursued. One approach to implementing “bad avoidance” is to a premium on a high security level, an extreme illustration of which is using the maximin decision rule.

¹⁴ Elster here uses “uncertainty” to refer to situations where probabilities of future states of the world are not known, but his minimalist conclusion applies to many situations of hard uncertainty.

Since the essence of a decision under hard uncertainties is that the decision-maker may be surprised by the outcome, preparations for surprises or outcomes that appear unlikely may be worthwhile. There are many ways to pursue such resilience. Contingency planning, war games, and rehearsal are methods of planning for circumstances whose possibility is foreseen. Unforeseen circumstances can be dealt with better when general response speed has been improved, and when a variety of skills, perspectives, and talents are available.

Strategies that have many possible times for adjustment are valuable in the face of serious, hard uncertainties, as Southern California Edison concluded in an example cited in section I.4. Plants that can be constructed in stages may be preferable to all-at-once plants, even if the average cost of such a plant is higher. Purchasing options, or getting permit acquisition out of the way before a project has been fully approved, are additional examples of multi-staged action.

Under hard uncertainties, many variables and parameters cannot be estimated as precisely as is desired. This can pose serious problems for quantitative methods, and it may be helpful to adopt qualitative methods, which are generally more tolerant of imprecise data. As Maynard Keynes put it, "it is better to be roughly right than precisely wrong."¹⁵ There are recent efforts to develop mathematical methods that are tolerant of imprecision, typically based on fuzzy sets or fuzzy logic.

It is not often emphasized in the literature on decision-making under

¹⁵ Quoted by Dror [1991:8]

uncertainty that any decision made under these circumstances amounts to a gamble. MEU calculations do not remove the gambling aspects at all; what MEU purports to do is to tell you which gamble is the best. Only by designing new alternatives that are less sensitive to the uncertainties or changing the values guiding evaluation of outcomes can the gambling aspects of a single decision under hard uncertainties be reduced.

III.3.2. Locus of uncertainty (data vs. model)

“Uncertainty” tends to be used in two different ways, depending on time horizon. Longer-range planners often use the term to refer to the whole environment, and think of themselves as “facing uncertainty”: the structure and dynamics of the situation are not known. Planners dealing with near-term considerations are more likely to use “uncertainty” to refer to a specific variable whose value cannot be precisely determined, such as the price of oil three years from now, or the prime interest rate available at the beginning of the next fiscal year. (“Uncertainty” in this work generally refers to a condition of the decision-maker, but it can be useful to distinguish between prevalent environmental patterns and specific events as causes of this condition.)

These two usages suggest a major difference in the mental models of the world that may be employed according to the planner’s time horizon. In the short term, the planner may have confidence in a model¹⁶ of the situation, and be uncertain only about the model’s data. In longer term

¹⁶ “Model” here is used very generally to refer to any form of representation of the actual situation, especially including mental models, in addition to more formal models, such as mathematical models.

cases, there may be major uncertainty about the model itself. In order to make it worthwhile to spend much time considering the range of values a particular “uncertainty” may take on, one must feel fairly confident that the uncertainty, or variable, will be relevant when the plans are implemented in much the same way that it is relevant now. That is, one must believe that the general structure of the world will not change all that much over the planning horizon, and that one understands it fairly well. If that assumption is warranted, then a model representing the current structure of the environment will also be useful for dealing with the environment being planned for.

Using the notation introduced in the GDM and recapitulated in figure 28, a represents our actions, w all exogenous influences (the state of the world), and c the outcome. The outcome is treated as a function of our actions and the state of the world, that is, $c = c(a, w)$. Data uncertainty indicates that we do not know w adequately; structure uncertainty means that the nature of $c(\cdot)$ is not known well enough.

Planners facing the longer-range, more general uncertainty may feel that the actual structure of the world in which their plans will be implemented is difficult or impossible to predict. Their problem is not that the future characteristics of elements of their model are unknown, but that the relationships among those elements, and even the set of relevant elements, are unknown. The entire “rules of the game” may be different in the future.

In most situations for which we engage in formal planning, we would be willing to accept the notion that there is some time frame that is short

enough that the fundamental structure of the situation is not likely to change.

Conversely, there is usually a time span large enough that we would

Figure 28. Locus — data vs. model

acknowledge that fundamental changes might very well take place. This distinction, easy enough to draw in this general discussion, is rarely distinct for planners, but is nonetheless useful for the purpose of identifying different kinds of tools and methods for planning for uncertainty.

I will refer to this dimension as the *locus* of the uncertainty — whether our uncertainty is mainly about *data*, or whether we are uncertain as to the very *structure* of the situation and therefore how to model it. If we believe that we have a sound understanding, i.e. we have confidence in our model, but there is a factor that is essentially exogenous that is uncertain, then we face *data* uncertainty. If, on the other hand, we are not sure how

to represent the system itself, or think that its structure may change significantly, then we face *model uncertainty*. In most planning contexts, long-term plans face *model uncertainty*, as the world is complex enough that major structural changes can occur within the horizon of a long-term plan.

Illustration. The energy industry in the U.S. is well-enough understood that the short-term (say, less than one year) impact of changes in the price of oil can be predicted fairly well. That price is uncertain, being affected by global events, especially by actions of OPEC. As long as the time horizon of our examination is short enough, the oil price uncertainty can be treated as a *data uncertainty*. Over the longer term, as little as five years, however, the impact of large changes in the price of oil can cause structural changes in the industry. This happened after the oil shocks of the 1970s when high oil prices drove up electricity prices, which spurred price-induced conservation that exceeded most forecasts and dramatically lowered national growth in electricity demand. For such longer horizons, oil price variability contributes to *model uncertainty*.

Implications of locus. Much of the importance of the distinction between model (or function) uncertainty and data uncertainty was discussed earlier in this chapter under “Subjects of Uncertainty”, but it is worth considering this distinction separately because it can help to quickly clarify whether or not modeling efforts have any chance of being helpful.

It is virtually impossible to make a decision other than randomly without employing models of some sort, however informal. When the structure of a situation is well known, it can be very useful to construct

formal models to better understand the implications of uncertainty in data. Mathematical modeling has benefited greatly from the explosive growth of computing power in the last quarter of the 20th century. Computerized models make it easy to use a great many alternative data assumptions and aggregate their results.

However, when the structure of the situation is not well known, time spent elaborating a particular model of the structure may be wasted, and may even obscure the uncertainty about the structure. The seductiveness of the quantitative output that computer models offer and the hope for improving that output provide a strong temptation to make the models a little better, a little bigger. This makes it harder to use alternative assumptions about the structure of the situation. As a formal model grows, more and more assumptions are built into it, meaning that developing an alternative model structure would involve either more modifications or a larger *de novo* construction effort.

Structural modeling is a field that helps to build alternative structures in order to compare each with what is known about the system of interest. It is the general behavior of the variables in structural models and their interrelationships that are examined — structural modeling tools are not intended to produce numerical outputs whose values represent the system of interest with any precision. This is not a real loss, since if the decision-maker is uncertain about the structure, then numerical outputs of any model ought to be taken very skeptically.

The distinction discussed here becomes especially blurry when a chaotic model is employed. Chaotic systems are well known for being very

sensitive to initial conditions. A very well-specified chaotic model can exhibit dramatically different behavior when started at two points that are very close together. The effect of this may be to sharply limit the value of reducing uncertainty in the data when a system is believed to be chaotic. It may be worthwhile to explore how sensitive the relevant behavior of the model is to changes in its formulation; in some cases it may be enough to know that chaotic behavior is likely, and further refinement may not contribute much to better decision-making.

III.3.3. Parametric vs. strategic

Elster's distinction between parametric rationality and strategic rationality (e.g. Elster 1984) can clarify some cases of uncertainty. A *strategic uncertainty* is one in which the uncertainty depends in a significant way on the choices and actions of other agents; a *parametric uncertainty* is one for which other agents choices are largely irrelevant.

For dealing with parametric uncertainty, forecasting processes as if they are outside the control of the decision-maker may be both legitimate and helpful. The future depends only on nature, which is presumed to be neutral towards the decision-maker. For strategic uncertainty, such forecasts may be quite deceptive, since other agents may actively attempt to help or harm the decision-maker, especially if they learn of the forecast. Forecasting their actions as if they were not susceptible to the decision-maker's influence is likely to obscure opportunities for mutually advantageous interaction as well as possibilities of deliberately harmful action by others.

As the number of other decision-makers increases, the strategic aspects

may diminish. When the actions of individuals are relatively unimportant, groups of people can often be characterized rather well by statistical methods.

Illustration. Weather is generally beyond our current ability to control, and ought to be treated as parametric uncertainty. The amount of electricity that consumers demand for their heaters or air conditioners is affected by weather, specifically how cold or hot a particular day turns out to be. The amount of water in the Columbia River that can be used to drive the turbines in the hydroelectric system depends on weather, specifically the amount of precipitation in the Columbia's drainage basin. Data has been collected on these variables, and it can be used to estimate the likelihood of demand and supply of electricity in the Pacific Northwest.

Other supply and demand factors are affected by human choice. The privately owned utilities in the Pacific Northwest may choose to meet additional demand from their customers by asking BPA to sell power to them. BPA would be required to ensure that it has sufficient power for such sales (under specified conditions). This uncertainty about the demand BPA must meet is not based on weather but on strategic decisions made by other players in the west coast utility arena.

There are similar uncertainties on the supply side of the picture. For example, British Columbia has a large hydro system that often has more hydroelectric fuel (water) than its own needs require. BPA may use forecasting techniques to estimate how likely it is that wet weather will produce such a surplus in the British Columbia system — a parametric uncertainty. Whether or not British Columbia makes any of that power

available to BPA is a strategic uncertainty. Regulations in Canada governing export of energy may change. As the provincial government of British Columbia changes leadership, its attitudes towards the degree of cooperation and interdependence with the Pacific Northwest of the United States also changes. Operational planning concerns in the British Columbia power system may call for storing more or less of their surplus in their own reservoirs.

Canada and the United States have made extensive agreements about how to operate the dams on the Columbia River. These agreements both reduce strategic uncertainty and reduce the importance of some parametric uncertainties. The agreements cannot affect the amount of precipitation or the weather-driven electric loads. They have, however, provided tools for dealing with the weather uncertainties that enable the two regions to reduce the magnitude of the effects of the weather uncertainties on the power systems, which has reduced the power uncertainties.

Implications of parametric vs. strategic. Different schools have developed to support decisions under strategic uncertainty and under parametric uncertainty. Game theory was developed to deal with strategic uncertainty. Decision analysis was developed for parametric uncertainty. Cases of parametric uncertainty are sometimes referred to as “games against nature”, but nature plays an unusual role, in that nature’s “moves” are not influenced by the decision-maker’s moves or values. Situations in which both chance and other agents play a major role are less tractable.

Two key contrasts between these two disciplines are in the use of

probabilities, and the treatment of worst-case outcomes. Generally, probabilities are very important in decision analysis, and much less important in game theory; worst-case outcomes figure prominently in game theory and receive much less attention in decision analysis.

Game-theoretic analyses often begin with an assumption that all actors will act rationally. From this it follows that if your opponent has a “most rational” move, she will take it. That assumption eliminates a potential uncertainty on your part about your opponent’s move. This amounts to basing your own move on worst-case assumptions, though in the context of games against rational opponents, what is worst for you and what your opponent is most likely to choose are often the same (and in two-person games of pure competition are always the same). In situations of parametric uncertainty, however, the worst outcome and most likely outcome may be very different, and probabilities are very often employed as weight in combining these two (and other) outcomes for comparing alternatives. The worst outcome may be very unlikely, and in expected-value calculations, its small probability will give it a small weight.

Probabilities are treated very differently by these two analytical traditions. In game theory, probabilities are not used except in the cast of *mixed strategies*. Mixed strategies are typically analyzed as candidates for your own action, but your opponent, assumed to be equally rational, may also select a mixed strategy. Mixed strategies are preferable when adopting any single strategy (a “pure” strategy) would allow your opponent to select an advantageous strategy, but by deliberately maintaining a set of strategies from which any one might be picked you can prevent your

opponent from faring as well. Here probabilities are not used as assessments of what another agent might do or which state of the world may eventuate, but to select one action from a set according to rationally chosen probabilities. This is one of the very few situations in which mainstream prescriptive rationality advocates making a choice randomly. Decision analysis, however, makes extensive use of probabilities. Instead of evaluating an alternative against the most likely, worst case, or most rational strategy of an opponent, a decision analysis will evaluate each alternative against all of the future states of the world in the feasible set. Then the utility of sets of individual outcomes associated with a particular alternative are aggregated into a single measure of the utility of the alternative by weighting them by the likelihoods of the states of the world.

Quantitative analysis in general, and calculating modes of decision-making, such as MEU machinery, is more suitable for parametric uncertainty, except for those strategic situations where the “rules” are clear enough that game theory can be applied. Strategic uncertainties depend on the *intentions* of others, and this offers the hope of understanding those intentions and getting a better sense of what the others might do. A further possibility is to find ways to alter those intentions through the use of incentives, threats, reasoning, persuasion, bargaining, and so on. On the other hand, the intentions of others may lead to their choosing alternatives that seem very unlikely to us, sometime just because they seem unlikely.

III.4. Sources of uncertainty

Uncertainty that a decision-maker experiences may come from many

different *sources*, or root causes. This section enumerates and illustrates the possible sources of uncertainty in decision-making. This is not meant to suggest that many situations are encountered in which there is a single source of uncertainty, or that determining the sources of a particular instance of uncertainty is trivial.

The model underlying this analysis is that to make a decision, a *decision-maker* has to

Figure 29. Three sources of uncertainty

process *information* about the *world*. This defines three major categories of sources of uncertainty: (1) the world — the “real” processes and objects that are the sources of data, and the space within which the consequences the decision-maker cares about can occur; (2) evidence — any type of information the decision-maker does or does not possess or acquire about the world; and (3) the decision-maker — characteristics of the agent or entity that is apprehending and processing the evidence and making inferences and decisions. (While we can make decisions only on the basis of evidence we have, a major assumption underlying the Western thinking is that there is something more fundamental than evidence “out there”, and that differences in the “out there” space matter.)

An ultimate answer to the question of whether or not the world is perfectly amenable to prediction by an omniscient Laplacian being is not needed in order to notice that the distinction between the world and our evidence about the world can be useful as we consider what to do about our

uncertainty. The uncertainty of not knowing how a coin about to be tossed will land differs deeply from the uncertainty of not trusting someone who tells us that a coin that has already landed fell tails-up. Similarly, using probabilities to describe the likelihood of heads on the next toss differs from using probabilities to represent our judgment of the likelihood that the observer in the second case is lying to us. In both cases, the probabilities may indeed be opinions, as the subjective probabilists assert, but they are clearly not the same kind of opinion.

In several cases, it is an interaction between two or more possible sources of uncertainty that determines how much uncertainty a decision-maker will face. These interactions are discussed after the sources are described.

III.4.1. *The world*

While we have no direct contact with the world, and all of our information about it suffers the potential uncertainties of evidence, discussed below, we often have good reason to believe that there are worldly processes that in themselves contain variability or unpredictability. The first distinction here is between the natural world, treated as non-sentient, and actors (or agents) in the world who exercise a power of choice.

Figure 30.
Uncertainty from *the world*

III.4.1.1. Natural world

Uncertainty arising from the natural world is largely a variability or irregularity of one or another kind: a completely unvarying world

contributes no uncertainty. That portion of the natural world relevant to a particular decision may have many degrees of regularity.

These range from a complete absence of regularity to perfect constancy over a specified time period.

Implications of Natural world.

Understanding and clarifying uncertainty generated by the natural

world is one of the strengths of the T

Perspective. Since much natural world uncertainty is irreducible, creating effective responses generally requires the O and P Perspectives as well.

Information gathering will not remove this sort of uncertainty. Incremental actions, such as refining and gathering more information or improving your processing capability, are of relatively little use in this case for actually reducing uncertainty, whereas they can be significant in reducing uncertainty caused by insufficient evidence or processing power. Uncertainty about the weather, for example, when the source of this uncertainty is variability in the atmospheric system, or even chaotic variability, cannot be reduced by data collection. Developing a technology to control the weather could reduce or even remove this uncertainty (if

Figure 31. Uncertainty from *natural world*

possible), but that goes far beyond gathering data. In everyday life, the toss of a coin is considered to be a random event, meaning that we do not think that data gathering can lead to predicting which side will be facing up. It does not seem impossible, in principle, to possess so much information about the wind currents and the precise movements of the mechanism propelling the coin (e.g. your thumb and forefinger) to be able to predict the entire course of the coin through the air, but this is an amount of information completely beyond the hope of most people who care about the results of coin-tossing — obtaining a little bit more information would not be sufficient to permit prediction of the coin toss. Highly coordinated people may be able to pluck the coin from the air at just the right time, or to turn it over surreptitiously, but this is clearly more a matter of replacing a random, physical process with a human motor skill performance task than reducing or eliminating the uncertainty generated by an aspect of the physical world.

The role of reason in this case is not to reduce the uncertainty but to clarify, to bound it as much as possible, reduce the uncertainty in utility resulting from the underlying uncertainty, and to devise defenses against the remaining uncertainty in utility.

III.4.1.1.1. Disorder

A world without any order, one of total randomness, devoid of natural laws, would not appear to permit rational decision-making any advantage over random behavior.

Illustration. This research has not uncovered any meaningful examples of complete lack of order in electric utility contexts, with the

possible exception of extraordinarily energetic physical phenomena, such as earthquakes.

Implications of Disorder. Genuine disorder would utterly frustrate attempts to predict what the world would be like in later moments, which would make it impossible to forecast the relative success of different actions. Success is likely to be very dependent on luck. Since in such circumstances it would not be possible to prevent failure, it would be important to try to take steps to see that failure would be as survivable or as inexpensive as possible. Instead of putting all of your eggs in one basket, put at least some eggs in as many baskets as possible. Strategies that have many different possibilities for success or for partial success, and that offer the most opportunities for changing course are promising. Research, being generally an attempt to discover regularities, would be unrewarding. (The question of how a situation could be determined to be one of true disorder is non-trivial: if attempts to discover regularities are unsuccessful, that could be due either to inadequate attempts or to fundamental disorder, and distinguishing between these two causes may be difficult.) Fortunately, there are not very many situations totally devoid of regularity.

III.4.1.1.2. Partial regularity

There are many forms and degrees of partial regularity, especially in situations composed of several factors, some of which are more regular than others. A specific type of partial regularity to gain recent prominence is chaos. In partial regularity, there is some order, and therefore some leverage for rational decision-making.

Illustration. There are indications that weather systems are chaotic (Casti 1990:108). This would mean that the changes in electric load that are due to weather phenomena, such as electric demand for home heating, will always be unpredictable to some extent. It may be possible to determine a band containing the impacts during particular time periods, but much of the weather-induced uncertainty is hard.

Implications of partial regularity. Partial regularity offers the hope of at least bounding the uncertainties, and making plans for best-case, worst-case and in-between situations. A chaotic system, as an example of a kind of partial regularity, may be amenable to mathematical modeling, which would facilitate determining the extreme conditions of the system, and perhaps of other features of the system's behavior. Possibilistic tools, as they are developed, may be useful.

Chaotic systems, and nonlinear systems in general, can exhibit wide ranges of behavior without any change in external conditions. While linear systems can be in stable, oscillating, or unstable states, nonlinear systems can additionally be in *chaotic* states, in which their behavior can appear to be random but actually be completely determined by a set of equations (Gordon and Greenspan 1994:49). Thus, an observer could find it difficult to determine whether stable behavior is a result of linear dynamics, or a product of a chaotic system that is temporarily stable. It can be similarly difficult to distinguish apparently random behavior from chaotic but deterministic behavior, to determine whether a change from one type of state (stable, oscillating, unstable, random) to another is due to a fundamental change in structure or simply a transition from one domain of

a chaotic system's behavior to another, or to find out if the erratic behavior of a system is due to the influence of external factors or a chaotic but purely internal relationships.

Naturally, the behavior of the stock market and other markets has attracted considerable interest, especially in the hope of being able to use the mathematics of chaos to forecast the performance of markets. That would appear to require being able at least to identify the behavior of the market as chaotic, and probably to learn something about the parameters of that particular chaos (though it would be great to be wrong about this). Douglas Patterson is not encouraging about the prospects for such identification:

Most researchers in the area have concluded that detecting deterministic chaos in economic-finance time series is probably impossible. (Patterson 1993:405)

Even if a system is identified as chaotic, there are severe limits to the sort of prediction that would be possible. Chaotic systems are sensitive enough to initial conditions or small deviations in inputs that prediction of exact values is unlikely to be possible. However, it might be possible to predict with great confidence that the system would be somewhere within a range of possibilities, and this might facilitate developing strategies that would be satisfactory for any possibility within that range (Gordon and Greenspan 1994:50).

Patterson reports that Brock, Hsieh and LeBaron do not provide a clear answer in their recent book *Nonlinear Dynamics, Chaos, and Instability*¹⁷

¹⁷ Brock, W.A., S.A. Hsieh, and B. LeBaron 1992. MIT Press, Cambridge.

to the question of whether security returns might be forecasted from the history of the returns (Patterson 1993:406). Although chaotic systems can be created whose behavior mimics that of the stock market, for example, it seems highly unlikely that the behavior of the stock market is determined solely by its previous behavior; surely there are significant external influences that need to be taken into account, which would make any chaotic modeling of the market quite difficult.

Still, there may be important reasons to consider the possibility that a system under observation is chaotic, and Theodore Gordon offers this advice:

Look for nonlinearities. Look for feedback paths. Observe whether speeding up the system leads to oscillations or apparently random behavior. Question whether externalities are always the cause of performance anomalies. A system that is noisy and fails to respond as anticipated is not necessarily inoperative; it may be operating in an unanticipated regime. (Gordon 1992:15)

The mere possibility that a system could be chaotic should induce uncertainty in an observer. Because a chaotic system can move quickly from relative stability to turbulence or vice versa, a system that might be chaotic cannot be safely assumed to continue to display the sort of behavior typical of its recent past (Gordon and Greenspan 1988:1).

Some approaches to reducing the severity of the behavior of chaotic systems are to slow it down — change the time scale, reduce the flow of information, or dampen rate changes in it; reduce feedback — change the ways components interact to reduce recursion or feedback, or limit particular flows of information; or actually decouple the system to prevent feedback (Gordon and Greenspan 1994:60). Sensitive monitoring may be

useful also, in that chaotic systems can exhibit turbulent behavior that is not unstable (i.e. does not decrease or increase without limit). If the system can be observed carefully enough that a shift into a different domain of chaotic behavior can be distinguished from blowing up, more freedom can be taken to experiment with control measures without running excessive risk of catastrophe.

The implications of believing a system to be chaotic are important enough both for avoiding wasted efforts at uncertainty reduction and for determining the range of system behavior for which preparations must be made that detection of chaos in sets of data has become an important topic. Market behavior in particular has been the subject of research that attempts to discern whether or not particular time series of data comes from a deterministically chaotic source or some other source, such as generalized autoregressive conditional heteroskedasticity process (GARCH)¹⁸.

III.4.1.1.3. Stochastic regularity.

Many physical processes can be usefully represented as stochastic systems with specific parameters which can sometimes be estimated empirically with great precision. The toss of a fair coin exemplifies this type: we believe the result of the toss to be unknowable, but we also believe

¹⁸ For example, Gilmore, C.G. 1993. "A New Test for Chaos", *Journal of Economic Behavior & Organization*, vol. 22, no. 2, Oct., pp. 209-37; Brorsen, B.W., & S.-R. Yang 1994. "Nonlinear Dynamics and the Distribution of Daily Stock Index Returns", *Journal of Financial Research*, vol. 17, no. 2, Summer, pp. 187-203; Willey, T. 1992. "Testing for Nonlinear Dependence in Daily Stock Indices", *Journal of Economics and Business*, vol. 44, no. 1, Feb., pp. 63-76; Blank, S.C. 1991. "'Chaos' in Futures Markets? A Nonlinear Dynamical Analysis", *Journal of Futures Markets*, vol. 11, no. 6, Dec., pp. 711-28.

we know the only two possible outcomes and their probabilities. The question of whether any particular coin is indeed a fair coin is an empirical one, and much less simple.

Illustration. Utilities deal with many mechanical processes for which stochastic representations are widely used. The rates of failure for many of the components in thermal power plants are well known. Utilities are able to calculate reasonably well the probability that a thermal plant will be unable to produce power when it is called upon to do so. (This is less true for nuclear plants, where the length of time required for some important stochastic processes to take place, e.g. deterioration of structural materials, may be of the same order of magnitude as the history of nuclear plant maintenance. The processes in a nuclear plant may be just as stochastically regular, but insufficient time has elapsed for collecting data, which reduces the efficacy of statistical tools.)

The total electric demand at any moment for use in electric water heaters can also be usefully modeled as a stochastic process. It is impossible for a utility to predict just when the water heater in an individual house will be drawing current, but utilities serve enough water heaters that the total demand for water heating is fairly predictable. This is an example of the potential power of aggregation, in which a large number of similar, stochastic processes can be represented as a single, larger, far more predictable process.

Implications of stochastic regularity. Stochastic regularity offers the most natural and powerful opportunities for probabilistic techniques, since observation or analysis is likely to permit the inference of frequency-

based probabilities. Frequency is the basis for probabilities that is most widely accepted as legitimate.¹⁹ Stochastically regular situations are also the most likely to provide repeated opportunities for essentially the same decision, which increases the legitimacy of the hope that the decision-maker's average results will approximate the mathematically expected results.

Stochastic regularity may be inferred for analytical reasons, as in the case of logical analysis of dice or card games, or empirical reasons, as in the case of observing the time between successive arrivals of a bank customer in a queue and then testing the "goodness of fit" between the observations and a representation of the phenomenon as an exponential random variable. The former case justifies beliefs about the long-run frequencies by declaring that the analysis explains the phenomenon and therefore the long-run behavior of the phenomenon ought to conform to the analysis; the latter by observing a long-enough run to claim that a longer run would merely be more of the same. However derived, frequency-based probabilities can be utilized in a wide variety of well-developed techniques, such as inventory theory, queuing theory, and Monte Carlo modeling and many others.

III.4.1.1.4. Complexity < interaction 1 >

Generalized Uncertainty Principle: Complex systems exhibit unexpected behavior.

John Gall

¹⁹ For example, von Neumann and Morgenstern declared that only long-run frequency probabilities ought to be used in the calculation of expected utility [von Neumann & Morgenstern 1953:19].

Even in the absence of variability or randomness in the world, the complexity of a phenomenon can engender uncertainty. While Laplace could imagine a being that could know all of nature's forces and the positions of all matter, and process information as quickly as needed²⁰, all earthly entities are constrained in their knowledge and processing ability. The complexity of a situation can overwhelm a decision-maker's ability to understand the situation and predict future states. Zadeh's Principle of Incompatibility states that an increase in the complexity of a system makes it more difficult to describe the system in ways that are both precise and relevant — and may make it truly impossible (Zadeh 1973; also see section II.3.10).

As noted earlier (Section II.3.10), Klir claims that the only the complexity of *systems* is potentially measurable, not the complexity of the *objects* the systems represent (Klir 1985:132). He does not deny that the complexity of objects is meaningful, however. Thus, our inability to measure objectively the complexity of objects does not prevent us from noting that objects differ in their complexity, and that complexity can hinder our understanding and cause us to be more uncertain. The impact of an object's complexity is co-determined by our ability to process information (see also "processing capability" below).

Complexity typically induces uncertainty when a process that may be deterministic is so complex that its outcome cannot be predicted with confidence. In the G.D.M, this corresponds to difficulty in predicting the

²⁰ Laplace is quoted in section II.3.8.1.

future state of the world, or to difficulty in predicting the consequences of a choice/world-state pair (model or function uncertainty).

Illustration. Von Neumann and Morgenstern prove that the outcome of a chess game is in principle completely determinable, but acknowledge that the complexity of the game renders the outcome uncertain:

This shows that if the theory of Chess were really fully known there would be nothing left to play. ... But our proof, which guarantees the validity of one (and only one) of these three alternatives [white wins, tie, or black wins], gives no practically usable method to determine the true one. This relative, human difficulty necessitates the use of those incomplete, heuristic methods of playing, which constitute “good” Chess; and without it there would be no element of “struggle” and “surprise” in that game. (von Neumann and Morgenstern 1953:125)

The paucity of practical significance of von Neumann and Morgenstern’s conclusion, one of substantive rationality, illustrates Simon’s assertion of the importance of procedural rationality, i.e. “executable algorithms” (see section II.3.3.1; Simon 1981:43). (*Interaction 1* below, section III.4.4.1, contains another illustration of complexity.)

Implications of complexity. The conflict between relevance and precision caused by complexity poses enormous difficulties for mathematical descriptions and models, where precision and internal consistency have traditionally been highly important. This problem is vividly illustrated by the varied acceptance that MEU models for decision-making receive. Such models are typically cast in the forms preferred by mathematicians and economists, i.e. they are quite precise, and few complain of problems due to *imprecision*. It is instead problems due to perceived *irrelevance* of the models that are most often and loudly decried. In most actual situations, the precise requirements of the models fail in

some degree to be met²¹. This leaves open the question of how directly the results of an MEU analysis ought to be applied to any given situation. In contrast, the concept of *satisficing* (see section II.3.3) is often perceived to be *relevant* to both individual and organizational decision-making, but it has not been a fertile source of *precise* implications for specific situations.

Statistical aggregation can be very helpful for dealing with *disorganized complexity* (see section II.3.10), i.e. when there are many similar events to deal with. This is at the core of strategies employed by insurance firms, casino operators, and opinion pollsters. Shackle goes as far as to say this can eliminate uncertainty:

Uncertainty is not present when the consequences of the action are actuarially certain. (Shackle 1957:107)

Organized complexity is more difficult to deal with. In some instances, reductionism or hierarchical decomposition can help divide up a situation of organized complexity into a set of sub-situations, each of which is more like *organized simplicity* and more amenable to analysis. This is such a powerful tool that it may be a virtual prerequisite for understanding the world:

The fact then that many complex systems have a nearly decomposable, hierarchic structure is a major facilitating factor enabling us to understand, describe, and even “see” such systems and their parts. Or perhaps the proposition should be put the other way round. If there are important systems in the world that are complex without being hierarchic, they may to a considerable extent escape our observation and understanding. (Simon 1985:219)

²¹ For example, a decision-maker may not actually have a complete preference order over all possible outcomes, or the decision-maker may have a distinct preference for or against gambling per se, which MEU axioms prohibit.

Klir suggests that fuzziness may be helpful in handling complexity:

Another way of dealing with very complex systems that possess the characteristics of organized complexity, perhaps the most significant one, is to allow imprecision in describing properly aggregated data. (Klir 1985:138)

An implication of Klir's assertion is that tools built around fuzzy logic and/or fuzzy sets may allow for an imprecision in inputs that permits more compact representation of input values and simpler logic in a representation of a complex object.

In order to capture the ever increasing richness of economic phenomena, we need increasingly complex models. When the required complexity for obtaining realistic models becomes unmanageable, we must simplify. In every simplification, unfortunately, we are doomed to lose something. When we insist that the simplified models give us precise predictions, we lose relevance of these predictions to the real world. We can preserve some of this relevance only by allowing some uncertainty in the models. That is, we can trade certainty for relevance under a given limit of acceptable complexity. (Klir 1992a:86)

Klir is suggesting a variant of Zadeh's Principle of Incompatibility — an increase in the complexity of a model makes it more difficult to retain predictive precision and the manageability of the model. Unless we allow more uncertainty into the model, the model is likely to become too slow, too difficult to maintain, modify, update, or learn, or too expensive.

Seen in the light of Klir's observation, statistical aggregation is another form of trading off one kind of certainty in order to obtain more of another kind, within limits of a given complexity and a desired level of relevance: it deliberately introduces uncertainty about individual cases (by omitting virtually all their details), so that the uncertainty about any one instance is maximal, while retaining as much precision as possible about the

aggregate.

Another approach in which uncertainty is tolerated for the sake of maintaining relevance is greater reliance on judgment and qualitative analysis. Judgmental analysis and intuition are less certain than more formal, rational analysis in the solidity of the grounds for proceeding from one step to the next; qualitative analysis lacks the potential numerical certainty (i.e. precision) that quantitative analyses at their best provide. However, judgment and qualitative analysis are better suited to situations where relevance is essential, and within the constraints that follow from that requirement, as much precision as possible is sought; formal and quantitative analysis often start with great precision and the struggle is to increase relevance.

Possibly useful responses to increasing complexity:

Aggregation: convert a series of uncertain events into a single, aggregate event that is certain, by permitting maximal uncertainty about individuals.

Simplification of models: with a given limit of complexity, trade away certainty (accept imprecision) to get relevance.

Reductionism or hierarchical decomposition: divide a single situation into several components, thereby trading away complexity to increase relevance and/or precision.

Qualitative analysis: as complexity increases, sacrifice precision and rigor to retain relevance and processing speed.

Unhelpful responses to increasing complexity:

“Formalism”: as complexity increases, continue to insist on precision and lose relevance.

Unrestrained model growth: as our understanding of an object being modeled becomes more complex, increase the complexity of the

model at the expense of its manageability.

III.4.1.1.5. *Fundamental dynamism*

The future is never concretely determined —
that is the principle of statistical mechanics —
although we cannot speak of the indeterminism
of the past.

Eberhard Knobloch

Decision-making requires prediction, since we care about the future consequences of the options available to us. A major source of uncertainty about future consequences is *fundamental dynamism* — the fact that things change over time. In contrast to stochastic regularity, where a variable of interest changes over time but according to certain laws, dynamism as used here produces “the kind of uncertainty associated with the unforeseen emergence of novelty” (Casti 1990:42). Stochastic processes require that a series of situations arise, each equivalent to the others for this purpose. There is a circularity to such processes, as they return in a sense to the same beginning each time. When each next situation is significantly new, the past is a less useful guide to the future, and events proceed more linearly, never returning to the same place.

The distinction between *variables* and *parameters* in models of processes, especially stochastic processes, is important.²² The processes

²² *Parameters* are part of a description of the nature of a system that distinguishes it from other systems; *variables* describe the different states the same system can be in at any one time. A system may be in many different states at different times; observations of the system at these times would yield different values of its *variables*. Its *parameters* would remain constant over these observations. For example, in a stochastic representation of a coin toss, the *variable* is the side of the coin showing, i.e. either heads or tails in any observation. The chance of a head coming up, e.g. 50% for a fair coin, is a *parameter* of the representation of a coin, and doesn't change from one toss to the next.

most suitable for stochastic treatment are those in which the parameters can reasonably be assumed to be constant over time, as we believe the probability of a coin landing heads-up is. In other cases, that assumption is dubious. When the future values of a system's variables cannot be predicted adequately because the system's parameters are changing, it is sometimes possible to identify a new system whose *variables* include the first system's *parameters*. This can permit predicting the course over time of the first system's parameters, and thence the pattern of the first system's variables.

Illustration. The total annual volume of water coursing through the Columbia River's system of dams and generators can illustrate the difference between *variables* and *parameters*. There is a record of over 100 years of data on this streamflow, and this permits detailed estimation of the probability of specific levels of annual run-off. Recent concerns about global warming, however, highlight the possibility that the basic climate patterns (parameters) which produce each year's weather (variables) may be changing. This casts doubts on the relevance of data on past streamflows when forecasting future flows: the parameter values about which we have information may not characterize the current climate/weather system, and we don't have enough data to know whether that is the case or not. Perhaps some time in the future we will have accurate global climate models. These may include as output variables the average annual precipitation in the Columbia River basin. This could then be used as a parameter in models simulating the annual operation of the PNW electrical supply system.

For another example, as BPA's energy planners evaluate alternative sources of electricity, the future spot market prices of natural gas and electricity are both dynamic variables. When evaluating the energy available a few months into the future, it is reasonable to expect that the west-coast energy markets will be much the same then as they are now, with largely predictable differences due to seasonal patterns of demand and supply. The system is expected to be stable (no changes in basic parameters), but the system's variables are still hard to predict exactly. The picture is very different when forecasting the costs of gas and electricity 5 years from now. 5 years is long enough that the system itself can change very significantly. A short-term phenomenon like a cold snap could temporarily drive prices up sharply without changing the system: the variables change, but the parameters don't. News that the North American supply of natural gas is either much larger or much smaller than previously believed could change the parameters of the energy market, causing a decrease or increase in prices that would not be temporary.

Implications of fundamental dynamism. The prospect that genuine novelty will emerge can be prepared for only in a generic manner. Clearly plans that are based on the assumption that the future has been predicted accurately are at much greater risk than those in which the possibility of being surprised has been incorporated. Environments differ both over time and over space in how often surprises occur. For example, the electric utility industry was relatively stable in the 1960s, but was upset by two large "oil shocks" in the next decade. More recently there have been increasingly numerous changes, none as wrenching as the oil shocks, but

with cumulative effects probably much greater. The overall rate of change appears to be much higher now, though it may not remain this high forever. When it is clear that changes are occurring rapidly, plans with a great many opportunities for mid-course corrections are advantageous. Pursuit of robustness instead of pursuit of optimality will be less vulnerable to the surprises that will inevitably occur. Times of relatively low rates of change permit more use of optimizing, but there is a very difficult problem of estimate the meta-rate of change: how soon might the rate of change itself increase?

If an organization or an individual spends too much time conducting what-if scenarios and planning for the unexpected during a stable period, a competitor who is optimizing for a future simply extrapolated from the present may have a performance advantage. On the other hand, an agent who spends too little effort preparing for surprises during a time when they are common is likely to be hard hit by some of the surprises. In a turbulent environment, luck may play a greater role, with more good plans going awry and more bad plans achieving success than in stable conditions. Even one surprise may prove fatal to a person or an enterprise, and though reason leads us to believe this is less *likely* to befall those who plan for the unexpected than those who do not, it may not be any less *possible*.

III.4.1.2. Actors

An especially important type of world-induced uncertainty is due to the

unpredictability of actors (or agents)²³. While the ultimate predictability of individual human actions is an open question, it is clearly beyond our current capabilities in so many situations that human choices represent a major source of uncertainty. On the other hand, there are situations pertaining to human actions, especially those of large numbers of people, where statistical prediction can be quite accurate, and other situations of lesser regularity where we are able to reduce uncertainty by ruling out some possibilities or by estimating the likelihoods of the possibilities.

One type of actor-induced uncertainty arises when an actor has not yet made a decision; a future act of will can resolve this uncertainty. A second type, differing from the first as

Figure 32. Uncertainty from *actors*

parameters differ from variables, is that in the time before an actor makes a decision, the actor may change, and in some way become a “different” actor. This can be due to increases in knowledge, experience that results in value changes, more or less regular changes over the course of a life, etc. A third kind of reason to be uncertain about actors is that there are times when it appears that no amount of information about the past could possible have enabled a behavior to be predicted — a genuinely surprising, unique action is performed, due to inspiration, irrationality, disease, or

²³ “Actors” here primarily means other people, but the essence of an actor is the possession of volition. Non-human entities may be deemed to be capable of *choosing* in some contexts and therefore be viewed as actors.

any of a myriad other causes.

Implications of actors. In dealing with strategic uncertainty stemming from the future behavior or choices of agents, probabilities, especially probability-based calculating methods such as MEU, are generally less often recommended than in dealing with parametric uncertainty, especially. Three exceptions are as a measure of confidence in a description or prediction in qualitative analyses, in randomizing your own behavior in certain games, and when aggregation makes it possible to characterize statistically pertinent aspects of groups of people or a series of similar interactions with the same actors.

III.4.1.2.1. *Un-made decisions*

Un-made decisions that would affect a decision-maker's current choice involve both *actions* and *values*. Actions yet to be taken may change the world and thereby affect the decision-maker's options or the consequences that will occur. Un-made value decisions by the decision-maker, or the party (parties) for whom the decision-

Figure 33. Uncertainty from *un-made decisions*

maker acts, can make it impossible to compare the desirabilities of outcomes and thence the desirabilities of options. When there is a need to aggregate the preferences over actions or values of more than one person, resolving the uncertainty is yet more complicated.

III.4.1.2.1.1. Actions

The decisions of others may be intentionally or accidentally relevant. Their actions can affect the feasibility of your options or the future world within which the consequences of your choices will occur, thereby changing some possible outcomes and thence the desirability of some options. Game Theory deals specifically with situations in which two or more actors know their decisions can affect each other according to public rules. Cooperative situations no less than competitive ones can produce uncertainty; simply knowing that another person will act cooperatively does not remove all uncertainty about the nature, extent, timing, etc. of the eventual cooperation. Many times others act in ways that affect us without having any intention of doing so; we can rarely be sure that such an inadvertent impact will not occur. Actions that the decision-maker will take in the future may be considered in this category also. For example, a decision about whether to re-paint or merely patch up the existing paint of your house is complicated by uncertainty from unmade decisions if you haven't yet made up your mind whether to move across town next summer or not.

Illustration. BPA has many wholesale customers, each of whom makes many decisions that can affect BPA. When they expect their own load to grow, they may choose to enter into long-term contracts to purchase additional power from BPA or to construct their own generating resources. Until they make their decision about how much power to purchase, BPA faces irreducible uncertainty about its load. Some actors who have an impact on BPA have no intention of doing so, such as a vandal who shoots out insulators on transmission lines purely for kicks, knowing little and

caring less about who owns the insulator or what impacts its destruction may have. The consequence of the loss of the insulator be loss of the use of the line until it is repaired, which could prevent BPA from serving a customer. There is a chronic uncertainty about transmission availability caused by the possibility of such vandalism

Implications of actions. Game theory can offer insights in how to deal with uncertainty arising from decisions yet to be made by other actors when the actors and the decision-maker are interdependent. When the other actors are either unaware of, or genuinely unconcerned about, the impacts of their actions on the decision-maker, game theory is less powerful.

Attempts to influence the decisions other actors make can make the situation more favorable to the decision-maker. This may result in reducing the uncertainty, but improving the situation does not always require uncertainty reduction. The purpose of negotiating agreements is to reduce uncertainty or increase security levels, and the purpose of various ways of making the agreements binding is to eliminate uncertainty over the actual carrying out of the agreements. The search for ways to influence the decisions of others is age-old, and has produced methods ranging from those of Dale Carnegie to covert intelligence operations, from advertising to blackmail and extortion.

Sometimes uncertainty about the future actions of others can be reduced by openly asking others what they plan to do. If there is a serious possibility that the others would lie about their plans, their answers may not reduce the uncertainty. Externally imposed sanctions for deliberate

deception may reduce this risk. For example, while rival utilities may have incentives to mislead their competitors, when they are required to file plans with regulatory agencies and justify any deviations from their plans, the plans provide much more substantial information about their future actions.

The presence of interdependencies is a key complication. When the decision-maker's actions will not affect the other actors, the others have little or no motivation to deceive the decision-maker. Potentially infinite regresses of second-guessing also require interdependence in order to arise (e.g. "What's he likely to do? What's he likely to do if he thinks I figured out what he was likely to do?").

If uncertainty from this source cannot be reduced, it can be treated with the methods discussed above under "hard uncertainty".

III.4.1.2.1.2. Values

Making decisions requires knowledge of values as well as predictions of future actions, and important value choices may not have been made yet. Often the decision-maker and the party (or parties) whose values will guide the decision are one and the same. However, there are other times when the decision-maker wants to use the values of others in making the decision, as when the party making the decision acts as a trustee for another party²⁴. Until the relative weight of different values has been

²⁴ There are also many times when someone making a decision attempts to justify a choice on the basis of what he claims others value, as when President Nixon invoked the unspoken (and therefore unknown) desires of the "silent majority". This is a significant kind of action, but not one of uncertainty: such a decision-maker may not know the values of the parties for whom he claims to act, but he does not *want* to know.

assessed, comparisons of the desirability of possible outcomes cannot be judged, and the options available to the decision-maker cannot be ranked.

Illustration. One way in which the electric utility world is changing is that wholesale electricity is more available now from a wider variety of power providers than it was a decade ago. As mentioned earlier, this means that BPA's wholesale customers do not have to purchase power from BPA — their choices are proliferating. A factor in their power purchase decisions, BPA's top management believes, is the *hassle factor*. This refers to the finding from surveys of managers of utilities that BPA's customers view BPA as a very difficult and frustrating business partner. It is hard for them to get a decision from BPA; they would like their services to be more individualized than they currently are; they perceive that BPA too often changes the terms of arrangements after they have been agreed to, and so on. Many of these utilities would be willing to pay a premium for power from another source to avoid experiencing this hassle factor.

But how large a premium would they want to pay? Assume for the sake of this example that a single person makes the power purchase decisions and also determines the value of the hassle factor for a utility. Until that person has decided how much money avoiding the hassle factor is worth, she faces uncertainty in power purchase decisions from her own un-made value decisions, and BPA faces a similar uncertainty when it tries to estimate how attractive its products will be in comparison to those of competing suppliers.

Implications of values. Value decisions that have not been made yet prevent complete evaluations of the relative importance of different

possible outcomes, but they do not substantially interfere with assessing what the world will be like. Values clarification techniques comprise one general approach that may be helpful — efforts to explore what values or principles are most fundamental to a stakeholder, whether that stakeholder is the decision-maker or someone with whom the decision-maker has an amicable relationship. Clarifying fundamental values can often illuminate the importance of secondary or derivative values.

Another way to treat value uncertainty is try to grant another person the status of trustee or guardian, empowered to identify values on another person's behalf, chiefly when someone is believed to be incapable of specifying their own values for any of several possible reasons. A related technique is to identify people or organizations that are similar in relevant ways to the pertinent stakeholders, and whose values have already been established. This clearly has major risks of mis-characterizing the stakeholders' values, but the need to proceed with the decision may make the risks worth taking.

A very different approach is to attempt to influence or even mold the value preferences of other. Political campaigns often aim at this kind of influence, either by trying to get voters to accept as their own the values espoused by the campaign, or by trying to persuade voters that the voters' values, presumed in this case to be fixed, ought to lead the voters to agree with the campaign's goals.

When a new issues comes up about which people have not yet determined their preferences, it can take a while for people to understand how the issue relates to other issues they care about and to develop their

own opinions. When the decision-maker genuinely cares that they develop their own views on a new issue, it may be helpful to seek attention for the issue before people will need to express their values so that they will begin the process of learning about the subject and start thinking of it in ways that will facilitate a later evaluation. When the decision-maker has a strong stake in just which way people come to feel about the issue, this approach has drawbacks: people will have more time to develop their own views, which may not coincide with the decision-maker's, and other people may have opportunities to influence the development of those views. How far the decision-maker may legitimately go to manipulate the information others obtain and to influence the values they develop is often ethically or legally constrained, especially for those who are perceived to have a responsibility to be impartial or who stand to gain significantly by particular outcomes.

Value uncertainty differs from action uncertainty, and poses a more difficult obstacle to the application of reason. When you don't know what someone else will do, it may make sense to devise a plan that would work out as well as possible under all possible actions that others could take, i.e. to pursue robustness. But when you are not sure which types of consequences are most desirable, it doesn't make sense to try to devise a plan that minimally satisfies all the values being considered, i.e. to pursue a kind of robustness across values. This amounts to treating values as random, and the meaning of value is incompatible with random selection of values. Under other circumstances, a compromise plan can be very important — when you know that two different values are equally

important to you, or when different stakeholders have different values but are each entitled to influence the decision. These conditions, however, involve certainty about conflicting or competing values, not uncertainty about values.

Value decisions are required as input into rational decision processes as rationality is ordinarily conceived. When they are not available, rationality cannot proceed. (Meta-rationality, however, can suggest ways to help obtain the value decisions.)

III.4.1.2.1.3. Group decisions and preference aggregation

Many of the decisions that may affect the decision-maker's choice are matters of collective, not individual, choice. In one sense, the uncertainty from these un-made choices is the same as that produced by un-made choices of individuals. Uncertainty from this source can, *in principle*, be completely resolved by acts of will. However, resolution of the uncertainty stemming from un-made collective decisions is *in practice* sufficiently more complicated that it amounts to another source of uncertainty and merits separate consideration. Disagreement among people or parties can cause uncertainty over their collective actions.

One major complication in group decisions is that while an individual who has not yet chosen an action or a prioritized set of values can under many circumstances be interrogated; it may be difficult or impossible to interrogate each member of a group in a similar fashion, especially when the number of members of the group is large. Even if all of the individual preferences were known exactly there might be uncertainty about the group preference, either because the method of aggregation has not been

specified, or the aggregation has not yet been performed. Once the group of stakeholders has been expanded beyond a single decision-making person, it can be difficult to determine the set of people whose interests are to be taken into account.

Illustration. BPA faced uncertainty due to un-made collective choice of action in an illustration from section III.2.9.2, when an initiative drive placed a ballot measure before the voters of Oregon forcing closure of the Trojan nuclear power plant. In the time after certification of the initiative and before the election, BPA knew that in a few months a plant contributing some 300 aMW to BPA and 1,000 aMW in total to the region either would or would not be available. This is a very large resource uncertainty for a region with an average load of around 20,000 aMW.

More complicated than uncertainty over the outcome of a single future decision to be made by a group is chronic uncertainty over how to value and rank the possible outcomes of a set of BPA decision options that affect many parties. There are several different views of who the stakeholders are whose values ought to guide BPA decision-making: ratepayers, those who purchase electricity from BPA's whole sale customers; only those ratepayers who live in the Pacific Northwest; all consumers of electricity in the Pacific Northwest; all people whose environment is affected by BPA decisions; unborn generations who will inherit a world partly shaped by BPA's decisions, etc. Even with a specific definition of the set of stakeholders whose values are to guide the decision-making it is very difficult to actually ascertain a confident assessment of their individual values, let alone their collective values.

Implications of group decisions and preference aggregation. An attribute of un-made group decisions that makes them easier to deal with is that the choice of any one individual is less important when it will be combined with many others — future decisions of groups, especially large groups, are less vulnerable to the effects of the changing of one person's mind. Statistical assessment and probabilistic forecasting may be quite helpful. Opinion polling is one variety of statistical sampling that is widely used.

On the other hand, simply asking a group what it will do is less likely to be successful in reducing uncertainty than asking an individual, since groups cannot be queried as if they were single entities. When samples of groups are selected for direct querying about their intentions, the results are less directly indicative of the ultimate decision than when an individual is queried, while still being prone to the risk that intentions will change before the decision is made. For smaller groups, querying each member can be feasible; this is often used in attempting to predict the outcome of an imminent vote by a legislature. It is also more difficult to reach an agreement with a group than with an individual about what it will do.

Another difficulty posed by un-made group decisions is that the process of reaching a decision is itself more complex and may be less decisive. When an individual says she has made a decision, it is reasonable to accept her statement as factual (though of course minds and circumstances may change): an individual is automatically assumed to have the right to speak for herself. It is often less clear who has the right to speak for a group. It

many cases, a specific process, such as the casting of ballots on the second Tuesday in November, has been agreed upon, but in other cases there may be no identified procedure, and thus no clear way to tell when a group has reached a “decision”. Attempts to provide a structure for such a decision can reduce uncertainty, or make a favorable result more likely.

Mobs are a special kind of group in which the social and interpersonal dynamics may make predicting the behavior of the aggregate more difficult than predicting individual actions. The characteristics of the members of a mob may be overwhelmed by the mood of the group, and the group can do things few or none of the members would do individually. Preventing the formation of a mob, or the transformation of a crowd into a mob, is probably a more effective way of dealing with the uncertainty over a mob’s behavior than attempting to influence or control the mob.

III.4.1.2.2. Actors may change

We expect that the laws of nature change slowly if at all (though our apprehension of them may change dramatically). The laws or other regularities governing the behavior of people and organizations, however, are subject to transformations of vast scope. An implicit assumption in most attempts to predict the behavior of other people is that the people will not change substantially. This assumption underlies any use of information about people’s prior behavior as a predictor of their future behavior. Thus, knowing that a person will or might undergo a significant change makes predicting their behavior more difficult. These fundamental changes amount to changes in human *parameters*, while predicting behavior corresponds to predicting a human *variable*. Models with well-

calibrated parameters are useful in forecasting future values of variables. forecasting the values of the parameters themselves generally requires a much more sophisticated and more comprehensive model whose variables are the parameters of the first model, or else enough experience with similar systems that heuristic descriptions of common patterns of parameter changes have been developed. An extreme case of this uncertainty is when some of the actors whose values are important have not yet been born. Their values will be shaped by experiences whose outline we can only guess at.

Organizations are similarly susceptible to transformations. Public institutions may be modified by the people via initiatives or legislatively by their representatives. Privately held firms can be purchased by new owners with new goals. Both public and private organizations can change course when new managers assume control.

Illustration. BPA and other organizations face major uncertainty about employee performance. Even when good information is available about performance in previous jobs, there is always the possibility that an employee will in a sense not be the same person in the future, for reasons that may or may not be job-related. Some employees respond to increased responsibility by developing greater maturity and judgment, some respond to it by becoming more egotistical and authoritarian. Appointment to the Supreme Court is regarded as an experience that can permanently change new justices. Employees may develop new skills, suffer health problems that impair their productivity, or undergo major shifts in their values, thus changing unpredictably the contribution the employee makes to the

organization.

Implications of actors may change. Some of the ways in which the nature or regularities of behavior of people and organizations change can be predicted to some degree. For example, engineering start-up firms will be dominated by engineers, and as these firms grow, the importance of acquiring managerial and marketing skill will increase; the nature of the demands made of the chief executive will become more organizational and less entrepreneurial. The impetuosity of youth is often tamed over time by growing sense of purpose and responsibility. The fact that these changes are fairly common developments in situations that occur frequently allows for preparing for their possibility, even though they cannot be predicted with certainty.

When possible changes do not produce early warning signs, only generic precautions can be taken. For example, while the balance sheet of a firm might indicate that it is an attractive take-over target, it will not indicate much about the corporate philosophy of potential new owners; individuals are likely to be changed by some types of experiences, such as personal crises or having their first child, but the nature of the change is hard to predict. Generic precautions include avoiding plans that depend critically on absence of change, and explicitly allowing for numerous options and other mid-course corrections.

III.4.1.2.3. Unpredictable behavior

Even when people do not change in fundamental ways, there is always a chance they will do something surprising. This might be due to a particular mood, lack of sleep, an obscure detail in a complex situation, a

creative brainstorm, a lapse of concentration, or many other factors.

An assumption virtually always adopted in game theoretic analyses is that other agents will act rationally (attempt to act rationally and succeed). This assumption provides a foundation for attempting to predict their behavior from knowledge of a situation and of rationality. Under typical field conditions of imperfect information about other agents' goals, resources, beliefs, etc., this is hard enough. In many situations, it is difficult enough to understand what "rational" would mean to a particular person or organization that predicting their behavior on the basis of rationality is a formidable challenge, may even obscure the depth of our ignorance about them. Decisions that depend on any condition that varies from moment to moment are harder to predict than those more influenced by enduring factors. Agents who attempt to act rationally may make errors. They can adopt approaches that are fundamentally non-rational, such as pure inspiration, or by interpreting scriptures. Even a much weaker assumption that agents' behavior, whether rational or not, will be consistent over time, thereby allowing prediction by extrapolation from past to future, is often unjustified.

Illustration. In 1975, a husband and wife bombed transmission towers 11 times in the Pacific Northwest, threatening further attacks and electrical blackouts if their extortion demands were not met. While they did not cause a blackout, they did disrupt the transmission system. The actions of individual agents such as these are not predictable, and inasmuch as they may have a significant impact on BPA's system, they contribute an irreducible element of uncertainty into decision-making.

Implications of unpredictable behavior. People do not always even attempt to act rationally; they may choose to be impulsive. In many other cases, people genuinely strive to act rationally, but do not appear to an observer to be acting rationally. Apparent nonrationality of other actors may be due to the actors' genuine nonrationality or to the observer's lack of information which would otherwise make the actors' behavior appear to be rational, but in either case the actors' behavior can be hard to predict. Not all nonrationality complicates prediction, of course; while an observer might think that it would be rational to put part of a bonus paycheck into a savings account, if a particular person always spends the entire bonus immediately, though that may appear nonrational, it is easy to predict how long the next bonus will last.

One of the main causes of uncertainty we can have about the future actions of other people, presumed now to be rational, is decisions those people face about which they are uncertain. The more uncertainty they face, the more likely it is that rationality cannot provide decisive recommendations to them, and therefore the decision will be made on a basis we cannot see and therefore cannot predict. Complexity of decisions others face can make their decisions hard to predict: while chess has clearly defined rules and a finite set of possible moves, the game is complex enough that a very rational move can still be surprising. In other situations, a stroke of genius may provide a decision-maker with a new option that we were not aware of — selecting the new option may be clearly rational while the existence of the option creates the surprise. Conversely, a decision-maker may surprise us when unaware of a

particular option — rational selection from an improperly restricted set of options may appear irrational to observers who see a superior option.

A particularly important phenomenon closely related to uncertainty due to possible error is false certainty about the future performance of others. A common failing of technology designers is to give insufficient attention to the fact that their systems will be operated by human beings who are liable to make mistakes (Perrow 1984). A system that is completely safe when operated correctly may be susceptible to catastrophic failure after a single operator error. The designers of systems should be more uncertain about the behavior of human beings, and should respond to that uncertainty by crafting systems to be less vulnerable in the event of operator error.

III.4.2. Evidence

The second major category of sources of uncertainty comprises evidence of the world (or of actors' actions and preferences), and evidence-like data that is itself fundamental, such as written laws. In the former case, the evidence is secondary, and only a representation of process or objects that are relevant (but only knowable through evidence, whether original or supplemental). In the latter case, the evidence, such as a written law, is itself the primary source of information, though a more fundamental source of data can sometimes be hypothesized.

In interpreting a law, for example, a court may attempt to gather information about what the *intent* of Congress was in passing the law. This distinction can be relevant in deciding how to attempt to reduce evidential uncertainty, since when the evidence is only a representation of a more

fundamental object, it may be possible to perform additional measurements.

When the evidence is itself the primary source, there is generally no additional measuring that can be performed, only additional attempts at

Figure 34. Uncertainty from *evidence*

interpreting the evidence. The impossibility of further assessment of the object also arises when the object is destroyed, when the author of a statement dies, if a required channel of communication is lost, etc. Both representative evidence and fundamental data give rise to uncertainty in the same ways, so they are both treated under the term “evidence” in as described next.

Even when the object or phenomenon of interest is completely determinate and has no inherent uncertainty, we can be uncertain because of characteristics of our data. The nature of human language begets additional uncertainty in linguistic evidence. Evidence from actors, whether from the decision-maker or from other actors, engenders uncertainty due to the possibility of error or deception.

III.4.2.1. Data, measurement

Properties of evidence can generate uncertainty in different ways. All measurement processes induce uncertainty. Secondly, decision-makers can be uncertain because they lack evidence that is potentially available. The third major uncertainty-causing property, conflicting sources of evidence,

is in a sense the opposite of incomplete evidence — excessive evidence.

III.4.2.1.1. *Imprecision*

The act of measuring objects and processes, bearing always some degree of *imprecision*, inescapably introduces uncertainty, as has long been noted:

Figure 35.
Uncertainty from *data*

For one thing, an empirical operation is always attended by error. Thus Lebesgue^[25], who strove so well to perfect the concept of mathematical measure, took explicit note that, in the assignment of number to a physical magnitude, precision can be pushed, as he said, “in actuality only up to a certain error. It never enables us,” he continued, “to discriminate between one number and all the numbers that are extremely close to it.” (Stevens 1968:850)

Many measurements carry with them an implicitly understood assessment of their own accuracy; indeed, it could be argued that any measurement that is accepted as worth quoting by more than one person has some commonly understood accuracy, even if the accuracy is not expressed in numerical terms. (Stigler 1987:287)

[Borel’s] crucial starting point was that the notion of the precise numerical value of a physical quantity is a mere mathematical abstraction that corresponds to no reality. In order to measure or to define a physical quantity it is necessary to give additional explanations. The greater the precision desired, the longer these explanations become. (Knobloch 1987:229)

As these quotations show, the errors that inevitably attend any measurement are well recognized as sources of uncertainty. The type of error considered here is a random error related to the precision or

²⁵ H. Lebesgue 1966. *Measure and the Integral*, K. O. May (ed.), Holden-Day, San Francisco

resolution of the measurement process. Systematic error, such as miscalibration of an instrument, does not cause uncertainty in the same way. Uncertainty only arises from systematic error when the *possibility* of systematic error is conceived of by the decision-maker; the error itself induces wrongness, not uncertainty.

Implications of imprecision. Characterization of imprecision is straight-forward in many contexts. Opinion poll results often include the margin of error; resistors are sold according to the tolerance around their specified value for which they have been tested, and so on. Numerical measures of imprecision can be used to statistically assess how much the imprecision weakens any conclusions that could be drawn from the measured data — essentially, a determination of the amount of uncertainty in conclusions arises from imprecision in the inputs. This is helped by having information about the distribution of the results of the measuring process (e.g. normally distributed with mean of zero and standard deviation of 1 inch).

Other approaches can be useful when this distributional information is lacking. For instance, if a value is described as being at least 400 and no more than 500, with no information about the likelihood of its being near the middle of that range, an analysis could be made using 450 as the value, followed by sensitivity analysis to determine if a change to 400 or to 500 would affect the conclusions. Methods based on interval arithmetic may be able to use ranges directly as inputs (Alefeld and Herzberger 1983). Additional methods are likely to come from the current vigorous research into *fuzzy logic* — ways to build algorithms and apparatuses that tolerate

imprecision (e.g. Kandel 1986; Kosko 1993).

If the imprecision results in unacceptably high uncertainty in the conclusions, it may be possible to use a larger sample size, or use the average of a longer series of measurements to produce a more precise value. Another route is to improve the measuring instruments or process.

Whether any of these approaches are worth pursuing depends in part on the other uncertainties present. In particular, it is not worth using scarce resources to reduce one imprecision when other uncertainties already have a larger impact on the conclusions. For instance, if a tank is being emptied at a rate of between 15 and 16 gallons per hour, and we want to know how many hours it will be until the contents are gone, it is probably not worthwhile to spend resources to reduce the range of uncertainty over the rate of emptying from 1.0 to 0.5 gallons per hour if all we can know about the current contents of the tank is that it is full, and it is either the 400 gallon model or the 500 gallon model, with no further information about which model it is. There is uncertainty in time remaining of 20-25% is due to not knowing which model the tank is, and only an uncertainty of 6-7% due to the uncertainty over the rate. The time to empty is in a range of 25 to 33.3 hours; if we knew that the rate was between 15 and 15.5 gallons per hour, the range of possible emptying times would still be 25.8 to 33.3, a reduction of less than 10% in the imprecision in conclusion for a reduction of 50% in the imprecision in one of the inputs.

Application of the Uncertainty Link Principle (see section V.2.2.3.2) is especially important in cases where some of the uncertainties stem from

measuring processes and some do not. Frequently the uncertainty in measured variables, or in quantitative factors in general, will be easier to reduce than uncertainty from other sources, even though the former uncertainty may have less impact on the certainty of the conclusions. It is tempting to reduce still further the measurement imprecision simply because it is feasible, but it is not necessarily helpful to do so.

III.4.2.1.2. Incomplete data < interaction 3 >

Decision-makers are often uncertain about an object or phenomenon because they do not have sufficient evidence. This can occur for sundry reasons, such as insufficient data collection, too little experience with a new phenomenon to develop adequate observations, or information loss during transmission of a message. The object of interest is knowable, but is not yet known. What would be needed to obtain more evidence varies across the different situations. Uncertainty about other persons' ranking of different values may be reduced by asking them (presuming that they know now or will decide upon being asked). Uncertainty about physical objects may be reduced by measurement, or by inventing better measuring instruments. One lawyer's uncertainty about the interpretation of a statute may be reduced by researching previous cases that have set precedents in that area.

Further evidence is not always obtainable. A person whose values or observations are desired may not be interrogatable for many reasons. Time may be too short to complete a search through libraries or databases. The evidence may be so expensive to obtain that the effort is not justified.

Illustration. Total potential fuel supply is an important variable to

estimate in energy fields. The total supply of oil, for example, depends on the quantity of oil existing underground and how much of it is economically recoverable. It seems reasonable to assume that the quantity is the same today as it will be a year from now, adjusting for extractions. In other words, uncertainty about the quantity is not due to variability in the phenomenon of interest, but due to the fact that we do not have enough evidence to know exactly how much oil there is. It is possible to obtain more evidence, through exploratory drilling, for example, but this is expensive, and estimates of the expected value of the additional evidence are used to determine how much drilling to perform.

Implications of incomplete data <Interaction 3 – value uncertainty>. The most satisfying way to deal with incompleteness of data is, of course, to acquire additional and sufficient data. This may be feasible through new research efforts, repetition or refinement of previous measurements, direct inquiries, or many other ways. When this is impossible — or not worth the expense, it may be necessary to modify the approach taken for treating the uncertainty caused by the incompleteness. In the most extreme cases, lack of data can prevent any quantitative or even any rational process from producing meaningful conclusions. Maximization of expected utility has demanding requirements for information, including probability estimates for the possible future states of the world. A strict interpretation of von Neumann and Morgenstern's work is that the lack of long-run frequency-based probabilities invalidates MEU techniques, though Savage and other subjectivists would use subjective probabilities and proceed with MEU. Further data

incompleteness, such as lacking information about how desirable each potential outcome is compared with the others, renders even the subjective version of MEU unworkable. However, other techniques can operate with less complete data. The maximin rule, for instance, neither requires nor takes advantage of the probabilities of future states of the world, and does not require a complete ordering of preferences over potential outcomes, and therefore may be useful when MEU cannot be.

One of the most important assessments to be made concerning incomplete data is the extent of the incompleteness and the magnitude of its impact on the certainty of conclusions. This is performed routinely in the analysis of experimental results, when statistical tests of significance are applied to determine the likelihood that observed results were due to chance rather than a causal mechanism under investigation. Sensitivity testing should be used to explore how much difference obtaining additional data could make.

Long-run strategies for reducing this problem include “growing” better information systems, focused on the type of uncertainty most troublesome for the individual or organization; research programs designed to increase the acquisition of new information; developing new and better information-gathering tools, such as measuring instruments or voter survey techniques; redesigning processes so that information that is difficult to obtain is not needed as critically; and improving the availability of previously acquired information (better indexes, search methods, intra-organizational communication, memory enhancement training). Short-run approaches include trying (or re-trying) existing methods of acquiring

information, relying less on quantitative techniques and more on subjective and judgmental methods.

III.4.2.1.3. Conflicting data

Data from different sources often conflict with each other, causing serious problems of aggregation and reconciliation for the decision-maker. If one person comes in from outside and reports that it is hot out, an observer experiences some uncertainty in interpreting the fuzzy adjective “hot”. That uncertainty is far less than the observer would face if a second person came in and reported that it was cool out there. Then a conflict between the two reports is apparent. Sometimes it is safe to assume that there is a single “out there”, with a “true” temperature, but often it is not safe. Perhaps the two people had come in from different sides of the building, one in a breezy shade, the other in still sunlight. Perhaps one was dressed much more warmly, or one prefers much warmer environs than the other. It may be important to examine the interpretive context of each of the conflicting data, and the reliability (accuracy and veracity) of each observer. Sometimes a new frame of reference can be developed within which the formerly conflicting data are seen to be compatible.

The conflict may arise between two statements from the same person but at different times. Economists often assume that people have stable patterns of preference; this implies that there should be some resolution possible of a conflict between different statements of preference — perhaps one was made with less care than another. A less mathematically convenient possibility is that the person’s preferences did actually change, and therefore each statement offers only an uncertain clue to the person’s

future preferences. Two conflicting statements may each be an original source, as in the case of laws or contracts. Then a resolution may come from deciding that one takes precedence over another, or that a new interpretation accommodating or superseding both must be developed.

Illustration. BPA's policies are guided and restricted by several major acts of Congress. The 1980 Pacific Northwest Electric Power Planning and Conservation Act (Northwest Power Act) required BPA to consider energy conservation along with generating resources when planning ways to meet additional load growth, and required that conservation measures be accorded a 10% cost advantage in numerical assessments. A previous law, the 1974 Federal Columbia River Transmission Act, requires that BPA's rates are to be "consistent with sound business principles". It is possible that there are some conservation policies and actions mandated by the Northwest Power Act that would produce higher expenses than revenues. Clearly, planning in such a way as to lose money is inconsistent with "sound business principles"; but how to resolve such a conflict between the two acts is not clear.

Implications of conflicting data. A natural reaction to noticing that two or more sources of data conflict to attempt to determine which one is right, and occasionally this can be helpful. One may be more authoritative, or one could have become corrupted. Frequently there will be no way to determine which is more right, and the most direct conclusion that can be drawn from the conflict itself is to note the extent of the uncertainty. For instance, Koopman describes a probability statement as an *asserted proposition*, not an *experimental proposition*, claiming there is no "single

crucial experiment” that can even in principle verify it (Koopman 1940:119). Subjective probabilities are, if anything, even less verifiable. If two consultants differ in how likely they believe Event E is to occur, that difference cannot be resolved by determining who is right, but ought to be taken as indicating that the uncertainty in the probability is at least as great as the difference between those two consultants.

Another response to a data conflict is to examine whether each source of data is illuminating a different aspect of the object under consideration. Someone viewing a cone from the side would describe it as triangular, while someone viewing it from above would describe it as circular; the three-dimensional cone contains both aspects, neither more correct than the other. Different mathematical models frequently will produce different answers when applied to the same problem: each has its own set of built-in assumptions, and is likely to be sensitive to some aspects of the problem and blind to others.

III.4.2.2. Linguistic evidence < interaction 2 >

Much of our evidence comprises statements in a human language. Linguistic constructs are prone to two types of uncertainty, ambiguity and fuzziness. Most of the ways in which a decision-maker might respond to these two types are similar, but the distinction has been a prominent one in the literature, especially with regard to the type of mathematical representation that is appropriate for each type. This work focuses on a

Figure 36. Uncertainty from *linguistic evidence*

broad look at different uncertainties, but for more detail on the representational distinctions, a reader could refer to Klir (1992b).

III.4.2.2.1. Ambiguity

A statement that is *ambiguous* is open to more than one interpretation. The interpretations differ qualitatively, not only quantitatively. As most individual words have more than one meaning, this is difficult to avoid in natural languages. The contexts within which different observers interpret linguistic evidence emphasize different connections to parts of a statement, and thus may occasion different interpretations. When the author is available for clarification, especially in face-to-face discussion, ambiguity can be sometimes be reduced by further communication. In other situations, such as the interpretation of a law passed by Congress scores of years ago, the author cannot provide additional clarity, and another method, such as clarification by a court, must be used to resolve the uncertainty. This does not reduce any uncertainty about the author's message, but does reduce the uncertainty about how that message will be interpreted henceforth.

Illustration. The rates BPA sets for its power products are of great concern to BPA's customers. It is widely repeated that the customers want BPA to offer "stable rates". This term is ambiguous, however. The customers may employ that term to mean "rates that don't change". BPA might interpret it to mean "rates that change only slowly and steadily". For instance, rates that increase at the rate of inflation are stable in the latter sense. Rates that always increased 3% per year, and are guaranteed to increase at that rate for the next twenty years are even more stable in

that sense, but may be unsatisfactory to customers. Another ambiguity in the phrase comes from the fact that rates can be described in either real or nominal dollars, and rates that are constant over a span of years in real dollars are increasing over that span in nominal dollars.

Implications of ambiguity. Ambiguity is more likely to be deliberately introduced in order to deceive another party than is fuzziness. It can be especially effective when the other party either has a predisposition to interpret a message one way — the “wrong” way — or can be led to interpret it that way. For this reason, the ability to simply detect ambiguity may be a valuable safeguard against purposeful deception.

III.4.2.2.2. Fuzziness

In contrast to ambiguity, *fuzziness* is a quantitative uncertainty. Fuzziness is uncertainty in *degree*, whereas ambiguity is an uncertainty in *kind*. English contains many terms that pertain to degree or amount, and many others that modify the meaning of the first type of words. All of these are imprecise. For instance, we might say that a man is “tall”, or “old”, or “fast”. We might modify these terms, and say the man is “very tall”, “kind of old”, or “not very fast”. These terms are all in common usage, and their meaning — in general — is well understood by native speakers, but their meaning is *fuzzy* with respect to, for example, how many inches in height a “tall man” would measure. In most contexts, linguistic fuzziness is generally helpful: greater precision would require much more effort, and would not usually be worthwhile — the extra time spent on it might actually degrade the communication.

Illustration. The amount of energy that can be generated from the

Columbia River power system and usefully applied to the regional electric load depends on the precipitation throughout the Columbia's drainage basin. Amounts that have been recorded vary from below 12,000 average megawatts (aMW) to over 19,000, with an average of about 15,500 aMW. In dry years, there is not enough hydroelectricity available to supply all of the power BPA's industrial customers want, and there may not be enough to provide the special flows that have been planned to help migratory fish. BPA's commitments to these and other programs can clearly not be allowed to suffer any more fuzziness than possible. If BPA agreed to provide power to industrial customers except in "dry years", or except in "extremely dry years", the comparison between the numerical estimates of the power available in the current year with those verbal terms would inevitably have to be made by a court, since the terms are simply not precise enough. Agreements of this nature would be much less problematic in interpretation if they were couched in less fuzzy, more measurable terms, such as "except in years at least as dry as the fifth-driest year in the 50-year historical record", or "except in years with less than 12,300 expected aMW of firm power".

Implications of fuzziness. Fuzziness and sureness of conclusion can be traded off for each other: "really hot" is a fuzzier description of the weather than "at least 100°", but in the absence of a thermometer, it is easier for people to agree with certainty that a particular day is "really hot" than to be sure that it is "at least 100°" outside.

Further implications of ambiguity and fuzziness. In most respects, ambiguity and fuzziness present the same challenges to a decision-maker

and can be treated in similar fashions.

The symbolic, intentional nature of linguistic evidence distinguishes it from other types of evidence. The way we view an explicit, straightforward attempt to communicate exemplifies this: the evidence, that is, the actual linguistic communication, is often taken to be an expression of an idea or intent of the sender of the message. What the receiver of the message wants to know is what the sender *intended* the receiver to know. When fuzziness or ambiguity prevents the receiver from inferring the intended message clearly enough, many avenues can offer hope of reducing the uncertainty.

The receiver could request another communication, or another longer communication, and could describe the uncertainty that she hopes will be reduced in the forthcoming message. Evidence about the sender's linguistic patterns may permit a clearer interpretation of previous or future messages. There may be other evidence that can be obtained that helps illuminate the sender's intention, such as previous decisions by the sender, or interpretations others have made of the sender's current or earlier messages.

Linguistic evidence, at least messages, differs in another way from measured evidence. Suppes (1975) argues that the approximating aspect of measurement comes not from imperfect measurement but from the fundamental nature of measurement. This denies the validity of conceiving of "true values" with regard to which measured values can be assessed. In the case of messages, however, it is often the case (though not always) that the sender does indeed have a clear, unambiguous meaning in mind, so

that the notion of the “correctness” of the meaning inferred by the receiver appears valid. However, since no direct comparison of the sender’s intention with the receiver’s inference can be made without being mediated by additional communication, it is difficult to apply this potential standard of accuracy with precision.

For longer-term problems of linguistic uncertainty in a series of communications, developing a mutual understanding of terms, a particular code, or shorthand expressions can help. Training in clear communication for the sender or in understanding idioms and nuances for the receiver may also help.

Contracts, laws and regulations may present other issues. There are occasions when they can be treated as communications from the drafters of the document to later receivers. For instance, some ambiguities in the treaty entered into by the U.S. and Canada over the benefits of some Canadian power projects on the Columbia river were resolved by referring to earlier drafts of the treaty in order to infer the intentions of the drafters. Courts may settle cases in part via attempts to infer the intent of Congress in phrasing a bill in a particular way. In other situations, however, the words in a document are required to stand by themselves, ambiguous or not. The uncertainty may be partially reducible then by specification of a particular linguistic context or “sub-language” to be used in interpreting the document, e.g. “as commonly used by native speakers”, or the habits of a particular profession. Actors may agree among themselves how they will henceforth interpret the wording (courts are also called upon to do this). This amounts to *designing* the meaning more than *discovering* it.

Linguistic ambiguity and fuzziness are by no means always avoided. They can be important to people with similar but not identical views who want for whatever reason to agree on a particular statement, such as a contract or a public announcement. Fuzziness and ambiguity in the common statement can let each of them interpret it slightly differently and therefore find it acceptable. If the ambiguity masks a difference that is more than slight, however, there could be trouble. Fuzziness and ambiguity are also used in attempts to deceive others.

III.4.2.3. Evidence from actors

When evidence is provided by actors or agents, it may be either inadvertently or deliberately misleading to the decision-maker. A decision-maker who recognizes this possibility will experience uncertainty over how much confidence to put in the evidence. These will be defined here as uncertainty due to possible error or possible deception

respectively, with the understanding that it is the awareness of the possible error or deception, not the possibility itself, that causes the decision-maker's uncertainty.

III.4.2.3.1. Possible error

Evidence, whether gathered by others or by oneself (or one's group-self) may contain *error*. The possible flaws are innumerable, including using an inappropriate or flawed data-gathering procedure or instrument, making mistakes while employing an appropriate procedure, not checking the

Figure 37. Uncertainty from evidence from actors

results, checking them improperly, bad recording or communicating of results, and so on. While it is often easier to feel confidence in one's own efforts or those of one's group, such efforts are vulnerable to virtually all of the same potentials for error. It is not the errors that may be made, but the decision-maker's apprehension of the *possibility of error* that induces uncertainty in this way.

Perhaps the most insidious form of error is the *assumption*. No examination or analysis can be performed with the making of many assumptions, and these often have two dangerous characteristics: 1) they have a large impact on the results of the analysis; and 2) they are largely invisible. While a decision-maker is still unaware of these assumptions, she would be falsely certain, or at least less uncertain than she should be if the assumptions were acknowledged. Experienced analysts often point out the significance of the assumptions that underlie any model or even perspective on the world, and would urge decision-makers to be *more* uncertain in the conclusions they draw if they have not taken explicit account of the assumptions on which any analysis rests (e.g. Morgan and Henrion 1990:38; Ascher and Overholt 1983:105; 257).

Illustration. The process of entering data into a computer record contains many possibilities for error, and most organizations have procedures for checking the accuracy of the data entry. A method often used by secretaries for checking the accuracy of their typing is to ask a second person to proofread the finished product. This guards against two types of error: the possibility that the first secretary mis-typed, and the possibility that he was so familiar with the material that his own double-

checking did not catch the error.

Implications of possible error. The possibility of human error in preparing evidence includes gross as well as minor errors and unlike measurement error, which is typically assumed to result in only small deviations from “true” values, the possibility of error includes the possibility that the wrong variable was measured, that an entire component of a measurement process was omitted, or that digits were transposed.

Prior experience with the specific actor can provide information about that actor’s reliability, or experience with a class of actors may yield general characterizations of the likelihood of error, sometimes including inferred probability distributions. Sometimes people who process a particular kind of data are required to be licensed, and any mistakes they make jeopardize re-licensing, so they have additional incentive to be accurate and their clients have more reason to be confident in the processing. Training programs can help reduce the possibility of error, as can guidelines about how long a duty shift a person is expected to perform with breaks or time off.

When this kind of doubt arises about specific information, it may be possible to re-check the processing; having a different person perform the re-checking further reduces the uncertainty by decreasing the likelihood of making the same error twice. If the possibility of error appears to be too high (according to whatever bases for judgment are available), it may be prudent to choose actions that depend less critically on the suspect information, or to devise contingency plans to deal with any consequences

of the information turning out to be erroneous.

For the decision-maker aware of the general importance of assumptions and the uncertainty that they ought to induce, *assumptional analysis* can be a useful way to uncover and clarify the key beliefs on which any analysis or disagreement rests (Mitroff and Linstone 1993). Will Rogers got it right (section I.4): it is better to be aware of your assumptions, and uncertain as a result, than to be ignorant of them, and falsely certain.

III.4.2.3.2. Possible deception

When some of the evidence comes from parties with clear interests in the outcome of a decision, it is possible that the other parties will attempt to deceive the decision-maker in order to induce her to make a particular choice. Anytime the decision-maker suspects some of her evidence may be deliberately flawed she experiences uncertainty due to possible deception. Deception can range from falsification and fabrication of facts to more subtle actions such as presenting un-fabricated data in a fashion that makes a specific misinterpretation likely, or withholding of some relevant information, and “unconscious” deception. Uncertainty from this source requires not only the possibility than another party may attempt deceive the decision-maker, but also that the decision-maker be aware of that possibility.

Illustration. In 1993 BPA defined a set of Strategic Business Objectives, one of which is to achieve high and continually improving customer satisfaction. BPA will measure this by asking its customers how well they think BPA is doing. Clearly these customers will always have an interest in seeing that BPA continue to improve its customer service, even

if it is already excellent. There will be a very natural pressure for them to report a degree of satisfaction somewhat lower than their actual satisfaction. While it is possible that some or all of BPA's customers will nonetheless resist this pressure and report accurately, the mere awareness of this pressure by BPA staff interpreting the survey results will cause a degree of uncertainty for the staff about the extent of the customers' (conscious or unconscious) deception. In the cases of some surveys, response errors can be expected to balance out, but in this case, all of the customers are likely to face the same pressure, and therefore the errors are likely to all be underestimates of BPA's performance. There is no reason to expect the total bias to be zero, nor indeed, any simple reason to expect it to take any other particular value — it remains a difficult factor to estimate.

Implications of possible deception. The possibility of deception differs from the possibility of error in that the intentions of another actor are directly involved. In the case of error, or random processes in general, we often feel justified in assuming that one situation is much like another, and therefore evidence about the first can be used to infer conclusions about the second. The involvement of another person's intentions, however, undermine this assumption. First there is the complication that another's intentions might differ quite markedly from one situation to the next, and secondly, intentions are invisible, so we may have no clues as to how others' intentions differ across situations. This greater relevance of factors unique to each situation make it much less likely that techniques relying on probabilities will be justifiable or useful than in the case of

error, which is in turn less suitable for probabilistic treatment than the case of uncertainty generated by physical processes.

Evidence can be checked for potential deception in various ways. Obtaining corroborating evidence does not completely rule out deception by multiple sources, but reduces the hazard. Similarly, finding conflicting evidence does not indicate which evidence is most valid (if any). Lack of consistency, either entirely within a body of evidence or between it and other evidence, justifies concern, but again may no more prove deception or error in any particular body of evidence than consistency guarantees accuracy.

The real-time clues that can be used to assess whether or not another person is lying appear to be subtle enough that it has not yet been feasible to construct sets of rules that would allow a naive observer to apply them effectively. It does appear that experienced and skilled negotiators, poker players, and criminal investigators can develop an above-average ability to detect lying, but this knowledge appears to be difficult to codify.

III.4.3. The Decision-maker

Decision-makers can be uncertain about well-measured and well-assessed deterministic phenomena when they are unable to process the evidence fast enough, when they have problems interpreting the evidence, or when the mental models they

Figure 38.
Uncertainty from *the decision-maker*

employ are inadequate for understanding relevant phenomena or predicting the future states of relevant variables. This is essentially the domain of *bounded rationality* discussed earlier in section II.3.3 (along with the fact that information is not free, so that resolving incomplete evidence, discussed above in section III.4.2.1.2, is non-trivial).

III.4.3.1. Processing capability < interaction 1 >

A key ability that a decision-maker must have in order to make rational decisions is that of processing the information he has obtained fast enough to be useful for the decisions.

Facts do not speak for themselves clearly enough in most cases; they must be assembled sensibly, perhaps added or multiplied together, their implications traced out to logical conclusions, and their meaning for the decision-maker inferred through judgment.

This processing takes time to complete, and depends on memory and various other resources.

Figure 39. Uncertainty from processing capability

III.4.3.1.1. Memory failure

A decision-maker may be uncertain because relevant information, previously collected, is not available. For individuals, this type of memory failure is usually a simple forgetting. The stored information may still in some sense “be there”, but efforts to retrieve it are not successful within the time required. Institutional memories are especially vulnerable to retrieval failure. For individual human beings, retrieval efforts are

generally intuitive, involving little or no prior preparation, although memory training techniques can sometimes dramatically strengthen associative chains during memory storage and enhance the later retrieval. Organizations often develop explicit mechanisms or procedures for storing and retrieving information. The challenge of ensuring that critical information existing somewhere in the organization will be available to those who need it, when they need it, typically requires a major allocation of internal resources.

Implications of memory failure. When previously acquired information cannot be retrieved, the short-term remedies are fairly limited, primarily comprising emergency measures. For individual memories, various techniques can be tried to rebuild an association with the needed memories, such as thinking of similar items, visualizing the setting in which the information was first acquired, attempting to “prime” the subconscious by trying very hard to remember and then moving the conscious mind on to other topics so a subconscious search can proceed. When the information has been stored physically, physical searches may turn up the original records or related information. Sometimes information that has been stored in computer disks that have crashed can be partially or wholly recovered by specialized software or hardware.

Long-term strategies to prevent this problem or to fix a continuing problem that has developed can be quite varied. Writing is perhaps the oldest and simplest strategy. When a decision is too complicated for a person to remember all of the factors at once, writing them down helps overcome the limitations of conscious memory. Organizing written records

provides many more strategies, from obtaining filing cabinets to creating sophisticated computer databases and search algorithms. Imposing an order on stored information is a key strategy; hierarchical order is especially useful in reducing the time or effort that later searches will require. It may be difficult to know just how much effort should be put into developing a better information-retrieval system, since most of the costs are incurred immediately and with certainty, and the benefits are reaped later and to an unpredictable degree.

III.4.3.1.2. Time, resource limitations

A key constraint on decision-makers' ability to resolve uncertainty is the time and other resources available for gathering and analyzing evidence. In some cases, increasing any one of several resource types could help reduce an uncertainty: more time could allow completion of an analysis that would answer a question, or assigning more people or purchasing a faster computer could speed up the analysis.

Illustration. Several species of Columbia River salmon have been listed as threatened or endangered, and alternative system operation strategies (sets of restrictions on when and how much water can be released from specific dams on the river) are being developed and considered. The effect of these restrictions on BPA's ability to produce electric power from the river is of great interest whenever a new system operation strategy is proposed. The only estimates of the power impacts that can be developed within hours are based solely on experienced intuition, and carry a high degree of uncertainty. With a few days of time available, estimates can be made using experienced judgment augmented

by spreadsheet analysis. If more time than that is available, mainframe computer models can be used to estimate the power impacts with smaller final uncertainty in the results. Some of the computer models require so many computations to be performed that upgrading to this year's version of a computer can cut days off the time required to produce an estimate of comparable uncertainty compared to last year's version of the same computer.

Implications of time, resource limitations. When the resource limitations cannot be ameliorated sufficiently and the residual uncertainty is still troublesome, the decision process may need to be adjusted. A process that would eventually completely resolve uncertainty but not until long after the decision opportunity has passed is of no use. If the process can be abridged, streamlined, simplified, or performed in an approximate fashion it might still be useful. While the pressures facing decision-makers within bureaucracies incline them to want thorough analyses that depend little on subjective judgment, judgments can often be made faster than analyses. Frequently the result of resource limits is that an uncertainty that is in principle, and even in practice at other times, reducible must be treated as if it were hard because it cannot be reduced in time or at low enough cost (i.e. it is a *virtually hard* uncertainty). This often requires adopting a fundamentally different process for reaching decisions, such as using scenario analysis instead of cost-benefit analysis, or using qualitative descriptions instead of quantitative descriptions, or using judgment instead of calculation, or even using a non-rational process for decision-making instead of relying on one in which it is important to

“know” what the right answer is.

In the long term, to prevent this type of problem from occurring, it may pay off to arrange to have more resources available, or more effective resources. This could involve developing faster methods of analysis or ways to incorporate fuzziness in data, or obtaining better computers and faster algorithms. It may be worthwhile to devise scalable processes provide the capability to be improved-upon when needed with much lower up-front costs (e.g. computer systems where additional cpu power can simply be added in without requiring major system adjustments, or staffs that have been trained in several disciplines or tasks so that moving them to a new project does not require as much retraining). Decomposable processes may be advantageous, so that a decision process can be divided into many, smaller decisions that can be tackled simultaneously. If the smaller decisions must be addressed serially, it will be harder to gain an advantage through decomposition.

III.4.3.2. Ability to interpret evidence < interaction 2 >

To yield insights into a decision, evidence must be interpreted and made meaningful. This requires both linguistic ability and contextual knowledge, two closely related concepts.

An insufficiency in either one can cause the decision-maker to be uncertain about linguistic evidence, even when the evidence itself is adequate (i.e. complete enough that other observers

Figure 40. Uncertainty from *ability to interpret evidence*

would not be uncertain in the same situation).

III.4.3.2.1. Linguistic ability

Understanding evidence presented linguistically requires familiarity with the language, and skill in using it. An extreme case is when a decision-maker confronts evidence in a language she cannot understand. In other cases, evidence, especially about complex situations, is difficult to understand, and demands considerable skill from the decision-maker. The process of extracting the meaning the decision-maker needs from the evidence can be a very active one. Organizations typically develop what amount to local dialects comprising acronyms, nicknames, and other kinds of verbal “short-hand” that speed up communication among those who know the dialect.

III.4.3.2.2. Contextual knowledge

Language cannot be made completely free of context. Indeed, this is a central challenge to people trying to send messages to other species: what context can we hope that we share with beings from other worlds? The answers must be drawn either from those aspects of the universe that are truly universal, such as the patterns of atomic structures or number theory, or aspects of our corner of the universe that others could observe, such as the structure of the solar system and Earth’s position in it. Many terms have more than one meaning, and the receiver of a message must use additional information not included in the message to know which meaning was intended. Without sufficient contextual knowledge, even perfect evidence about completely determined phenomena will leave the decision-maker uncertain.

Illustration. BPA publishes many documents each year for public consumption. Each is written by people who spend 40 hours a week thinking and talking about electric utility issues, using the specialized language of the industry. Members of the public often do not have a comparable familiarity with the industry's terminology, and they may be much more uncertain about the intended meaning of documents that can appear straightforward and clear to others with more experience.

Implications of ability to interpret evidence <interaction 2 – linguistic uncertainty>. The same methods can help reduce uncertainty arising from either inadequate linguistic ability or insufficient contextual knowledge. In the short term, the decision-maker can supplement his ability or knowledge by seeking assistance from others. Reference books may be useful. Comparisons with other passages that have already been interpreted may help, as would discovering another Rosetta stone.

In the long term, the decision-maker can seek to acquire additional skills and knowledge himself. Training can increase a person's ability to detect ambiguity. Further training can help a person to infer a "most likely" interpretation, when appropriate, or to highlight ambiguous passages in order to have other parties clarify their intents. For organizations, skill-building can take the form of developing training programs for employees, or nurturing pockets of expertise of a particular nature, or providing standardized glossaries of terms that have specialized meanings within the organization's field. Computer translators are becoming sophisticated enough to usefully augment the decision-maker's skills in some situations, and computers are likely to continue to grow in

their power to enhance or complement people's linguistic skills.

III.4.3.3. Mental models

A mental model is a set of assumptions or images that represent the decision-maker's beliefs about some aspect of the world or about how the world works. The assumptions may be explicit, formulated as homilies or mathematical rules, for example, or implicit, outside of the awareness of the decision-maker, for example, a subconscious belief held by an individual, or the composite of rarely-stated beliefs in a firm's corporate culture. "Model" here does not imply mathematics or formalism of any sort, merely that decision-makers act on the basis of their representations, or models, of the world, not on the basis of the "true nature" of the world.

Uncertainty can arise specifically from the use of mental models when the models are incomplete, incorrect, or when there are two or more mental models that conflict with each other, even with an abundance of unambiguous evidence about the immediate situation. Mental models often are used in the answering of "what if" questions, i.e. to predict the consequences of a particular choice and world state (the role of function c in the GDM).

Figure 41. Uncertainty from mental models

Implications of mental models. A factor that complicates attempting to reduce uncertainty generated by problems with our mental models is that these models often act as perceptual filters, and it is difficult to

ascertain what aspects of the world your own filters are leaving out. A similar problem occurs when your mental models do not explain a phenomenon well enough to permit the accuracy of prediction you would like — how do you know if there is any other mental model that would explain it better without having the other mental model to test? (That is to say, how do you know that a phenomenon that appears to be random is actually random and not just poorly understood?) Referring to deterministic, chaotic processes, Casti notes that

The observed output of a dynamical process can be generated by a fixed, purely deterministic rule, yet be indistinguishable on the basis of any standard statistical test from a completely random process. (Casti 1990:243)

As causes of uncertainty, incompleteness and incorrectness of models are quite similar in effect and in the difficulty of identifying the model as the cause of uncertainty. There must be some reason for suspecting that the object being modeled is more regular than current models suggest. This doesn't mean that an identification of the model as the problem must be made before attempting to improve the model — it is not uncommon to work hard at improving or replacing a model on the mere hope that the phenomenon of interest is more regular than current models suggest. Clues about the possibility of creating better models can be gotten from statistical tests of randomness, or from tests such as the *correlation dimension* and the *K-entropy* (Kolmogorov entropy) that can point to deterministic but chaotic causal factors (Casti 1990:248).

It is often very difficult to distinguish incompleteness from incorrectness in a specific case. The reason for attempting to distinguish them is that the ideal resolution of uncertainty caused by the former is to

elaborate and complete a current model while the ideal resolution in the case of the latter is to find (and complete) an alternative model. Often the only practical approach in the long run is to attempt to do both and compare the results; the short-term approach is to treat the uncertainty as virtually hard. One of the benefits of recognizing that incorrectness or incompleteness of models is a potential source of uncertainty is that it prompts examination of the models and assumptions on which the “normal” view of the decision-maker’s situation is based, and can thus help to reduce false certainty.

III.4.3.3.1. Incomplete models

Incomplete models may represent much of the relevant world, but simply not cover the situation facing the decision-maker, and therefore leaving the decision-maker uncertain about which choice to make, like a map with a hole in it. Novel situations often lie outside the domain of relevant mental models. Teenagers’ models of how to behave around members of the opposite sex, and their models of what would please members of the opposite sex, are often very incomplete at the time of their first dates. Members of the team that built America’s first atomic bomb had incomplete models of the explosive potential of the atmosphere itself: they knew quite a bit about the chemistry and physics of combustion, and could calculate the approximate amount of energy that would be released by the bomb, but their representations of the world did not indicate whether or not the atmosphere would catch fire, and this concerned many of them deeply.

Another important way in which mental models or other models can be

incomplete is to be simplifications of the objects they represent. Often this simplicity is a major advantage, for instance by making the most important dynamics easier to understand. Nonetheless, a result of simplification is uncertainty, or lack of precision, in the output.

Illustration. The impact on human health of exposure to electromagnetic fields (EMF) has recently become a topic of widespread concern. Some studies appear to show that people who have experienced more exposure to particular types of EMF have suffered from higher incidence of health problems. We would generally say that we “understand” these health impacts when we have a model of how the field affect human cells that is both consistent with other knowledge of electromagnetism and human biology and consistent with the our observational data. We do not have such model now. For example, some studies seem to associate larger or more frequent health impacts with *smaller* levels of EMF strength, a finding that is at odds with other accepted models of causes of disease such as carcinogens. Decisions in which EMF could be important are made more difficult by the current uncertainty from our lack of a model explaining how different electromagnetic fields can affect human bodies.

Purposeful incompleteness — simplification — is an advantage of some of the energy planning models used by BPA. “ISAAC”, for example, divides the year into a few seasons, while some of the other models in use divide the year into 14 periods (twelve months, two divided in half). ISAAC’s simpler representation of time means that there is considerable uncertainty in its results as to the month-by-month distribution of effects.

The advantage the simplification offers is a reduction in the detail required for its inputs, and faster running times than the more chronologically detailed models. ISAAC's developers have been aware of the trade-offs of complexity and accuracy versus simplicity and speed.

Implications of incomplete models. [Many of the implications were covered above under Implications of mental models.] The primary effect of incompleteness of models is uncertainty. If an observer can know that a more complete model is potentially knowable, she could describe this uncertainty as a false uncertainty. Decision-makers are not always able to characterize their own situation in such a manner, and even doing so might do little to help them, except to encourage them in their efforts to improve their mental models. When a model is deliberately simplified, a decision-maker is aware that more complete models are possible, but prefers the simplified model, and the uncertainty that the simplification causes ought to be described.

III.4.3.3.2. Incorrect models (false certainty; false uncertainty)

Mental models that are incorrect can cause not only uncertainty (both true and false uncertainty²⁶), but also false certainty and error. This is analogous to a map that is complete but out of date.

Illustration. During the 1960s, demand for electricity in the PNW grew at a very steady rate from year to year. Forecasters in the region shared a mental model of this load growth that included a number of

²⁶ See section I.4.

factors. Demand (D) was a function of population (P), the number of electrical devices employed (N), the efficiency of the devices consuming electricity (E), and other factors (O):

$$D = f (P, N, E, O).$$

This model did not include the price of electricity — it was thought that demand was fairly inelastic. As a result of the two oil shocks of the 1970, the price of electricity rose sharply, and it became apparent that consumers did in fact decrease their consumption when their cost increased. This showed that the model forecasters had used was incorrect²⁷, and that cost (C) should have been included:

$$D = f (P, N, E, C, O).$$

Implications of incorrect models. [Many of the implications were covered above under Implications of mental models.] Our models of the motion of the planets let us predict phenomena such as solar eclipses with great precision. Earlier models based on what we now call superstition could not have done so: people employing such a model would have been far more uncertain (or falsely certain) about the occurrence of planetary events. But how would they come to recognize that they had the wrong model and that it was causing them to be more uncertain than otherwise necessary? “We believe the world acts according to such and such, and therefore we cannot predict what will happen.” And the result is that they indeed do not predict it. Only evidence of a greater regularity than their model ascribes to the world could suggest the inadequacy of the current

²⁷ Incorrect instead of incomplete, since the model did represent the situation at hand, just incorrectly.

model. This evidence could come from an alternative model which demonstrates greater predictability and therefore greater regularity in the world, or from data where “regularity” itself can be observed or suspected without necessarily knowing the specific pattern. (Evidence might appear that indicates a consequence of using the wrong model: the model may imply that the world is *more predictable* than subsequent evidence shows. In this case, the incorrectness of the model would be causing *false certainty* or *error* rather than unnecessary *uncertainty*, which fall largely outside the scope of this work.)

III.4.3.3.3. Conflicting models

When multiple models of the world conflict it may be because one is fundamentally better, i.e. more useful for making the predictions asked of it, as Copernicus’s model of planetary motion was eventually shown to be more useful than the model of epicycles it replaced, or because each model represents only some of important features, but is “right” in that aspect. An example of this sort is provided by the wave theory and particle theory of light. Each theory captures aspects of the behavior of light that the other misses. Resolution of the conflict in the first case comes from demonstrating which model is more right, perhaps by designing experiments that could refute each model. Resolution in the second case comes from constructing a more comprehensive model that incorporates the truths of each of the conflicting models, or perhaps a meta-model that guides the selection from a set of models a representation to be used in particular circumstances.

Illustration. A great many of the recent electrical generating plants

burn natural gas — it appears to be the technology of choice for the 1990s. Consequently, forecasting world-wide supplies of this fuel is of increasing importance to electric utilities. How much gas can be economically extracted over the next 30 or 50 or 100 years affects the desirability of constructing additional gas-burning generators, and of permitting the electric utility industry to become more dependent on natural gas.

The traditional explanatory model that underlies forecasting efforts suggests that natural gas and oil are both the decomposed remains of ancient life. Over half of the known recoverable deposits lie in a very small region in the Middle East. While any model representing a supply mechanism, such as this one, that cannot be measured with precision induces uncertainty, another uncertainty in the supply arises from an alternative model of the origins of Earth's natural gas. Thomas Gold has described a different model which posits that a completely separate process is responsible for the creation of many natural gas deposits (Gold 1987). He proposes that most gas is from inorganic, stellar origins, and that incredibly large quantities of this gas lie deep in the earth. If his theory is correct, then the world supply of gas is far greater than estimated by the traditional model. Each model leads to forecasts with their own uncertainties, but the uncertainty arising from not knowing which model to use is of a different nature.

Implications of conflicting models. Unlike the case of incomplete or incorrect models, the case of conflicting models can at least be recognized. A very natural and common question to arise immediately is “which one is right?”, and it is very convenient when that can be determined. Sometimes

there is room for compromise, as in many political discussions. For example, it currently appears that congressional Republicans have a model of society in which harsher punishment for criminals is the most effective way to reduce the rate of crime, while the model held by Democrats suggests that reducing the intensity of the social causes of crime is the most effective. The result is that Congress passes some laws making penalties more severe and some laws that attempt to modify conditions the Democrats believe lead to crime. Congress doesn't settle the issue of which model is the most valid, but it is able to get on with its business.

In other situations, compromise is less likely, especially when the question about which model is "right" is interpreted by at least one party in an argument as referring to a moral sense of right in contrast to rightness as accuracy, predictive utility, or simply which model will be adopted for the time being. Firmly ideological opinions also make it more difficult to test and improve models, since ideologues have an incentive to overlook data that threatens their model even if it could lead to changes that improve the model's accuracy or utility. An open-minded attitude that welcomes all new evidence that challenges existing models is likely to be helpful in resolving conflicts.

In a great many situations, none of the models is "right", as in the case of light qua wave versus light qua particle just mentioned. When each model captures part of the nature of the object of concern their joint use will inform us better than either model could alone. (Some would say that both models of crime are valid, but the more strident members of Congress do not typically admit that their opponents' model has validity; they

compromise for practical reasons, not for the technical reason that a combined model might be more accurate.) The existence of conflicting models may indicate that none is satisfactory, and that a better one needs to be found, possible by combining existing ones into a more comprehensive model.

When it is not possible to resolve a conflict by testing, improvement, combination, or other methods, the decision-maker may face uncertainty not only over what consequences *would* occur if he adopts a particular alternative, but over what consequences *could* occur. This not only makes probability tools inappropriate, but makes application of decision rules that do not rely on probabilities, such as the maximin rule, difficult to apply. Finding a robust alternative, one that appears attractive if not optimal no matter which model is used to evaluate it, is the best that can be done when the conflict cannot be resolved rationally. Barring that, the decision-maker could look to non-rational means for selecting a model to use in a rational decision process, or use non-rational means for the decision itself.

III.4.4. Interactions

III.4.4.1. Interaction 1: complexity and processing power

Clearly there is an important relationship between the complexity of a situation and the processing capability of a decision-maker. A situation that might leave a less capable decision-maker uncertain about the future state of a situation may be entirely predictable to a more capable one, or even to the first decision-maker if more time were available. It is the complexity of a system relative to the processing power of the decision-

maker that is relevant: a less complex object or better information processing for the decision-maker could reduce the uncertainty. When either the decision-maker's information gathering or processing capabilities insufficient in a particular setting, she will be uncertain (or else falsely certain). To paraphrase Arthur C. Clarke:

Any sufficiently complex deterministic process is indistinguishable from randomness.²⁸

The fact that human beings are a part of or are inside of the system of the universe may imply ultimate limits to how much of the universe we can understand.

Illustration. The control center for BPA's generation and transmission system receives large amounts of data every second on the state of the system. The output of the system's generators must be matched to the system load moment by moment. The pattern of instantaneous variations is too complex for unaided human brains to respond to quickly enough, but human brains augmented by computers are able to do this. The uncertainty an unassisted human being would face in controlling generation levels can be reduced to a tolerable level with sufficient processing power.

Implications of the interaction of complexity and processing power. Because the degree to which limitations of processing capability may cause a decision-maker to be uncertain clearly depends on the complexity of the decision situation, and vice versa, it is worth considering

²⁸ Clarke: "Any sufficiently advanced technology is indistinguishable from magic."

responses to both sources of uncertainty whenever either is noticed. That is, when processing power is inadequate, the responses described above for complexity may be helpful, such as decomposition in simpler sub-problems; and when complexity is seen as a problem, the responses described above for limitations in processing power may be helpful, such as memory enhancement or increasing computer speed.

III.4.4.2. Interaction 2: linguistic uncertainty and linguistic ability

The impacts on a decision-maker's uncertainty of the ambiguity and fuzziness of linguistic evidence depend on more than the details of a linguistic structure. Some people may be better able than others to interpret a structure and derive a specific meaning. Expressions specific to a sub-culture may be far clearer to members of the sub-culture than to outsiders. Some people have larger vocabularies, or more experience with a variety of sub-cultures, which lets them recognize phrases as belonging to a particular sub-culture and requiring an interpretation within that context. The linguistic skill of an observer can offset some linguistic uncertainty.

Illustration. The contracts BPA executes with its customers often have many provisions and contingencies. It can be very difficult for a layperson to understand what BPA is or is not permitted within a contract to do in a particular situation. To lawyers and others experienced at interpreting the contracts, understanding the nuances of phrasing and the meaning of the technical wording, the permissible actions are far less uncertain. They have come to collective understandings of the terminology, based on common utility and legal practices and court rulings. Thus,

phrases that may have a wide range of possible meanings to the lay public have come to be understood much more narrowly and specifically to the initiated.

Implications of the interaction of linguistic uncertainty and linguistic ability. While ambiguity and fuzziness in linguistic evidence and linguistic knowledge and skill are different causes of uncertainty about the meaning of linguistic evidence, in many cases the uncertainty could be reduced by improving either the clarity of the evidence or the skill of the interpreter of the evidence, regardless of which cause is deemed to be more fundamental.

III.4.4.3. Interaction 3: incomplete evidence and decisions not yet made

While economists often posit that each of us possesses a utility function which describes the relative preferability of various possible consequences, there are many cases in which it makes sense to speak of needing to decide which of two or more values is more important, or needing to decide what you really want. Some such questions may never have been posed before, so an actor may never have been prompted to make an assessment of the relative importance of a particular set of values. An economist might suggest that the uncertainty is due to our not having yet uncovered the true set of values (whether the relevant values are those of the decision-maker or of other stakeholder); others would say the set has not been chosen yet, depending on whether they believe the utility function is pre-existing or not.

Illustration. After several species of salmon were officially listed as threatened or endangered, plans to operate the dams on the Columbia

River system have been developed to enhance the survival of juvenile fish migrating to the ocean and the survival of adult fish returning upstream to spawn. While the Endangered Species Act (ESA) requires mitigation efforts and does not prescribe balancing their benefits against their costs, plans have also been developed for other species of fish for which there are no ESA-mandated preservation efforts. These plans reduce the power-producing capability of the river, and the region incurs increased costs as it replaces this lost power. Is the fish enhancement worth the cost? Decision-makers who manage the river system do not know. There are many uncertainties about the benefits and costs of the enhancement measures, but a key type of knowledge that is missing is the relative preference for fish enhancement compared to cost-saving. Whether this is considered to be incomplete evidence about part of the world, i.e. the preference structure of the stakeholders who are deemed relevant, or an un-made decision, i.e. how much weight to give to each of these two values, depends on whether one believes utility functions already exist and only need to be discovered or whether they are sometimes only potential, still needing to be created.

Implications of the interaction of incomplete evidence about values and un-made value decisions. The previous two interactions each deal with an uncertainty that is a joint function of two different sources. This interaction deals rather with an uncertainty that could be attributed to either of two sources according to the observer's untestable assumptions about whether values are fundamentally part of a person's being, and need to be *discovered*, or whether they are choices that need to

be *decided upon*. When the former model is adopted, resolution of uncertainty about the values that relevant stakeholders will have is seen to be reducible through elicitation of responses that indicate what the stakeholder's values are (and have been all along). From the perspective of the latter model, reducing this uncertainty should be approached by using techniques that invite and facilitate the stakeholder's act of will.

Why does this matter? It seems intuitively obvious that the manner in which a person is invited to make a choice qua act of will may affect the particular choice made. This makes the particulars of how another person's values are investigated quite important. In the other case, where we hold to a belief that the values are preexisting, the values ought to appear the same through the lens of any unbiased elicitation procedure. Much of the evidence to date indicates that elicitation procedures do exert a strong influence on the results (McCord and de Neufville 1983; Munier 1988:16).

The same considerations apply more or less to a decision-maker whose own values are not clear — she could engage in activities designed to explore her fundamental values so that they can illuminate the ones she needs to know for her current decision, or she could engage in activities whose purpose is to facilitate her making up her mind. This is not to say that values-clarification activities cannot facilitate the *choosing* of values, just that not in all cases is it an explicit purpose to precipitate an act of will.

How would these issues affect uncertainty about the values of a group, such as an organization or an electorate? In some cases it may be clear that many or all of the members of the group have clear values, and that

the challenge is to construct a synthesis of these individual values. In other cases, the individuals involved need to develop an understanding of their values these can be synthesized into group values. For example, when a manager asks a board of directors to make a decision on an issue with which the board has a total lack of experience, he is more likely to get a thoughtful decision if he sets the stage by providing background information and suggests linkages to previously enunciated values, making it easier for the board to construct their position on the new issue.

III.4.5. Summary of sources

The distinctions among sources of uncertainty described above are depicted together in figure 42 below

Summary of implications. I have defined three general sources of uncertainty: the world, evidence, and the decision-maker, with a further subdivision of “world” into “natural world” and “actors”. Uncertainty stemming from variability in the natural world must generally be accepted: it can be clarified and better understood, but often cannot be reduced. It is generally hard, parametric uncertainty. Basic responses to uncertainty caused by the natural world are clarification, devising better alternatives that may not be as sensitive to the uncertainty in the world, and in the end, gambling — taking an action whose outcome is not known. Decision analysis has been developed for this source of uncertainty, and will generally be more successful with it than with uncertainty from the other sources.

Uncertainty generated by actors in the world is the province of game theory, politics, bargaining, and studies in communication. These

uncertainties may be reducible by gathering more information from or

Figure 42. Uncertainty from all sources

about the other actors, or by negotiating with them to produce a change in the world and its uncertainties. These strategic uncertainties can be hard, virtually hard, malleable, or soft.

Uncertainty generated by evidence may be reduced by research that produces better evidence. This might involve first attempts at gathering a particular sort of data, developing new instruments, extending the spatial or temporal scope of observation, or improving the interpretation of existing evidence. Unless better evidence is in principle unobtainable (which would suggest that the uncertainty stems from the world), these uncertainties are generally soft or, when improving the evidence is either not currently possible or not cost-effective, virtually hard. There are some situations when the world could be changed in such a way as to make an evidence-caused uncertainty irrelevant.

Uncertainty caused by characteristics of the decision-maker may be reducible by “self-improvement” — memory training, faster computers, improving language skills, acquiring more contextual knowledge, inventing better mental or computer models.

A further distinction that can be drawn among sources of uncertainty concerns the severity of error that they can cause. Some uncertainty, such as that coming from data imprecision, linguistic fuzziness, or in some of the cases of other sources, generally implies the possibility of relatively minor error: the uncertainty is a matter of *degree*, of the exact magnitude of a variable. Other uncertainty, such as that arising from conflicting data, ambiguity, possible error or deception, fuzzy data in a chaotic system, for example, carries with it the potential for very large errors: the uncertainty

is, or can be, a matter of *kind*, not only degree. A fuzzy answer is likely to be “close”, but an erroneous or deceptive answer may be far off the mark.

Technological choices can have a similar impact on the potential magnitude of errors. For example, an error in reading an analog clock is likely to be an error of minutes, and not too many at that. An error of several hours is much less likely. Errors in reading a digital clock, however, are just as likely to affect the reading of the digit on the extreme left, the most significant, as they are to affect the reading of the digit on the extreme right, the least significant.

III.5. Decision context

The context of the decisions we face can have a major influence on the kind of responses to uncertainty that we can make. Several significant contextual characteristics are discussed next.

III.5.1. Unique vs. repeated decisions

The distinction between decisions that will be faced over and over again and decisions that are unique is important. When virtually the same decision will be faced repeatedly, the decision-maker has rational grounds for believing that her long-run average results will be close to the mathematically expected results. When a decision is unique, it may be impossible for the mathematically expected results to eventuate. For instance, when two people bet each other \$1 on the toss of a fair coin, each has an expected gain or loss of zero, but that is not one of the only two possible outcomes — one will lose a dollar, and one will gain a dollar. On the other hand, if the two people make the same bet every day for a year, it is very likely that neither will be very far ahead, even though a dollar will

change hands 365 times.

Uncertainty is not present when the consequences of the action are actuarially certain. (Shackle 1957:107)

Shackle argued vehemently that the application of objective probabilities to unique decisions is illegitimate (Ford 1983:55).

Most gambling games can be played identically many times (assuming the player does not go broke during the process). Most planning decisions are essentially unique: while similar decisions must be made again and again, the circumstances of each, the options available, or the potential consequences are likely to be different, rendering each decision an essentially new task.

Illustration. When an employee of an organization such as BPA retires, she often has the option of receiving her benefits in a lump sum or taking it as an annuity. Valuing the lump sum option is complicated by the uncertainty she has about how long she will live, which affects how fast she would want to consume the money. If she does not live very long, she doesn't want to die with much of the money unspent — she could have used it faster and enjoyed it more. But she surely does not want to use it up before she dies and have nothing left to live on. For her, this is a unique situation. For the organization, however, it is merely one of a great many very similar situations. The organization can agree to pay her a fixed annuity for the rest of her life, no matter how long she lives. Some employees will outlast their expected lives, and receive more money than the organization calculated, but others will live less long than expected. As long as the organization's actuarial tables are reasonably accurate, it can calculate quite well its future outlays, and provide its retirees a far more

certain support during retirement than each could provide for himself.

Implications of unique vs. repeated decisions. A main argument for using expected utility as a criterion in repeated choices is that the decision-maker can expect her average results over the long run to approximate the expected results. How long “the long run” must be can’t always be known very well, and the argument actually only says that over the long run it is less and less probable, though still possible, that the average result will differ significantly from the expected result. This is far short of a guarantee, but sufficient for the successful operation of casinos and insurance companies. In one-shot decisions, it is not infrequent that the expected result *cannot* occur. An argument could be advanced that a decision-maker who made all decisions on the basis of expected utility would be better off in the long run, even if no two decisions were similar, but this does not seem a very powerful justification. In a sense, the particular outcome of one uncertain event in a series of like events does not much matter, while the results of a one-time decision matter greatly. The expectation of the result is a plausible proxy for the actual results only in the case of repeated decisions. The possibility, raised in section III.2.9, that a person might have a preference over security level, hope level, or regret, further weakens the importance of the expectation in a one-shot decision.

Lopes (1988:413) argues, in fact, that the acceptance of expectation as a (as *the*) criterion by what she calls “the American school” illustrates its central difference from “the French school”. The American school is built upon “mathematical sophistication”, while the French school has instead

based its perspective on “psychological intuition”, and the two disagree utterly on the relevance of mathematical expectation for one-time decisions, as expressed by a leader of the French school:

There is a yawning gulf between the conduct that is rational in relation to a random choice related to a non-repeatable event, and the conduct that is rational when the event will recur very often and in similar circumstances. (Allais 1979:490)

Shackle’s claim that there is no uncertainty when there is actuarial certainty suggests that converting instances of one-time decisions into a series can be an effective way of eliminating uncertainty (Shackle 1957:107; see section II.5.1). When the consequences of an uncertain event can be translated into money, insurance may be an effective way to gain the benefits of actuarial certainty even for one-time events.

When similar if not identical decisions will be faced, it may be more cost-effective to design decision rules for the whole set of decisions rather than to analyze the consequences of particular decisions. (Elster 1989a:25)

III.5.2. Equilibrium

It can be important to compare the rate of change of the *parameters* of a dynamic system and the rate of change of the system *variables*.

Many inferences about dynamic systems depend on the assumption that the systems are in equilibrium. For this assumption to be reasonable, the state variables must be capable of change that is fast relative to the rate of change of a system’s parameters. Thom’s Catastrophe Theory (Zeeman 1976), for example, describes systems whose equilibrium points undergo jumpy changes as a result of smooth changes in their parameters, but this requires that the state variables change enough faster than the

parameters that the system can follow and remain at equilibrium while the equilibrium point moves. Economic theories of markets and firms also frequently rely on the assumption of equilibrium.

There are at least two ways in which a system can systematically fail to reach equilibrium while constantly moving towards it. In the first way, the system's parameters are changing fast enough that the location of the equilibrium changes, and though the system moves towards equilibrium, the moving target is never reached. In the second, there are frequent enough perturbations or shocks that the system never settles down into equilibrium, though the equilibrium point itself is unmoving. In the first case, not just the state of the system but the nature of the system is changing, and this can preclude being able to calibrate a model's parameters well enough to predict future states without a high degree of uncertainty. In the second case, valid models may be made, but the perturbations to which the system is subject introduce sufficient variation that predicting future states of the system may again be very difficult.

Another possible reason for a system to exhibit non-equilibrium behavior is chaos. Chaotic systems may be temporarily stable or periodic, but they are best known for irregular, non-periodic behavior even when completely determined by simple nonlinear relationships. These systems do not always have equilibria.

Implications of equilibrium. The main reason for examining whether a system under investigation can be presumed to be in a state of equilibrium is that most of the tools developed in classical and neo-classical economics have assumed equilibrium (Nelson and Winter

1974:886; Simon 1988/89:3), and thus the implications and tools of economics may not apply to disequilibrium situations:

The problem ... is that optimality and equilibrium go hand-in-hand. Under conventional methodology, if we require our agents to behave optimally, then we are restricted (at the market or economy level) to analysing equilibrium: whether it be equilibrium states, equilibrium adjustment processes, equilibrium adjustment of adjustment processes, or whatever. True disequilibrium analysis requires us to abandon optimality.

To my mind, this appears to indicate what we should be doing next. Rather than extend indefinitely optimality analysis to increasingly complex environments, we should be modelling true disequilibrium through the use of *reasonable* 'rules of thumb'. I stress the word 'reasonable', as I believe actual behavior is generally reasonable: people do not go around using stupid 'rules of thumb' — certainly, not for very long! Naturally, I would like to say what 'reasonable' means, and how one models 'reasonable' behaviour; unfortunately, my thoughts along these lines are insufficiently well-developed to commit to print. (Hey 1981:242)

In an environment not in equilibrium, then, the very notion of optimizing loses its meaning, and cannot be made operational. This makes striving for optimality dangerous in two ways: 1) the actions taken in that effort may not actually be helpful; and 2) decision-makers may be falsely comforted that they are doing the best they can. If optimizing is proscribed, satisficing, i.e. the avoidance of options that are not good enough, is probably the most useful objective, supported by the use of heuristics that are "roughly right" in Keynes's phrase. This may also be impossible to achieve, but unlike optimization, it has a chance.

It should not be surprising that there has been more enthusiasm in technical literatures to date for attempting to find out if chaos theory can lead to better predictions of the behavior of market indices than for examining the implications of chaos for equilibrium-based economic

theory. On the one hand, the behavior of markets has been historically difficult to forecast, and any new tool with a hint of potential is eagerly tested. On the other hand, chaos, with its potential for sudden and large shifts in system behavior, represents a considerable threat to the machinery of mainstream economic theory which assumes not only that systems have equilibria, unlike chaotic systems, but that the systems are at the equilibrium points.

III.5.3. Surprise-prone environments

Another important context for decisions is one that is “surprise-prone” (Dror 1986b): there is a high probability of a low-probability outcome (even the “most likely” outcome has only a small chance of occurring). An environment that is characterized by increasing, pervasive uncertainty (about many factors) presents problems for planning that go beyond the problems of a single decision.

An environment can be surprise-prone for several reasons. If you know nothing about it, anything will be a surprise. Some environments are roughly stable in the sense of not undergoing major structural change, but are nonetheless turbulent. Even if you know a lot about its current state, if the fundamental state is changing rapidly, surprises are likely. Situations of great complexity may have so many possible states that no single state is very likely. Or they may be complex enough that prediction is very difficult. Surprises can also be generated deliberately by other actors. This is especially likely when the actions of individuals can cause large changes, or when the actions actors might take are particularly hard to predict, perhaps because either some of the actors or their circumstances are

unusual.

Implications of surprise-prone environments. In the case of the previous paragraph's first cause of surprise proneness, simply not knowing much about the environment, attempting to learn about it rapidly is clearly to be recommended. Courses of action that maximize learning should be adopted, such as experimentation, and encouragement and tracking of many decentralized efforts in lieu of a single centralized effort.

The second reason is environments that are "roughly stable" but turbulent; these offer the hope that boundary conditions can be discerned. If the set of possible conditions can be learned, and responses developed for each of them, the main challenge then is to learn to recognize which condition holds as quickly as possible, and to initiate the appropriate response quickly and capably. Each response may be made reflexively rather than by invoking an optimizing process, but it may be possible to optimize the strategy comprising the set of responses and their triggering conditions. Systems that are relatively insensitive to small perturbations, or that can be made more so by creating shock-absorbing mechanisms, are helpful.

The third possible cause is that the fundamental nature of the situation is changing rapidly. It may be very difficult to even infer the boundaries of what is possible, unless it is possible to understand something about the nature of the changes the situation is undergoing. It is far more difficult to assemble in advance a set of responses when the possible conditions cannot be bounded. Optimizing is clearly impossible; even satisficing may be quite difficult. Since novel conditions are more likely to be encountered,

creativity in responses will be more valuable than if the environment is roughly stable. Providing a wide diversity of skills, perspectives and knowledge bases will facilitate creation of a varied mix of possible responses, though increasing the diversity in options or in personality types will exacerbate the challenge of responding quickly. Describing circumstances with less detail — fewer significant digits, fewer but larger categories, coarser resolution, etc. — may improve response time, and make it less likely that people will be seduced into attempting to be more accurate than the situation warrants.

In the fourth case, when the surprising nature of the environment is due to its complexity, simplified representations of it may omit so much detail that prediction is not enhanced at all. More powerful modes of representation and prediction are needed — better mental models, faster computers, better ways to link together the insights of many people or many perspectives. Hierarchical representation is helpful for systems that are “nearly decomposable”.

There are two kinds of implications of surprise-prone environments caused by the actions of other agents, the fifth reason. The first is the adverse impact on the decision-maker’s options and response time, which is similar to the impact of surprises due to other causes, and can be dealt with similarly. The second is that the unsettling impact on the decision-maker may be due to deliberate choice by other parties, and may be providing advantages to these or other parties. In such circumstances, it may be helpful for the decision-maker to become less predictable herself, in order to disrupt the planning and decision-making of the other parties.

This would at the least make it more difficult for the others to devise strategies based on sound predictions of the decision-maker's behavior, and at best will prevent others from generating surprises that interfere with the decision-maker's planning and decision-making.

A tactic that can help no matter what the cause of surprises is to become more used to surprises so that they are less jarring. Simulations and increased experience can people become less shaken by surprises. A related approach is to use training and rehearsals to improve performance even when people become unsettled, and to put crisis management plans in place and ensure that they are suitable for being carried out by people who may be shaken.

All of the responses described above require time to implement. When there is not enough time for such strategic responses designed to decrease the amount of surprise, the decision-maker is left with trying to cope with surprises. The ability to respond quickly is clearly paramount. Responses that can be prepared ahead of time can help to save time when they are needed; the more general these are, the more likely they are to be useful in surprise-prone conditions. Judgmental responses are often faster than calculating responses. Procedures that are cumbersome are detrimental, not merely inconvenient, e.g. those requiring lengthy analysis, multiple levels of approval, or high standards of proof or justification. In extreme cases, the most rational mode of response may be pure improvisation.

A particular regimen recommended by Dror (1986a:257) begins by generating a list of potential surprises, perhaps by brainstorming, which is to be updated periodically. Each possibility is then characterized to

whatever extent is feasible in terms of its likelihood, the magnitude of its effects (both good and bad), and the practicality and cost-effectiveness of preparing for it. The list is then sorted, with highest priority going to those events with high probabilities, or high impacts, or high ease of preparation. Low-probability, low-impact events are ignored, while low-probability, high-impact events are divided again, so that those that can be easily prepared for are dealt with, and those that are difficult to prepare for are monitored and fed into crisis management preparations.

III.5.4. Low-probability, high-impact events

A special category of risky decisions comprises situations where there are possible outcomes that are fairly improbable, but whose impacts are severe enough to warrant attention. These low-probability, high-impact events present special challenges and merit special treatment, and the field of risk analysis and assessment has emerged out of that need. Some considerations to take into account are described next, but this is an established area of specialization, and a work of the broad scope aimed at here cannot do justice to this growing literature.

The cognitive difficulties that human beings have in dealing with very small or very large numbers vastly complicate risk analysis where such numbers are common for likelihoods and consequence magnitudes respectively. For example, Kahneman and Tversky (1979) assert that subjects often both over-estimate the probability of unlikely events and give excessive weight in their decision-making to their estimates (using MEU calculations as the standard of reference). Engineers are able to perform expected value calculations but are not very successful at

persuading others that their calculations are meaningful.

Events that can recur but do so only infrequently provide us with few observations on which to base estimates of their frequency, which leaves us with a large standard error for these estimates (Fischhoff et al. 1981:16). When lack of observations prevent satisfactory estimation of frequencies but the favored analytical tools require probabilities, risk analysts often substitute subjective probability estimates.

Aside from problems of estimation of the magnitude of consequence or the likelihood of low-probability events, the interpretation of these estimates is often difficult:

In fact, we have difficulty conceptualizing very high and very low numbers. This applies to astronomic distances at one end of the spectrum and to subatomic particles and very low probabilities at the other. While we can clearly distinguish between probabilities of 0.50 and 0.10, it is almost impossible for us mentally to distinguish 0.004 from 0.00004, despite the fact that the former is 100 times greater than the latter. (Linstone 1994:265)

Similar difficulties are presented by comparisons of dollars of cost, or millions of gallons of oil spilled, when the number run into 6 or 9 digits.

The possibility that the decision-maker may die or go bankrupt is a particularly important possibility to consider when it exists, and difficult to compare against more "normal" consequences, such as a change in relative prosperity. Risks to human health and to human life generally pose serious problems of comparability when they must be balanced against mundane considerations such as expense. With large potential impacts, and often high costs of prevention, it can be both psychologically and politically difficult to decide how to prepare for risks, especially when the estimate of likelihood is subjective and perhaps controversial.

Implications of low-probability, high-impact events. Many authors recommend dealing with low-probability, high-impact possibilities with techniques other than those used for “normal” possibilities (e.g. Dror 1988a:257), and this typically implies less reliance on strictly technical approaches.

Three major conclusions of Fischhoff et al. in their study, *Acceptable Risk* (1981), are that the level of risk that is acceptable depends critically on values, there is no “right” way to determine the appropriate level of risk, and decisions about acceptable risk change over time. Decisions about how much risk to accept cannot be settled by purely technical methods; they must be resolved through personal, organizational and political processes.

Perrow sorts high-risk systems into three categories: systems that are inherently very risky, and difficult or impossible to make less risky; systems whose benefits are essential and whose risks can be better mitigated, though perhaps with difficulty; and systems that are “self-correcting”, even though risky (Perrow 1984:304). He recommends that systems in the first category, such as nuclear power and nuclear weapons, should be eliminated, because their risks are too high to be justified by their usefulness. Systems of the second type are important enough that they should be kept, but greater effort should be undertaken to make them safer.

Perrow’s recommendation concerning the first type of high-risk system does not match that of most technical “experts” who believe such systems can be operated safely. Perrow attributes much of the difference in

estimation of riskiness between experts and citizens in general not to different perceptions of the characteristics of a hazard, such as the “dread” associated with it, the invisibility of the risk, or how voluntary the exposure to the risk is, but to the fact that experts generally do not attach as large a negative value to dread, invisibility, or involuntariness as citizens do. There is some potential for reconciling what Perrow calls the different “rationalities” of the two groups if they can understand each other better and understand the basis of their evaluation differences.

Bowonder and Linstone (1987:198) provide additional ideas that can be used to reduce the uncertainty in the management of high-risk systems. One category of tools is to reduce the likelihood of disaster. This can be done via technical means (e.g. decouple key linkages to make systems “safe-fail”²⁹), organizational means (e.g. provide better training, better manuals, or build crisis management and risk assessment into the regular procedures of the organization), or personal means (e.g. make it easy for individuals to draw attention to potential problems, and include a wide variety of viewpoints different from the standard technical and organizational perspectives). Systems can be redesigned; typically they were built originally to control technical sources of danger, but adding features to reduce the risk from non-technical sources, such as cultural differences in how people would operate the system in different countries. Research efforts may produce new technologies that are inherently less risky.

²⁹ This may be especially effective with chaotic systems [Gordon & Greenspan 1994:60].

III.5.5. Types of environments

Emery and Trist (1965) describe four types of environments, and the types of uncertainties that each can induce are significantly different (see also section II.4.7). These descriptions relate to either the state of the world or the function mapping actions and world states into consequences. In Type 1 (Random), uncertainty arises from the random distribution of goals and noxians. There is no regularity to use in search or prediction. In Type 2 (Placid), the environment is regular, but may simply not be well enough known. Type 3 (Disturbed Reactive) is characterized by the interaction of multiple agents, and the unpredictability of other agents' behavior introduces a new source of uncertainty. In the Type 4 (Turbulent Field) environment, the patterns of regularity change. Even if the current structure of the environment is well known, that knowledge is flawed by the possibility that the structure will change within relevant time horizons to make the knowledge obsolete. Summarizing these sources of environmental uncertainty:

Type 1. The environment is random and somewhat unknowable.

Type 2. The environment is knowable, but imperfectly known.

Type 3. The future behavior of other agents is not known.

Type 4. The environment changes faster than we can acquire knowledge.

Implications of types of environments. Types 1 and 2 below comprise the domain of parametric rationality (see section II.3.3.5); Type 3 is more or less the domain of strategic rationality, and in Type 4, the efficacy of rationality is questionable.

Type 1 -- Random: When the uncertainty is caused by the random distribution of “goals and noxiants”, better search strategies and better mental models will not be especially useful. Probabilistic characterization of the random distributions may be quite useful and appropriate (according to both the subjectivist and objectivist views of probability). It may be feasible to bound the possibilities, and to devise canned responses in advance for use in different external conditions.

Type 2 -- Placid: In Emery and Trist’s *placid* environment, the decision-maker’s uncertainty stems from inadequate knowledge, and an obvious remedy is better knowledge. Observation, programs of research, better facilities for storing, comparing, combining and retrieving information are good prospects. Probabilities may be useful, but as measures of the confidence the decision-maker places in estimations, rather than as measures of the frequency with which certain conditions are found. (Subjectivists would endorse use of MEU techniques with this kind of probabilistic data, but objectivists would not.)

Type 3 -- Disturbed Reactive: In a *disturbed reactive* environment, other agents exercise decision-making discretion, causing uncertainty for the decision-maker. This interdependence is the natural domain of strategic rationality and game theory, and both zero-sum and positive-sum games may be useful models. Probabilities as estimates of the actions of others are not recommended; if the decision-maker can calculate them, another agent could also, and could possibly benefit from deliberately changing his behavior to frustrate such predictions.

Type 4 -- Turbulent Field: In *turbulent fields*, while the behavior of

other agents can pose challenges for a decision-maker, the possibility of fundamental changes in the arena within which all of the actors operate decreases the usefulness of game theory, since the rules themselves are subject to significant change. Any approaches that rely on calculating may be handicapped by the possibility of fundamental structural change. The decision-maker may be more concerned about survival than about gaining a competitive advantage, which would ordinarily be a primary goal in a Type 3 environment. This distinction could increase the motivation for cooperative behavior among actors. Optimization here is difficult to define, and this environment more than the others calls for contingency planning and improvisation, pursuit of robustness, and avoidance of ruin.

III.5.6. The Role of time

III.5.6.1. Uncertainty as a function of time

One clear relationship between uncertainty and time is that uncertainty generally increases with space/time distance from here/now (Simon 1981:178 ff.) While uncertainty increases with the distance in the future, so does the prospect for gaining information before that future time is reached.

In general, the further a commitment to a line of action extends into the future, the larger the proportion of all the information entering into the final decision will be that is news coming in over the course of implementation of the decision. (Stinchcombe 1990:4)

A reasonable person would expect major differences in techniques for dealing with problems depending on whether their consequences are predominantly near-term or long-term, whether they are well-known or poorly glimpsed, whether there will be additional information available

after the impending decision is made, and whether there will be future opportunities for adjusting or correcting decisions made in the present.

Implications of uncertainty as a function of time. One method for incorporating the increased uncertainty accompanying assessments of events further into the future is to use a discount rate to decrease the weight given to benefits anticipated in later years. This is especially attractive when the costs are either known very well or are paid up front. When an important stream of future costs is uncertainty, this technique is harder to use. A negative discount rate could be used to increase the weight given to future costs, but this has the effect of giving a greater impact to the more distant and less well-known costs, which is not as intuitively sensible as giving less weight to future benefits that may or may not materialize.

Since the uncertainty in future consequences exposes the decision-maker to risk, explicitly including risk as a decision criterion is one way to compare alternatives whose future consequences differ in the precision with which we can know them. Other things equal, alternatives with longer-lasting consequences or commitments clearly present greater risks than those with only short-term impacts.

III.5.6.2. Speed of our learning and acting vs. rate of change in world

How fast we can learn relevant details about the state of the world compared to how fast the world can change affects the actions that we ought to take. If the world changes too fast for our knowledge to “catch up”, we need to use decision methods that do not require extremely high-quality information. There are other situations where the most effective

action is to increase the speed of our information acquisition so as to improve the timeliness of our information rather than increase its detail, precision, or accuracy.

Equally important is the relationship between how long it takes us to make decisions and how fast events develop. When events occur very quickly, a very quick method of making decisions is demanded. Many of the normal procedures of American government are modified during wartime for this reason.

Illustration. After the two oil shocks of the 1970s, Americans made significant changes in their patterns of consumption of electricity, and the rate of growth of electrical load dropped more sharply than utility planners had expected. The industry had previously believed that the price of electricity had only minor effect on consumer conservation, but price-induced conservation was part of the reason the rate of load growth dropped so sharply. It took several years for utilities to become sure of this phenomenon — and in roughly that same time span, the relative price of oil had dropped enough that price-induced conservation was no longer as important as it had been in the years immediately after the second oil shock. A major factor affecting electric load growth had arisen and then significantly retreated in the amount of time it took the utility industry to gather sufficient knowledge about the factor to incorporate it in its planning.

Implications of speed of our learning and acting vs. rate of change in world. When the rate of change of the world is high relative to the speed with which we can learn about the world, or compared to the

speed of our decision-making, efforts that take a lot of time are less useful and can even be potentially dangerous. In the former case, research is less likely to be helpful, since there is a risk that by the time we learn something, it will no longer be true. Fast modes of learning are clearly called for. This might mean using a larger number of smaller experiments that can be performed simultaneously, each looking at fewer variables, for instance by encouraging state governments to launch independent trial projects could produce more information in the same time than the same amount of money and other resources directed in a single Federal project. It is very likely to imply that more uncertainty will have to be accepted at the time of making a choice, and that generally argues for alternatives that include more options, contingencies, and fall-back plans.

Rapid change in the world relative to decision-making speed puts all calculating methods of decision support at a disadvantage. Computer modeling may simply take too much time, unless existing models can be quickly adapted to new uses. Much simpler, more easily modified models are needed, such as one-page visual models or informal mental models. Maintaining a set of conceptual models that can be quickly reexamined may be helpful. In extreme cases, where action is imperative and time minimal, it may be more rational to decide by improvisation or intuitive reaction than to attempt to decide rationally and risk being too late.

III.5.6.3. Un-simplifying the passage of time

As Robin Pope (1988) points out, the standard model of von Neumann and Morgenstern assumes that there are two points in time, with no actual passing of time involved in the transition from one to the next. The first is

the choice point, at which the decision-maker deliberates and chooses. Making the choice immediately triggers the second and last point, an outcome point where the consequences and utilities are all assumed to become known. Pope points out that in the case of a Bayesian view of statistical testing of hypotheses, there is no final outcome point and thus there is no post-outcome period: additional evidence may always be presented to be combined with the prior probabilities of the falsity of the hypotheses.

There are generally two periods of time other than the outcome point in which consequences may be experienced. The outcome “point” may in reality be an extended period of time during which a great many successive consequences occur. Frequently, the details of the consequences are not known clearly at the time of the initial “outcome”, and they may be subject to numerous additional uncertainties. Even when streams of future consequences can be foreseen, uncertainty in their significance generated by uncertainty over an appropriate weighting or discounting may be very large relative to the magnitude of the significance.

The other period lies between the decision and outcome points. Options may vary considerably with regard to how long the post-choice, pre-outcome period lasts, and this can be important to decision-makers. The standard model is obviously inadequate for a psychologically realistic treatment of gambling, among many other activities, for the gambler's experience in a post-choice, pre-outcome period is very important: this is the period where most of the excitement of gambling is felt, and where people experience “hope”. People do not gamble only in genuine efforts to

increase their wealth — indeed, gambling is generally recognized to be an abysmally poor way to do that. It is, though, for some people an effective way of getting very excited precisely by the uncertainty that exists only during the time before the outcome is learned.

For many organizations, the post-choice, pre-outcome period has real costs that we can refer to as the *cost of uncertainty*, which is distinct from risk. Suppose that a firm has excess factory capacity for manufacturing chemicals, and can choose to make a food product or a food additive. The firm will not know until it seeks approval of the food additive whether it will be regulated by the Department of Agriculture or the Food and Drug Administration, while it is clear to the firm that the food product would have to meet only Agriculture's rules. In either case, the additive is expected to be more profitable than the food product. However, the uncertainty about the which regulations would apply to the food additive means that the firm would have to take one of two measures: 1) wait until the regulatory status has been resolved after applying for the needed permits, or 2) begin preparing to comply with both sets of guidelines at once (instead of just one set). In either case, the firm bears additional cost due to the uncertainty about which rules will be operative. Thus, we need to consider not only the utility of the post-outcome period(s), but also the utility of the pre-outcome period(s).

In fact, expanding the scope to include pre- and post-outcome periods still serious over-simplifies the matter. The "veil of uncertainty" is not always lifted at once, not always lifted completely, not always lifted at all. The events that characterize different states of the world, and resulting

consequences, may occur without our finding out about them, or some may occur long before others. Options may differ in how long the decision-maker will remain uncertain about the outcome, and this can be an attribute of great interest to a decision-maker.

Consequences may begin to occur long before any resolution of uncertainty follows a decision, indeed can begin before the end of the decision process. A decision may be rushed because anxiety from uncertainty about what to do is great. A decision may be delayed because the uncertainty provided by the absence of a decision gives additional license to the decision-maker.

Characterizing decisions as having only two artificial points in time obscures consequences that occur after a decision and before the uncertainty is resolved, and obscures the fact that the uncertainty may never be completely resolved. Consequences may be experienced before the resolution that are quite different in nature from those that are experienced only after total resolution of the uncertainty of a gamble.

Figure 43 illustrates how the representation of time in the SDM needs to be expanded. The gray bars indicate possible ranges of time for several processes that the SDM assumes for convenience to take no time. The complexities hidden by the SDM include possibilities that new alternatives will be generated during the decision process, the potential importance of hope or fear (as an attribute of a specific decision option, not merely as a function of the decision-maker's whole situation), and the possibility that resolution of outcome uncertainty and the occurrence of consequences might both take place over extended periods of time, and may never

terminate.

Figure 43. Un-simplified passage of time

Implications of un-simplifying the passage of time. Perhaps the most straight-forward way to un-simplify the representation of the passage of time is to consider directly whether the consequences will all occur at a single point, as assumed by MEU (von Neumann and Morgenstern 1953:19), or whether there will be significant consequences before or after the "outcome point". There may be many decisions in which the errors caused by this simplification can be judged to be small enough to not merit further consideration.

When this is not the case, some way of assessing the importance of consequences occurring at different times will be needed. It may be possible to directly compare consequences occurring after the "outcome point" with each other. They may be of the same nature, merely occurring

at different times, as would generally be the case with payments or receipts of money. However, some of the consequences that can occur before the “outcome point”, such as hope or fear, are quite different from those that occur after the “outcome point”, and such comparisons may be much more difficult. It can be helpful to imagine the after-effects of consequences — those that would probably be judged to be significant even after they are experienced may be worth more weight in the decision, and those whose impacts are intense but brief, judged less important after the fact, could perhaps be given smaller weight (see Elster’s discussion of *akrasia*, i.e. weakness of will, *Ulysses and the Sirens* (1984)).

III.5.7. Once-and-for-all vs. correctable decisions

Some options will provide opportunities for later course corrections while others are made once and for all.

There are two kinds of decisions: those that are expensive to change and those that are not.

A decision to build the Edsel or Mustang (or locate your factory in Orlando or Yokohama) shouldn’t be made hastily...

But the common or garden-variety decision — like when to have the cafeteria open or what brand of pencil to buy — should be made fast. No point in taking three weeks to make a decision that can be made in three seconds — and corrected inexpensively later if wrong. The whole organization may be out of business while you oscillate between baby-blue or buffalo-brown coffee cups. (Townsend 1970:45)

The future monitoring and decision-making may be a cost which is undesirable as long as an uncorrected option works well. When the future environment is difficult to predict, flexibility may be highly advantageous. Should surprises occur and a plan meet with a state of the world that begins to lead to an undesired consequence, high flexibility would facilitate

additional action that more vigorously forestalls the unwanted outcome.

Implications of once-and-for-all vs. Correctable decisions. When the consequences of a decision will not all be evident immediately, the possibility of being able to take additional steps in the future to adjust or correct the original decision can be very valuable. Without this advantage, it is much more difficult to ensure that today's decision will not have some kind of highly adverse impact. This argues for a more defensive approach to making decisions that cannot be corrected later — placing a greater emphasis on what Dror calls “bad avoidance” and less on “good achievement” (Dror 1991:27). Thus, robustness becomes more important than optimality as a goal for the design and selection of decision options. An additional implication is that when it will be possible to correct the decision in the future, it is not worth spending as much time and effort on the decision now. Utility theory in general is less suitable for correctable decisions, because “[an] essential feature of von Neumann-Morgenstern utility theory [is] that it conceives of choice as being a once-and-for-all affair.” (Hey 1979:30)

III.5.8. Second-order consequences: anticipating the next decision

The real result of our actions is to establish initial conditions for the next succeeding stage of action. (Simon 1981:187)

The previous section suggests that future decisions can improve the outcomes from today's decision; this section describes another perspective on that link, in which today's decision can improve the prospects for future decisions. The outcome from one decision generally has an impact on both the set of states of the world and the set of available actions for later

decisions. One T Perspective on planning and how it differs from decision-making is that the goal of decision-making is to find and select the action whose (set of possible) consequences are the most desirable, while the goal of planning is to find and select decision-producing processes that yield a desirable stream of outcomes resulting from a long sequence of future decisions. This leads naturally to the inference that one of the more important kinds of consequence of any decision is the effect on the sets of actions available for future choices, and the range in desirability of the consequences that are associated with those actions. It also suggests an explicit role in planning for information gathering. Indeed, Michael (1973) and De Geus (1988) identify facilitation of learning as the chief function of societal and organizational planning. The goal in rational decision-making is to choose actions whose consequences are desirable. The goal in rational planning is two-fold: to ensure that actions with desirable consequences will be available throughout the future, and to ensure that good decision-making will identify, select and implement these desirable actions.

Implications of second-order consequences. A premise of MEU is that the decision-maker has a complete preference over *outcomes*. Indeed, the adjective “final” is often used (e.g. Kahneman and Tversky state that decision analysis recommends that decisions be coded in terms of “final assets” rather than relative gains and losses (Kahneman and Tversky 1979)). When the result of an uncertain decision is viewed as the initial setting for future decisions, it relates perhaps as much to the *means* of future decisions as it does to the *ends* of the decision at hand. It seems unrealistic to require that all rational decision-makers have complete

preferences orders over intermediate outcomes, if these are based on their relative utility for future decisions (whose results are, in turn, the initial conditions for yet further decisions). But without a complete preference order over the intermediate outcomes, MEU cannot be used. Heuristic analysis can help to indicate the desirability of alternative future positions or intermediate outcomes, but usually with enough more uncertainty about the exact utility as to render inconclusive decision processes dependent on calculation, such as MEU. Judgmental analysis and subjective evaluation must be used for all such situations except the most predictable.

For example, the game of chess, being a well-defined, deterministic two-person game, is completely solvable, in principle (von Neumann and Morgenstern 1953:125). It is reasonable to assume that each player has a preference over the possible outcomes, which are win, lose, and draw. However, those are not the outcomes of a particular *move*. The direct consequence of each move is to help define the initial conditions for the player's next move. The utility of this outcome can be measured only heuristically by assessing the strength or desirability of the player's position, and this is currently an inexact art. Lacking a preference order over these intermediate positions, the expected value of a particular move cannot be calculated, and MEU techniques cannot help.

IV. IMPLICATIONS OF UNCERTAINTY

The fundamental cause of trouble in the world is that the stupid are cocksure while the intelligent are full of doubt.

Bertrand Russell

IV.1. Introduction

The previous chapter identified ways in which uncertainties can vary. This chapter looks at the implications of those differences for rationality and methods supporting rational decision-making. A series of topics is examined so see what light is shed on them by the expanded range of uncertainties developed earlier. First, limitations imposed by uncertainties on the possible usefulness of rationality are considered. Then implications of uncertainty for the typical methods of rationality, including measurement, optimization, and maximization of expected utility theory are discussed.

The mainstream conception of normative decision theory and the types of uncertainty it addresses are then examined. Part of this examination is a review of a 4-stage model of the historical unfolding of attitudes towards uncertainty first applied by Smithson to developments within mathematics, and applied here to decision theory.

The chapter concludes with miscellaneous results of the current study, some ideas for future research, and some illustrations of the kinds of insights application of the taxonomy may facilitate.

IV.1. An assessment of rationality in light of uncertainty

I can stand brute force, but brute reason is quite unbearable. There is something unfair about its use. It is hitting below the intellect.

Oscar Wilde

When is it rational not to choose rationally? To start with, it may be *non-rational* but it is hardly *irrational* to decide to choose via non-rational processes. More than that, we can state that it is irrational to choose rationally:

- when it would be more rational to improve your options (and you shouldn't choose at all yet); or
- when the cost of rationality (resources, time) is greater than the benefit (e.g. and you need to choose faster or cheaper).

In the third case below, it may not be so much that it is irrational to “choose rationally” as it would be dishonest to claim that basing a selection on an indeterminate rational selection process actually constitutes “choosing rationally”:

- when the procedures invoked for attempted rational choice turn out to be indeterminate, and additional selection processes are needed.

When this happens, the decision may be made partly by rational processes, partly by other processes.

IV.1.1. Uncertainty as a threat to rationality

Uncertainty poses a serious threat to the effectiveness of rationality as a foundation for choice:

The failures of indeterminacy [of rational choice theory] appear to me to be quite serious. One way of assessing the power of the theory is to distinguish choice situations by two criteria: the importance of the problem and the number of agents involved.

'Small' problems, that is, problems in which the options do not differ much in value from one another, do not lend themselves to the rational approach. Either the options are equally good or it is not clear that it would pay to find out which is the better, or pay to find out whether it will pay to find out. 'Large' problems, in which the choice can be expected to have wide-ranging consequences, also tend to fall outside the scope of the theory. Preference rankings over big chunks of life tend to be incomplete, and subjective probabilities over events in the distant future tend to be unreliable. The theory is more powerful when applied to medium-sized problems like the purchase of a car or a house, but even here the question of optimal search is largely indeterminate.

Other things being equal, decision problems with one agent or with many agents are more likely to yield determinate solutions than problems with a small number of agents. By definition, one-agent problems have no strategic indeterminacy. With many sellers and many buyers, competition forces a unique set of equilibrium prices. With one seller and one buyer, there is often a large range of mutually acceptable outcomes and much indeterminacy concerning which outcome will be realized. A rough conclusion is that rational-choice theory is applicable mainly to one-agent and many-agent problems of intermediate size. Although precise quantification is impossible, indeterminacy is not a marginal problem that can be assimilated to 'friction' or 'noise'. (Elster 1989a:26-7)

Elster judges, then, that rationality is most powerful when used in "medium" problems with either one or very many actors, and that outside these limits problems of indeterminacy are very significant.

[In some cases the] "infinite regresses involved in games without solution and in the problem of the optimal amount of information prevent the rational-choice model from even getting off the ground. This does not mean that people in such cases cannot make rational choices, only that there is no choice which is *the* rational one. The situation is not a 'single-exit' one ... but this does not mean that all exits are equally good. (Elster 1984:156)

In such "multiple-exit" situations, a rational decision process fails to be decisive, and some other means of making the decision among the "exits" is needed. (As described in section II.3.4, some of the alternatives may have

been eliminated by rational considerations, even though the selection has not been narrowed down to a single alternative as desired.) Another decision process must be invoked. Pushing rational processes to identify a single answer beyond the point where they are able to draw distinctions on rational grounds produces an answer that should not be characterized as a “most rational” answer.

Elster is describing his perspective on rationality in general, not on specific rational approaches, such as rational choice theory, or maximization of expected utility, and it is thus not likely that future innovations can get around the problem he sees.

Stinchcombe points out that unavoidable defects in predictions of future events prevent make it impossible to be perfectly rational (in the substantive sense):

Rationality necessarily involves an analysis of the future, because the consequences that give purpose to acts are necessarily in the future. Thus all rationality is based on predictions of one kind or another, not on knowledge. Assuming that actors are perfectly rational, of course, implies that they are certain what the future holds, that in all relevant respects our notions about the future constitute knowledge. (Stinchcombe 1990:1)

This dependence on prediction imposes a severe limit on the power of rationality.

Klir (1985) recapitulates Warren Weaver’s three classes of complexity: organized simplicity, disorganized complexity, and organized complexity. He believes analytical methods, such as Newtonian mechanics, are often quite successful in dealing with organized simplicity, statistical methods can be powerful in dealing with disorganized simplicity, but that organized complexity is not well handled by either type of approach. This is not as

strong a conclusion as Elster's, since it allows for the possibility that new methods might handle organized complexity more powerfully than current methods, but does point out areas where currently rationality is more likely to be decisive.

Virtually all conceptions of rationality in the technical literature exclude from its scope the matters of identifying problems and opportunities, generating options, and holding appropriate values (e.g. selfish versus altruistic, impulse gratification versus delay of gratification, or the relative importance of consequences occurring at different times). These issues are very important to the people and organizations confronted with the need to decide under uncertainty, and can have a great bearing on how well their actions work out over time. Even if rationality is not directly useful in these matters as a guiding principle, considerations of meta-rationality (section II.3.3.6) can indicate how important they are and the consequences of paying them too little attention. It is irrational (meta-irrational) to focus too much on those aspects of a decision situation that are most suitable for rational solution.

IV.1.2. "Null Hypothesis": assess significance of uncertainty

A key step whenever uncertainty is present is to assess the significance of the uncertainty for decisions at hand. One straightforward way to do this is to adopt the use of a Null Hypothesis which states that no option is superior to the rest. If the Null Hypothesis can be rejected with the desired level of significance, then the uncertainty does not cause excessive indeterminacy in the rational analysis of the decision. If it can't be, then the results do not show that one option is significantly better than others,

and any apparent difference may be due purely to happenstance and have nothing to do with the actual desirability of the options. As the purpose of the analysis is to assist in selection of options that are better, an important question is, are they *really* better? When the Null Hypothesis cannot be rejected, the rational decision selection process is indeterminate, and another process needs to be adopted in order to produce a decision, i.e. a selection of a single option.

For example, suppose that a gamble is to be evaluated using subjective probabilities and MEU, and that all aspects of the assessment are perfectly certain (!) except the probabilities, each of which has a “subjective margin of error” of $\pm 20\%$. Suppose that two options have expected utilities of 2.45 and 2.15 respectively. The uncertainty surrounding the probabilities should be propagated through the calculation of expected utility so that the uncertainty surrounding the expected utility values. Then a test should be made to see whether it is possible to reject at the desired level of significance the hypothesis that the difference of 0.30 in expected utility is due to chance. If it is not, then the conclusion ought to be that the rational procedure adopted in this case does not declare one option to be a more rational choice than the other. This does not mean that the first option should not be chosen if a choice must be made — only that the selection should not be justified on the basis that one option is more rational than the other.

Results of intermediate clarity may still be useful. A number of null sub-hypotheses can be developed, such as Option A is not superior to Option B, Option A is not superior to Option C, and so on. Even when

much of the data is fuzzy, it may be possible to reject many of these, and conclude, for example, that Option A is superior to all other options except B, and similarly Option B is superior to all except Option A. Then the choice has been narrowed, and the task remaining for a subsequent selection process is smaller.

Techniques for assessing and processing measures of uncertainty or fuzziness are not as well developed as numerical methods for working with single-point (non-fuzzy) values. Ideally, the uncertainty in all input data would be allowed to propagate throughout the analysis so that the uncertainty in the output would be visible. Fuzziness could be assessed via ranges, margins of error, standard errors, probability distributions, possibility distributions, or perhaps many other methods not yet known.

If there are no measures of superiority that indicate that one option is *significantly* better than another, but there are measures that indicate slight (non-significant) superiority, can these be used to make a decision? Yes, with two cautions. The first is that there is a possibility that a better strategy would be to attempt to reach a significant conclusion, by generating a new and better option, or by performing another, more decisive analysis, and there is a risk that this possibility will be overlooked if the non-significant difference is adopted as the decision criterion. The second is that it is likely to be difficult to remember that the decision made on the basis of the non-significant criterion was decided on a basis no sounder than the flip of a coin. Typically individuals and organizations are likely to be far happier to have a decision made via rational analysis, even inconclusive rational analysis, than to have it made by flipping a coin, so

there may be P Perspective or O Perspective advantages to this, even though there is a T Perspective risk that the metaknowledge that the decision was made virtually randomly will be lost.

Relying on a rational selection process that does not appear to be determinate has advantages and disadvantages. The advantage is that it may be possible to use the process to produce a decision even though the assessments of the merits of the options do not show one to be significantly superior. When no obvious choice for the sort of selection process to be used next is apparent, this may be a big practical advantage. The disadvantage is that the result of the *rational* selection process was indeterminate, and pushing the process further to show a unique answer can produce a “winner” but cannot produce a “most rational winner”. Thus, the process, pushed that far, may be determinate but is no longer rational. The outcome depends so much on the particular ways in which the process is forced to a conclusion that its result is as likely to represent differences in technique as differences in the merits of the options. Analysts with integrity who are asked to do this become cynical or will leave; decisions that are made in this way and are described as representing a rational kind of “best” are open to serious challenge by others.

IV.1.3. Summary of when rationality is likely to be useful

Rationality is more likely to be indecisive (i.e. indeterminate) in situations with more than one but not many actors, and in situations of organized complexity. It has better prospects when zero, one or two or very many actors are involved, when the situation is one of organized simplicity or disorganized complexity, and when information is readily available.

While science can tolerate open-ended questions, decision-making often cannot, and ways to proceed must be found.

Policy-reasoning must face the question what to do when evidence is inadequate for judgment, even after validity criteria are relaxed to meet policy-reasoning realities, as will often be the case. The answer is to introduce an additional element of explicated and recognized uncertainty into policy-reasoning. This approach leads to the broader view of policymaking as fuzzy gambling, which takes into account also uncertainty in respect to evidence in the conventional sense of that term. (Dror 1988b:145)

Dror's view is that a choice taken without knowing its consequences is a gamble, usually a fuzzy gamble (section II.7.3), and that often that is the best we can do. Meta-rationality then suggests that if we must make fuzzy gambles, we learn to be the best fuzzy gamblers we can be.

Uncertainty plays a critical role in rationality: without any uncertainty, rationality in decision-making is unimportant, because decisions are trivial; and yet also uncertainty also invalidates rationality when it is extreme.

IV.1.4. Subjective rationality

The unavoidable subjective aspect of any implementation of rationality is of the same nature and at least as important as the role of subjectivity in probability. A conclusion of the analysis here is that

- especially under uncertainty, rationality is subjective, and rational conclusions are opinions.

One supporting argument follows directly from the views of De Finetti, Savage, and other subjectivists who hold that probabilities are opinions. If so, and if probabilities are needed for determining which actions are rational, then rationality is necessarily subjective.

But a more fundamental point is that determination of a rational action or choice requires employment of a procedure. There is no such thing as “the rational alternative” until we have employed a particular procedure that takes account of some aspects of the situation under consideration and neglects others. Implementation of any scheme for rational decision requires that many choices need to be made about which factors to consider, what evidence to seek and to use, and how to go about assessing the lot; all of these decisions require human judgment and therefore subjectivity.

In most cases, there is virtually an unlimited number of ways one can interact with an object. As a consequence, the interaction is almost never complete. It is based on a limited (and, usually, rather small) number of attributes that the observer is capable of distinguishing on the object and that are relevant to his interests. These attributes are not available to the observer directly, but only in terms of their abstract images, which are the results of perception or some specific measurement procedures. (Klir 1985:132)

It is important to stress the subjective nature of the choice situation. The fact that options are objectively available to an agent cannot enter into the explanation of his behaviour if he has no rational grounds for believing that they are available, nor can the fact that certain options will objectively lead to certain outcomes if he has no reason to think that they will. This point is sometimes overlooked. (Elster 1986:4)

“The true rational choice” is as impossible to operationalize as “the true probability” is. There is no conceivable experiment that would prove that a probability estimate, or a rational choice, is “true”. There are candidates that can be disproved, largely on the grounds of internal inconsistency or being dominated by another candidate, but none can be proven true. We can, of course, gather evidence, and a great many of us may happen to agree on what our judgment of the probability or the rational action is, but

that agreement does not constitute proof.

Even the matter of interpreting evidence used to inform a rational decision is subjective:

Coming to have suitable degrees of belief in response to experience is a matter of training — a *skill* which we begin acquiring in early childhood, and are never quite done polishing. The skill consists not only in coming to have appropriate degrees of belief in appropriate propositions under paradigmatically good conditions of observation, but also in coming to have appropriate degrees of belief between zero and one when conditions are less than ideal. (Jeffrey 1968:237)

Stinchcombe employs the concept of rationality as a principle of organizational growth instead of a principle for making decisions, but reaches a similar conclusion about the non-objective nature of rationality:

In this book, the postulate of rationality is what connects conditions in the environment to the structure of parts of organizations.

But the concrete substance of rationality depends on what part of the total uncertainty a particular part of the organization has to deal with. Thus, rather than postulating one rationality for the whole organization (such as “maximize net present value”) and assuming that the rest of the problems will be dealt with by an opaque “production function,” we postulate many different rationalities depending on what uncertainty has to be dealt with to get a given subtask done. (Stinchcombe 1990:242)

Stinchcombe’s terminology is different from mine: where he refers to “multiple rationalities”, I would refer to different sets of values (and possibly different sets of beliefs and information) as the primary explanation of differences in how rationality appears in different situations or as applied by different people or organizations. What his point highlights is that there isn’t a single rationality (or in the terms of this work, a single procedural implementation of rationality), and this

inevitable variability or subjectivity of rationality again constrains the ultimate potential of rationality as the be-all and end-all of decision-making.

Several other authors have characterized different rationalities. Perrow (1984:310ff.), for instance, discusses *absolute* rationality, *bounded* rationality, and *social* rationality, and those would be seen in the current view as a mainly substantive rationality, procedural rationality, and procedural rationality with a particular set of (social) values. Goulet (1986:302) distinguishes a slightly different three rationalities: *technological*, *political*, and *ethical*. Again, those could be translated into the more general usage here as naively pure substantive rationality, rationality with goals of maintenance of power that cannot be overtly stated, and rationality in the service of non-material, more “moral” values, such as freedom, justice, peace, etc.

Others use different terminology for describing similar kinds of differences. Linstone’s multiple perspectives are similar to the notion of multiple rationalities, though Linstone associates the term “rationality” primarily with the technical perspective. He suggests that logic and rationality form the ethical basis of the technical perspective, justice and fairness the ethical basis of the organizational perspective, and morality the ethical basis of the personal perspective (Linstone et al. 1994:110). Thus, Linstone’s ethical basis of the technical perspective corresponds closely to Goulet’s technological rationality, which is based on instrumental rationality and “*hard* logic”; and Linstone’s ethical basis for the personal perspective corresponds fairly well to one of the two

characteristics of Goulet's ethical rationality, the primacy of a "holistic meaning or belief system". The second characteristic of ethical rationality is the ascribing of high importance to the impacts of events and actions on people who possess little worldly power. There are many organization or political perspectives or rationalities, and Linstone and Goulet have chosen differently. There is little connection between justice and fairness, Linstone's ethical basis for the organizational perspective, and Goulet's political rationality, in which the primary motivation for political actors is the maintenance of their own power.

My own view of Goulet's rationalities attempts to relate them by means of examining the values adopted in each. The operative value in technological rationality is whatever goal has been set for the particular application. Mainstream conceptions of rationality generally exclude the choice of values from the province of rationality, with the result that sometimes rationality is invoked without any explicit choice of values, and any concrete goals can then be perhaps inadvertently elevated to the status of values. Thus, while Linstone describes logic and rationality as the ethical basis of the technical perspective, I view logic and rationality as tools that are independent of ethics and values, and would characterize a rationality guided only by instrumental goals (Goulet's terms) or a perspective ethically grounded in logic as incomplete in that it is missing a set or system of values. A feature that makes Goulet's political rationality distinct from his other two rationalities is that political actors' true value is the maintenance of their own power, but they must publicly claim other values or suffer a loss of legitimacy. This partly explains why political

actors are so much more willing to compromise their (public) goals and values than technological or ethical actors: the goals the politicians compromise are not their true goal. When viewed as if they are actually most interested in maintaining their own power, their behavior often makes very good sense, even from a coldly rational point of view. The behavior of Goulet's ethical actors also makes rational sense when their values are taken into account, especially in light of Goulet's assertion that they view "asserting their dignity as persons" as more important than achieving concrete goals or preserving existing power bases (Goulet 1986:303). It is not obvious someone trained in technological rationality would behave all that differently from a politician, if he were convinced maintenance of his own power was the most important goal, or all that differently from an ethical actor, were he convinced that promoting a specific value system and the dignity of the powerless were the most important (though a technocrat might not be as talented in these efforts as a skilled politician or community activist). It may also be that political and ethical actors are more comfortable at acknowledging times when rationality is indecisive and then using other procedures for devising and choosing actions.

Both discussing different rationalities and emphasizing the crucial role of the values adopted within a single, broader rationality can serve to clarify the widely divergent ways in which people may use the term "rationality", and either approach can be used to help different groups understand how others may be reaching different conclusions and to find common ground.

Under conditions of uncertainty, we cannot often hope to prove a particular action the most rational or the most right, even when the ends are clearly defined, and we must then build our confidence to act on a base other than a belief in the absolute, unique rightness of our choice. This must be our confidence in the integrity of the process that leads us to act in a particular manner, recognizing that there may be other processes with comparable integrity which could perhaps lead us to take different actions.

IV.2. An assessment of rational methods in light of uncertainty

IV.2.1. Mathematical approaches in general

Yet it is precisely by way of the proper and judicious joining of the schematic with the empirical that we achieve our beneficial and effective mappings of the universe — the schemapiric mappings known as science. The chronic danger lies in our failure to note the distinction between the map and the terrain, between the simulation and the simulated. The map is an analog, a schema, a model, a theory. Each of those words has a separate flavor, but they all share a common core of meaning. “Contrary to general belief,” wrote Simon and Newell ⁽¹⁾, “there is no fundamental, ‘in principle,’ difference between theories and analogies. All theories are analogies, and all analogies are theories.

... And lest the physical sciences should seem immune to what Schwartz ⁽²⁾ called “the pernicious influence of mathematics,” consider his diagnosis: “Thus, in its relations with science, mathematics depends on an intellectual effort outside of mathematics for the crucial specification of the approximation which mathematics is to take literally. Give a mathematician a situation which is the least bit ill-defined — he will first of all make it well defined. Perhaps appropriately, but perhaps also inappropriately ... That form of wisdom which is the opposite of single-mindedness, the

¹ Simon, H. A., and A. Newell. 1956. in *The State of the Social Sciences*, L. D. White (ed.), Univ. of Chicago Press, Chicago, pp. 66-83.

² Schwartz, J. 1962. in *Logic, Methodology and Philosophy of Science*, E. Nagel et al. (eds.), Stanford U. Press, Stanford, pp. 356-360

ability to keep many threads in hand, to draw for an argument from many disparate sources, is quite foreign to mathematics. ... Quite typically, science leaps ahead and mathematics plods behind.

... In the behavioral sciences, where the discernment for nonsense is perhaps less sharply honed than in the physical sciences, the vigil must remain especially alert against the intrusion of a defective theory merely because it carries a mathematical visa. An absurdity in full formalized attire may be more seductive than an absurdity undressed.

... The illusion of objectivity seems to preserve itself despite the admitted necessity for the investigator to make improbable assumptions, and to pluck off the top of his head a figure for the level of probability that he will consider significant. His argument that convention has already chosen the level that he will use does not quite absolve him. (Stevens 1968:850ff.)

Streeten (1972:280): mathematics permits rigour, though possibly at the expense of relevance;

... [it] lulls its practitioners into a false sense of certainty. (Streeten 1972:280)

The passages just quoted do not invalidate mathematics as a set of tools to use when uncertainty is present, but they do constitute a warning of some hazards which always attend the use of mathematics and sometimes cause serious problems.

IV.2.2. Measurement is fundamental

Measurement is fundamental to any process of comparison of options. While fuzziness complicates comparisons, the issues surrounding imprecision have received much attention outside of decision support.

To a statistician, the whole of the history of science revolves around measurement: all sciences (and the social sciences are no exception) require at least the possibility of quantitative measurement. But measurement alone is not enough — it is essential that the measurements be susceptible to comparison, and that, in turn, requires some understanding of the accuracy of the measurements. (Stigler 1987:287)

The historical tendency since the work of Simpson, Lagrange, and Laplace in the eighteenth century is to treat the approximate character of measurement procedures as due to *errors* of measurement arising either from 'systematic' or 'random' effects. This concept of error matches up with Laplacean determinism to give weight to a theory of exact results of measurement being in principle, if not in practice, possible. The fairy-tale quality of this view of measurement in physics should be evident enough two hundred years after its first telling. Even if it is not, I shall not attempt to make it evident here, but restrict myself to the more special arena of personal beliefs and decisions and their measurement

It is surely surprising if not literally paradoxical that the view of ideal measurement characteristic of classical determinism has been carried over almost untouched to modern theories of partial belief and decisions. It is, for example, a characteristic requirement of Savage's (1954) theory of decisions that the axioms on preferences between decisions be strong enough to lead to a numerical representation in terms of a *unique* probability distribution on beliefs and a utility function on consequences unique up to a linear transformation.

As in the case of classical physics there is a mythology of ideal measurement leading to unique numerical results which is widespread in contemporary theories of belief and decision. But the most casual common-sense consideration of actual beliefs and decisions generates almost immediate skepticism about any numerical uniqueness claim. The desire for uniqueness I view as a remnant of past hopes for an orderly, thoroughly deterministic universe. (Suppes 1975:153)

By the end of the nineteenth century the official philosophy was that you understood something only when you could measure it. (Hacking 1987:48)

I am not saying, in the spirit of dominant paradigm of the late nineteenth century, that we cannot understand uncertainty unless we can measure it, but that we do need to improve our ability to measure it: we need to be able to assess the uncertainty attending other measurements and estimations, e.g. of probabilities, in order to be able to make valid comparisons and interpret the meaning of the comparisons. We may need

several different units of measurement for different types of uncertainties.

IV.2.3. *Optimality (and satisficing)*

The problem, as we have already argued, is that optimality and equilibrium go hand-in-hand. Under conventional methodology, if we require our agents to behave optimally, then we are restricted (at the market or economy level) to analysing equilibrium: whether it be equilibrium states, equilibrium adjustment processes, equilibrium adjustment of adjustment processes, or whatever. True disequilibrium analysis requires us to abandon optimality. (Hey 1981:242)

As Winter (1988/89) and Simon (1988/89) argue, economists have spent more time elaborating the implications of equilibrium in economic systems than investigating whether the systems we live with and work with show evidence of being in equilibrium.

A second and perhaps more fundamental weakness of optimizing is that it chooses the best of the lot, when improving the lot may be more beneficial, as discussed above in section IV.1.

IV.2.3.1. "Optimizing behavior" and fuzziness

Seeking optimal values can be seriously handicapped by the presence of uncertainty. As Miller and Starr point out, "It is always questionable whether the optimum procedure is to search for *the* optimum value." (Miller and Starr 1967:51) Fleurbaey maintains that optimizing is so dependent on good information that uncertainty can make it non-optimal to attempt optimizing:

For the notion of optimal decision to be meaningful, it seems necessary that the context of the decision-making process should involve the following elements: A, the set of the actions among which the agent can choose (taking into account the constraints); C, the set of the possible consequences or outcomes of the actions of A; $\gamma: A \rightarrow C$, an application which to each action $a \in A$ associates its consequence $c = \gamma(a) \in C$; $\succeq$, a binary relation ("is preferred to"),

defined on C , which denotes the agent's preferences relative to the various possible outcomes, $\succeq$ is assumed to be a complete preorder on $\gamma(A) \subset C$.

In such a context, one can define an action $a \in A$ as optimal for the agent if it satisfies:

$$\forall a' \in A, \gamma(a) \succeq \gamma(a')$$

Let α^* denote the set of optimal actions. α^* is built with constituents which are all of an objective kind, even the preorder $\succeq$. In order to define what we will refer to as the optimizing behaviour in the sequel, we must on the contrary consider the subjective elements of information that the agent possesses: A^e , the perceived set of possible actions; C^e , the set of perceived relevant outcomes, $\gamma^e: A^e \rightarrow C^e$, the expectation function of outcomes; $\succeq^e$, a preorder, complete on $\gamma^e(A^e) \subset C^e$, supposed to represent the agent's preferences.

We will say that the agent adopts an optimizing behaviour if he has in mind all these elements, and chooses his action in the set:

$$\alpha^e = \{a \in A^e / \forall a' \in A^e, \gamma^e(a) \succeq^e \gamma^e(a')\}.$$

It clearly appears that as soon as one of the elements of information A^e , C^e , γ^e or $\succeq^e$ differs from its true equivalent, the sets α^e and α^* have in general no reason to coincide. This allows us to state that in a situation of incomplete information, an optimizing behaviour is not optimal in general. (Fleurbaey 1988:461)

Fleurbaey's criticism of the utility of optimizing behavior does not even touch on probabilities or expectation, just on the effects of incompleteness of information. An argument could be made that if the incompleteness of information results in "small" or "localized" errors in A^e , C^e , γ^e or $\succeq^e$, so that none is "far" from the true values, then α^e will probably be "near" α^* , but that argument rests on shaky assumptions about the meaning of "near" and "far" and general linearity of interactions among the key variables. "Sensitivity to initial conditions" is now popularly associated with chaos, but Poincaré's warning may apply to a great many situations:

A very small cause which escapes our notice determines a considerable effect that we cannot fail to see, and then we say that the effect is due to chance. If we knew exactly the laws of nature and the situation of the universe at the initial moment, we could predict

exactly the situation of that same universe at a succeeding moment. But even if it were the case that the natural laws had no longer any secret for us, we could still know the situation approximately. If that enabled us to predict the succeeding situation with the same approximation, that is all we require, and we should say that the phenomenon had been predicted, that it is governed by the laws. But it is not always so; it may happen that small differences in the initial conditions produce very great ones in the final phenomena. A small error in the former will produce an enormous error in the latter. Prediction becomes impossible. (Poincaré, *Science and Method*, cited in Gleick (1987:321))

IV.2.3.2. "The best" is nonexistent

The term 'optimization' and its underlying way of thinking should be avoided in most of the philosophy of policy-reasoning, because [it] assumes that a definite optimum exists, which is doubtful, both in respects to policy contents and with respect to policy-reasoning. (Dror 1988b:143)

The problem is that even use of the term "best" creates some problems, as in the simple question, "What is the best thing for me to do?" A psycholinguistic requirement for processing that question upon hearing or seeing it is to presuppose that the concept named in the question, "the best thing", is meaningful, implying that, though its identity may currently be unknown, a single best thing exists. Much of the current work supports the claim that conditions of uncertainty can keep it from having any operational meaning. Furthermore, efforts to identify the referent of that phrase will lead less directly to helpful actions than pursuit of merely "good things for me to do". (Eventually, a decision-maker would meta-optimize, and ask, "What's the next best thing for me to do if I can't identify the best thing for me to do?" A reasonable answer would be, "Find a very good, or a very non-bad, thing to do" — in other words, satisfice.)

Remember Satisficing Rule No. 1 (section II.3.3.3): always optimize

when you can. The hard part is to know when you can.

In the short run, assessing a particular decision with only the evidence already on hand, you can attempt to reject the Null Hypothesis, which states that there exists no alternative that is preferable to all others; failing to reject it, you attempt to determine whether at least some alternatives are demonstrably inferior and can be ruled out. With assessments of the uncertainty of the factors considered, even if the assessments are subjective, these two determinations can be made.

In the long run, knowing in advance when it will be rewarding to set out to optimize is much more difficult, requiring a prediction of how strong a conclusion you will be able to draw at a later time.

IV.2.3.3. When is optimizing promising?

When a great deal of information is available about the feasibility of options, which values are most important, what consequences will occur if specific options are chosen, etc., optimization is the best strategy. Under conditions of increasing uncertainty, it becomes less and less justifiable as a strategy.

An optimizer facing uncertainty is not led automatically to take the steps that this analysis of uncertainty recommends (especially section V.5.1). In contrast, a satisficer confronted with uncertainty does tend to take the appropriate actions. Suppose strategy X has been selected, and is thought to provide for a “good enough” outcome. If uncertainty increases, the challenge to the selection of strategy X is “Is X still good enough?” If the decision-maker, due to additional uncertainty, is no longer sure that X is good enough, she will take steps to bolster it, to improve its weak points.

A strategy that is considered optimal (before the increase in uncertainty) is harder to bolster — how do you improve upon the best? Indeed, it may still be the best of the lot — just no longer good enough!

The nature of satisficing as “good enough” is similar to “not too bad”, which is what is needed when information about how any option might work out is insufficient. The essence of optimization is what Dror calls *good achievement* (section II.7.6), while satisficing emphasizes *bad avoidance*. Both aims are important; uncertainty tends to elevate the importance of bad avoidance by make good achievement less attainable.

Satisficing may have lower cognitive burdens, which can be important if cognitive resources are strained by uncertainty. Options can be developed and evaluated one at a time, and no more are developed than necessary to find a satisfactory option.

Optimization will never be a completely sufficient principle, since it can only be applied once problems are found and framed. There is no way to calculate the optimal amount of information to seek:

At some level of analysis, all goal seeking behavior is satisficing behavior. There must be limits to the range of possibilities explored, and those limits must be arbitrary in the sense that the decision maker cannot know that they are optimal. (Winter 1964:264)

IV.2.4. Probability

A reasonable probability is the only certainty.
Edgar Watson Howe

Howe means, I believe, that we can get no closer to knowing something with certainty than to know it with “a reasonable probability”. We do know, with certainty, that some uses of probabilities have been very controversial, which brings up the question of just what is a reasonable

probability, and a reasonable use of probabilities?

Morgan and Henrion believe that probabilities should only be used in descriptions of *empirical variables*:

To be empirical, variables must be measurable, at least in principle, either now or at some time in the past or future. ... it is generally not appropriate to model uncertainty about [value parameters] probabilistically. (Morgan and Henrion 1990:69)

Nor, they continue, should probabilities be used when there is uncertainty over the proper form of a model (function uncertainty). They would also, then, rule out use of probabilities for uncertainty about the decision-maker's authority, the feasibility of options, or for measuring the uncertainty in probability estimates. Few advocate the use of probabilities for predicting the future behavior of other actors, unless the actors' behaviors are important only in the mass and they are so numerous that statistical measures of their behavior are relevant (Weaver's disorganized complexity, section II.3.10). When the actions of individual actors are important, game theory is recommended, which does not rely much on probabilistic assessment of other actors. Probability also

is ill-suited for capturing the full scope of uncertainty associated with phenomena characterized by Weaver as phenomena of organized complexity. (Klir 1992:74)

Klir suggests that probabilities are appropriate when class frequencies are well known, or when there is evidence from a series of random experiments (1990b:258).

IV.2.4.1. Objective, subjective and fuzzy probabilities

One of the main debates in decision support has been over the admissibility of subjective probabilities when objective (frequency-based)

probabilities are not available. Dror (1986b§24.1.1) and Michael (1973:119) and others assert that substituting subjective probabilities into analytical frameworks built for objective probabilities is an dangerous practice, essentially providing the illusion of greater certainty than actually exists (especially if MEU has been adopted as a decision selection process). Elster is willing to consider subjective probabilities, but thinks they are very prone to hysteresis in their elicitation:

How can we be sure that the final probability p is independent of the path that has led up to it; independent, that is, of the particular values of p_1 and e that were chosen? (Elster 1984:129)

His point is that we should want our beliefs, including probabilities, to be based only on causal factors relevant to the beliefs, and should want them to be independent of *belief-irrelevant* factors such as the particular method used to elicit subjective probabilities.

I think the main concern of Dror, Michael and Elster is that the use of subjective probabilities may hide important uncertainty. Klir shares this concern, formalizing it in his Maximum Uncertainty Principle (section V.2.2.3.4). A method that can guard against this danger is to be sure that uncertainty attending the estimation of the probabilities be assessed and propagated through any analysis. This would in any event be required when a Null Hypothesis approach was employed, as suggested earlier. We would expect that there would generally be higher uncertainty surrounding subjective probabilities than objective probabilities, and that for the same degree of regularity of object or process under consideration, there would be higher uncertainty surrounding objective probabilities estimated from sparse data than from extensive data.

This matter is illuminated to some extent by distinguishing between sources of uncertainty — in particular, uncertainty traced to variability in the physical world is very different from uncertainty arising from insufficient evidence or from characteristics of the decision-maker. In the case of physical variability with ample observations, we may have a very high degree of confidence in an assertion about the likelihoods of possible events. In the latter case, we may have only very low confidence in any estimates we make at all, and should take care to ensure that the fuzziness around those estimates is conserved throughout all analyses. Elster's concern over hysteresis could be partially addressed by eliciting subjective probabilities via multiple methods, and taking the differences in results as a minimum measure of the uncertainty that should be ascribed to the probability values.

A potential remaining problem is that the uncertainty of subjective probabilities would generally have to be estimated by the same people who make the estimates, and there would often be pressure on them to make the estimates appear more certain by underestimating their fuzziness. Resisting this and other pressures for convenience and expedience that threaten the validity of their work is one of the main functions of analysts' professional integrity.

Assessing the uncertainty of all probabilities should go far in defusing the debate over the use of subjective probabilities. A consequence, I expect, would be that many analyses currently using subjective probabilities would not show that one option is significantly superior to others when the uncertainty in the probabilities is accounted for. The weakened

decisiveness of analyses would be unwelcome to decision-makers who want to be able to lean on analytical efforts to justify their decisions.

IV.2.4.2. Distributional nature of probabilities

Another problem many have pointed out with the use of probabilities is that probability functions are distributional (section IV.2.5.1):

Any probabilistic model of uncertainty requires either a full enumeration of all relevant events and possible system outcomes or the use of a 'miscellaneous' category for unforeseen outcomes. Turner (1978: 169) points out that many practical decision making contexts require decision makers to specify the range of foreseeable outcomes, and that imposing the further requirement that the probabilities of these outcomes sum to 1 implicitly assumes that the probability of unforeseen outcomes is 0. ...

Some of the more sophisticated PRA [probabilistic risk assessment] experts acknowledge this problem but almost invariably place it in the 'too hard' basket. In an article on the estimation of probabilities for risk assessment, Winkler expresses the conventional tendency to bracket this concern by declaring it irrelevant: "Although the identification of possible outcomes is an important problem, the concern in this paper is with the probabilistic aspects of risk assessment" ([Winkler] 1982:352). (Smithson 1988:87)

This clearly has the potential to be a very serious problem when the possible outcomes cannot be enumerated. One of the primary reasons why this work argues for the display and clarification of uncertainty is specifically to help keep decision-makers from being surprised, and mathematical frameworks that imply that there is no chance of surprise outcomes may significantly worsen the danger. Unless this risk can be ruled out or adequately mitigated, probabilities ought to be avoided. Possibilities do not expose decision-makers to the same danger, and may be preferable when the probability that ought to be reserved for surprises cannot be estimated well.

IV.2.4.3. Summary of probabilities

Probability measures are well-suited for representing statistical regularity, though so when fewer observations are available. They should not be used when surprise outcomes are significantly possible, nor with value uncertainty, nor with function uncertainty, nor when the actions of other individual actors are significant sources of uncertainty, especially when they know their potential impact on the decision-maker, nor with authority or feasibility uncertainty. Probabilities are also poorly suited for uncertainty due to complexity, limitations of time or other resources. The use of a technique for estimating the fuzziness around probability estimates significantly mitigates most of the hazards associated with their use that are due to small sample sizes and subjectivity, though it would not help much to compensate for the distributional nature of probabilities. Two such techniques discussed by Good (1980) are *upper probabilities* and *lower probabilities*, which could be used to define a range for probability estimates, and *Type II probabilities*, which are probability distributions over the values probabilities could take.

IV.2.5. Possibility Theory

Possibility theory appears to be stimulating a great deal of thought and research. Practical methods based on possibility theory for making decisions, or more generally for providing decision support, do not yet appear to have entered the mainstream, though there is great promise.

IV.2.5.1. Probability and possibility compared

The degree to which possibility and probability represent essentially different types of uncertainty is not yet settled. As the quotation from

Lindley indicates (section II.3.8.5), some probabilists believe no other form of representation is needed or even valid. Possibilists argue that probability cannot represent nonspecificity or total ignorance (e.g. Klir 1989:356, 361). Klir refers to probabilistic and possibilistic bodies of evidence (Klir 1989:359-60), as if the nature of the body of evidence determines which type of formalization can be utilized. On the other hand, Klir goes to great lengths to demonstrate a particular mathematical translation between probability and possibility (Klir 1990b:258-69), which would not seem meaningful unless the two measures were capable of capturing the same aspects of an uncertain situation.

Various measures of types of uncertainty have been developed. For example, Klir (1992b) defines four such measures, which he says quantify fuzziness, nonspecificity, dissonance, and confusion. Probability theory is capable of representing only dissonance; possibility theory can represent both fuzziness and nonspecificity (Klir 1992b:83).

One rationale for the use of possibility in place of probability is the case of *nondistributional likelihood* (e.g. Ford 1983). The concern here arises from the requirement for probability measures that the probabilities of an exhaustive set of mutually exclusive events sum to 1.0. In cases where the possible events cannot be enumerated, this is impossible to achieve. Suppose three possible events are known, and considered to be equally likely. Suppose next that the feasibility of a fourth event is discovered. It seems mandatory and yet unreasonable that learning of a new possibility should reduce the numerical *probability* of the first three events.

Possibility measures, however, as noted above are super-additive³, and are thus not required to change when a new event is added to the set of feasible events.

Another problem with distributional likelihood, noted by Dror (1986b:26) and Smithson (1988:87), among others, is that the probability of an unforeseen event is implicitly defined to be 0. In many situations, that level of confidence in our ability to enumerate the possibilities is not reasonable.

Kahneman and Tversky (1982b:519) argue that

When uncertainty is assessed in terms of propensities, arguments, or confidence, it is less obvious that the probabilities should add up to unity — even if it is known with certainty that one of the alternatives is correct.

Dubois and Prade (1988:13) suggest that probabilities represent precise but disperse evidence (e.g. precise evidence from a series of observations), while possibilities represent imprecise but coherent evidence (e.g. fuzzy, subjective estimations from a single subject). Klir proposes that possibilities are better at representing predictions about the future (1990b:258), or for representing total ignorance (1989:361). He also suggests that the arithmetic processes typically used in convolving possibilities, $\max()$ and $\min()$, are less likely to accumulate and magnify errors than the main arithmetic processes in probability convolutions, the sum and product functions (Klir 1989:362). Zadeh proposes that possibility theory has good potential when the *meaning*, not the measure, of

³ Superadditive: the sum of the *possibilities* of all possible, mutually exclusive events is equal to or greater than 1.0. In contrast, the *probabilities* of such a set must sum to 1.0, and the sum of the *necessities* is sub-additive, being equal to or less than 1.0.

information is important (Zadeh 1978:4).

While these theoretical considerations of possibility theory do not seem to have been translated into practical applications of much note, they do indicate where applications of possibility theory are most promising. The mainstream of rational decision-making literature is still based on probabilities, as are the commonly used tools and methods.

IV.2.6. Maximization of Expected Utility (MEU)

Maximization of expected utility has been so thoroughly accepted by the American technical mainstream as the foundation for decision-making under uncertainty that it deserves closer attention here. MEU serves as both the framework for normative work and as the standard against which observations of subjects' behavior is compared. Many specific challenges have been leveled against its use in decision-making. 13 of these challenges and their implications are reviewed, and finally summarized.

IV.2.6.1. Whole-gamble features

MEU ignores whole-gamble features, e.g. security level, hope level, and preferences for or against risk. This exclusion of whole-gamble features is implied by the axioms of expected utility theory. For example, Axiom 5, Application of Probability Rule (section II.3.7.2),

removes the possibility of any pleasure or preference being derived from the way the actually outcome of any choice is arrived at — for example, it assumes away any pleasure derived from gambling at cards or on the roulette table. (Hey 1979:30)

Considerations that are not simply functions of a final outcome but of a combination of possible outcomes or of comparisons among options, are sometimes called “distributional variables” or here, whole-gamble features.

When they are important, MEU is inapplicable. For example, someone who valued security level above everything else would want to use the maximin rule, and someone who valued a high hope level might choose the maximax rule and select the option whose best outcome was higher than the best of any other outcome.

Implications. An important preference for or against any distributional characteristic of an option or a gamble makes the standard applications of MEU very awkward or inappropriate because these preferences are not well represented.

IV.2.6.2. Subjective probabilities and probability estimation

Though von Neumann and Morgenstern argued that only long-run frequency probabilities should be used in MEU calculations, Savage and others have countered that all probabilities are opinions in any event, and that subjective probabilities represent our knowledge of what will happen as well as it can be represented. Still others have argued that when we are highly uncertain about our subjective probabilities, we should use techniques that do not require probabilities at all (e.g. Elster 1984:131), or else we can seriously underestimate the risk of surprise outcomes. As shown in Tables III-1 and III-2, even small changes in probabilities can cause very significant changes in the results of expectation calculations. (In that example, changing the probabilities by no more than 3% for each of four possible futures completely reversed the order of preference of four alternatives.)

Implications. The fuzziness associated with all probability estimates should be assessed and propagated through all expectation calculations,

even if the only way to assess the fuzziness of subjective probabilities is again subjectively. Whether a definite conclusion can be reached in spite of the fuzziness should then be assessed by attempting to reject a Null Hypothesis. (section IV.1.2.)

IV.2.6.3. The distributional aspect of probability

As discussed above (section IV.2.4.2.), the distributional nature of probability requires that the total of 100% of likelihood be allocated to the set of possible states of the world. In cases of strategic uncertainty and possibility uncertainty, it is often impossible to know that all of the possible future states of the world have been identified, or to know how many of them there are or what the total probability that an unknown state would occur. If no probability is reserved for unknown states, there is an implicit assignment of 0% to the possibility that an unknown state (a surprise) will occur. An alternative is to assign a non-zero probability to a surprise category without specifying what the possible surprise states are. It seems extremely likely that this assignment would have a high degree of fuzziness.

Implications. When even a subjective probability distribution cannot be developed without running the risk of badly misallocating subjective probabilities to the unknown but possible future states of the world, MEU ought not to be used. If it is possible to estimate how likely it is that the future will be a surprise, probability can be assigned to that possibility and the remaining probability allocated over the identified possible states; the uncertainty in the assessment of the probability of surprise should be not be ignored, and its implications for the certainty of any conclusions should

be explicitly assessed.

IV.2.6.4. Use of mathematical expectation for one-shot choices

Authors such as Shackle and Allais have maintained consistently that expectation is neither an appropriate foundation for decision-making in unique situations, nor an accurate representation of the way people actually decide. However, Lopes suggests that the specific intent of von Neumann and Morgenstern's theory "was to justify expectation maximization for single plays" (1988:410). They appear to have succeeded, at least in the U.S.; the legitimacy of expectation as the guide for one-shot choices is not discussed much here. Most of the complaints are not that MEU should never be used in one-shot decisions, just that such situations are significantly different from repeated decisions, and that it is wrong to assume that MEU is the only appropriate decision rule, or the only rational one.

Implications. The axioms of expected utility theory are fairly restrictive; when a decision-maker actually believes they apply to his situation, MEU is one reasonable approach to adopt, even in unique decisions. This means that he must have no significant preferences over whole-gamble features. Additionally, the Null Hypothesis should be applied, which in many cases will prevent MEU calculations from being decisive due to uncertainties in the measurement of variables (especially probability measurements). However, many decisions are complex enough, and fit the axioms of MEU poorly enough, that the case for claiming that mathematical expectations should govern all one-shot decisions seems extremely weak, and other decision rules are often at least as appropriate.

IV.2.6.5. The meaning of utility

The very notion of *utility* is controversial. It is central to the working of a great deal of decision analysis and economic theory, and supported by little empirical testing. Utility is “whatever it is that decision-makers maximize;” a decision-maker ascribes more utility to one choice than to another if and only if she prefers the first choice to the second.

Utilities are not generally considered to be interpersonally comparable. Even comparisons for the same decision-maker at different times may not be appropriate:

It is unrealistic to suppose that utility functions are given and remain fixed. (Simon 1981:57)

Implications. The attractiveness of a tool that can combine many different evaluations of outcomes into arithmetically precise values presents a strong temptation to ignore problems of interpersonal or intertemporal comparisons. When these issues are important, they should be explored before MEU or any mathematical approach covers them up

IV.2.6.6. Hysteresis and the non-uniqueness of utility functions

By giving the subject a series of choices between risky options, it is possible to define concepts that express the cardinal utilities associated by him with certain events as well as the subjective probability with which he believes they will occur ... The question, of course, is whether this operation really measures what it purports to measure; whether there are distortions inherent in the procedure that prevent the techniques from being valid. (Elster 1984:128)

The axioms of MEU permit the derivation of a utility function for the decision-maker, and, as Elster notes above, results can almost certainly be obtained. However, such a utility function is neither unique for that decision-maker nor useful for making comparisons between different

decision-makers (Hey 1979:36). The MEU axioms permit an infinite number of utility functions to be derived, as any linear transformation of a utility function will remain consistent the MEU conditions. This itself is generally not a great problem. What is much more problematic is that the utility function is quite sensitive to the specific procedure used to derive it. Cohen and Jaffray (1988), McCord and de Neufville (1983), and Bouyssou and Vansnick (1988) report on studies showing that different methods of assessing utility functions result in different shapes of the functions that cannot be reconciled. They interpret these findings as very serious challenges to the application of MEU. McCord and Leotsarakos note that realistic decision problems have more than one dimension of outcome, and the additional parameters required for assessing multi-dimensional utilities also have unavoidable impact on the assessment values (McCord and Leotsarakos 1988).

Implications. Whether one believes that there is a *true utility function* within the decision-maker that must be revealed or merely that constructing one as an artifice will be convenient, some process must be employed to determine the characteristics of the decision-maker's utility function. The findings noted above indicate that there is no unique way to assess this function, and therefore the results of any procedure that depends on utility assessments should be carefully examined to find out how sensitive its results are to uncertainty in the utilities; employing the Null Hypothesis would be one way to do that.

The problem of how to estimate the uncertainty around a particular elicited utility function remains. Comparing two methods chosen to be

likely to produce as divergent results as possible will provide a minimum estimation of the uncertainty, which can be checked to see whether the result of MEU workings is sensitive to this difference. However, a reasonable measure of the uncertainty may be far greater than the discrepancy between any two specific elicitations.

IV.2.6.7. MEU does not clarify uncertainty

The main effect of employing an expected utility calculation is to *hide* rather than to *clarify* uncertainty. The typical output is a set of numerical scores, one of which is best. None of the uncertainty about the evidence used, nor about which state of the world will eventuate, is generally displayed in that set of scores.

Implications. When MEU is chosen to be guide the selection of an option, care should be taken to avoid gearing the whole analysis around the MEU procedure. If the focus of the work is prematurely narrowed to MEU, it will be more difficult to explore the nature of any uncertainties, possible ways to reduce them, or to design better options. Other method of displaying analytical results should be included, since MEU tends to emphasize the precision of expectations rather than the extent of uncertainties.

IV.2.6.8. MEU does not facilitate learning

Adopting MEU as a decision selection procedure imposes quite a bit of structure on a decision problem, and defines many of the major questions that will be pursued during an analysis. Some decision situations may not fit this structure well. Acceptable risk problems, for example, typically are at least as dependent on accommodating conflicting views by developing

political agreements as on the technical assessment of the factors MEU requires (states of the world and their probabilities, etc.). The mathematical emphasis that MEU will add to an analysis will make it more difficult to incorporate “softer”, less quantifiable concerns, though these may often be the areas in which the most fruitful learning might be possible.

Implications. MEU is primarily a *selection* process, and does not facilitate *exploration*. If MEU is adopted, efforts should be made to encourage exploratory efforts at the same time.

IV.2.6.9. Information requirements

In their exposition of expected utility theory (section II.3.7.2), von Neumann and Morgenstern describe a very well-informed decision-maker, one who has perfect information and unlimited processing power and speed. To the degree that any decision-maker has imperfect information on any component of her decision or any limitations on processing ability, her situation falls outside the territory in which the expected utility hypothesis has been validated.

On the basis of their postulates, von Neumann and Morgenstern assert that the “maximum utility which can be obtained in this situation ... is a well-defined quantity.” (1953:30) It is obvious, however, that many decision-makers have nowhere near the perfection of information required, which frequently has the result of keeping the maximum from being a “well-defined quantity”. The Null Hypothesis proposal can be used to assess whether or not the maximum is well-enough defined to be decisive.

High cost of information, or of information processing, could also

invalidate MEU. When these costs are high, it may be possible to determine the point of optimum utility, but it may not be optimal after deducting the search costs. When search or evaluation costs are high, sequential evaluation and satisficing may be preferable due to being able to avoid some of the search and some of the evaluations.

Implications. Since the justification for the use of maximization of expected utility rests on the assumption of a quality of information rarely attainable in practice, it is rarely possible to show convincingly that MEU ought to be used for the decision rule.

IV.2.6.10. "Small worlds"

An important assumption underlying MEU applications is that it is meaningful to consider with precision decision situations that have been abstracted out of a person's or an organization's life into what Savage (1954) calls "small worlds". This, too, is controversial; Good (1980) believes, for example, that estimations of probabilities depend on entire bodies of beliefs, not on those atomic beliefs comprising the body of evidence that bear most directly on the event under consideration. Others dispute the validity of small worlds:

All our every-day choices involve risk and uncertainty and we are always involved in several different types of decisions at the same time. We must always take this into consideration and make our choices depend not only on the outcomes of the current problem, but also on how they combine with the prospects of other future or simultaneous situations. (Hansson 1975:179)

Jungermann complains that the use of "small worlds" excludes consideration of many other goals of experimental subjects, and thus predisposes the experimenter to believe that any perceived deviations from

the MEU normative standard are due to cognitive errors rather than to complicated motivations of the subjects (Jungermann 1983:76). The emphasis in the American tradition has been on decision rules, especially MEU, rather than on motivation and interpretation, but

task representation may be of more importance in defining errors than the rules they use within that representation. (Einhorn and Hogarth 1981:60)

Implications. The implication here for experimenters is that careful assessment of the way subjects interpret the tasks they are asked to solve is important if understanding how well their behavior corresponds to the MEU model is a goal; subjects presented with the same idealized task may bring enough of their own goal structure to the task that each solves a different problem. The implication for decision-makers and decision support is that the universe of values and goals that are relevant to any decision may be quite large, and attempts should be made to avoid settling on a too-simple representation, and to keep in mind that it is likely that there will be values and goals omitted from the representation that may become important before the decision has been complete.

IV.2.6.11. Completeness of preferences

A requirement imposed by the axioms of EUT is that the decision-maker have a complete preference order over all possible consequences.

Many authors object to this:

Having complete preferences in no part of what it means to be rational. On the contrary, to insist that preferences must be complete and all pairs of alternatives be comparable can be a form of hyperrationality — that is, of irrationality. (Elster 1989a:9)

Loomes interprets experimental evidence as indicating that, whether

they *should* or not, individuals *do* not have complete preference orders:

If individuals had easy access to a clear ordering of preferences over certain and uncertain actions, we might have expected a) that valuations would have been expressed to a greater degree of precision and b) that the distributions of valuations would not have been greatly affected by the particular method of elicitation.

However, the evidence of considerable rounding suggests that even though the actions involved simple probabilities and sums of money well within participants' normal experiences, most did not find it easy to be very precise. This tends to support the view expressed by von Neumann and Morgenstern (1953, p. 20) that "the common individual, whose behavior one wants to describe, does not measure his utilities exactly, but rather conducts his economic activities in a sphere of considerable haziness." (Loomes 1988:51)

This haziness can be the reason for a decision-maker not knowing which of two outcomes is preferable — i.e. for not having a complete preference order. This is especially likely to be the case outside of the experimenter's "small worlds" when outcomes can be conveniently defined to have only a single dimension, often money. In "large worlds", there are often multiple dimensions of consequence that matter, many of which are themselves hazy (or fuzzy), and frequently more than one stakeholder's evaluation of the consequences must be taken into account.

Implications. It seems inevitable that there will be times, many times in my view, when decision-makers do not have complete preference orders. To assert that those decision-makers cannot act rationally because a favored tool, MEU, assumes a complete preference order in its axioms is to relegate rationality to a minor role in human affairs. It seems far preferable to decide instead that when the incompleteness of a decision-maker's preference order is significant MEU ought not to be used. Otherwise we end up with a conception of rationality that is always precise

and nearly always irrelevant.

IV.2.6.12. Complexity of MEU assessments

MEU presupposes that

knowledge of an individual's choices among some prospects each containing only one or two possible incomes would permit prediction of his choices among all other prospects, no matter how complex. (Friedman and Savage 1952:470)

But assessing utility is actually quite tricky, especially when multi-dimensional or interpersonal complications are important; in order to assess the individual's utility accurately enough it may take a great many experiments. If the goal is just to assist in a single decision, performing a complete enough set of experiments to derive a utility curve may be more work than directly comparing the options from the actual decision, thus becoming guilty of Elster's hyperrationality:

Because the total number of relevant events is usually very large, the number of probability judgments is likely to make it impossible to do a thorough job of handling the uncertainty in the environment. (Bass 1983:73-74)

Implications. In those cases when MEU is legitimately applicable, before attempts are made to derive the decision-maker's utility curve, the magnitude of the effort likely to be required should be compared to the anticipated gain in decision-making speed, clarity or quality.

IV.2.6.13. MEU versus satisficing

When searching for, inventing, or evaluating alternatives is expensive or must be performed one alternative at a time, a satisficing approach that stops as soon as an alternative is found that exceeds a particular threshold may be more desirable. For MEU, or any other form of optimization, all the alternatives must be found, clarified, and evaluated, while a satisficing

approach will only evaluate all the alternatives when either (1) only the last to be evaluated is satisfactory, or (2) none of the alternatives are satisfactory. In the latter case, it would seem to be more rational to try to develop better alternatives than to choose the least unsatisfactory alternative. Except in very simple decision problems, it is uncommon that all of the alternatives can be described in a fully-developed fashion without expending what could be considerable resources. Determining the optimal level of such expenditure is generally not possible, as this determination again requires information, the optimal level of which to pursue must be determined (an infinite regress mentioned earlier in section II.3.3.4).

On the other hand, a pure satisficing approach runs the risk of defining a standard for good enough that is too easily satisfied: a hasty acceptance of an early alternative prevents consideration of other options that may be better.

Implications. Some judgmental balance of optimizing and satisficing is required: if none of the current options is satisfactory, finding or designing better options could be very valuable. If it is not possible to do this, nor to avoid the decision, picking the best of the lot is better than picking an option worse than the best of the lot. Searching for and devising alternatives also must be guided by both optimizing and judgment; when it is impossible to know how much better an undiscovered option may be, it is still necessary to choose a level of effort in searching for new options, hoping not to set the level too high or too low.

IV.2.6.14. Summary

MEU as a model is very attractive, and the *concept* of utility can be

helpful in clarifying options for the decision-maker. As a *procedure*, it has flaws that have not been gotten round yet. In particular, there is strong evidence that different utility assessment procedures produce different utility curves, opening the possibility that they would lead to different decisions in the same circumstances. This type of arbitrariness would be very undesirable, as described later in the RAG (section V.2.2.3.3).

Capturing accurately what the decision-maker really cares about appears to be the crux of the matter. The mechanics of MEU are trivial in comparison to this problem; there is a risk that someone who starts by setting up an MEU analysis will be seduced by its elegant structure and neglect to pay sufficient attention to the details of evaluating consequences. Starting with the utility or preference assessment issue may make it possible to proceed directly to a decision without ever detouring into arcane mathematical computations.

The fundamental problem with MEU is that once you leave the lecture hall and step outside the “small worlds” that experimenters impose on their subjects, decisions are seen to be quite complicated, especially in the ways they depend on preferences of the stakeholders. MEU then runs afoul of Zadeh’s Principle of Incompatibility, when decision complexity makes precision and relevance virtually mutually exclusive. MEU is exquisitely precise (when provided with all the right data). But its relevance to particular decisions, given the narrowness of its axioms, is often very weak.

There are four major dangers associated with the use of MEU:

1. Whole-gamble features may make it inappropriate;

2. It can be indecisive in specific cases;
3. Although indecisive (i.e. its conclusions are not shown to be significant), people will take its precise results as being decisive; and
4. Adopting it too early or giving it too much emphasis can reduce the incentive for learning about the environment and exploring other options, and increase the risk of buried assumptions determining the outcome.

Given that alternatives with comparable mathematical convenience have not appeared, the question with regard to MEU is not so much whether it should be used, but (a) under what conditions might the axioms of MEU be so poorly met that it would be inappropriate to use MEU; (b) if it is not inappropriate, when is it likely to be helpful; and (c) what safeguards ought to be taken with MEU?

IV.2.6.14.1. When Is MEU inappropriate?

One case when MEU is inappropriate is when it is not yet time to be making a selection, i.e. when it is more useful to be clarifying authority or designing better options. If there may be feasible options not yet described that are superior to those known, using MEU before exploring options further could prevent selecting the best option. Focusing all analytical efforts on MEU procedures will be detrimental when the decision-maker needs to be learning more about the environment. In general, the less is known about the environment, the less appropriate a pure selection procedure, such as MEU, will be except in conjunction with vigorous efforts to learn more about the environment.

When the consequences of the alternatives under various states of the world cannot be estimated accurately, the weighting scheme of MEU is not

appropriate. This includes cases of conflicting models or conflicting data, unless any of the possible models or data would yield the same MEU conclusion, and consequence uncertainty in general. When the set of possible future states of the world cannot be described well, and there is a significant but unknown probability of a surprise, i.e. that the state of the world will come from outside the set, MEU should not be used.

Any sort of value uncertainty, whether that of a single party evaluating a single dimension of consequence, or multiple parties, or multiple dimensions of consequence, needs to be resolved before using MEU; MEU does not contribute to this resolution.

Another time when MEU should not be used is when the decision-maker has important preferences that conflict with the axioms of MEU, that is, when the decision-maker has preferences for or against gambling per se, preferences over security level, hope level, or regret.

The appropriateness of the use of mathematical expectation in one-shot decisions is unresolved, with the American mainstream generally endorsing it and Europeans much less sanguine about its use. In specific cases, it may help to judge this issue by exploring whether security level and hope level are important, with MEU being more appropriate when they are not important.

IV.2.6.14.2. When Is MEU likely to be helpful

Hey describes when he thinks MEU is likely to be helpful:

von Neumann-Morgenstern utility theory is eminently suitable for simple decision problems — that is, those in which the computational and other decision costs are essentially ignorable — but for more complex problems, it is more realistic to suppose that the decision-maker adopts some simple ‘rules of thumb’. (Hey

1981:241)

High costs of information or processing do not automatically make MEU inappropriate, but they make it more likely that any benefit from its use will be outweighed by the costs of the process.

Similarly, greater fuzziness does not by itself rule out MEU, but it makes it more likely that the uncertainty or fuzziness about an MEU conclusion will be large enough that the result will not be judged to be significant. Fuzziness that can cause this includes uncertainty about the magnitudes of consequences or the value of those consequences, or how to combine multidimensional or multiparty preferences, or fuzziness about the probabilities over states of the world, whether objective or subjective

IV.2.6.14.3. Safeguards

The most important safeguard in the use of MEU is to recognize that MEU, and all rational decision processes, may be indecisive, and to test for this using the Null Hypothesis approach described earlier. Many decision-makers will not want to recognize this potential indecisiveness, since it clearly implies that they will often have to gamble, and deprives them of a technical justification for selecting a particular option, leaving them more vulnerable to second-guessing. In many cases, it is likely that some of the alternatives can be ruled out even if a single best option cannot be identified. Then the decision-maker must select from the remaining alternatives using a different decision process.

Another safeguard is to perform MEU calculations with estimates of the utility of consequences rather than with the simple magnitudes of the consequences. If it can be safely assumed that utility is a linear function of

the consequence this is not important, but generally utility functions are non-linear.

The difficulty of obtaining unique utility assessments and therefore the unavoidable uncertainty over utility evaluations emphasizes the importance of utility assessment, and increases the importance of not having this assessment buried in an analysis. Techniques that leave much of the utility assessment to the end of the process where it can be performed more visibly run less risk of overlooking the possible influence of the assessment process on the outcome of the analysis.

IV.3. A Perspective on Normative Decision Theory

IV.3.1. Relative emphasis

IV.3.1.1. Over-emphasis on selection

Normative decision theorists have given the majority of their attention to the selection phase of decision-making. In contrast, Mintzberg et al. in their study of how strategic decisions were actually made within organizations (section II.4.3) suggest that selection is often a mere “trimming” on processes primarily consisting of identifying and refining options, a finding supported by other authors (e.g. Stinchcombe 1990:4, Wilson and Wilson 1970:203, Dror 1987:40).

Personally, I feel that this quest for a “scientific” and “mathematically objective” [decision] rule is all wrong! When there is a paucity of evidence at hand, we require a methodology that brings information, however vague and imprecise, into the analysis, rather than a methodology that suppresses information in the name of scientific objectivity. We must be wary of objective panaceas that grind out answers by throwing away information that is not easily quantified. (Raiffa 1968:155)

Over 25 years later, much of Raiffa’s criticism is still warranted. Too

much of what is called decision support is actually an attempt to automate decision-making by providing models that produce decisions — newer, easier-to-use models into which the decision-maker's world is plugged. Usually such tools are prepared to handle a few kinds of uncertainties, but in other aspects they anticipate perfect, point-valued information. This can make it much harder to clarify both what is known and what the limits of that knowledge are. As Raiffa urges, we should be making it easier, not harder, to develop the sort of information that is most difficult to obtain, which in many cases is information about the desirability of consequences and how to balance interpersonal, intertemporal, and multiple objectives. For the most part, this is not sexy work: faster computers do not yet contribute much to helping a person understand his basic values.

Raiffa also points out the importance of attending to less-readily quantified sorts of information, which has been made more difficult by over-emphasis of MEU.

IV.3.1.2. Under-emphasis on preferences

With MEU largely accepted as *the* rational selection process, the axioms of MEU have acted as filters that have kept researchers from paying much attention to kinds and patterns of preference that are incompatible with those axioms. Some preferences that are incompatible with MEU are legitimate, such as whole-gamble preferences, but there are other preferences that, while natural and understandable, are internally inconsistent or flawed in other ways that decision-makers would acknowledge once aware of them. More attention needs to be given to questions of the kinds of preferences people actually have, how to clarify

what those preferences are, and how to “improve” them when they are flawed from the perspective of the decision-maker and/or a decision theorist, and how to account for them in decision support if they are legitimate even though awkward for MEU. It is irrational to concentrate so exclusively on the calculating aspects of rationality at the expense of the valuing aspects.

IV.3.1.3. Over-emphasis on state-of-the-world uncertainty

As the Taxonomy and the Generalized Decision Model presented here show, decision-makers can be uncertain about aspects of a decision than just the state of the world that will eventuate. Normative decision theory, though, has concentrated on this one uncertainty, usually on the still narrower case when a probability distribution is known (or assumed) over the possible states of the world. That has meant that the impact of these uncertainties on decision-making has not been examined enough, and few aids have been developed for decision-makers experiencing other types of uncertainty.

IV.3.1.4. The role of decision support

Good decision support will facilitate high-quality decision-making, not just faster or less-conflicted decision-making. As Elster and Dror have argued persuasively, decisions made under uncertainty cannot be evaluated simply by *results*; the quality of decision-making *processes* must be assessed, perhaps using criteria such as those developed by Janis and Mann or Fishchoff et al. (section V.2.1.2). Decision support efforts, then, need to be evaluated by the degree to which they promote good decision-making, e.g. by facilitating full exploration of options and clear explication

of uncertainties.

An important role for experts in any subject is to help their clients understand not only the strengths but also the weaknesses of their specialties. In the case of rational decision support, a main limitation is that any rational decision selection process may fail to be decisive, and decision support experts need to help their clients understand what that means for them. Client decision-makers may not be able to determine what their best choice is, even after receiving excellent support help, due to not only the gambling aspect of a decision (e.g. state of the world uncertainty) but also due to currently irreducible fuzziness.

A main function of systems-analysis units is to educate decision-makers to understand correctly their function as fuzzy decision gamblers. (Dror 1988:263)

Aside from the fact that this educational process may be difficult when expert and client are both enthusiastic about it, two motivational problems further complicate the issue. Decision-makers, especially bureaucratic ones, will seldom be happy to be taught that their decisions often amount to gambles for which they may nevertheless be held responsible. Secondly, decision support experts may be disinclined to advertise the limitations of their product and realistically advise that the best they can do may be to clarify the parameters of a truly difficult choice instead of being able to offer their clients some relief from their burden of responsibility by identifying the “best” option for them.

IV.3.2. Better tools needed

IV.3.2.1. Tools for fuzzy data

The emphasis on substantive rationality in the mainstream of the

technical literature has downplayed the importance of measurement as a key issue for rational decision-making. Though March, Simon and others have described the bounded nature of human rationality, and described the differences between substantive and procedural rationality (section II.3.3), techniques for rational decision support have largely been based only on the concerns of substantive rationality. The issue of how the fuzziness in any aspect of the data will affect the certainty of a conclusion is mainly an issue of the procedural impacts of measurement uncertainty.

Klir's Maximum Uncertainty Principle (section V.2.2.3.4) urges that all uncertainty be preserved. While input data is often fuzzy, methods for processing it (combining and analyzing) typically require non-fuzzy data. De-fuzzifying the data removes uncertainty, which prevents any later assessment of the importance of the fuzziness. A great help in this area would be more tools that are capable of accepting input data with all of its uncertainty and processing both the data and the uncertainty. An example of such a tool is FuzziCalc, a microcomputer tool very much like a spreadsheet but with significant extensions. Whereas typical spreadsheet cells contain non-fuzzy numbers, FuzziCalc cells can contain ranges, and output cells can thus also have ranges, indicating the extent of uncertainty about an output variable.

In a recent project at BPA, for example, executives wanted staff to provide answers to the key questions in the form of ranges, recognizing that the attendant uncertainty was so great that single-point answers would give a false impression of sureness. The staff, however, had no tools for working directly with ranges. All the tools available — computer

models, spreadsheets, etc. — operated on point values. The only way to use those tools to produce ranges of answers would be to run one case with a high value representing the top of the range, and one with a low value. Unfortunately, treating a second input variable in the same way would again double the number of cases: there were at least four critical variables and parameters to treat in this way, and that would require assessing at least $2^4 = 16$ cases. The problem of aggregating and displaying the results meaningfully is similarly compounded.

IV.3.2.2. Displaying uncertainty

Various ways of displaying fuzzy data and fuzzy conclusions are possible, though few are nearly as common as typical displays of non-fuzzy data such as bar charts, line charts, pie charts, or tables of numbers. Greater experience with producing displays that indicate not only values but the uncertainty surrounding them would be very useful — as would more experience on the part of decision-makers and readers in general in interpreting such displays.

IV.3.2.3. Possibility Theory and Fuzzy Set Theory

Both possibility theory and fuzzy set theory offer the hope of spawning new tools and methods. Both may be more effective at producing tools for the pursuit of robustness instead of optimality, since they both work with ranges of futures rather than with single futures. Mathematical tools that can accept ranges, such as interval arithmetic, could help to make possibilistic analyses more computationally feasible.

A perverse development noted by Smithson (1988:303) is that much of the literature on fuzzy set theory, which was launched by Zadeh in an

attempt to provide a framework that could handle approximate data and imprecision, is no freer of precision and rigor than the previous mathematical approaches from which fuzzy set theory seeks to distinguish itself. We have not yet seen many tools built on fuzzy set theory that are able to actually help us accept imprecision and fuzziness without forcing it into one form of precise representation or another.

IV.3.2.4. Chaos Theory

Systems that can be identified as chaotic are likely to be both predictable and unpredictable: we may be certain their behavior will not vary outside the domain of their attractor, and may also know that we must remain uncertain about precisely where the system will be within that domain.

How widely relevant tools based on chaos theory will be does not seem knowable yet; it depends in large part on how common chaotic systems are in our world. As a model of how non-periodic variation can occur it is powerful, but it may be that there are few systems that are purely chaotic in the sense of their variability coming only from their own non-linear dynamics. Systems that are affected by external disturbances in addition to being non-linear internally will probably be more difficult to analyze. Advances in non-linear analysis in general, not only within the field of chaos, would be very helpful.

IV.3.2.5. Tools for evaluating consequences

A major cause for rational processes being indecisive is significant uncertainty about how to compare one set of possible consequences with another. This can arise from uncertainty about many factors: the

desirability of a particular consequence or the relative desirability of different magnitudes of a consequence, how to compare different dimensions of consequence with each other, the relative importance of consequences that would occur at different times, how to combine or compare the views of different stakeholders, or basic values about making risky choices. Uncertainty over these preferences can be resolved only by clarifying one's values, either actively by making a decision about the preferences one will hold, or less actively by "discovering" previously implicit preferences.

When the utility of consequences is uncertain, the Uncertainty Link Principle implies that it is wasteful and can be misleading to over-analyze the more readily calculated aspects of the decision. It may be possible to resolve the general evaluation uncertainty and then to proceed with a rational analysis. If not, it may be helpful to provide some analytical inputs to a political or psychological process for making a decision, while recognizing that the analysis cannot demonstrate which option is the most rational, and that it ought not to be pressed for such an answer.

The technology for eliciting and clarifying preferences does not appear to be as well developed as technology for performing calculations on measures of values once they have been obtained. Better support for resolving these issues of evaluation would be very useful. A large special case that is especially in need of attention and better tools are when two or more parties need to reach an agreement on an evaluation, and when more than one dimension of consequence or value is important. These are areas where "soft technology" is likely to be more helpful than new computer

algorithms, though the display capabilities of computers may be very helpful in framing issues so that decision-makers are better able to understand their values and to share theirs with other stakeholders. This is not a matter so much of mere assessment as of forming opinions and forging consensus.

The description of values is another area where the ability to work on data described in ranges would be helpful. Von Neumann and Morgenstern themselves believed that people are hazy about their utilities (1953:20), and ways of accepting hazy data without forcing it into a spuriously precise form would be useful. When data is available in the form of points or shaped distributions other than simple ranges, that additional information should be helpful.

IV.3.2.6. Complexity

Organized complexity poses very difficult challenges for rational analysis. Probabilistic tools work well for disorganized complexity, analytical methods often work well for organized simplicity (except when human agency or other sources of uncertainty are important). Disorganized complexity is frequently still intractable. Better tools would be very helpful. Complex models offer considerable hope, though they present additional complications such as the difficulty in finding and fixing errors, keeping assumptions from becoming obscured, and taking a lot of resources to maintain and a lot of time to modify. New ways of structuring complex problems that preserve the uncertainties in the data and are as transparent as possible would be useful.

IV.3.3. Smithson's four stages

Section II.6.2 examined Smithson's description of four stages through which the treatment of uncertainty in mathematics has gone:

1. Exclusion (or banishment of uncertainty);
2. Inclusion (control, often through reductionism);
3. Second-order strategies;
4. Suburbanization (subdivision, relativization).

Rational decision-making approaches have also changed over time, and can be related to Smithson's four stages as follows.

IV.3.3.1. Stage 1: exclusion and banishment

A major strategy employed by economists and decision analysts is to ignore many of the major uncertainties that can complexify decisions. Use of the term "uncertainty" has generally been closely identified with probabilities. Lindley, for example, quoted in section II.3.8.5, asserts that probability is satisfactory for representing all uncertainties and that nothing else is appropriate. Clearly, someone who accepted Lindley's perspective would have to deny the existence of any uncertainties that could not be fit into a probabilistic framework. This strategy restricts the definition of uncertainty to a small set of uncertainties — the possible future states of the world with a probability distribution over the set of possible states.

A second, more explicit strategy of banishment is the normative endorsement of maximizing expected value as the only criterion for rational choice. The essence of this standard can be described simply, and it was postulated that the implied model was relevant to all decision

problems. This provides the appearance of eliminating the uncertainty in making the decision, and of rescuing the hope of achieving certainty.

IV.3.3.2. Stage 2: inclusion and reductionism

At least two inclusionary strategies are identifiable in the mainstream of decision analysis. The first grew from the problem that there were many possible events for which objective probabilities could not be estimated from the observed frequency of previous occurrences for any of a number of reasons. The Subjectivists then developed the view that probabilities are essentially opinions that can not be proven wrong or right. This permits broadening the domain of usage of probability from only those cases where historical data sufficient in quantity and regularity are available to any situation where opinions can be held.

The second major initiative has been the construction of axiomatic systems of utility, notably MEU, and the attendant elicitation processes. These methods can be applied to nearly any situation, and will produce numbers in the form that MEU needs for its calculations, allowing its use to spread. This again preserves the hope that while certainty of outcome cannot be obtained, at least certainty of the rightness of one of the choices can be.

IV.3.3.3. Stage 3: second-order strategies

Problems have arisen in the actual use of utility elicitation procedures, and it now appears to be impossible to assure the uniqueness of their results, which MEU requires. One second-order response has been to study how such procedures differ from each other.

Another key second-order strategy has been the explosion of interest in

fuzzy set theory and its offshoots on the one hand, and Dempster-Shafer evidence theory on the other. Both of these fields offer ways to assess the second-order uncertainty of first-order measures of uncertainty. One goal of second-order strategies might be to attempt to define with precision and certainty when certainty would be possible to achieve. Klir (1993) contributes to second-order attempts by producing mathematically precise formalisms (e.g. probability theory, possibility theory, and Dempster-Shafer evidence theory) and relating them to different kinds of uncertainty (discord or strife, nonspecificity, and fuzziness). Klir acknowledges that “different theories may be appropriate at different stages of a problem-solving process or for different purposes at the same stage” (1993:312), but stops short of heading fully into suburbanizing uncertainty. Instead, he attempts to provide a standard for unifying the treatments of uncertainty, and states that

every uncertainty theory, even the least general one, is capable of characterizing (or approximating, if you like) the uncertainty of every situation. This characterization may not be as natural as its counterparts in other more adequate theories due to the constraints of the theory. I claim, however, that such a characterization does always exist. (Klir 1993:312ff.)

He describes an *Uncertainty Invariance Principle*, which requires that translations from one mode of characterization of uncertainty to another preserve both the amount of uncertainty and the degrees of belief. His claim, just quoted, implies that he believes transformations that meet the requirements of the Uncertainty Invariance Principle are possible, whereas a genuine suburbanite would assert that in at least significantly many situations such conversions cannot be made, and that uncertainties

vary enough that no grand translation scheme is possible. For instance, it is hard to see how the systemic, inter-institutional uncertainty that Thompson and Warburton (1985) describe as preventing even the formulation of a problem could ever be characterized by probability theory.

Both Smithson's work and this dissertation are at least in part second-order efforts in that they take uncertainty as a main subject, not merely a complication of a primary subject.

IV.3.3.4. Stage 4: suburbanization and relativism

Smithson's last stage (as of 1988, at least) is suburbanization, in which the hope for a single approach to uncertainty is abandoned, and sub-fields emerge that each deal with uncertainty in their own ways. This might mean that rather than attempting to define a societal standard for certainty, there would be one standard for criminal trials, one for civil trials, a third for engineering calculations, a fourth for the conduct of police officers, and there would be acknowledgment not only that these standards might be different but that it might not even be possible to compare them meaningfully. Or there may be different calculi developed for handling subjective and objective probabilities, or clearer rules of thumb for when to use possibilities instead of probabilities and how to use either for clarifying decisions under uncertainty.

An excellent example of a suburbanizing approach is Thompson and Warburton's "Decision Making Under Contradictory Certainties: How to Save the Himalayas When You Can't Find Out What's Wrong with Them" (1985). The authors focus on a widely accepted belief that a deforestation problem exists in the Himalayas, so their analysis is not of uncertainty per

se. Their conclusions, however, are that a solution process neatly fitting Smithson's definition of suburbanization is required. Thompson and Warburton find that climate, culture, agriculture, trade and perceptions of the "problem" vary so much from one locale to another that no single solution will work. Instead, they argue that the details of each locale are a rich source of potential micro-solutions. The focus of their recommended problem-solving has shifted from determining which solutions are right and which are wrong to which ones are appropriate to each set of circumstances. As a result, they do not proffer a "solution" as much as describe a process that they believe can be used to generate a multitude of solutions. Were I as clever as they, I would describe a process decision-makers could undertake that would lead to their discovering for themselves what steps would be appropriate for each of their particular circumstances. I have made do with a broad survey of the territory of uncertainty, pointing out many kinds of obstacles to be watched for, without being able to offer as neat a description of how to navigate between obstacles as I would like.

A primary conclusion of this dissertation is that certainty of decision-making is unobtainable except in rare, lucky instances. Furthermore, second-order certainty — for example, certainty of choice of methods for decision-making under uncertainty — is also unobtainable. Uncertainties and decisions vary so widely that very few absolute rules can be devised. Each situation is different, and judgment will have to be used in selecting and applying methods.

For the future, it is likely that Smithson's fourth stage will continue to

unfold, bringing a larger and more diverse collection of approaches and methods for handling uncertainties, with better information about the relative strengths, weaknesses, and most suitable applications for each of them, but with considerable uncertainty and debate still attending their use.

IV.3.3.5. Tension: suburbanization vs. comprehensive view

I experience an odd tension when reviewing Smithson's discussion of the 4th stage of suburbanization of the study of uncertainty: on one hand, I believe that a comprehensive framework is important, and I have been attempting to build one. On the other hand, Smithson's description of the suburbanization and separate handling of different uncertainties makes sense to me also, which might imply that no comprehensive treatment is possible, not even a structuring framework based on a taxonomy of uncertainties.

Perhaps the resolution is to aim for a comprehensive framework which is able to provide a loose organization for many different distinctions and different treatments, but without requiring it to reconcile them all precisely, or provide an automated package that can give precise recommendations to any uncertain decision-maker. My effort here could then be viewed as a survey of the property in preparation for a full-blown subdivision. Individual tracts will be developed by different parties, at different times, and may be related to each other more by a kind of geographical proximity than by a shared set of axioms or procedures. For instance, it seems very likely that there will be edifices constructed around possibility theory, though exactly where they are located, and how much of

the ground they cover isn't known yet. Chaos theory may give rise to buildings in this tract also, though probably not as extensive in this realm of decision-making under uncertainty.

Eventually nicely made maps will be available. Maps do not impose order or structure upon their territory; they display some of the aspects of what structure may exist in different ways. You can buy a map of the West Coast in which physical relationships are preserved as much as possible, so that straight lines on the map correspond to sight lines on the land; but you can also buy maps where freeways are represented as straight lines no matter what their physical shape really is, which can make identifying sequences of exits easier and faster. Different maps, drawn from different perspectives, will be needed for different purposes in this field also.

IV.4. Miscellaneous

IV.4.1. The Law of Requisite Variety

Ashby's Law of Requisite Variety (Ashby 1956:206-18) is formulated in terms of disturbances (in the environment) that may be counteracted by a regulator in order to keep a system in a desirable state. The Law states that the variety or uncertainty in outcome (undesirable) will be at least as great as the variety of the disturbances, unless it is reduced by variety in the behavior of the regulator. Furthermore, the variety in outcome is kept to a minimum when the regulator's response to the disturbance is deterministic. Any single, predetermined action or plan has no variety in it, and will therefore not reduce the (undesirable) variety in outcome. A set of responses from which one is chosen according to the particular disturbance is required, or a plan with a set of optional components; the

greatest reduction can be achieved when the regulator's response is a deterministic function of the disturbance. Accurate information about the disturbance is required to inform a deterministic response. This argues for good information systems about potential disturbances.

Of all the possible outcomes, the optimal outcomes are a subset of the non-bad outcomes, and therefore it generally requires greater ability to control the outcome to ensure that an optimal outcome eventuates than to ensure that a non-bad outcome occurs. If the decision-maker has a great deal of variety of response compared to the environmental uncertainty, optimizing makes sense. This may occur either because the environment is very predictable, or because the decision-maker has very numerous responses to make at the right time, such as being able to make mid-course corrections as new data is received, or having extensive contingency plans and enough time to put them into play, etc. When the decision-maker's responses have relatively small variety compared to the environmental uncertainty, pursuing optimality may not be feasible, and pursuit of robustness and avoidance of ills may be more realistic.

IV.4.2. Information, attention and the cost of certainty

While information is a multiplier positive-sum resource, attention is a zero-sum nonrenewable resource. ... In an information-dense world, therefore, most people cannot afford to indulge in an unselective consumption of information. (Smithson 1988:233)

Smithson asserts that ignorance increases faster than knowledge as people become more specialized and because new questions are raised faster than old ones are answered. Thus, there will never be a shortage of

ignorance to be reduced, but each person has finite attention, and must be sparing with it. Simon (quoted in section II.5.6) agrees that limits on attention mean that “information may be an expensive luxury”.

This suggests that uncertainty is essential for permitting a limited amount of attention to comprehend and deal with a world that contains a virtually unlimited amount of potential information (unlimited complexity? unlimited detail?). We cannot afford to be certain! The challenge is to allow in the right uncertainties, in the right amounts — to be uncertain in the least harmful or most helpful ways.

IV.4.3. Total uncertainty and sub-type trade-offs — silver lining?

Without additional information, you can’t honestly reduce the uncertainty in a representation; but you can shift it around. For example, the confidence in a statement and its precision can be traded for each other: you might be 99% sure that John is “tall”, but only 75% sure that he is “very tall”. Precision and relevance might be traded off, such as “alcohol might be bad for women” versus “alcohol is definitely bad for pregnant women” (the latter is more precise but relevant to fewer women).

Given a set of items of information, the opposition between precision and uncertainty is expressed by the fact that in rendering the content of a proposition more precise, one will tend to increase its uncertainty. Conversely, the uncertain character of an item of information will in general confer a certain imprecision on the conclusions to be drawn from it. (Dubois and Prade 1988:4)

Aggregation of input data, for example, which increases the uncertainty about individual cases, may permit more certainty in concluding whether or not a general relationship is supported by the data; or it may permit a simplified model to be developed that is feasible to verify when more

complex models would be too large to verify within the time available.

IV.4.4. Probable knowledge

As Jeffrey (1968) asserts, even knowledge which we treat as certain is merely probable. One reason we can so often overlook this is that we must take so much “knowledge” on faith:

we live in second-hand worlds in which most of what we consider knowledge never gets tested or experienced in a first-hand manner. However, because reality-testing is instrumental and intrusive, it is highly circumscribed and usually reserved for high-status members of society with approved authority, qualifications, competence, and reliability. (Smithson 1988:235)

Along with describing clearly the uncertainties we know we face, we must acknowledge more openly the less-than-certain nature of much of what we take for truth. This undermines efforts to eliminate and control uncertainty, and makes anticipatory planning more difficult (Wildavsky 1985). However, it is probably necessary in order to proceed with better handling of uncertainty by acknowledging its constant role in all aspects of life, rather than considering it to be an intrusive complication, a contamination rather than constituent of our ordinary reality.

IV.4.5. Is social science too hard?

I would hardly be the first to observe that the production and acceptance of general laws and theorems appears to proceed much more rapidly in the physical sciences than in the social sciences. It is possible that this is partly due to greater difficulty in defining systems with natural boundaries that have limited complexity, so that the total of the knowledge and complexity associated with the systems is not too vast. The objects of social science investigation do not appear fit Simon’s notion of “nearly

decomposable” as well as those of physical science. Hierarchical decomposition can prevent uncertainty about one aspect of a system from impeding the development of information about another. For instance, the structure of matter has been decomposed hierarchically with good success, and uncertainty about the constituents of electrons does not impair our understanding of the behavior of different molecules. In contrast, however, the behavior of organizations does appear to be significantly affected by the internal workings of human beings within the organizations.

Whether the approach of abstraction and generalization can succeed depends upon the nature of the subject. If there are 10 equally important factors in a system, including 9 or even fewer of them may well be sufficient. If there are an infinite number of factors, at least several possibilities exist. It may be that the second is only half as important as the first, whose importance we can call a unit (1.0) of importance. If each succeeding factor is again only half as important as the previous, the total importance of all the factors is only 2.0. Including the first three factors will then account for 87.5% of the total importance, and abstraction is very likely to succeed.

If, on the other hand, the second factor has $1/2$ the importance of the first, the third factor has $1/3$ unit importance, the fourth $1/4$ unit, and so on, the total importance is unlimited. Any finite combination of factors can account for only a negligible fraction of the total importance of the situation, leaving enormous uncertainty, and abstraction far less likely to be helpful.

Fortunately for chemical science, the relevant characteristics of

individual atoms are fairly few, and one iron atom behaves enough like any other iron atom that the differences are usually unimportant.

Unfortunately for social scientists, human beings appear to vary in a great many ways, and many of the relevant characteristics are difficult to observe without the use of highly intrusive techniques.

IV.5. Future research

IV.5.1. *Uncertainty = f (Complexity – Knowledge) ??*

For any topic, there may be a relationship between the sum of knowledge and uncertainty and a function of the complexity of the topic:

$$\text{knowledge} + \text{uncertainty} = f(\text{complexity}) ?? \text{ or}$$

$$\text{uncertainty} = f(\text{complexity} - \text{knowledge}) ??$$

That is, you can be both more knowledgeable *and* more uncertain about a subject that is highly complex. This may hold better if the complexity is assessed on a *system* in Klir's terminology, that is, a defined representation, instead of an *object* (the thing a system represents). For any system, there actually is a lowest level of detail knowable. In the real world, there may be no limit to the fineness of detail that could be assessed, and therefore no limit to the complexity ascribable to a situation.

Smithson argues that uncertainty, in fact any kind of ignorance, is socially created. He does not see uncertainty as pre-existing and infinite, to be reduced by each piece of knowledge. Then how should we construe the transmission of information in communication? The usual approach is to measure the information in bits, where one bit is the amount of information necessary to answer a yes/no or true/false question. But how does the *question* get transmitted?

In my optimistic moods, this dissertation appears to have quite a bit of information in it. In pessimistic moods, though, I am dismayed by how much uncertainty there is. Then at other times, both views seem correct, as I appreciate the complexity of this territory.

It seems that some sort of boundary is implicitly defined by a document, and the total territory is composed of knowledge plus uncertainty. It is not simply a matter of length; a 100-page document of random letters would have virtually no information, and while technically the uncertainty about which letter would be the next in the document would be high, in another sense there would be virtually no uncertainty in such a document either (no “significant uncertainty”? no “uncertainty of meaning”?). Perhaps uncertainty requires not just the potential for information but some actual presence of information.

What would be the maximum amount of information a 100-page document could contain? If the meaning of each place for a character is pre-established, then using each letter individually to answer a 26-way question would maximize the information content. But what if the questions themselves had to be contained in the document — then what would the maximum be? (Since any use of language presupposes a great deal of knowledge, part of the burden of transmitting the question is carried by the requirement of knowing the language in which the document is written.)

Which is more uncertain: (1) Chris is a tall man; or (2) Chris is 74.25 inches tall? Which is more informative? (1) gives you information that (2) does not, such as the fact that Chris is a man, but the height information is

not absolute, only relative to a population of men. If you have substantial contextual information about that population, then you may be able to interpret it fairly well. (2) does not tell you what kind of an object Chris is, but it tells you that it is equal in height to 74.25 inches (whatever inches are), and the precision is unknown, probably to the nearest .25 inches or to the nearest .05 or the nearest .01 inches. A height of 74.23 would be less fuzzy, and 70 would be much more fuzzy (actually, that is meta-fuzzy: it's not that 70 is necessarily less precise, but that the degree of precision is harder to infer). Or consider (3) The object is spherical with a diameter of 5.5 inches, or (4) The ball is somewhat larger than a baseball, but smoother and maybe denser. I might claim that (4) contains more information and more uncertainty than (3)!

This suggests that uncertainty is not only a function of the amount of information but also of the implied context — a larger amount of information within a much larger context may come with greater uncertainty. Therefore, one way to reduce uncertainty when the total amount of information that can be obtained or transmitted would be to reduce the context.

IV.5.2. Testing risk aversion vs. security level

The importance of a preference for security level compared to the importance of risk aversion could be tested by asking subjects to choose between sets of lotteries in which a lottery with a higher security level was also riskier than another (a mathematical definition of risk is given above in section III.2.9.1). One reason that confirming the importance to at least some decision-makers of the security level is that MEU theory has

difficulty in modeling such preference. This would be evidence that as a descriptive theory it needs to be extended, and for at least some people it could not serve as an adequate normative approach.

In the example shown in table 18, there are two options and four possible futures, with given probabilities of occurrence. The two options have the same expected value, and the same hope level, but they differ in riskiness and security level. Alternative 2 has a higher security level than Alternative 1, but it is also more than 50% more risky (as measured by the variance of options).

Table 18. Risk vs. security level

Probabilities of Futures				
Future	A	B	C	D
Prob.	0.20	0.30	0.30	0.20

	Utilities of Future/Alternative Pairs				Exp. Value	Risk (Var.)	Sec. Level	Hope Level
	A	B	C	D				
Alt. 1	80	54	46	10	48	506	10	80
Alt. 2	69.5	80	15	28	48	806	15	80

Subjects could be tested for their preference for one or the other type of gamble to determine how common a preference for security level is, and further, what conditions (e.g. duration of the post-decision, pre-outcome period), what personal characteristics, make preference for security level more likely.

IV.5.3. "Optimizing behavior"

If optimizing procedures cannot guarantee identifying the optimum (due to imperfect information, see Fleurbaey on optimizing procedures, section IV.2.3.1 above), then we have an open question: what procedure is likely to get us closest to the optimum under imperfect information? What

conditions indicate an advantage for one approach over another?

IV.6. Potential applications of this work to theory

IV.6.1. *T Perspective help for O Perspectives*

For the T perspective, which generally attempts to find conclusions that are as sure as possible, uncertainty that can prevent reaching such sureness thwarts a major goal. The T perspective typically attempts to reduce or eliminate uncertainty by discovering additional data, which can work well when the source of the uncertainty is a physical object or process that is researchable, or by establishing a “right way” to do things, as the mechanisms of MEU have been accepted as the right way to make a risky choice. However, T is frustrated when uncertainty cannot be reduced or eliminated.

The O and P perspectives emerge from worldly contexts in which the banishment of uncertainty is not even dreamed of. They may not be comfortable with uncertainty, or with some types of uncertainty, and its presence may be suppressed, but O and P have many ways of tolerating its existence. These methods frequently involve *design* more than *discovery*. For example when a court establishes the meaning of a phrase in a contract or a law (responding to linguistic uncertainty), or negotiators reach an agreement, or a decision-maker chooses values or designs new alternatives, uncertainty is reduced by doing rather than by learning. Uncertainty is accepted as a given by O and P, perhaps inconvenient, perhaps advantageous. In contrast, uncertainty often constitutes a serious threat to the fundamental validity of T, which is based on knowing, not guessing.

A researcher from the T perspective with a passion for service, recognizing this, might set about in a serious, task-oriented way to help develop 4th-stage tools for organizations to use in “suburbanizing” their uncertainties without trying to banish them.

IV.6.2. Alienation

Recalling Crozier’s assertion that rationalization of business tasks leads to a more rule-oriented workplace with fewer opportunities for power (section II.4.8), we could infer that workers would encounter fewer uncertain situations, or situations with lessened uncertainty. If they have less uncertainty about the consequences of their actions, then it is generally true that in comparison to the “expected” or mean outcome of their actions (however estimated), the best possible outcome, the “hope level”, is lower, and the worst possible outcome, the “fear level”, is higher. They will experience less hope and less fear. A cluster of possibly testable hypotheses grows from this line of thinking. One is that lower hope and lower fear lead to greater feelings of alienation. (I am not recommending increasing fear as an antidote to alienation.) Socialistic policies may have a similar effect: high tax rates may reduce the level of hope for material success that people experience, and social safety nets (e.g. guaranteed minimum incomes) can reduce the fear of material failure. If higher levels of hope can help to combat alienation, then lotteries would be more popular when alienation is common. (High correlation between low hope and high alienation could complicate this assessment.)

Crozier’s perspective suggests that as work is rationalized and more rules are created, the total amount of freedom within an organization will

decrease, and therefore the total amount of interpersonal power will also decrease. Workers and supervisors alike may find that they gradually have less freedom and less power, and at the same time are less subject to the personal power of others. Within the larger context of society, as the people and institutions with whom an individual interacts care less about his specific characteristics or are better able to predict his behavior, the individual will have less power over them to obtain the goods he desires. Conversely, he will be less subject to their power. As mutual irrelevance replaces uncertainty and interpersonal power struggles, feelings of alienation are quite likely.

IV.6.3. Gamblers and natural selection in the organization

If Dror is right, and most important decision-making is fuzzy gambling (Dror 1986b), then good gamblers are needed at the helms of organizations. However, there is often a strong inducement within the ranks of organizations for employees to avoid risks, and for the organization to promote those employees who successfully avoid risks — they are the ones to whom no taint of failure is attached. Risk aversion leads to more inaction in times of uncertainty. An organization that does not reward risk-taking, one whose native perspective does not focus on gambling as a skill, may not be able to distinguish good gamblers from reckless gamblers, only gamblers from non-gamblers. Even if the need for executive gambling skills were recognized, such a group may be inexperienced enough with the characteristics of skillful gambling that it cannot select the responsible gamblers it needs.

V. GUIDELINES

Reason is and ought to be the slave of the passions and can never pretend to any other office than to serve and obey them.

David Hume

V.1. Introduction

The purpose of this chapter is to present a set of guidelines for uncertain decision-makers that emerge from the preceding taxonomy and implications. Before the exposition of the guidelines, a set of Rational Advice Guidelines (RAG) is developed in order to evaluate proposed methods for aiding uncertain decision-makers, i.e. to evaluate decision support tools. I assume for this analysis that the decision-maker wants to act rationally, but that how to do so is not obvious due to uncertainty. Many discussions of rationality deal with rationality in the abstract, and attempt to characterize the ideally rational decision-maker. This work, however, begins by assuming a human decision-maker in the murk of the real world, and will attempt to characterize behavior that would be rational for that decision-maker. The RAG will operate as criteria for evaluating the suitability of specific decision approaches in specific circumstances, i.e. for evaluating advice that might be given. The RAG outlines the purposes of rational decision support under uncertainty that are assumed for this work. The product of the subsequent analysis amounts to advice to decision-makers facing uncertainties — not so much advice on which choices to make, but on the process they use to go about

making their choices.

The guidelines are grouped into sections on problems caused by fuzziness and those caused by gambles, the special difficulties imposed by fuzzy gambles, and meta-rationality concerns, essentially issues of how to rationally manage processes intended to assist rational choosing.

Building upon a taxonomy in which decisions are divided and subdivided in first this way and then that, this analysis may seem to take formalization to such an extreme that it loses touch with actual decisions. It is intended to be an extreme view, dominated by the technical perspective, which must eventually be tempered in application with judgment, which is always so much harder to describe analytically and precisely than taxonomies. Speaking for judgment:

In human decision making there can be no separation of means and ends, alternatives and objectives. Both sides of the decision process—analysis of what is available and analysis of what is desirable—are equally important, interdependent, and interactive. Neither aspect is primary.

...Neither means nor ends are determined independently of each other. Objectives evolve on the basis of available alternatives which, in turn, are adjusted and generated in accordance with existing objectives. Modeling of decision processes and economic behavior must take into account the fact that the mathematically convenient separation of means and ends does not accurately describe the way in which people typically approach decisions. (Zeleny 1982:102)

I've heard it said that "God is in the details". Then again, others say "The Devil is in the details." A range like that certainly indicates the importance of details. Details, however, are often what the technical perspective not only loses but *must* lose in its quest for general truths. The richness of any particular organizational or personal perspective comes from just those details that technical eyes skip over in their search for

common themes. Any generalization ignores particulars that are important in specific cases, and aiming for a broad scope has forced me to ignore especially many. What this means is that there may be exceptions to all the guidelines given here when the political, personal and organizational realities of specific cases are taken account. What I can hope for is that parts of this technical elaboration of uncertainties can be useful when adopted by decision-makers according to their unique needs and modified by their judgment. Even the best decision support still leaves the decision-maker(s) with the final responsibility to *make* a decision, to create movement in one direction and not in other even though maps and compasses are flawed. The ubiquity of uncertainty and the difficulties it poses for decision-makers ensure that those who can face uncertainties openly without becoming paralyzed or reckless and still choose intelligently will be highly appreciated.

V.2. The context of the Rational Advice Guidelines

To know one's ignorance is the best part of
knowledge.

— Lao Tzu, *The Tao*, no. 71

V.2.1. *Purpose of decision support*

Broadly speaking, the purpose of decision support is to assist a decision-maker. Decision-makers have many kinds of needs, as mentioned in section II.10.7 (Ascher and Overholt 1983:44ff.). They have political needs to be convincing and to be well-supported by their information, to be right, or at least not badly wrong, or at least not to be blamed, and they need to maintain their ability to make further decisions. They have

psychological needs for certainty (to reduce anxiety), for simplicity, and to have their personal views of the world confirmed (in order to maintain their self-confidence). These two classes of needs belong to the O and P Perspectives described earlier. Decision-makers also have rational needs, corresponding to the T Perspective, however, and the focus of the field of decision support has tended to be on some of these. As the current work is a study of how to support rational decision-making, it too will focus on these rational needs along with some rational considerations of the non-rational needs. The political need to avoid having any mistakes attributed to the decision-maker is not discussed in the technical decision support literature, nor is the psychological need of the decision-maker to feel less distressed by uncertainty. Analysis of these needs could be quite interesting, raising, among other issues, ethical questions of how much attention could be given to them before running the risk of encouraging dishonesty or selfish gamesmanship, in the case of organizational needs, or merely tranquilizing the decision-maker, in the case of psychological needs, but that analysis will be left for later works.

V.2.1.1. Rational decision support

The purpose of rational decision support is to help the decision-maker make better decisions, chiefly by meeting her rational needs for information, but also by reasoned consideration of organizational and personal realities. This can include recommending specific methods, implementing the methods, gathering data, interpreting it, highlighting risks, helping to develop new options, considering the context of the decision, and other efforts that can lead to better, or lead away from worse,

decisions. Rational decision support is aimed at helping the decision-maker decide better, not feel better — and this can create serious conflicts for the honest analyst supporting a decision-maker known to dislike bad news (e.g. inconclusive results).

V.2.1.2. Assessing the quality of decisions made under uncertainty

The assumption common in this field is that decisions are better, i.e. more rational, when they select means that produce more of the ends the decision-maker desires (“means-ends rationality”, as noted in section II.3.2). The problem here is that when a decision is being made under uncertainty, the results of a potential choice cannot be foretold, so that the desirability of the outcome may depend as much on the luck of the decision-maker as on the quality of his decision-making. Therefore, the quality of the decision-making must be assessed prospectively — before the luck — not retrospectively (Elster 1989a:117, Dror 1988b:157).

Dror suggests that grand policy analysis (which is what he calls broad-scope decision support) should be based on philosophy of judgment more than philosophy of science:

Science strives for truth, as defined from time to time by consensus in different disciplines. But the mission of grand policy analysis is different though correlated, namely to improve decisions, in the sense of arriving at better ones than would otherwise be obtained...

Grand policy analysis accepts the adage of Maynard Keynes “it is better to be roughly right than precisely wrong”. (Dror 1991:7)

Janis and Mann have proposed that good decisions share seven characteristics that can be used in *ex ante* evaluations:

From the extensive literature on effective decision making ...¹ we have extracted seven major criteria that can be used to determine whether decision-making procedures are of high quality. Although systematic data are not yet available, it seems plausible to assume that decisions satisfying these seven "ideal" procedural criteria have a better chance than others of attaining the decision maker's objectives and of being adhered to in the long run.

The decision maker, to the best of his ability and within his information-processing capabilities

1. thoroughly canvasses a wide range of alternative courses of action;
 2. surveys the full range of objectives to be fulfilled and the values implicated by the choice;
 3. carefully weighs whatever he knows about the costs and risks of negative consequences, as well as the positive consequences, that could flow from each alternative;
 4. intensively searches for new information relevant to further evaluation of the alternatives;
 5. correctly assimilates and takes account of any new information or expert judgment to which he is exposed, even when the information or expert judgment does not support the course of action he initially prefers;
 6. reexamines the positive and negative consequences of all known alternatives, including those originally regarded as unacceptable, before making a final choice;
 7. makes detailed provisions for implementing or executing the chosen course of action, with special attention to contingency plans that might be required if various known risks were to materialize.
- (Janis and Mann 1977:11)

Janis and Mann's criteria are clearly dominated by the T Perspective, except for a brief nod to the importance of objectives and values.

Fischhoff et al. developed a similar list of requirements for approaches to "acceptable risk problems" which would serve well in evaluating any decision process (Fischhoff et al. 1981:54ff.). Concluding that acceptable risk problems should not be treated as strictly technical issues, they give

¹ Eleven references in original omitted here.

roughly equal weight to organizational and personal perspectives as well. Their criteria are more sensitive to the context of decisions. (The letters following each criterion indicate which perspective or perspectives each is based on.)

1. Comprehensive: “an approach should accommodate a comprehensive problem definition, reflect the uncertainty surrounding technical issues, acknowledge the labile or conflicted nature of social values, realistically appraise the human failings confronting the decision-making and -implementing processes, and assess the quality of its own conclusions ... and be flexible enough to accommodate new information.” (T+O+P)

2. Logically sound: an analysis should provide a clear conclusion reached by a decision rule that is sensitive to nuances, reliable or replicable, justifiable by either sound theory or prior success, suitable to the particular task, and unbiased. (T)

3. Practical: “it must be possible to implement the approach with real problems, real people, and real resource constraints.” (O+P)

4. Open to evaluation: the assumptions should be made explicit, the boundaries between facts, judgments and values should be clear, and the total uncertainty in conclusions should be described. (T)

5. Politically acceptable: an approach must be capable enough not to be vulnerable to attacks, must not be both so capable and so unwelcome in its recommendation that opponents elect to attempt to negate the whole process, nor can it ignore the need to bring constituents along with it in its process. (O)

6. Compatible with institutions: major institutions cannot be changed

over night, but if they need to change, the approach should facilitate it. (O)

7. Conducive to learning: “an approach should educate its participants, eliminate opportunities for obstructionism ... Somehow, society should become better or wiser for its adoption”, perhaps through “(a) leaving a clear record of deliberations and assumptions ... (b) affording two-way communication between scientists and decision-makers ... (c) educating lay observers ... and (d) having enough scope to be used on many problems.” (P)

Clearly, then, one standard of good decision support is that it facilitate good decision-making along the lines of the above criteria.

Incidentally, while some hope that market competition will reward those organizations that have better decision-making, Winter thinks that uncertainty prevents this variety of *ex post* evaluation of decision-making quality as well:

Although the decisions are of the utmost importance, the problems of prediction involved in making them are so large that the uncertainties may swamp any systematic superiority of one firm over another. (Winter 1964:266)

Therefore, better decision making may not be rewarded with competitive advantage.

V.2.1.3. Decision-makers' rational need for information

The most basic need is for information about available options, their likely consequences, the relative importance of those consequences and so on — the “subjects of uncertainty” described earlier comprising the Generalized Decision Model (GDM) in section III.2. But decision-makers have additional rational needs for information:

1. Knowledge of the *uncertainty* of the information...
2. Knowledge of where the factual information ends and the recommendations begin. Recommendations disguised as facts ... can mislead the policy maker and limit his range of choice.
3. Knowledge of the *context* from which the specific information was derived, so that the policy maker can infer the full meaning of the specifics...
4. Knowledge of the *assumptions* underlying the analysis ...
5. Knowledge of which factors are *relevant* for adequate understanding. (Ascher and Overholt 1983:44)

These five factors differ from basic knowledge about the components of the GDM in that they are information *about information*, or

metaknowledge: an appreciation of what we do know and what we do not know. Normally, we define knowledge as consisting of all the facts, concepts, relationships, theories, and so on that we have accumulated over time. Metaknowledge concerns a higher level of expertise: understanding the nature, scope, and limits of our basic, or *primary knowledge*. Metaknowledge includes the uncertainty inherent in our premises and world views. (Russo and Schoemaker 1992:8)

One main criterion for methods recommended is that they increase, or at least not decrease, the decision-maker's *metaknowledge*.

[Decision makers] are better off with an accurate picture of how extensive their ignorance is than a falsely precise assessment. (Smithson 1988:262)

The primary rational purposes of this type of decision-support is to inform the decision-maker. She needs to know as much as possible, within the constraints of available time and attention, about the options, their consequences, etc. She also needs metaknowledge — information about how to interpret the first type of knowledge and how much confidence to put in any of these descriptions.

V.2.2. The Rational Advice Guidelines

Arguments are put forth here for criteria for evaluating decision

support under uncertainty, based on the belief that effective decision support should promote good decision-making, discussed a bit earlier. Since these guidelines are intended to provide a framework for evaluating *procedures* that decision-makers could adopt, it follows that procedural rationality will be the source of most of these principles. (Substantive rationality, procedural rationality, and meta-rationality are discussed in section II.3.5.)

V.2.2.1. Meta-rationality guidelines

The main consideration arising from meta-rationality is that we must be alert for possibility that rationality and procedures based on reason may plausibly fail to be decisive. For any of a number of reasons, there may not be enough information for a decisive rational conclusion to be drawn. There can be ties, or virtual ties, in which there is no clearly best option, or cases in which there clearly is no single best option. Slight differences may not matter: a difference, to be a difference, must make a difference — a significant difference.

Additional reasons for the muteness of rationality on particular issues come from the O and P Perspectives. Rationality alone cannot settle political disputes, though it may be helpful in illuminating them or clarifying possible resolutions. Rationality cannot point the way to specific options when the decision-maker is not sure of which consequences he prefers; again, reason may be useful in highlighting this as an issue needing to be settled, but most treatments of rational decision-making explicitly define the choosing of values as exogenous to rationality.

V.2.2.2. Substantive rationality guidelines

Decision support is substantively rational when it informs the decision-maker, both on the level of basic information (about options and values, etc.) and on the level of metaknowledge (about the basic knowledge). Merely flooding the decision-maker with data does not inform him — appropriate scaling of level of detail and of quantity of information is necessary. Another aspect of informing a decision-maker is to avoid misleading the decision-maker by providing deceptive or erroneous material.

Another substantive goal is motivated by the recognition that the decision-maker will make other decisions later. Therefore, one goal of decision support is to leave the decision-maker in a better position to make good choices in subsequent decisions. A chief way to achieve this is through learning from current experiences, for example, by learning effective strategies for reaching decisions, acquiring more information about the relevant environment or the values of relevant stakeholders, and by strengthening memory and communication processes so to facilitate making later use of what is learned.

V.2.2.3. Procedural rationality guidelines

V.2.2.3.1. Practicality

Since the fundamental premise of this work assumes the presence of significant uncertainty in decision-making, the most critical procedural criterion is that any method for assisting the decision-making must be feasible in conditions of insufficient information, i.e. under uncertainty. Methods that require, for example, a level of detail about future

consequences that cannot be predicted or that require knowing more about the preferences of each of a large number of stakeholders are not helpful.

A closely related principle is that methods must be ones that can actually be implemented by the decision-maker. In some cases, it may be worth while for the decision-maker to learn new skills in order to utilize a technique; in many cases, recommendation of a technique has to take into account the limitations imposed by the decision-maker's current skills and abilities.

V.2.2.3.2. Uncertaintest Link Principle

The Uncertaintest Link Principle:

The conclusions of an analysis can generally be no more certain than the most uncertain component in the analysis.

There are two immediate implications of this principle. The first is that to reduce uncertainty in the conclusion, effort spent on reducing the uncertainty in the uncertaintest link is likely to be the most successful. The second is that if the most uncertain link cannot be made more certain, there may be little reward for reducing the uncertainty in the other links.

Which link is in fact the most uncertain depends not on an isolated measure of uncertainty in that link but rather on the amount of uncertainty in the conclusion propagated by the link. For example, if the total dollar costs of all the options were known only to within 50%, that would ordinarily appear to be a high level of uncertainty. But if the dollar costs were not especially important, perhaps because all estimates were deemed to be easily affordable, a measure of the effectiveness of the options might be the uncertaintest link, even if the variance of the estimates were much lower than that of the dollar costs.

To take another example, it is often argued that when we do not have statistical data on which to base estimates of the likelihood of occurrence of significant possible events, we ought to use judgmental or subjective estimates of those likelihood, since that is the best information available. This is a persuasive argument, I believe, as far as it goes. It does not compel us, however, to treat the subjective estimates the same way we would treat estimates derived from historical data, and it does not absolve us from the responsibility to assess the uncertainty in those estimates themselves and the consequent uncertainty in any conclusion about which option is best that is induced by the uncertainty in likelihoods. In many expected-value analyses conducted within organizations, there will be very small investments of time and other resources in the creation of subjective probabilities compared to the overall investment in the analyses. When the uncertainty in the probability estimates is large, the resulting uncertainty in the expected value of each option must also be large, severely limiting how much could be gained by obtaining highly precise estimates for other parts of the analysis.

Testing the sensitivity of the uncertainty in the conclusions to reduction in particular uncertainties should be a prerequisite to major investments in uncertainty reduction.

V.2.2.3.3. Non-arbitrariness

Concern for the efficacy of reason suggests a fundamental criterion: a procedure for deciding should not be arbitrary, that is, it should not point to one option instead of another due to the influence of factors that have little or no bearing on the utility of the options. While this may make any

given decision procedure less decisive, it will prevent a kind of outcome that I think we can all agree is not desirable. For instance, if our data is poor, and generally has a margin of error of $\pm 50\%$, we would not want to accept a difference of 1% in estimated utility as decisive: we would say that a difference that small in comparison with the much greater uncertainty in the data is not significant. To decide on such a basis would be committing an error of arbitrariness.

This condition will, I think, strike most people as reasonable. A procedure for deciding ought not to claim decisiveness of result when any apparent decisiveness is an artifact of the procedure instead of a reflection of real differences in the desirability of the choices. If the procedure does not find a distinct advantage for one alternative, the result should be indeterminate. It is not rational to choose one option over another, for example, merely because of the order in which various attributes were examined, or because of a difference in the 16th decimal place in some kind of desirability score.

Non-arbitrariness is a condition that is much more important in assessing procedures than it is in discussing general, abstract principles such as maximization of expected utility (MEU), because implementation is such a rich source of arbitrariness. Any procedure for implementing maximization, for example, ought to meet the condition of non-arbitrariness. The value of MEU as a principle is not diminished by the determination that a particular implementation of it is arbitrary; there may be far better implementations available. On the other hand, if there are decision contexts in which it appears that *no* implementation *could*

avoid arbitrariness, then the *principle* of MEU is not *feasible* for those contexts—it isn't workable, and other principles must be invoked. There are other problem contexts in which a principle that is successful in theory cannot be made practical, many of which require computations that cannot be completed with current methods. The game of chess, for example, is in principle solvable by analyzing all possible sequences of moves. While this exhaustive search of possibilities can be shown to be finite, it is nonetheless so large that the "in-principle solvability" is of no use to human players, and is of little use in practice for computer chess-playing programs at this time.

V.2.2.3.4. Maximum Uncertainty Principle

Klir proposes a *maximum uncertainty principle* dealing with this matter:

maximize the appropriate uncertainty function (or functions) subject to constraints that represent available information. ... in any ampliative² inference, we should use all but no more information than available. That is, we require that conclusions resulting from any ampliative inference maximize the relevant uncertainty within the constraints representing the premises. (Klir 1992b:81)

If this principle is followed, the risk of arbitrariness will be minimized, and the decision-maker will not end up believing the results are more conclusive than they really are. This requires careful presentation of the uncertainty of any conclusions, and may require modesty and self-restraint from analysts. The distinctions between what is known and what is merely assumed for any analysis must be made clear.

² Enlarging or extending (type of analysis).

This is an area where decision support can usefully borrow from experimental science. Just as the results of an experiment are explicitly tested for the quantitative significance of their implications for the truth of hypotheses, the results of a decision analysis can be tested for their significance regarding the superiority of one or more of the options. This Null Hypothesis approach was described above in section IV.1.2.

An assessment of the impact on the uncertainty of the conclusion that each input fuzziness produces would also highlight which are the most important, and would be the best candidates for fuzziness reduction so as to reduce output uncertainty sufficiently to produce a unique recommendation of an option. Conversely, if the uncertainty stemming from one factor is shown to be both irreducible and large enough to render the rational selection process indeterminate, it makes no sense (for this decision) to attempt to reduce any of the other uncertainties. Another selection process will definitely be needed.

V.2.2.4. Summary of the RAG

A good decision support method should:

- Provide information: illuminate the consequences of choices, trade-offs among different values, etc.
- Provide metaknowledge: clarify what the decision-maker knows and doesn't know, what she can and cannot know.
- Improve decision-making: result in making better decisions now and in the future.
- Be practical: Must be usable when information is limited, must be feasible for the decision-maker, must not depend on unobtainable information.
- Avoid arbitrariness: represent the results as being as no more

certain than warranted by what is known, avoid false impression of precision, avoid false impression of certainty.

V.2.3. Diverging from economic rationality

The perspective employed here differs from the typical perspective of normative decision theories. One difference is the virtually automatic adoption in the literature of MEU as a decision principle; another is the literature's requirement that decision-makers have consistent preferences.

V.2.3.1. Adoption of MEU

Means-ends rationality is accepted both here and in the literature, but the implementation of means/ends rationality remains a major issue.

David Gauthier (1982) has "claimed that in our society the received conception [of rationality], which most persons do accept initially given their socialization, does identify rationality with maximizing activity."

While accepting means/ends rationality, I will not automatically accept its identification with a procedural implementation of MEU.

MEU is not actually employed as a criterion for evaluating advice to decision-makers — it is used as an *ex ante* criterion for the rationality of decisions themselves. Use of MEU as a criterion presupposes that MEU has already been selected as the method for reaching a decision; this work looks at an early stage in the process of adopting a method for making a decision, and compares alternative procedures, including MEU, against a set of criteria and the particulars of each decision situation.

V.2.3.2. Consistency of preferences

There are usually two criteria defined for rationality (see section II.3.2.2. Economic Rationality): 1) means/ends rationality, and 2) consistency of preference. I suggest here that for evaluating decision

support methods, it is important to prohibit arbitrariness. Furthermore, preference consistency is not a characteristic of support methods or of decisions, but of decision-makers, and therefore should not be adopted here as a criterion.

What do we lose by giving up the consistency criterion? We lose analytic power: permitting only consistent preferences makes it more likely that the situations we do consider can be resolved by rationality alone. The huge downside to that is that there could then be many decision-makers that we could not even assist. Omitting it, we are faced with the possibility that some decisions cannot be decided solely by rationality — a result that should not seem counter-intuitive to anyone with much experience of the real-world contexts in which human beings need to make decisions.

Many theoretical discussions of rationality do include a restriction on preferences that rules out preferring A to B, B to C, and yet preferring C to A. Such a preference structure simply seems irrational to many observers. It has long been noted, however, that a similar intransitivity can arise in the Voters' Paradox³ with three voters, three choices, and pair-wise balloting. In fact, if we consider that the consequences of the choices open to an individual may have three components, and that the chooser's utility is "fuzzy", a plausible argument can be made that cyclic preferences may occur quite naturally. Suppose that aspect 1 of choice A, called A1, is

³ There are, say, three voters ("1", "2", & "3"), who are to select one of three candidates ("A", "B", and "C"). Voter 1 ranks the candidates A, B, C; voter 2 prefers in the order of B, C, A; and voter 3 prefers C to A to B. Any pair-wise vote is won by a 2-1 margin, with A defeating B, B defeating C, and C defeating A. The ultimate winner depends on the order in which the comparisons are made.

slightly preferred to B1, which is slightly preferred to C1: the chooser may be nearly indifferent to an A1 - B1 choice, or to B1 - C1, but feel that the difference between A1 and C1 is indeed significant. If similar preference orderings, properly out of phase, are assumed for the second and third aspects, cyclic preference results⁴. It is certainly convenient to be able to exclude the awkward case of cyclic preferences, but if one cannot be sure that an otherwise reasonable person (or organization) would never possess such a preference structure, then a comprehensive normative guide needs to allow for that possibility. It may be that in such cases the decision-maker has pair-wise preferences but is indifferent to the choice among all three alternatives, and that rationality cannot point to a "best" choice. Another guideline needs to be invoked to break the tie.

Therefore, since it has been known that group preferences can be cyclic (insofar as groups can be said to have preferences), and an argument can be made that cyclic preferences are plausible for individuals with preferences over more than one dimension of consequences, demanding consistency of preference from decision-makers does not appear to be helpful in advising decision-maker on how to proceed when they are

⁴ B2 is slightly preferred to C2 which is slightly preferred to A2, but B2 is strongly preferred to A2; and C3 slightly preferred to A3 which is slightly preferred to B3, but C3 is strongly preferred to B3: A is preferred to C, due to strong preference on aspect 1; C is preferred to B, due to strong preference on aspect 3; B is preferred to A, due to strong preference on aspect 2. The decision-maker may feel that the cost of caring about a "slight" preference is as great as a slight difference in utility.

Another source of cyclic preference: assume the same preference orderings as before, but make all preferences of equal (or indistinguishable) intensity. Thus, A1 is preferred to B1 which is preferred to C1, but the preference for A1 over C1 is not perceived to be of any greater magnitude or importance than for A1 over B1 or B1 over C1. This is essentially another version of the Voters' Paradox, with the three voters replaced by three aspects of the preferences of a single decision-maker.

uncertain.

V.3. The problems of uncertainty

The difficulty in life is the choice.

George Moore

Uncertainty is not always a problem, of course. As Jon Elster notes (section II.7.1), an overly certain life may be repellent, and the popularity of gambling, a deliberately chosen uncertainty, remains high in spite of strong social sanctions. Furthermore, in competitive situations, causing uncertainty for an adversary can be very helpful. When there are possible futures for which planning or preparation can be helpful it is generally better to be aware of them, and thus to be uncertain, than to be ignorant of them, and thus falsely certain.

But in the context of decision-making on the basis of rationality, uncertainty is associated with problems, which are of three general sorts. The first is that uncertainty about the value of a relevant factor can make it difficult to determine which option is the best to choose. The second is that there are often possible outcomes that are undesirable but not certain. In this case, it is the undesirability that is the main problem; the uncertainty is actually a virtue, since an uncertain ill is generally preferable to a sure one. The uncertainty of the ill complicates the task of designing an appropriate remedy, especially since there may be many possible harms which therefore cannot all be prepared for. The third sort of problem associated with uncertainty is that the gamble type of uncertainty can itself have costs when it persists through time, delaying other actions that require resolution of the uncertainty. Of these three sorts of problems, the first two are the most important; the third will be addressed below as

part of the second.

Dror's term *fuzzy gambling*, which he believes applies to virtually all policy decisions, neatly illustrates the distinction to be made here (see section II.7.3). Fuzziness by itself can make a decision among even completely deterministic options very difficult. A choice between two piles of money may be hard if estimates of the total value of each pile have large margins of error. At the other extreme, a choice between two perfectly specified options is hard when one option carries a risk of a bad outcome (relative to other possible outcomes). Choosing such an option constitutes a gamble, whether we want to admit it or not (and frequently there are no non-gambling options).

Often improving the possible outcomes of one option will make the selection easier. Reducing uncertainty is not always the best approach to faring well under uncertainty; the end should be improvement in the decision-maker's general prospects, to which uncertainty reduction can often but not always contribute.

V.3.1. What can be recommended?

When in doubt, win the trick.

Edmond Hoyle

Throughout the Taxonomy chapter, various suggestions for coping with uncertainty about different factors in a decision, or uncertainty stemming from different sources or arising in different circumstances were proffered. These are summarized next in several sections. In the first section, methods for coping with problems of difficult decision-making relating primarily to fuzziness are described. In the next, I discuss approaches for dealing with problems caused principally by the gambling aspect of a

choice, including the cost of uncertainty. The third section deals with problems that involve both fuzziness and gambling aspects. The fourth section considers bounded rationality — aspects of the decision-maker and his pragmatic constraints. The fifth section, on meta-rationality, treats the question of how to be rational about the first four sets of issues.

The material covered here is broad, and many others have considered particular areas within here in much more detail than I ever can. This is intended to be a framework, an organizational schema, with a some new distinctions I did not find in any literature. All of the specific suggestions here should be supplemented by the consultation of a more focused treatment. Risk analysis, to take one topic, is too broad and deep to be covered properly by any single work, let alone by a few paragraphs here; this framework may help someone determine that risk analysis is relevant to their needs, but it cannot meet their needs in detail.

V.4. Advice for problems of fuzziness

The chief problem caused by fuzziness is indeterminacy — it is not possible to identify clearly which option is the best or most rational, and a decision selection process fails to yield a unique recommendation. Indeterminacy is not always a serious problem. For example, if two options are both very good, it may be impossible to tell which is better and therefore is more rational — but it may also be unimportant. If you know that they are both great, then you can't go far wrong picking either one, and it is not worth wasting much effort on the selection. However, if the outcome is not that two are tied for best, but that the result of the comparison is fuzzy, then it does appear to be possible to go wrong, and

more effort may be warranted.

In the pure case of fuzziness-induced indeterminacy, there is no fundamental uncertainty about whether Future X or Future Y will occur, or what would result from making a particular choice. Instead, there is fuzziness about the precise values of relevant factors that makes selecting a “best” option difficult. The typical factors that can cause this are fuzziness about the state of the world (uncertainty of *degree*, not uncertainty of *kind*), fuzziness in the function mapping a state of the world and an option into a set of consequences, and fuzziness about the value of consequences (fuzziness about a specific utility function, the methods of combining evaluations of multiple parties, multiple attribute fuzziness, or fuzziness about the relative importance of consequences that will occur at different times). The ideal “rational” solution for this uncertainty problem is to reduce the fuzziness to such an extent that it is possible to identify a “best” option. When the fuzziness is due to uncertainty about collective preferences, in lieu of attempting to collect preference information and then determine which option is the “best”, a social choice process can be adopted so that the information, which is often difficult or expensive to pursue, need not be explicitly obtained; the information acquisition and subsequent decision-making are collapsed into a single step.

Approaches to dealing with indeterminate decisions due to fuzziness generally consist of obtaining more information or using a social choice selection process. (While it may not make the task of deciding any easier, searching for new options that are better but perhaps no less fuzzy may pay large dividends.)

V.4.1. Reduce fuzziness, obtain more information

There are no great insights to be offered here: obtaining more information, or more precise (i.e. non-fuzzy) information is a major human enterprise with many subfields to which this work can add little. The approach to seeking information often depends on the type of information that is needed.

V.4.1.1. Better measurement

Since the problem encountered here is largely one of difficulty in making decisive comparisons, improving measurement capabilities can help. More powerful or more sensitive instruments can reduce imprecision.

V.4.1.2. Experiment

Conducting experiments can provide a great deal of information. Multiple, parallel experiments may provide faster learning about complex systems than pure analysis or than fewer, larger experiments. For example, learning about the effectiveness of interventions in social problems may be faster with multiple, decentralized trial programs than with unified Federal Government trials.

V.4.1.3. Function uncertainty

Fuzziness about predictions of what consequences will follow from pairings of options and future states of the world may be reducible by gathering data about the world to refine theories about how the relevant portions of the world work. (Uncertainty due to belief that the consequence function is either non-deterministic or deterministically chaotic generally puts the decision-maker in the position of gambling, described later.)

Uncertainty due to the existence of more than one model of how the world

works may be reducible by research; if not, the decision-maker is left with a gamble.

V.4.1.4. Values clarification

Imprecision in estimation of values may be amenable to reduction, but it is important to remember the conclusion of McCord and de Neufville (section II.5.4.6) that different methods of eliciting information about individuals' utility curves produce incompatible results. Thus, while it is likely that a procedure can be designed that produces non-fuzzy estimates of values, some uncertainty or margin of error ought to be ascribed to the results.

Von Neumann and Morgenstern noted that people typically experience considerable haziness about their utility evaluations (section II.3.7.2). Nonetheless, people may be able to reduce that haziness by *deciding* more clearly on their values, through explicit exercises, more vivid renderings of consequences to stimulate evaluations, or other methods of values clarification. Open-minded debate and exploration can help surface more factors that bear on people's assessments of value. Comparing differences in both relative and absolute utility⁵ across options may highlight preferences across options.

When the decision-maker is uncertain about the values of others who for some reason cannot speak for themselves, designating trustees to speak for them can facilitate decision-making.

⁵ "Relative" here means assessing differences between the current level of utility and prospective levels; "absolute" means making comparisons solely on the basis of prospective, "final" utility levels. Either may be useful.

V.4.1.5. Fuzzy time comparisons

Fuzziness in comparisons of the relative importance of consequences occurring at different times is a fuzziness of value, and not soluble by purely rational means. Individuals may be able to clarify their feelings on this by imagining greater temporal distance from the issue, for instance by imagining looking back on the decision from their death bed, or by trying to predict how their descendants would view the decision much later.

Organizations often adopt policies governing such comparisons, such as an agreed-upon discount rate. Another sort of policy is to decide that you must leave your children a better world than you inherited. The advantage of a policy is that until you decide you need to change it, you have a guide for making intertemporal comparisons. When parties with different value systems attempt to make a decision together, assessments of the relative merit of options may be very contentious if time comparison issues remain unresolved, and adopting in advance a process that will produce a social or political agreement on time comparison may facilitate reaching agreement on the evaluations. The disadvantage of adopting a policy is that it has costs in time, effort, or for individuals, possible feelings of unnaturalness, and additional relevant factors may turn up after the policy has been adopted. Without a policy, the time comparison decision can be made implicitly as options are compared with each other. The more often similar situations arise, and the more similar they are, the more convenient a policy will be.

V.4.1.6. Multiple objective evaluation uncertainty

Uncertainty over comparisons of multiple dimensions of value shares

some of the characteristics of the problem of choosing among gambles. It is not risk of a bad outcome under one possible future for one option balanced perhaps against the risk that a different option would turn out badly under a different state of the world, but the fact that one option is better under one criterion than another, while the second is better than the first under another. The fundamental similarities with choosing among gambles are that (1) additional information may not help at all, and (2) the best outcome for the decision-maker would be to develop a new option that does not have the liabilities of the current options. In this case, that means to develop a new option that is at least as good as any current option on the first dimension of value, and also at least as good as any current option on the second dimension of value, and so on.

Assuming, however, that new options have been explored, choosing may still be difficult. As described in section III.2.6, two general routes decision-makers can take are to develop multi-dimensional utility curves explicitly, which can then be applied to a set of options and consequences, or to use techniques that operate more directly on the set of available options under consideration, such as trade-off methods. Optimization requires adopting a single criterion, and the method by which multiple objectives are aggregated into a single objective clearly has a great influence on the final choice, and therefore on the claims that can be made for the optimality of the choice.

Among the most common methods of developing an explicit utility curve the translation of all consequences into dollars (in many analyses, it is then usually assumed that utility is linear with dollars). This can make

the comparisons of fundamentally different consequences more palatable (but can also mask trade-off issues that are not yet legitimately resolved). If numerical parameters are used in the comparison, such as weights, the sensitivity of the result to slight changes in the parameters should be investigated to reduce the risk that the result is driven by the details of the process and not by the fundamentals of the decision. If all consequences can be expressed in a single measure, optimization might make sense; if the attributes are incommensurable, maximin combined with tradeoff may make more sense.

The more the options differ from each other on any of the dimensions of value, or the more dimensions are relevant, the more likely it is that fuzziness on any one dimension or in the balancing of dimensions will make the comparison indecisive, and require adopting another selection process, perhaps after screening out some of the options.

V.4.1.7. More powerful comparison methods

There is some potential for increasing the power of the comparison to reduce the fuzziness or indeterminacy in output even when input fuzziness cannot be reduced further. For example, multiple measurements with the same technology can produce a more precise answer, as can increasing the sample size in statistical measurements, or changing the way data is aggregated to retain more of its original information may increase the comparison's decisiveness.

V.4.2. Social choice processes

V.4.2.1. Agree on process

When uncertainty in the evaluations of outcomes and options depends

on the preferences of two or more people, an alternative to attempting to ascertain their preferences individually and then adopt a procedure for combining their preferences into a group decision is to adopt a procedure that accepts data from each person and directly produces a decision.

Majority voting is such a procedure: the complete preference structures of the individuals are not elicited, and there is no formal method for combining rankings. Rather, each individual expresses a more limited preference, e.g. “yea” or “nay”, which is used in a simple way to produce the decision. This is more a process of creating the decision than of discovering the values and then simply making a selection according to the discovered values.

V.4.2.2. Agree on representative or on government

Democratic governments represent other methods for reaching group decisions; adopting a form of government enables a large set of decisions to be settled. A direct democracy will refer many decisions to the voters themselves (as do the initiative and referendum options in many American state governments); a representative democracy empowers representatives of the voters to make many decisions on behalf of the people (as is typically done by American governments with the exception of referendum-like processes).

An alternative would be to use statistical or other assessment of the preferences of the members of the group to guide the decision-making of an executive, and if all leaders could be trusted to carry out the will of their constituencies as accurately as possible, this might be a more popular form of governance.

V.4.3. Choosing among fuzzy options

V.4.3.1. Adopt Null Hypothesis

Assess the significance of the input fuzziness for the output of the selection process. If one option is significantly better, given the uncertainties, choose it.

V.4.3.2. Reduce fuzziness

If a significantly better option cannot be identified, try to reduce the output fuzziness by targeting the key sources of fuzziness, essentially by acquiring more information. If this does not succeed in reducing the uncertainty in the output of the selection process sufficiently to identify a single answer, an additional selection process will have to be employed.

V.5. Advice for problems of gambling

The chief problem caused by the gambling aspect of choices is that there is a risk of a bad outcome. The uncertainty may or may not be a problem; a 50-50 gamble between two equally great outcomes contains uncertainty, but no risk of coming out badly. It is when one of the possible outcomes is significantly more desirable than another that the uncertainty in outcome really complicates the choosing.

The case of pure gambling is the one for which the SDM was developed — all factors are known with as much precision as is desired, and the consequences are all of the same nature, but the future state of the world cannot be known and is treated as a random variable with known probabilities. Casino games at low stakes exemplify this case: the rules are known, the odds are known, the outcomes are all monetary, and at low enough stakes the utility of the outcomes is nearly linear with the dollar

amounts, etc., but the actual result cannot be known ahead of time for the gambler. The ideal “rational” solution is to craft an option that assures that the best possible outcome will obtain with certainty.

Gambling uncertainty arises most directly from statistical regularity in the natural world. Decision-makers face fairly similar risks of possibly great uncertainty in the consequence of selecting a particular option when dealing with conflicting data, incomplete or conflicting mental models, uncertainty about the feasibility of options or their authority to implement particular options if they are chosen, or uncertainty about the consequence function. However, these factors can rarely be apprehended without significant fuzziness, and are therefore treated in the section below on Fuzzy Gambling.

There are basically two categories of actions to take when facing gambles. The first is to attempt to reduce or better yet eliminate the risk or the gambling aspect of the options. If that is not successful, then you cannot avoid gambling.

V.5.1. Reduce gambling aspect of options (de-gamble)

Several types of approaches are available to reducing the gambling aspect of options. One way is to improve the options; another is to reduce risks; a third is to improve response capabilities; and another is to aggregate decisions to balance or combine risks.

V.5.1.1. Improve options

While all of the approaches for mitigating a gamble could be described as improving your options, some ways to do that are more clearly specifically intended to improve individual options. Perhaps two or more

options that each have attractive features could be combined into a new option superior to any of its ancestors. Other routes are to attempt to increase the likelihood of desired outcomes for an option, or to increase the desirability of a good outcome. These do not necessarily make an option less of a gamble, but they can make it a more attractive gamble.

V.5.1.2. Mitigate risk

Two major techniques for mitigating specific risks are to reduce the likelihood of the risk, and to reduce the magnitude of the harm if the risk should occur. Slowing a risky system down may make it safer, or at least provide more time for dealing with problems that arise. In general, the specific risk dictates much of what can be done to reduce its likelihood or magnitude.

V.5.1.2.1. Reduce likelihood of harm

At its extreme, this is the *fail-safe* approach: attempt to prevent failure altogether. Frequently there are steps that can reduce but not eliminate the likelihood that a specific unwanted outcome will occur. When systems or situations have great risks that cannot be reduced, avoiding them entirely may be worthwhile. Systems that have many components which are both critical and susceptible to failure are more vulnerable than systems whose success can be obtained through more than one route.

V.5.1.2.2. Reduce magnitude of harms

Quite different from the previous strategy, this *safe-fail* approach intends to make a failure or bad outcome as tolerable as possible. Decoupling a systems or compartmentalizing it may prevent harms that do occur from propagating to other parts of the system.

V.5.1.3. Improve response capability

Instead of mitigating a risk, increasing your ability to respond effectively when hazards come about can help to neutralize their impacts, as suggested by the following categories.

V.5.1.3.1. Contingency planning

Planning for specific contingencies makes it possible to have measures in place that can be implemented rapidly when needed. This requires developing lists of possible states of the world or events that are worth responding to, and devising in advance responses to be taken. War games and crisis management are variants of this approach. Crisis management may be more general than contingency planning, in that the type of crisis may not be identified, only that one exists and that generic changes in operations will be made.

V.5.1.3.2. Plan for correcting course later

A more abstract or technical practice is to be sure that your plans for the future include options and opportunities for changing course as the world changes or as you receive new information about the world. This calls for designing projects that can be implemented in stages, or even reversed if conditions change sufficiently, and for preserving options.

V.5.1.3.3. Develop early warning systems, get news faster

A key to responding effectively is to obtain the information necessary for response as quickly and accurately as possible. Advance notice of events for which contingency plans have been prepared provides more time for implementing the plans. It is useful for individuals or organizations to develop information acquisition systems that are especially sensitive to the

kinds of information currently most important in distinguishing the best of the possible actions from the rest.

V.5.1.3.4. Flexibility and speed of response

Formulating contingency plans is a way to prepare for particular possibilities; improving your flexibility and ability to respond quickly prepares for dealing with sudden changes of any sort. This is especially useful when surprises are likely, or when there are very many different possible events to prepare for. Turbulent environments or ones about which little is known, or when diverse actors can each exert significant, different influence on the world may be so tumultuous that improvisation is more likely to be helpful than planning or analyzing, and it may be possible to plan ahead to make improvisation easier and more likely to succeed.

V.5.1.3.5. Compensation for harms

When harms are not prevented, it is sometimes possible to arrange compensation for them. Purchasing insurance is a widely endorsed way to do this; suing another party for damages is another version of this that is widely thought to be overused in America today.

V.5.1.4. Aggregation of decisions

The net effect of a combination of decisions may be very different from a simple sum of their effects.

V.5.1.4.1. Combining unique decisions into series

When one-shot risky decisions can be aggregated properly, the risk can be reduced or even eliminated and replaced by a statistical certainty. Casinos and insurance companies enter into so many similar risky

transactions with large numbers of customers that the risks are aggregated very effectively, and cease to be very risky. Individuals can form groups to pool their risks and reduce individual exposure.

V.5.1.4.2. Hedging

Hedging is essentially the aggregation of disparate risks into a package in such a way that the risks cancel each other out. The resulting package may simply not have the risks that the components did. Events could unfold in the same way, but the consequences for the decision-maker are very different.

V.5.2. Choosing among risky options — smart gambling

If it hasn't been possible to devise an option that preserves benefits and eliminates risk, the decision-maker must gamble. Even choosing an option whose outcome is known constitutes a kind of gamble if means declining other options that have a possibility of superior outcomes; the decision-maker will have foregone this possibility, and therefore incurs the risk of having made a choice that he or others will regret. This may or may not matter much to a particular decision-maker.

Given that gambling is not avoidable, alert gambling is better than blind gambling.

V.5.2.1. Elimination of dominated options

A first screening step is to rule out options that are dominated by others. Option A dominates (nearly dominates) Option B if Option A is at least as good as (not significantly worse than) Option B under all possible future states of the world, and is better than Option B under at least one. This analysis may be very useful in pruning a very large set of options

down to a smaller set, but is less likely to single out a unique dominating option.

V.5.2.2. Consider all relevant dimensions of difference across options

Determine whether hope level or security level is important. Establish whether unresolved uncertainty will have any cost (positive or negative) through time (for instance, by preventing further planning, or by providing a sense of excitement for a horse player). Consider whether psychological or strategic rejoicing or regret (section III.2.9.4) would be important potential consequences. (Experience from previous decisions by the same decision-maker may illuminate the propensity for psychological regret/rejoicing. In the case of potential strategic regret, it may be possible to estimate which parties may have opportunities for second-guessing that would matter.) Statistical measures of risk may be useful indicators. Whether some of the options contain more potential for mid-course corrections may be important. Consider not only the first-order effects, the immediate utility of future consequences, but also how well positioned you would be for the next decision — the value of options that may or may not be preserved by all alternative choices.

V.5.2.3. Decision rules and how to decide

Von Neumann and Morgenstern's Maximization of Expected Utility is a decision rule that is widely held to identify the "most rational" choice to make, that is, the most rational gamble to take. Von Neumann and Morgenstern defined several limitations on its applicability which rule it out when the decision-maker's information (including information about the decision-maker's values) is too fuzzy, when objective probabilities over

future states of the world are not known, or when the decision-maker has a preference for or against risk or certainty. Additionally, preferences for high hope levels or high security levels, or significant costs or benefits arising from the persistence of unresolved uncertainty make MEU difficult or inappropriate to apply. These and other concerns are discussed in more detail below in the Maximization of Expected Utility section.

Modeling can help illuminate the gamble that the decision-maker faces. Monte Carlo modeling, for example, can simulate hundreds or even hundreds of thousands of repetitions of the decision at hand (assuming that there is only state-of-the-world uncertainty, not function uncertainty). While the output most typically observed is averages over the entire run, which is what MEU needs, maxima and minima can as easily be observed, and can be used to highlight options or states of the world or combinations where option improvement could be promising.

Alternative decision rules include maximin, recommended when security level is the most important characteristic of an option; maximax, appropriate for maximizing hope level; and minimax regret, when the decision-maker or others will be second-guessing the decision. Except when the narrow conditions for adopting MEU apply, there are enough subjective considerations in selecting a decision rule that a single answer cannot be provided. When it is very difficult to determine what the right way to decide is, it may be that there is no single way that is right; there may be several that are each not-wrong or even right. Elster's conclusion is that under genuine uncertainty, rational choice theory can say little about what to do (Elster 1986:6). What reason can help us to do is to consider as

many of the relevant facts and values as possible before relying on our own fallible judgment in choosing a decision rule or choosing an option.

V.6. Advice for problems of fuzzy gambling

Few decisions are beset by only fuzzy uncertainty or only by gambling uncertainty; most suffer from both. Fuzziness about any of the parameters of a gamble turn it into a fuzzy gamble. A fuzzy decision in which the future state of the world is not known with certainty, whether due to variability in the natural world, un-made decisions of actors, or lack of information, is a fuzzy gamble. All of the suggestions for fuzziness and for gambling are potentially helpful for fuzzy gambles. Additional suggestions are given below.

Two strategies that may improve the decision-maker's outlook in nearly any fuzzy gamble are to gain more information and to improve the set of options available to her. Additional information may reduce fuzziness and thereby make some comparisons among options clearer, and it may also facilitate improving the options. Devising better options and reducing risks can not only improve her prospects, but may also make comparison and choosing among options easier. The significance of any fuzziness on the decision(s) at hand should be kept in mind. In a great many cases, neither the fuzziness nor the risk can be eliminated, and the decision-maker must ultimately gamble, perhaps without knowing clearly what the odds or payoffs really are.

In the case of many fuzzy gambles, especially those where the actions of other agents are or will be important, fuzziness and gambling are mixed closely together, and the twin solutions of reducing fuzziness and

improving the gamble are also mixed. For instance, when the feasibility of an option is uncertain, it may be worthwhile both to attempt to clarify the feasibility (reduce a fuzzy uncertainty) and also to attempt to increase the feasibility (improve the gamble).

V.6.1. Factors nearly always causing fuzzy gambling

Fuzziness about any aspect of a gamble renders it a fuzzy gamble. Fuzziness about how to value consequences is very common, blurring evaluations made by a single party, multiple parties, perhaps over multiple objectives, or across different time periods. The determination of values is outside the domain usually accorded rationality, but the implications for rational decision of imprecise evaluations are important. Estimations of probabilities are nearly always fuzzy, and in the case of subjective probabilities, always fuzzy. Low-probability, high-impact risks virtually always involve fuzziness about both the probabilities (due to infrequent data) and valuations of the consequences (which are outside the range of typical consequence evaluations).

Situations of conflicting data and conflicting or known-to-be incomplete mental models usually possess both fuzziness and a gambling character. Surprise-prone environments generally cause both the risk of bad outcomes and fuzziness about the likelihoods of the outcomes and perhaps additional fuzziness. Uncertainty over the feasibility of options or the authority of the decision-maker to choose one of the options induces fuzzy gambling. Much of the uncertainty over the future actions of other actors, such as their possible impact on the state of the world, the consequence function, or the possibility of error or deception produces both fuzziness

and gambling.

V.6.2. Improve options

Uncertainty about the feasibility of options or the decision-maker's authority to put them in place induces fuzzy gambles. The ideal situation for the decision-maker is to have complete authority, for each of her options to be completely feasible, and for her to know this. In any other situation, it would generally be good to move closer to the ideal, both by clarifying the extent of her authority and the degree of feasibility of her options, and by increasing her authority and the feasibility of her options.

V.6.2.1. Feasibility of options

Feasibility can be clarified by obtaining more information about potential constraints and barriers, including constraints on political feasibility, not just technical feasibility. The same information that helps to clarify feasibility can be used to begin efforts to increase feasibility. Such efforts could include describing options more attractively, making more powerful arguments for them, or redesigning them. Bargaining and consensus-building can also increase their feasibility.

V.6.2.2. Authority of decision-maker

Like option feasibility, the decision-maker's authority may be clarifiable through investigation of constraints or specific grants of authority. Experimentation may be useful: the decision-maker can act as if he has the authority to make a decision, and observe the results (contingency planning may be a useful adjunct!). This may have the effect of increasing the decision-maker's authority, if his actions are seen as successful, or could result in restrictions on authority if others become

sufficiently upset.

Authority can be increased by cashing in favors, acquiring more resources (personnel, legal, financial, physical, military, etc.), increasing or capitalizing on charisma, and so on. Seeking to define the terms of a debate may make a particular position appear to be more obviously the right one, and in extreme cases can even deny others the terminological ability to state their cases.

V.6.3. Mitigate risks

Reducing the likelihood or magnitude of possible harms is complicated in fuzzy gambles because the fuzziness can make it difficult to assess the severity of a risk or compare it to other risks.

Systems in which people play key roles, such as in the operation of nuclear power plants, can be redesigned to protect against non-technical sources of danger that may be less obvious to technically-oriented engineers than to industrial psychologists or consumer advocates. Better training or clearer, more complete manuals can reduce the chance of people causing accidents. Organizational efforts to facilitate problem-spotting efforts, or encourage whistle-blowers, can help to highlight potential problems before they become more serious.

Many of the rules, laws, and social norms that people adopt have the effect of reducing the uncertainty surrounding the behavior of other actors, though their primary ostensible purpose is to reduce the likelihood of undesired behavior. Nonetheless, proliferating social sanctions have the effect of reducing individual's freedom of action. Parents may find that the sheer variety of behavior of their kids is a problem, aside from whether

any specific act is undesirable, and attempt to directly reduce the uncertainty of the kids' behavior, which may not be seen as legitimate by their kids.

V.6.4. Prediction — state-of-the-world and function uncertainty

Uncertainty over the future state of the world and over the function that maps states of the world and options into consequences causes problems of prediction. When the consequence function is uncertain, the utility of mathematical modeling is sharply limited. Uncertainty about the behavior of actors very commonly produces this effect. When the state of the world cannot be predicted but the consequence function is reasonably well known, modeling can be important, though estimating the reduction in the certainty that ought to be ascribed to the model output that is caused by uncertainty over the probabilities used as inputs is often overlooked, and should not be. The fuzziness attending the probabilities, typically quite high when logical or frequency-based probabilities are not available, should be propagated through any analysis to see its effects on the certainty of final outputs.

Except for the case mentioned above when the consequence function is well known and the state of the world can be treated as one or more random variables, uncertainty about the state of the world or the consequence function pose the same problems for the decision-maker, and available remedies are substantially the same also, primarily acquiring more knowledge about the world, acting in order to make it more predictable, or improving options and response capabilities.

V.6.4.1. Uncertainty due to actors

Uncertainty about what actors will do can be reduced by information gathering, a relatively passive approach, or more active measures designed to restrict either the potential or likely variability in behavior. Your own behavior may be treated this way as well, and you may take steps, as Ulysses did when instructing his crew to tie him up before passing the Sirens (Elster 1984), to prevent yourself from taking an action which you now do not want but might find hard to resist in other circumstances. (This is not restricted to situations of uncertainty; Ulysses would have done the same thing whether he had been uncertain about his ability to resist the Sirens or had been completely certain of his inability to do so.)

Information gathering approaches include obtaining information about specific actors, on human nature in general, on specific organizations, on organizations in general, inquiring directly of actors or organizations what their plans are, and so on. Considering the decisions the uncertainties and contingencies they face can help understand what they are likely to do. Game Theory can be very effective in this area when the rules governing behavior and results are clear and the number of contingencies is relatively small.

Not all uncertainty arising from the relevance of actors comes from choices yet to be made; over longer periods of time, individuals and organizations may undergo fundamental changes. Some of these can be foreseen to some extent, such as maturational changes in children or engineering-focuses start-up enterprises, and understanding the nature of the people and organizations can help predict them. In other cases, only

the possibility of major change is foreseeable, and rather than taking steps based on particular possibilities it makes more sense to prepare for the generic possibility of change, e.g. by avoiding actions overly dependent on current assumptions and continuance of the status quo, increasing efforts to detect changes, and increasing response capabilities.

More active steps to reduce the uncertainty from others' actions include forming associations with others that adopt rules, such as laws or social norms, entering into contracts with others (implying that the reduction in your own freedom is more than compensated for by your increased certainty in the behavior of others). Various types of inducement can be tried, ranging from suggestion and persuasion, through marketing and lobbying on to use of threats and or employing physical or military force.

V.6.4.2. Awareness of possible error or possible deception

The result of awareness of possible error or deception is reduced confidence in the evidence at hand. Unlike fuzziness, which is largely a matter of *degree* (although in extreme cases, less so), reduced *confidence* in the evidence often induces an uncertainty of *kind*, and a fuzzy gamble is created, since the state of the world if the evidence is valid may be qualitatively different than if it is not valid. Once these possibilities have been raised, it is helpful to both gain information about how likely error or deception is, and also to reduce that likelihood.

Understanding more of the history of error or deception of the individuals or organizations involved, or their current motivations, may help. Seeking corroboration of the evidence from other sources, or seeking indications of error or deception, such as internal inconsistencies in the

evidence or non-verbal cues, can help in this assessment. Someone familiar with the frequent comments in the literature on the importance of the recency and accuracy of major assumptions will be aware of this possible source of error, and may undertake assumptional analysis (Mitroff and Linstone 1993).

Methods for reducing these possibilities include improving skills, redesigning process, and changing the contingencies others face. Skill-building can improve the abilities of individuals or organizations to perform tasks without (or with few) errors and for communicating their information unambiguously. Process redesign may be able to eliminate steps that are especially susceptible to error, to reduce the ability of one party to commit undetected falsehoods, or to build in error-checking. Many ways exist to change the contingencies others face, such as bonuses for low-error work or for identifying errors, emphasizing moral aspects of truth-telling, penalties for perjury, et cetera.

V.6.4.3. Conflicting data

Another threat to the confidence that can be placed in evidence is a conflict between different pieces of evidence. An ideal solution is to acquire sufficient information to become confident of the evidence, by determining the “truth”, or determining which source of data is the most authoritative. Sometimes more than one source is correct, perhaps about a different facet of the system under consideration, and a new, more encompassing view is needed. In some cases, compromises are possible that partially accommodate all data. When it isn’t possible within the limits of time and resources to resolve the conflict, the extent of the discrepancy represents a

minimum level of uncertainty about that factor (though the uncertainty may be substantially greater than that one discrepancy, perhaps as indicated by previous discrepancies).

V.6.4.4. Conflicting, incomplete or improvable models

Models — mental models, computer models, metaphors, archetypes, even folk sayings — are a fundamental way for us to codify our understandings about how the world works. When are uncertain about the future state of the world or about what consequences will follow from our choices, improving our models is a broadly applicable strategy that often works well over time. It works because our world is regular in many ways, and models that reflect those regularities enable us to predict future developments. Aside from limitations on our own abilities, the primary limits that we have encountered so far are processes that are or appear to us to be random, cases of deterministic chaos, and the actions of other actors and organizations. Science and engineering have proceeded on the assumption that the world is highly modelable, and in the hard sciences this approach has worked well.

While we often acquire information by chance and incorporate it into models without having planned to do so, explicitly attempting to construct more powerful hypotheses and models, and to seek the specific data needed to guide refinement of our models is usually faster. Models from other fields or arenas may be useful if borrowed. When our understanding of a system is only rudimentary, structural modeling can help to test and clarify what we think we know about major relationships within a system. Building multiple models and comparing them not simply for accuracy but

for the particular types of accuracy or error each exhibits can accelerate learning compared to refinement of a single model.

Sometimes we might have reason to suspect that there is regularity that we have not yet identified, and attempting to model it may help. Other models are constructed in the hope that the world is regular and not merely random. Statistical tests on observations may indicate whether or not the data are likely to have been produced by a stochastic process, or by a chaotic one, pointing to particular ways to look for regularity. Incorporating wider varieties of data or devising more detailed outputs may increase the predictive utility of a model, reducing uncertainty about future states. Most models are significant simplifications of the systems they represent, however, and there is always the possibility or even likelihood that the world contains sources of variation and surprise not depicted in the model.

Resolving uncertainty caused by conflicting results from multiple models can be more difficult than uncertainty from conflicting data. Models usually contain numerous assumptions which are first, difficult to identify, and second, very difficult to reconcile with the assumptions on which alternative models are based. It may be possible to find evidence that refutes some of the hypotheses or models. Each model may be capturing one aspect of the system, which implies at least two possible resolutions. One is to determine the conditions and uses for which each model is most suited. Another is to construct a meta-model in which each original model is a special case or subset.

When the discrepancies cannot be resolved at the level of the models,

accommodations outside the operation of the models can be made. The results can be incorporated into a single judgmental analysis. If two or more models each appear to have some validity, the decision-maker may have to proceed as if any of the results predicted by any of the models could occur, which may imply a large degree of uncertainty. Options can be sought that appear to be attractive under any of the futures predicted by the models. Ultimately, if two models, or any two alternative implementations of rationality, conflict irreconcilably, rationality is indecisive, and non-rational means for making a choice are required.

V.6.4.5. Probabilities

Probabilities nearly always contain the implication of fuzzy gambles — they usually apply to sets of states of the world, implying gambles, and they are hardly ever non-fuzzy. It is important to assess their fuzziness, since this fuzziness may well have more impact on the fuzziness of an expected value than fuzziness in other variables. If objective probabilities are not available, subjective probabilities can be useful substitutes, but subjective probabilities will virtually always have high fuzziness, and the impact of that fuzziness on the certainty of any calculations should be displayed, and in many cases may make numerical comparisons so fuzzy as to be indecisive. Seeking more information about the current conditions, the future state of the world, or news of trends may reduce this fuzziness.

V.6.4.6. Complexity

High complexity can make prediction of the future state of the world, or of the consequences of particular choices, very difficult. One approach that has served physical science very well is hierarchical decomposition, or

reductionism: divide complex systems into smaller systems that are individually easier to analyze. This works when the dynamics *within* each subsystem are quite important relative to the dynamics *between* subsystems. Another is to gather more data and develop more complex models and theories that are able to take into account a larger set of relevant factors.

However, as Zadeh's Principle of Incompatibility asserts (section II.3.10), complexity can make simultaneous precision and relevance very difficult to obtain. Two of the approaches just mentioned are to decrease the complexity of the model (reductionism), or to tackle this difficulty head-on via more complex models. Sacrificing precision may be quite helpful, by aggregating data into less-precise units, or producing or displaying output with less precision. Another approach would be to sacrifice relevance. Though this violates the RAG and would probably not be helpful to someone facing a decision right now, it may make possible models that can improve our understanding of complex systems and enable us to gain insights that will be valuable in the future. It may be easier to emphasize relevance instead of precision by increasing the use of judgmental analyses, either by reducing the role of quantitative analysis or subordinating it to final judgmental interpretation.

V.6.4.7. Fundamental dynamism

Fundamental dynamism implies the possibility of genuine novelty in the world you must deal with. Technological innovation and human creativity are two of the main factors that make this difficult to rule out in situation in which human beings play important roles. Planning for

particular contingencies is less helpful than improving generic response capabilities such as speed and flexibility, and keeping your eyes open to avoid major surprises.

One of the most serious dangers here is that buried assumptions will be made obsolete by changes in the world, leaving the decision-maker more vulnerable to serious surprises; techniques that are not as dependent on assumptions or where the assumptions are more visible carry less risk of this. Encouragement of efforts to challenge assumptions can reduce the danger of assumption ossification.

V.6.5. Surprise-proneness

Surprise-proneness can be dealt with in some generic ways that do not depend on the cause of the surprises, such as gaining more experience with the environment or simulations of it to learn to be less shaken by surprises, increasing response capability, and using simpler, faster procedures, perhaps ones that are less quantitative, or have less rigorous standards of what constitutes a good enough reason to act.

Generating lists of surprises and preparing canned responses makes sense when the environment is regular enough and well-enough known that a set of likely surprises can be developed. Rehearsing the responses or rehearsing responses to surprise conditions may help, especially in reducing the debilitating effect of being surprised on people's ability to function effectively.

When an actor's environment is surprise-prone because not much is known about it, learning about it quickly is clearly promising. Gathering information passively by simple observation or more actively by

experimentation, especially parallel experimentation, keeping track of the information, and making effective use of it can all pay off. Better models and more powerful modes of representation may help, especially where complexity complicates understanding what may happen.

It may be possible to decrease the frequency, speed or variety of surprises by reducing information flows, buffering perturbations to localize their effects, or increasing safety margins and shock absorbers.

When fundamental changes appear likely it is important to encourage creativity, maintain a diversity of skills and perspectives, either within individuals or in specialized parts of organizations. Facilitating open and frank discussions can facilitate fast learning, but may hamper efforts to respond quickly; the time scales of preparatory actions need to be matched to the speed of the possible environmental changes.

Human interactions can cause surprise-proneness, and active methods for reducing it can come from nearly any branch of sociology, psychology, politics, or negotiation techniques and interpersonal skill methods like the Dale Carnegie School's.

V.6.6. Low-probability, high-impact

These cases need special treatment, as the existence of the distinct field of risk analysis attests. Nearly all cases of high-impact risks involve not only the possibility of loss or harm but also considerable fuzziness about the importance of the harm relative to possible preparations and relative to other possible harms, and cases of low-probability risks typically involve fuzziness about likelihood. Experts may be better able than laypeople to estimate likelihoods, though their judgments are not immune to

systematic errors, but evaluating the utility of the consequence is a personal, social or political task that should not be reserved for experts. The importance of the evaluation of consequences, often underemphasized, reduces the potential for settling such issues solely via technical or calculating methods.

V.6.7. Modify values

Deliberate attempts to modify the values of those who evaluate consequences may be undertaken simply to reduce variations in evaluations or to effect major shifts in the evaluations. Efforts to build consensus could be viewed as ways to help diverse evaluations come closer together without attempting to change the center of gravity of the whole set of evaluations. Marketing efforts usually are trying to change the center of gravity of evaluations, not reduce uncertainty in them.

Modifying either your own or others' values can improve a decision prospect. If you can reframe risk so that the consequence is no longer important, it ceases to be a risk. For example, a firm decision to evaluate your decisions only from an *ex ante* perspective and not to second-guess yourself would eliminate the danger of serious regret; persuading yourself that any rain that fell on you would dry soon without serious harm and could even be fun might eliminate the weather risk of taking a stroll in cloudy weather.

V.6.8. Choosing among fuzzy gambles

Often facing fuzzy gambles is so frustrating that there is a strong temptation to overlook important factors because they complicate rather than simplify an already burdensome decision. It can be helpful to make

simplifying assumptions, but these always have costs, at least in terms of obscuring some information, and they constitute a liability in that circumstances may change so that the same simplifying assumptions, now forgotten, would not be made anew, but still operate outside of awareness. The more slowly the person's or organization's environment, means, and ends change, the greater the power and smaller the drawbacks of simplifying assumptions.

Which simplifications offer the most help and the fewest problems varies with the values and circumstances of the decision-maker. For example, when an uncertainty will remain unresolved for a long time, experiences during the pre-outcome period are more important, and adopting maximin to maximize security level when peace of mind is paramount, or adopting maximax when the hope is more important may help.

Any of the recommendations given here could turn out to be the most relevant under particular circumstances, but this work, being general, is blind to the specific circumstances of any one decision. Judgment, informed by those specifics, must be used to guide adoption of particular remedies.

V.6.8.1. Assess gambles as accurately as possible

When avenues for improving the decision-maker's prospects have been explored, the remaining gambles should be assessed honestly, and uncertainties in those assessments noted. Use the Null Hypothesis to assess whether any apparent differences in desirability are significant.

V.6.8.2. In the end, you must gamble

A fuzzy gamble remains a gamble, and ultimately the decision-maker

must gamble. He doesn't have to like it; our frequent desire for more certainty is natural, but often doomed to frustration. I would count it as progress if we learn to recognize when hope for greater certainty is legitimate and when it is illusory.

V.7. Bounded, meta- and hyper-rationality advice

V.7.1. Bounded rationality, limitations of the decision-maker

It is sometimes possible to reduce limitations imposed by characteristics of the decision-maker sufficiently to have an impact on the uncertainty of conclusions. Reducing these limitations will have costs, though, which need to be compared in some fashion with an estimate of the potential benefits.

V.7.1.1. Improve processing capability

Better processing capability may let the decision-maker reduce uncertainty, or obtain the same level of certainty when dealing with more complex issues, or respond faster with the same confidence. Infinitely many steps can facilitate this, including codifying knowledge as decision rules for faster use later; better computing power, more powerful means of representation or mental models, clearer understanding of fundamental values, firmer agreement on the processes to be used in making an organizational decision, etc. Memory limitations may be ameliorated by augmenting psychological or institutional skills, better storage and retrieval methods.

V.7.1.2. Time and resource limits

When time or resource constraints pose a problem, and neither more

time nor more resources can be obtained, redesigning the processes or adopting different processes may help. Processes that are scalable are easier to adjust to the resource availabilities of each situation; some can be broken into modules performed in parallel. Quantitative processes, especially those relying on large computer models, are often very difficult to accelerate; decisions about what trade-offs to make, for example between speed and thoroughness, or speed and accuracy, may be easier to make in processes relying more on judgment.

V.7.2. Considerations of time

The amount of time available may limit the variety of approaches available, but it is difficult to say a lot about specific limitations without knowing the details of specific situations. Some illustrations of the effects time constraints can impose are given next.

V.7.2.1. Longer term approaches

When a lot of time is available, any of the techniques mentioned here may be feasible, such as “growing” better information systems and early warning systems to produce better streams of information for future decisions (Stinchcombe 1990). This can include improving acquisition, processing, storage, or retrieval stages of the information processes. Long-term plans can be improved by designing them to take best advantage of information that is likely to become available in the future. Attempt to avoid situations attended by problematic uncertainty, or change the world to make it less problematic.

V.7.2.2. Some time is available

In the short run, adjust decision-making techniques to accommodate

uncertainty that it is too late to reduce. Clarifying uncertainty is more likely to be feasible than modifying it; gathering more information more likely to be manageable than developing better mental models. Rehearsal of responses, mental preparation for being surprised, and buying insurance may be feasible.

V.7.2.3. No time available

When there is virtually no time available, all a decision-maker can do is to gamble, perhaps choosing the gamble by a non-rational decision process.

V.7.2.4. One-shot vs. correctable decisions

When a decision cannot be corrected or adjusted later, avoiding bad outcomes is more important. This may indicate use of satisficing instead of optimizing, and pursuit of robustness instead of optimality.

V.7.2.5. Unique vs. repeated decisions

The methods, even the notion of rationality, appropriate for repeated decisions may not necessarily be the appropriate for unique decisions. Hope level, regret, and security level are likely to be more important in unique decisions.

V.7.2.6. Delay decision until uncertainty is resolved

One valid way of dealing with uncertainty is to delay action until more information becomes available, though this may have unacceptably high costs.

V.7.2.7. World changes fast compared to us

When the rate of change in the environment is high relative to our own ability to acquire and assimilate information and make decisions, processes of analysis and decision that would otherwise be attractive may

not be workable. Assumption ossification is a bigger danger. Precision may be an luxury. More decisions may have to be made before the decision-maker is able to become comfortable with her assessment of the options. Having a variety of models or even solutions on the shelf would be convenient. More effort may have to be expended to aggressively seek information instead of waiting for it to appear.

V.7.3. Meta-rationality advice

V.7.3.1. Focus on the big picture

Uncertainty is not necessarily a problem. When it is undesirable, it is usually as a complication to pursuing a good, rather than an ill in itself, by complicating a selection task, making improvement of options more difficult, or because a risk of relatively bad outcomes is present.

V.7.3.2. Allocate efforts sensibly

Don't work harder at a possible treatment than the task can be worth, e.g. if a reduction in one uncertainty will not improve your general prospects enough to be worthwhile (perhaps because other uncertainties remain large).

V.7.3.3. False certainty

This work focuses on what do to about being uncertain, but a danger clearly implied by this framework is that there are times when you *should* be uncertain but are not. There appear to be few easy ways to detect times when you don't really know what you believe you know. Being on guard for this possibility is probably the best single protection. From this follows an interest in making sure that uncertainty is not hidden; that major assumptions are highlighted and periodically examined, such as

assumptions that systems are in equilibrium, actors have perfect information, and human beings will operate systems the way the system's designers want; that more readily quantified factors are not emphasized at the expense of less quantifiable factors, such as esthetics, equity, and other value issues, no matter how attractive quantitative frameworks are; etc.

V.7.3.4. When uncertainty is high

High levels of uncertainty decrease the likelihood that analytical methods and quantitative methods, such as optimization, and MEU, will be decisive. In times of great uncertainty, decision-makers may have to fundamentally change the processes through which they define and make decisions. Pursuit of efficiency, for example, often attempts to reduce redundancy, duplication, overlap, and diversity, though these are qualities that Thompson and Warburton maintain are essential for avoiding fragility under high uncertainty (1985:31). Decision-makers may need to emphasize non-quantitative tools such as scenarios, sensitivity testing without probabilities, listing surprise possibilities, favoring resilience over risk reduction, seeking robustness instead of optimality, seeking bad-avoidance as well as good-achievement, using less-detailed tools with more highly aggregated data, or greater reliance on qualitative analysis. When uncertainty over facts is so large that it completely swamps technical views of the problem, there may be political or institutional facts that are much easier to ascertain, and these can be useful grounds for decision-making (Thompson and Warburton 1985:22).

V.7.4. Hardness and softness of uncertainties

The hardness of an uncertainty (section III.3.1) does not determine how

to respond to it by itself, but there are some generalizations about response categories that can be made. (Actions that can improve the options available to the decision-maker, as described in the above section on Gambling are worth remembering.)

V.7.4.1. Soft uncertainty

Soft uncertainties can be reduced or eliminated; the major questions are

1. How much reduction in a given uncertainty would be required to make a significant difference, considering any other uncertainties present?
2. What ways of reducing or modifying the uncertainty are available?
3. What level of commitment of time and resources is worth making, given the potential for improving the decision-maker's prospects?

V.7.4.2. Hard uncertainty

An uncertainty known to be irreducible must be accepted and coped with. It represents a genuine limit to how much we can know. Better options, better response capabilities, and better means of compensating for harms may help. Analytical processes or decision processes that depend on resolution of this uncertainty should be avoided.

V.7.4.3. Virtually hard uncertainty

Uncertainties that are essentially hard at this time have to be treated as if they are hard. There is a possibility that they will become reducible later due to changes in circumstances, such as new technology, more experience with a process, or better understanding of values and priorities, better theories or mental models, faster computers, new regulations, shifts

in political arrangements, or transformations in the dynamics of relevant systems. Thus, virtually hard uncertainties that are important to future or on-going decisions should be monitored for developments that render them reducible.

V.7.4.4. Malleable uncertainty

Some uncertainties present possibilities for changing the world to improve the decision-maker's outlook. These should not be overlooked by focusing solely on more passive approaches of information acquisition.

V.7.5. Hyperrationality

Be alert for signs indicating the limits of rationality, such as uncertainties known to be hard, or certainty about options that are tied for best.

If a rational decision selection process is indecisive, improve it, adopt a different rational decision selection process, or recognize that rationality is mute and employ a non-rational decision selection process.

Don't attempt to optimize, for instance, when you don't have enough information to do it. Be prepared for the possibility that rationality in all its forms will fail to be decisive, and that your decision will have to be made on other grounds.

V.7.5.1. Employ a Null Hypothesis

Test the certainty of the conclusions of any rational process when uncertainties are important. One way to do this is to borrow the Null Hypothesis from inductive statistics, as described above in section IV.1.2. The point of that approach is to avoid presenting as a rational conclusion a result that could have been arrived at through chance or methodological

artifact, and therefore ought to be attributed to the fundamental facts of the situation.

V.7.6. Other considerations

When objective probabilities are not available, subjective probabilities can always be derived, but their fuzziness is very important to assess. Alternatively, methods that do not require probabilities could be used, such as maximin.

Decisions that are sensitive to the evaluation of the importance of consequences cannot be decided solely by technical means; personal or political processes for evaluation are essential.

When alternative actions are developed serially and cannot be assessed simultaneously, satisficing may be more efficient than optimizing.

V.7.6.1. Utility value of consequence

In many cases it is important to estimate the utility of consequences and use that value rather than the simple magnitude of consequences. This is especially so when the actual magnitudes are obscured in mathematical calculations or models, such as MEU methods, or when some of the magnitudes of consequences are extreme.

V.7.7. Clarify extent of uncertainty for informed decision-making

All of the above are easier if there is honest appraisal of the relevant uncertainties, unless a shaky sense of confidence is at risk or there are extreme pressures of time.

Display uncertainty in assessments, especially for hidden assumptions and when probabilities have been estimated subjectively. An inescapable need for judgment becomes obvious here concerning how much effort

highlighting uncertainties ought to receive. Too much attention to uncertainties can make it appear that we can know nothing of importance with any surety at all, and we may become demoralized and paralyzed. If too little attention is given to the limits of our current knowledge, we may act with foolish blindness to real dangers. These two risks must be balanced against each other by a hand sensitive to the nuances of the specific context.

V.8. Schematic diagram (overview guide)

V.8.1. *Metarational considerations*

Figure 44 below shows schematically the meta-rational considerations described above. The apparently sequential nature of these processes is an artifact of putting them into a diagram; many iterations may occur, stage may occur before other stages shown later in this picture.

1. *Decide Rationally?* There is no reason why decision-makers must decide rationally (invoking *reason* presupposes rationality, but reasons only of paramount importance after adopting rationality, not when considering its adoption). Even for those decision-makers who would like to decide rationally, there are situations when the time or resources required by a rational decision selection process are too high compared to the added value rationality could provide.

2. *Choose a Non-rational Decision Selection Process.* When rationality is either not desired or not worthwhile, other bases for making a selection need to be adopted (e.g. coin flipping, “log-rolling”).

3. *Well-formed Decision?* Rational choice requires a reasonably well-formed decision, meaning that a choice will be possible, the decision-

Figure 44. Meta-rational guide to rational choice

maker(s) has authority to make a decision, there are some options that are feasible and available, there are ways to value and rank the possible consequences, about which something must be known.

4. & 5. Improve Well-formedness. If the decision is not well formed, then it may be possible to improve the formulation by clarifying whether a choice will be possible, assessing option feasibility or acting to increase feasibility, assessing or increasing the decision-maker's authority to act, etc.

6. Exit. If the decision cannot be made well-formed enough, then rational choice will not be very helpful. A decision may "emerge" from an ill-defined process, or better definition may become possible over time.

7. Adopt a Rational Decision Selection Process. Conclusions about which option should be taken on the basis of rationality can be obtained only through the exercise of particular processes, and many choices and assumptions are involved in choosing, designing and implementing decision selection processes. More than one process may be adopted at different times, for example, a quick screening may be performed first, followed by more detailed examination of only a few of the original options.

8. Is the Process Decisive? The rational decision selection process as implemented may fail to be decisive (section II.3.4), largely due to uncertainty, especially fuzziness. If is not the case that one option can be shown to be significantly superior, taking into account the uncertainties in the selection process, then the process is indecisive (Null Hypothesis, section IV.1.2). If there appears to be a reasonable chance of improving the process or adopting a new one in some way, it may be worth the effort,

though not always.

9. Attempt to Increase the Decisiveness of the Process. Many of the ways to improve the decisiveness of a selection process involve reducing uncertainty, either by gathering data about the world or stakeholders' values, or especially when the uncertainty stems from the relevance of other actors, by taking steps to change the world (e.g. by making deals). Improving options may also improve the decisiveness of the selection process, especially if two or more options present widely different trade-offs of multiple objectives or widely different risks. Such decisions could be made clearer by values clarification or utility curve elicitation, but improving the options would be more rewarding, if possible. Improving the decision-maker's processing capability may also help. Sections V.5 and V.6 and especially V.4 provide more details on this topic. Another schematic that illustrates treatments for uncertainties from different sources is given below in section V.8.2.

10. Hope of Better Options? If there is a reasonable chance of being able to find or create better options at a low enough cost, this possibility should be pursued before making a selection. If none of the options are considered to be good enough, even more effort should be made to find better options. At some point, however, shrinking time, resources, or the prospects of improving the options will indicate abandoning these efforts.

11. Improve Options. Multiple objective decisions and gambles are quite likely to be cases where none of the options is good enough due to unattractive trade-offs or risks associated with each option that the decision-maker does not want to take. Sections V.5 and V.6 discuss some

methods for improving options more fully.

12. *Exit.* If the result of the process is decisive, there is a single option that is shown to be significantly better than the others (the Null Hypothesis is rejected). If this option is good enough, it is a happy outcome; if not, the decision-maker can at least feel secure in having found the best of a bad lot.

If there is a failure of unicity, more than one option has been shown to be superior. Another selection process must be invoked to break the tie, but the decision-maker knows she cannot go too far wrong in choosing between two (or more) best options. Generally this result can occur only with very high quality information, good enough to be certain that two or more options are tied.

If there is an existence failure, then none of the options meets the criteria of the selection process. Another selection process must be invoked, but without the assurance in the case of unicity failure that the decision-maker cannot go too far wrong. However, it may be possible to show that some of the options are inferior, reducing the set of options that the next selection process needs to sort out.

V.8.2. Treatments of uncertainties from different sources

Figure 45 below illustrates some of key aspects of uncertainties arising from different sources. Its structure corresponds to that of figure 40, and summarizes material from chapter III. The methods for treating uncertainties described below might be used in step 9 of figure 44 above.

V.9. An irony

There is an irony here. The literature has focused on uncertainty

related to not knowing which of the possible states of the world will eventuate, and has also focused on the selection of options as a task made more difficult by this type of uncertainty. In this work, I have attempted to show that decision-making complications arise from a much wider array of

Figure 45. Treatments of uncertainties from different sources

types of uncertainty, and also that a much wider array of tasks than mere selection of options is worth considering as actions to be taken when facing uncertainty. The irony is that the primary problem caused by not knowing which state of the world will eventuate is that some states are usually much more desirable than others, and the best solution for that problem is not good selection, but the improvement of options for the purpose of coming out better.

On the other hand, the primary problem caused by the additional types of uncertainty described here is difficulty in selection, and tackling this directly by increasing your knowledge and improving the power of your selection process is often a very useful approach. The focus of much of the technical literature on selection is an emphasis more appropriate for the types of uncertainties ignored by that literature, and the best approaches to the problem highlighted in the literature is option design, which is ignored by that same literature.

VI. CONCLUSIONS

To know that we know what we know and that we do not know what we do not know, that is true knowledge.

— Confucius

In this chapter I will describe my primary conclusions, divided into three categories: background, taxonomy, and implications. I will then describe the contributions I believe this work makes, and its limitations. A review of the initial hypotheses follows. In the final section I discuss directions for extending this work.

VI.1. Background

VI.1.1. *Definition of uncertainty*

For this research, I have defined uncertainty to require an awareness by a decision-maker of the insufficiency of the available information (see section I.4). If enough information is available, and we know it, we are in a state of *true certainty*; if enough information is available, but we are not certain (for instance, because we fail to recognize the adequacy of the information), we are in a state of *false uncertainty*. Similarly, if we do not possess enough information for the purposes at hand, and we are aware of that, we are in *true uncertainty*; if we do not have enough information but are sure in spite of that, we are in a condition of *false certainty*.

I leave aside the problems of error (believing something to be true that is not, whether or not we have sufficient information), even though they can be serious, because under error we lack the *awareness* of the error that

would motivate us to seek tools or assistance. Error can preclude the possibility of making a rational decision, but it is not a form of uncertainty.

When used as a noun, an uncertainty is something about which we are uncertain.

VI.1.2. Decision selection processes can fail to be decisive in two ways

A process adopted for selecting one alternative from those that are available to a decision-maker is here called a *decision selection process* (see section II.3.4). Such a process can fail to be decisive — that is, it can fail to identify a single option — in two ways. The first type of failure is *non-uniqueness*: more than one option is identified. The second type is *non-existence*: no option is identified.

In both cases, the selection process is incomplete and must be augmented with another process, either a strengthened rational process or a non-rational process, to narrow the choice down to a single option. However, in the context of the decision-maker's prospects in general, the non-uniqueness failure is far less harmful, for the decision-maker has the assurance of believing that all of the identified options are unsurpassed. Even a coin toss could then guarantee the selection of an option that is as good as any possible. Non-existence failures, on the other hand, leave the decision-maker with the unavoidable risk of choosing an option that actually is inferior to another, as might be revealed by information obtained later, or when better information processing is available, etc.

VI.2. TAXONOMY

VI.2.1. Subjects of uncertainty

The literature on rational decision-making under uncertainty has

focused on uncertainty about what the future state of the world will be, generally assuming that the set of possible future states is known and typically assuming a probability distribution over those states (see the Standard Decision Model in section II.3.7). Many of the tools that have been developed for assisting uncertain decision-makers have been based on this model, which assumes that a great deal of information is available about factors in the decision other than which future state of the world will occur. When the decision-maker has good information about these factors, tools based on the SDM can be very powerful.

Often, however, a decision-maker is uncertain about many factors other than which state of the world will eventuate. A decision-maker may be uncertain about who has authority to make or block the decision, the membership and feasibility of the set of available options, which future states of the world are possible and how probable or possible each is, or what would happen if a particular option were chosen and a particular state of the world eventuated.

There are many issues of values about which a decision-maker can be uncertain. A decision-maker could be uncertain about whose values ought to be used to guide the decision or how values of multiple parties ought to be balanced, how conflicting values ought to be traded off, or how consequences occurring at different times ought to be weighed against each other (including consequences experienced before the outcome of decision is known). Evaluation of gambles requires additional judgments about which the decision-maker may be uncertain, such as how much preference the decision-maker has for or against risk, hope level, security level,

potential regret, or which gamble evaluation rule to choose. (The GDM and these factors are discussed in section III.2.10.) Uncertainty about these factors can render an SDM-based tool unhelpful or even deceptive in that it may encourage false certainty.

VI.2.2. Characteristics of uncertainty

Uncertainties can be distinguished by several characteristics to illuminate their impact and possible mitigation.

A key characteristic of an uncertainty that affects how it ought to be treated is its *hardness*. A *hard* uncertainty is one that is known to be irreducible. A *virtually hard* uncertainty may be reducible, but not within given constraints of time, labor, and other resources, and it must be treated as if it were hard. A *soft* uncertainty can be reduced, e.g. by investigation. A *malleable* uncertainty can be modified usefully by changing the world and thereby improving the decision-maker's prospects.

When a decision-maker is uncertain about the consequences of an action combined with a particular possible future state of the world, it may be important to determine whether the locus of the uncertainty is in the *data* or the *model* (e.g. mental model) the decision-maker employs. When the dynamics of a situation are well-understood, modeling may be very helpful. However, a single model is not sufficient to characterize the situation when the decision-maker is uncertain about the dynamics, e.g. about how some elements affect other elements. In this situation, a model is simply a codification of assumptions about the dynamics. Multiple models or structural modeling may help, but the value of ordinary computer modeling is more limited, since any "ordinary" computer model

embodies only one set of structural assumptions, while the decision-maker is uncertain about the structure.

Elster's distinction between *parametric* and *strategic rationality* is useful (e.g. Elster 1984). Strategic rationality is required when the future actions of other agents must be taken into account; this can be a significant complication. The simpler parametric rationality is appropriate when the actions of others are unimportant.

VI.2.3. Sources of uncertainty

The root cause of the decision-maker's uncertainty is often relevant to the question of if and how the uncertainty can be reduced, characterized, and handled. In section III.4, I defined three large categories of root causes: variability and unpredictability in the world; the decision-maker's evidence about the world; and characteristics of the decision-maker.

Uncertainty can be generated by processes in the world that are inconstant or unpredictable, whether that is due to disorder or a type of partial order, such as chaotic or stochastic regularity, or to fundamentally novel situations where lack of history prevents prediction. Some processes are sufficiently complex that, though fully determined, they are not predictable by a particular decision-maker with current processing aids. The presence of other decision-making agents can be a potent source of uncertainty. The decision-maker's choice may depend on choices of action or value that have yet to be made, perhaps by the decision-maker herself. Relevant agents may change between the time of the current decision and the decisions the others agents will make later, and that possibility makes predicting their future choices more difficult.

Whether or not the potential future state of the world is knowable, the decision-maker's evidence about the future may be imprecise, incomplete, or conflicting. Evidence obtained in linguistic form can cause uncertainty if it is fuzzy or ambiguous. When evidence comes from other agents, the decision-maker's awareness that agents *could have* made mistakes or *might have* attempted to deceive the decision-maker may cause uncertainty about the reliability of the evidence.

Whether or not the decision-maker's information is sufficient, he may be uncertain because his own characteristics prevent him from drawing all of the inferences the evidence supports. Such limiting characteristics include his information processing capabilities (memory, speed, time available, etc.), linguistic ability and knowledge of contexts, and the mental models he uses to explain and relate phenomena to each other.

VI.3. Implications

VI.3.1. Uncertainty can thwart rationality as a selecting process

A rational decision selection process compares the rationality of choices; that is, their effectiveness at promoting the decision-maker's desired ends. Clearly, uncertainty can complicate such a comparison or even render it indecisive (inconclusive).

It should be remembered that although uncertainty may be a very unwelcome complication for rational choice, it is not always bad for the decision-maker, given the alternatives. For example, the risk of a bad outcome is generally preferable to the certainty of the same bad outcome. It is also better to be uncertain than to be falsely certain, (though often much less comfortable).

Uncertainty can pose two different problems for rational decision-makers. One is that rational decision selection processes may fail to yield unique recommendations due to uncertainty, such as fuzziness in input data. This can occur even if all possible outcomes are good. The second problem has to do with the desirability of the decision-maker's prospects rather than with the decision-making process itself. Risky choices, or gambles, hold the potential for bad outcomes. Multiple-objective choices with conflicting trade-offs pose a problem structurally similar to that of gambles: they may offer only options that are bad on a least one dimension of preference in the case of multiple-objective choices. The second problem is not one of indecisiveness; it is that none of the options is good enough.

The key characteristics of these two problems call for different approaches to solutions. When fuzziness complicates a rational choice, better information or a more powerful decision selection process could solve the problem. In the case of gambles or trade-offs, efforts to improve the decision-maker's prospects appear to be more promising, such as developing options that do not have a risk of a bad outcome or a bad outcome on one dimension of preference.

Much of the technical literature has focused on risky decisions, or those of the second type, and has aimed at improving decision-making techniques, which is a solution to the first type of problem. Potent research opportunities remain for exploring more effective ways of making comparisons among choices when fuzziness is a problem and for devising better options when gambles or trade-offs appear.

VI.3.2. *Meta-rationality and the Null Hypothesis*

Since uncertainty has been shown to impose potential limits on the power of rationality as a decision-making approach, it is important to understand those limits — to be rational about our use of rationality, that is, to be *meta-rational*. Meta-rationality demands caution that a procedure for comparing the rationality of several choices may produce results that are more an artifact of the process itself than a reflection of the characteristics of the choices. One approach is to formulate a Null Hypothesis, stating that none of the choices is more rational than the others. A level of significance is adopted. Then the rational procedure attempts to reject the Null Hypothesis by showing that one (or, less desirably, more than one) of the choices is *significantly* more rational than the others. If the Null Hypothesis cannot be rejected, that rational procedure is judged to have failed to distinguish one choice for acceptance, and thus some other process, perhaps a non-rational process, must be implemented.

If it is not possible to determine which choice is optimal, it may still be possible to determine that some choices are significantly inferior.

The Uncertaintest Link Principle (see section V.2.2.3.2) suggests that when several uncertainties complicate an analysis, the conclusion cannot be more certain than the most uncertain link in the analytical chain, and therefore reducing uncertainty in one of the other links may have no benefit. (Which link is most uncertain must be measured by their impact on the certainty of the conclusion.)

VI.3.3. Rationality is inescapably subjective

VI.3.3.1. "Objective rationality" is a myth

No matter how objectively the principles of substantive rationality are stated, procedural rationality cannot avoid subjectivity. The most obvious way in which subjectivity enters the rational scene is that all evaluations of consequences are based on values. These values may be explicit or implicit, but they are not "objective" in either case. The choice of particular procedures for implementing a rational analysis equally depends on judgment. Judgment is also needed in choosing the dimensions deemed important enough to be represented, in deciding how the factors will be measured, and so on.

VI.3.3.2. Rational choice of rational methods is subjective

The combined import of the subjectivity of rationality and the potential indecisiveness of any rational procedure is that in many cases it cannot be demonstrated with certainty which approach to handling uncertainty ought to be taken, or how much emphasis each of several approaches ought to be given. Some guidance can be provided from a general perspective, such as this one, but the details of each situation cannot all be anticipated.

VI.3.4. "Decision under certainty" doesn't exist

In the traditional terminology, "decision under certainty" refers to a decision where the state of the world is known. In light of the expanded taxonomy of uncertainties described earlier, world-state indeterminism is only one possible source of uncertainty. Deterministic decisions suffer from most of the types of uncertainty described here. For instance, we can be uncertain about the past, where indeterminism has absolutely no role to

play. In the absence of perfect information, and perfect ability to process that information, there are no decisions under true certainty.

VI.3.4.1. Virtually all decisions are fuzzy gambles

As Dror observes, virtually all decisions are fuzzy gambles, though there may be a few exceptions, such as the decision to play or not to play a nickel slot machine. Assuming the potential player knows the pay-off odds for that machine, and knows the relative value of a nickel on hand versus a possible pot full of change (how big a pot?), and knows the dollar equivalent of the pleasure experienced during the slot machine's whirring deliberations before the outcome is known, and...

VI.3.4.2. Some choices may be clear nonetheless

While there may be no such thing as a decision under certainty, there may nonetheless be decisions where the right choice is clear. In the case of indeterminacy of world state, one risky option can *dominate* another, resulting in a very clear choice. In the case of fuzziness, if the difference between the measured attractiveness of two options is large compared to the amount of fuzziness in the measures, the choice may again be very clear. This emphasizes the need for a way of assessing the importance of the uncertainties in any comparison process, such as the Null Hypothesis test described above.

It is one thing to be uncertain about what the best course is; it is another to be uncertain about the course one will choose. The former may not be resolvable, but if so, it is often resolved through analysis; the latter is ultimately resolved through an act of will, or the *making* of a decision.

VI.3.5. Measurement

VI.3.5.1. Measurement is central

A fundamental theme in many of the conclusions presented here is the centrality of measurement. Measurement, or comparison, is essential to the efficacy of procedural rationality. Even when an analytical treatment is possible and clear, as in the case of determining the odds of throwing a die and having a particular number show, there is still the potential for any particular object to diverge from the ideal for its class of objects. Only with measurement can the degree of correspondence between the object's characteristics and those of its archetype be determined.

Judgment is required in deciding what characteristics to measure, how to measure them, and how to evaluate the results. For many types of measurement, assessment of the uncertainty associated with a particular result has long been part of the measurement process. It is important for assessing the contribution that any measurement's uncertainty makes to the uncertainty of the result of a decision selection process. This is no less true for assessments of probabilities, whether the assessments are objective or subjective.

The utility of possible consequences should be measured or estimated for all the relevant stakeholders, along with the uncertainty of those measurements and estimates. This part of a rational decision assessment is often taken for granted, but should not be overlooked; the uncertainty in value and utility assessments may be the uncertainty that has the largest impact on the certainty of the conclusions.

VI.3.5.2. "The best choice" does not exist

When the question, "What is the best choice?" is asked, it is linguistically natural to assume that the referent of the term "best choice" exists. This is not actually so. There is no inherent "bestness" property of options that is waiting to be disclosed. Instead, "best choice" is a potential human creation, the associating of that term with a particular option. Innumerable possible processes could be followed in attempts to *create* a correspondence between that term and a particular option. These processes may or may not identify a single choice, and different processes may identify different choices. Such processes either rely on measurements or assume results that substitute for actual measurements.

VI.4. Contributions

The fundamental contribution of this work is a conceptual scheme, the Taxonomy, which has emerged from a survey of the territory of uncertainty. The Taxonomy identifies many ways in which uncertainties can differ from each other. These distinctions can be useful to people who are building tools for use by decision-makers, to decision-makers as they deliberate, and to others who investigate phenomena that are affected by human beings and the different uncertainties they experience. Additional insight into the limits of rationality as a decision-making principle emerged from this research.

VI.4.1. Fundamental contribution: the Taxonomy, a conceptual scheme

I propose that the Taxonomy can function as part of the "concept package" called for by Dror (1986b; see section II.7.1). I have found no other taxonomies or typologies of uncertainty that deal with more than a

couple of aspects of uncertainty. The distinctions defined here can provide a language for further work in this area. The Eskimo language contains many words for the range of phenomena we call “snow”, and this allows for a fuller and more effective communication about these phenomena. Similarly, this conceptual scheme should help people communicate more clearly about the diverse phenomena we call “uncertainty”.

VI.4.2. This work as a survey of the territory of uncertainty

In section IV.3.3, I discussed uncertainty in light of Smithson’s four-stage model (1988:28-40) of evolution of intellectual attitudes towards uncertainty:

1. Exclusion (or banishment of uncertainty);
2. Inclusion (control, often through reductionism);
3. Second-order strategies (explicit analysis of uncertainty);
4. Suburbanization (subdivision, relativization).

Just as Smithson believes that mathematical approaches based on the first three stages were unable to sufficiently resolve problems relating to uncertainty, thus requiring a suburbanizing approach, I suggest that decision-making treatments based on the first three stages have left significant problems unresolved, and that suburbanization is needed in this area too. In the world of real estate, before a subdivision can be made, a survey must be completed. I believe the largest contribution this work makes is a survey of the territory of uncertainty so that an future developments in the territory can be planned appropriately and to aid the construction of relationships among current and future developments.

Figure 46 presents one preliminary visual rendering of two of the major dimensions of uncertainty defined here. Much work remains to be done before specific developments, such as utility theory, decision analysis, probability theory or possibility theory, and specific disciplines, such as psychology, political

science, or moral philosophy, can be located neatly on this map. Nonetheless, it illustrates that the territory is vast and suggests that no single development could suffice for all of uncertainty. J.D. Freund's assertion that statistics includes "the entire science of decision making in the face of uncertainty" (Freund 1962) is overstated.

A great deal of the attention given to uncertainty in decision-making has been focused

Figure 46. A partial map of the territory of uncertainty

on the lowest cell in figure 46, where the world as source of uncertainty and the future state of the world as subject of uncertainty intersect. Despite dissent from the objectivists, proponents of subjective probability have extended the use of probabilistic tools to other cells in the future state of the world column.

Although the map's axes do not represent continuums, the areas farther from the lowest cell are associated more with the organizational and personal perspectives and less with the technical perspective. A map such as this one can help integrate tools and techniques from the technical perspective with the more sensitive and detailed appreciation of decision-makers' needs that can be provided by personal and organizational perspectives.

VI.4.3. Useful for tool builders

Understanding the differences among uncertainties is necessary to continue progress in the technical study of uncertainty, probably leading towards what Smithson termed "suburbanization", when fields are divided into relatively independent sub-fields that are no longer required to be consistent with each other. This would facilitate developing methods that are tailored to the particular requirements of a decision-maker's tasks and environment and to the technical aspects of what information is available. A wide variety of uncertainties calls for a wide assortment of tools, and for careful matching of the tools to the specific task and uncertainty.

VI.4.4. Useful for decision-makers

The Taxonomy can provide critical assistance to decision-makers by clarifying the diversity of uncertainties encompassed by a decision.

Awareness of this diversity will aid selection of tools specifically suited to particular tasks and will facilitate effective use of the tools chosen. It can help decision-makers avoid using ill-suited tools which may waste time and energy. Understanding the full range of uncertainties may bring important problems to light and allow the decision-maker to anticipate and prepare for possibilities that would otherwise have been unexpected.

Major practical contributions in this work for decision-makers come from the Null Hypothesis and the Most Uncertain Link Principle. The Null Hypothesis is a feasible method for testing whether uncertainty makes it impossible to rely on rationality to identify which option to choose. If the Null Hypothesis is not rejected, unless a better rational selection process can be designed, rationality will not suffice for making the decision, and other decision-making means must also be used.

The Most Uncertain Link Principle can help allocate efforts in the most useful way when multiple uncertainties are present. When this principle is used in conjunction with the Taxonomy, which indicates how wide a variety of uncertainties may be present, limited time and analytical resources can be directed towards the uncertainties that pose the most serious problems for the decision-maker, regardless of whether they are the ones for which the most convenient tools are available.

Vi.4.5. Useful to other researchers

The distinctions among uncertainties identified here might be applied to other fields in several ways as part of explanatory or methodological efforts. They can also facilitate the adoption of technical treatments by workers in areas where the potential relevance of technical methods has

been underappreciated. Better understanding by workers in other fields of the ways in which meta-rationality and rationality can be extended beyond the Standard Decision Model and better discussion of the limits of rational choice by researchers in normative decision theory should facilitate communication and collaboration between the groups.

VI.4.6. Potential applications

VI.4.6.1. Uncertainty sub-type trade-offs

Zadeh's Incompatibility Principle (quoted in section II.3.10) states that as the complexity of a system increases, it becomes more and more difficult for any description to be both precise and relevant. In other words, there is an increasingly restrictive trade-off between certainty and relevance.

Trade-offs among subjects and sub-types of uncertainty are possible. Not all uncertainties in a formulation are equally harmful; it may be possible to decrease the most inconvenient uncertainties by allowing less important ones to increase. Alternatively, it may be possible to simplify a model to make it faster or more manageable by reducing the precision in an area where it is not really needed. Perhaps different ways of budgeting or allocating uncertainty within a project can be shown to impose significantly different limitations on goal attainment.

VI.4.6.2. Risk aversion and security level preference

People who prefer a certain gain over a gamble with a higher expected value are said to be risk averse; their risk aversion is often put forth as an explanation of why they prefer the certain but smaller payoff. The concept of security level offers an alternative explanation of their behavior. In section IV.5.2 I described a possible test of risk aversion versus preference

for a high security level. The security level explanation takes into account preferences relating to experience before the uncertainty is resolved, while the risk aversion explanation relies on preferences regarding final outcomes.

VI.4.6.3. Complexity as a function of information and uncertainty

While working on this project, I frequently experienced the uncomfortable feeling that although I was marshaling a great deal of information, each part of this dissertation was remaining far less certain and less specific than I would like. This suggests a possible relationship between uncertainty, information, and complexity. A *system* as defined by Klir (1985:132) is one particular abstraction among virtually infinitely many possible abstractions of a worldly object. Klir would measure its complexity by the amount of information needed to specify it, or to reduce any uncertainty regarding it. Klir's approach begins with the contemplation of a defined system, but the procedure could be worked the other way round. If this dissertation contains a large quantity of information and yet is highly uncertain, one might infer that the breadth of scope and level of detail presented represents a very complex system.

VI.4.7. Insights into a limit of rationality

As attractive as rationality is as a guide for making choices, specific procedures must be carried out to make assessments of the rationality of the options in any particular decision. The assessment in section II.3 shows that these procedures will not always be capable of identifying which option is the most rational to choose, that is they can fail to be decisive, and uncertainty is often the cause of such failures.

Many others have identified ways in which rationality may be of limited relevance in specific cases because of difficulties of incorporating some pertinent considerations, especially factors from O and P perspectives. The limit described here is different, in that it may be encountered in the purest of T perspectives. It is not an issue of what is able to be considered but one of possible failure to be decisive on the basis of those factors that are considered. Uncertainties of any of the types described here can directly impose this limitation.

The Null Hypothesis is suggested as a practical test for determining whether this limit has been encountered during the implementation of a rational decision selection procedure. When it has been reached and a procedure is indecisive, or when it can be predicted that no rational procedure will be decisive within the time and resource constraints, whether due to uncertainty or to other causes, it is not rational to attempt to rely on rational procedures. Rather, it is rational to adopt other, non-rational (but not *irrational*) procedures, such as improvisation or voting. It is important to be aware of this as a possible outcome of any application of rational analysis, and to be rational about what to do when this limitation of rationality is encountered.

VI.5. Limitations of this work

VI.5.1. *It is a first cut*

This dissertation represents one person's effort, which may be found later to need amplification or modification. The dimensions I have chosen to use in subdividing the territory of uncertainty may be superseded by more useful ones, or augmented with additional dimensions. I expect that

refinements would come fairly naturally, but that they would not invalidate the distinctions drawn here.

VI.5.2. It is untested

I believe this work can be useful in decision-making for both tool-makers and tool-users. I also believe it can be useful to others who find that uncertainty plays an important part of explanatory or predictive models in other fields.

These beliefs are currently untested. If sincere efforts to use this work to improve tools and to improve their utilization, and to weave technical, organizational and personal aspects of uncertainty into other work, fail to bear fruit, my assertions would be disproved.

VI.5.3. Design limitations

The deliberate concentration on uncertainty here has largely excluded considerations of two closely related and very important topics — false certainty and error! In neither of those two cases is a person aware of the problem, while I have premised my definition of uncertainty on that awareness. The issue that the three subjects together raise is how people ought to derive confidence, or conversely feelings of uncertainty, from the evidence they possess. This is important in order to avoid false uncertainty and an inability to act on the one hand, and yet avoid on the other hand coming to believe assertions that are either unsupported (false certainty) or false (error).

Both rationality and uncertainty are considered here in the context of an idealized decision, which cannot display the richness of the roles they play in organizational and personal arenas, such as the ways chronic

uncertainty affects morale, or helps shape the structural evolution of an organization.

VI.5.4. Technical focus

The Organizational and Personal Perspectives highlighted concerns about uncertainty that have been a major motivation for this work. These concerns include uncertainty about who has the authority to make a decision, how to balance different dimensions of value against each other in evaluating consequences, possible preference for avoiding feelings of regret, and others. However, once these concerns were identified, they have been treated within the mainly technical framework of the Taxonomy and subsequent chapters. Any use of this work in a specific decision context would require enriching its framework with the details illuminated by personal and organization views of the decision.

VI.6. Review of initial hypotheses

Have the hypotheses developed in section I.3.4 been substantiated? The hypotheses are recapitulated here with brief assessments of this work's success in evaluating them.

1. Distinguishing among the various ways in which uncertainty can complicate planning and decision-making can help to identify approaches that are useful to the planner or decision-maker facing a particular type of uncertainty.

The taxonomy of sources of uncertainty shows how the origin of an uncertainty can help a decision-maker determine how to go about reducing the uncertainty. The other parts of the taxonomy show different ways that decisions are complicated by uncertainties, and demonstrate that

uncertainties that cannot be reduced can be coped with in different ways.

Established.

2. Approaches that are helpful in some circumstances can be unhelpful or harmful if used in other circumstances.

A clear example is MEU. In the case of repeated decisions with known probabilities, virtually no one disputes MEU's legitimacy as the proper decision rule. However, when a decision-maker has a strong preference for a high security level, or a strong preference for risk, MEU is ruled out by its own axioms. Attempting to employ it will make it more difficult or even impossible to bring out the distinctions among options that are most important to the decision-maker. **Established.**

3. The definition of uncertainty usually assumed in technical treatments of uncertainty can be expanded to include more of the types of uncertainties that managers encounter in their work.

Fuzziness, a major kind of uncertainty in everyday personal and organizational life, can render rationality indeterminate. **Established.**

4. The technical approach can be shown by rational arguments to be inadequate if used alone.

The importance of value judgments shows that technical work is not enough to identify "best" or "most rational" options. Value judgments are necessary for all comparisons of the desirability of outcomes, as well as for making decisions where the values of more than one stakeholder must be considered or when more than one dimension of consequence or value is important. Value determinations cannot be made solely by technical means. In addition, rationality itself has been shown to have the potential

to be mute or inadequately decisive on specific decisions due to a wide variety of uncertainties. **Established.**

VI.7. Future work to extend this research

There are two branches of work that are very attractive candidates for extensions of this work. In the first, the surveying of the territory of uncertainty would be refined. In the second, particular regions within this territory would be targeted for development.

VI.7.1. Continue refining the survey of the territory of uncertainty

The matrix shown in figure 46 above represents one attempt at visually portraying some of the territory circumscribed by the expanded interpretation of uncertainty I have constructed. Other representations are possible, and ought to be explored. An obvious function for such representations is to serve as a vehicle for relating disciplines, developments, decision needs and varieties of uncertainties to each other. This continued survey work can help to clarify the strengths and limitations of tools for different purposes, and identify sub-regions whose development has lagged behind other sub-regions and may therefore offer great potential for further development work.

For example, probability theory is used both for representing likelihoods of future states of well-observed stochastic processes (as the objectivists demand), and for representing strengths of belief in particular assertions (encouraged by the subjectivists). The taxonomy presented earlier suggests that these uses are very different. In the first case, the source of the uncertainty is a variable process in the world, and the probability is treated as a characteristic of the process; in the second case,

the probability is used as a measure of the uncertainty in the conclusion of whatever process was used to infer the assertion, and no particular attribution of the uncertainty to a particular source has been made.

To suggest another example, uncertainty over the future state of the world has been described using both probabilistic and possibilistic concepts. Clearly, it would be useful to better understand the relative strengths and weaknesses of these two concepts for this purpose.

Inquiry could be made into the distinction between incomplete evidence and un-made decisions as possible sources of uncertainty about a stakeholder's values. If the values are unclear now because evidence is insufficient, different methods for obtaining that evidence might be hypothesized to produce similar results. On the other hand, if the stakeholder's values are unclear now because key value decisions have not yet been made, it is reasonable to predict that different elicitation methods could bias the stakeholder to make different value decisions.

VI.7.2. Identify specific needs and attempt to develop better tools

A second major direction for extending this research heads towards tool development. One area where current tools are unsatisfying is non-probabilistic uncertainty over future states of the world. When the set of possible states cannot be known well, the distributional nature of probability recommends against its use. Lack of relevant experience with a projected situation also weakens the help probabilities can provide. Both possibility theory and chaos theory are currently attracting considerable attention, and may help in this region.

The earlier discussion of uncertainty in the form of fuzziness about

other decision factors (e.g. knowledge of values and value tradeoffs) highlights a need for ways of performing rational analyses when basic evidence used as input is imprecise, and for measuring the impact that imprecision in the input has on the certainty of analytical conclusions.

As noted earlier, much of the work to date, especially the quantitative work, has been performed in areas important to the Technical Perspective. Areas less important to T but vital to the Organization and Personal perspectives, such as knowledge of values, clarity of decision-making authority, limitations of processing power, and the usefulness of better options, have seen significant developments. However, these developments have remained somewhat out of the technical mainstream, even though they deal with uncertainties that may sometimes be more important than the uncertainties treated by technical tools such as expected utility analysis. Additional work incorporating all three perspectives could provide valuable understandings of the multidimensional nature of uncertainty.

There is no such thing as absolute certainty, but there is assurance sufficient for the purposes of human life.

John Stuart Mill

The territory of uncertainty is vast, and only small areas of it have been charted. My Taxonomy helps define where and how uncertainty can affect a decision selection process and limit the power of rational choice. For scholars, this definition and my efforts at mapping the territory of uncertainty contribute to a theoretical understanding of how decision

selection processes function and how they might be improved. For decision-makers themselves, my definitions and mapping of the land of uncertainty provide guidance for selecting decision-making criteria and processes and for evaluating the appropriateness of rational decision-making procedures in specific situations. In this way, making advances at both theoretical and applied levels, our uncertainty about the territory of uncertainty is gradually reduced, and an increasingly detailed and useful map can be constructed.

VII. REFERENCES CONSULTED

VII.1. Cited works

- Aiginger, Karl 1987. *Production and Decision Theory under Uncertainty*, Basil Blackwell Ltd., Oxford, UK.
- Alefeld, G. and J. Herzberger 1983. *Introduction to Interval Computations*, Academic Press, New York.
- Allais, Maurice 1979. "The So-Called Allais Paradox And Rational Decisions Under Uncertainty", in M. Allais and O. Hagen (eds.), *Expected Utility Hypotheses and the Allais Paradox*, Reidel, Dordrecht, Holland.
- 1988. "The General Theory of Random Choices in Relation to the Invariant Cardinal Utility Function and the Specific Probability Function", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, pp. 231-90.
- Allais, M., and O. Hagen (eds.) 1979. *Expected Utility Hypotheses and the Allais Paradox—Contemporary Discussion of Decisions Under Uncertainty with Allais' Rejoinder*, D. Reidel, Dordrecht.
- Allison, Graham 1968. "Conceptual Models and the Cuban Missile Crisis," *The American Political Science Review* 63 (3), 689 - 718.
- 1971. *Essence of Decision: Explaining the Cuban Missile Crisis*, Little, Brown, and Co., Boston.
- Anderson, D. F. 1977. *Mathematical Models and Decision Making in Bureaucracies: A Case Story Told from Three Points of View*, Ph. D. Dissertation, MIT, Cambridge, Mass.
- Arrow, Kenneth J. 1982. "Current Developments in the Theory of Social Choice", in Brian Barry and Russell Hardin (eds.), *Rational Man and Irrational Society?*, Sage Publications, Beverly Hills.
- Ascher, William, and William H. Overholt 1983. *Strategic Planning and Forecasting*, John Wiley and Sons, New York.
- Ashby, W. Ross 1954. *An Introduction to Cybernetics*, Chapman and Hall, London.

- Baird, Bruce F. 1989. *Managerial Decisions Under Uncertainty*, John Wiley and Sons, New York.
- Barry, Brian, and Russell Hardin (eds.) 1982. *Rational Man and Irrational Society?*, Sage Publications, Beverly Hills.
- Bass, Bernard M. 1983. *Organizational Decision Making*, Richard D. Irwin Inc., Homewood, Illinois.
- Bateson, Gregory 1979. *Mind and Nature*, Bantam Books, Toronto.
- Becker, Henk A., and Joseph W. M. van Doorn 1987. "Scenarios in an Organizational Context", *Futures*, Dec., pp. 669-677.
- Bell, D. 1982. "Disappointment in Decision Making Under Uncertainty", *Operations Research*, 30, pp. 961-81.
- Bernoulli, Jacob 1713. *Ars Conjectandi*, Basel.
- Blair, Douglas, and Robert Pollak 1983. "Rational Collective Choice", *Scientific American*, August, pp. 88-95.
- Blau, P. 1963. *The Dynamics of Bureaucracy: A Study of Interpersonal Relations in Two Government Agencies*, University of Chicago Press, Chicago.
- Blumm, Michael C. 1983. "Risk Management and Northwest Electric Power Planning: Some Lessons From the Rearview Mirror", *Environmental Law*, vol. 13:739-773.
- Boudon, Raymond 1982. *The Unintended Consequences of Social Action*, St. Martin's Press, New York.
- Bouyssou, Denis, and Jean-Claude Vansnick 1988. "A Note on the Relationship Between Utility and Value Functions, in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, 1988, pp. 103-14.
- Bowman, E. H. 1980. "A Risk/Return Paradox for Strategic Management", *Sloan Management Review*, 23, pp. 17-31.
- Bowonder, B. and H.A. Linstone, 1987. "Notes on the Bhopal Accident: Risk Analysis and Multiple Perspectives", *Technological Forecasting and Social Change*, 32, pp. 183-202.
- Brown, Robert G. 1978. "Inventory Control", in Moder, Joseph J., and Salah E. Elmaghraby (eds.) *Handbook of Operations Research: Models*

- and Applications*, Van Nostrand Reinhold, New York, pp. 173-212.
- Burke, W.J., H.M. Merrill, F.C. Schweppe, B.E. Lovell, M.F. McCoy, S.A. Monohon 1988. "Trade-Off Methods in System Planning", *IEEE Transactions on Power Systems*, Vol. 3, No. 3, August, 1988, pp. 1284-90.
- Burns, T., and G. M. Stalker 1961. *The Management of Innovation*, Tavistock, London.
- Carter, C. F., and J. L. Ford (eds.) 1972. *Uncertainty and Expectations in Economics*, (Essays in Honour of G. L. S. Shackle) Augustus M. Kelley, New Jersey.
- Casti, John 1990. *Searching for Certainty: What Scientists Can Know About the Future*, William Morrow and Company, New York.
- Churchman, C. West 1979. *The Systems Approach and Its Enemies*, Basic Books, New York.
- Cohen, Michèle and Jean-Yves Jaffray, 1988. "Preponderance of the Certainty Effect Over Probability Distortion in Decision Making Under Risk", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, 1988, pp. 173-87.
- Crossley, David J., and C. A. Hooker 1987. "Retrospective Scenario Analysis and Government Energy Planning: Toward an Effective Synergy", *Energy Systems and Policy*, Vol. 10, No. 2, pp. 93-113.
- Crozier, Michael 1964. *The Bureaucratic Phenomenon*, University of Chicago Press, Chicago.
- Cyert, R. M., and J. G. March 1963. *A Behavioral Theory of the Firm*, Prentice-Hall, Englewood Cliffs, N. J.
- Daston, Lorraine J. 1987a. "The Domestication of Risk: Mathematical Probability and Insurance 1650-1830", in Lorenz Krüger, Lorraine J. Daston and Michael Heidelberger (eds.), *The Probabilistic Revolution*, The MIT Press, Cambridge, Mass., 1987, pp. 237-260.
- 1987b. "Rational Individuals versus Laws of Society: From Probability to Statistics", in Lorenz Krüger, Lorraine J. Daston and Michael Heidelberger (eds.), *The Probabilistic Revolution*, The MIT Press, Cambridge, Mass., 1987, pp. 295-304.
- De Finetti, Bruno 1937. "Foresight: Its Logical Laws, Its Subjective Sources", in H. E. Kyburg and H. E. Smokler (eds.) *Studies in*

- Subjective Probability*, John Wiley and Sons, NY, 2nd ed., 1980, pp. 53-118 (originally published in *Annales de l'Institut Henri Poincaré*, vol. 7 (1937)).
- 1976. "Probability: Beware of Falsifications!", in Henry E. Kyburg and Howard E. Smokler (eds.), *Studies in Subjective Probability*, 2nd Ed., Robert E. Krieger, Huntington, NY., 1980, pp. 193-224 (originally published in *Scientia*, 111, 1976, pp. 283-303).
- Dempster, A. 1967. "Upper and lower probabilities induced by multi-valued mapping," *Ann. Math. Statist.*, 38, pp. 325-39.
- Douglas, Mary, and Aaron Wildavsky 1982. *Risk and Culture: An Essay on the Selection of Technical and Environmental Dangers*, University of California Press, Berkeley.
- Dror, Yehezkel 1971. *Crazy States: A Counterconventional Strategic Problem*, Kraus Reprint, Millwood, New York.
- 1986a. *Policymaking Under Adversity*, Transaction Books, New Brunswick, NJ.
- 1986b. "Planning as Fuzzy Gambling: A Radical Perspective on Coping with Uncertainty", in David Morley and Arie Shachar (eds.), *Planning in Turbulence*, The Magnes Press, Jerusalem.
- 1987. "Planning as a Mode of Policy-Reasoning", in Leonard Guelke and Richard Preston (eds.), *Abstract Thoughts: Concrete Solutions*, University of Waterloo, Department of Geography Publications Series, Waterloo, Canada, pp. 25-47.
- 1988a. "Uncertainty: Coping with It and with Political Feasibility", in Hugh J. Miser and Edward S. Quade (eds.), *Handbook of Systems Analysis: Craft Issues and Procedural Choices*, North-Holland, New York, pp. 247-281.
- 1988b. "Notes Towards a Philosophy of Policy-Reasoning", in Mirian Campanella (ed.), *Between Rationality and Cognition: Policy-Making Under Conditions of Uncertainty, Complexity, and Turbulence*, Albert Meynier, Torino, Italy, pp. 117-171.
- 1991. "Grand Policy Analysis", Presidential Address Paper for Annual Meeting of the Policy Study Association, August, 1991, Washington, DC.
- Dubois, Didier 1988. "Possibility Theory: Searching for Normative Foundations", in B. R. Munier (ed.), *Risk, Decision and Rationality*,

- Reidel, Dordrecht, pp. 601-14.
- Edwards, Ward 1982. "Conservatism in Human Information Processing", in Kahneman, Daniel, Paul Slovic, and Amos Tversky, (eds.), *Judgments under uncertainty: Heuristics and biases*, Cambridge University Press, Cambridge.
- Einhorn, H.J., and R.M. Hogarth, 1981. "Behavioral Decision Theory: Processes of Judgment and Choice", *Annual Review of Psychology*, 32, 53-88.
- Elster, Jon 1984. *Ulysses and the Sirens: Studies in Rationality and Irrationality*, Cambridge University Press, Cambridge, 1979 (revised edition 1984).
- 1986. "Introduction", in Elster, Jon (ed.), *Rational Choice*, New York University Press, New York.
- 1989a. *Solomonic Judgements: Studies in the Limitations of Rationality*, Cambridge University Press, Cambridge.
- 1989b. *The Cement of Society: A Study of Social Order*, Cambridge University Press, Cambridge.
- Emery, F. E., and E. L. Trist 1965. "The Causal Texture of Organizational Environments", *Human Relations*, vol. 18, pp. 21-32.
- Fiegenbaum, A., and H. Thomas 1988. "Attitudes Toward Risk and the Risk/Return Paradox: Prospect Theory Explanations", *Academy of Management Journal*, 31, pp. 85-106.
- Fischhoff, Baruch, Sarah Lichtenstein, Paul Slovic, Stephen L. Derby, and Ralph L. Keeney 1981. *Acceptable Risk*, Cambridge University Press, Cambridge.
- Fischhoff, Baruch 1982. "Debiasing", in Kahneman, Daniel, Paul Slovic, and Amos Tversky, (eds.), *Judgments under uncertainty: Heuristics and biases*, Cambridge University Press, Cambridge.
- Fishburn, Peter 1982. "Non-Transitive Measurable Utility", *Journal of Mathematical Psychology*, 26, pp. 31-67.
- Fleurbaey, Marc 1988. "Rational Behavior and Adaptation", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, pp. 459-481.
- Ford, J. L. 1983. *Choice, Expectation and Uncertainty*, Barnes and Noble,

Totowa, NJ.

- Freund, J. D. 1962. *Mathematical Statistics*, Prentice-Hall, Englewood Cliffs, NJ.
- Friend, J., and W. Jessop 1969. *Local Government and Strategic Choice*, Tavistock Publications, London.
- Gall, John 1975. *Systemantics*, New York Times Book Co., New York.
- Gauthier, David 1975. "Reason and Maximization", original publication in *Canadian Journal of Philosophy*, 4: 411-433, 1975.
- 1982. "Reason and Maximization", excerpts in Brian Barry and Russell Hardin (eds.) *Rational Man and Irrational Society?*, Sage Publications, Beverly Hills.
- Gillwald, Katrin 1991. "Herman Kahn and the Developments in Germany: An Essay on Futures Research", *Technological Forecasting and Social Change*, Vol. 40, No. 3, Nov. 1991, pp. 289-96.
- Gleick, James 1987. *Chaos: Making a New Science*, Viking, New York.
- Gold, Thomas 1987. *Power from the Earth: Deep Earth Gas—Energy for the Future*, J. M. Dent, London.
- Good, I. J. 1980. "Subjective Probability as the Measure of a Non-measurable Set", in Henry E. Kyburg and Howard E. Smokler (eds.), *Studies in Subjective Probability*, 2nd Ed., Robert E. Krieger, Huntington, NY., 1980, pp. 133-146 (originally published in Ernest Nagel, Patrick Suppes, and Alfred Tarski (eds.), *Logic, Methodology and Philosophy of Science*, Stanford University Press, Palo Alto, 1962).
- Gordon, Theodore J., 1992. "Chaos in Social Systems", *Technological Forecasting and Social Change*, 42, pp. 1-15.
- Gordon, T.J., and D. Greenspan 1988. "Chaos and Fractals: New Tools for Technological and Social Forecasting", *Technological Forecasting and Social Change*, 34, pp. 1-25.
- Gordon, Theodore, and David Greenspan 1994. "The Management of Chaotic Systems", *Technological Forecasting and Social Change*, 47, pp. 49-62.
- Goulet, Denis 1986. "Three Rationalities in Development Decision-Making", *World Development*, vol. 14, no. 2, pp. 301-17.

- Grayson, C. J. 1960. *Decisions Under Uncertainty: Drilling Decisions by Oil and Gas Operators*, Harvard Business School, Division of Research, Boston.
- Hacking, Ian 1987. "The Erosion of Determinism", in Lorenz Krüger, Lorraine J. Daston and Michael Heidelberger (eds.), *The Probabilistic Revolution*, The MIT Press, Cambridge, Mass.
- Hamburger, Henry 1979. *Games as Models of Social Phenomena*, W. H. Freeman, San Francisco.
- Hansson, Bengt 1975. "The Appropriateness of the Expected Utility Model", *Erkenntnis*, 9, 175-93.
- Harsanyi, John C. 1988. "Assessing Other People's Utilities", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, 1988, pp. 127-38.
- Heidelberger, Michael 1987. "Fechner's Indeterminism: From Freedom to Laws of Chance", in Lorenz Krüger, Lorraine J. Daston and Michael Heidelberger (eds.), *The Probabilistic Revolution*, The MIT Press, Cambridge, Mass., 1987, pp. 117-156.
- Hey, John D. 1979. *Uncertainty in Microeconomics*, New York University Press, New York.
- 1981. *Economics in Disequilibrium*, New York University Press, New York.
- 1988. "Experimental Investigations into Economic Behaviour under Uncertainty", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, pp. 147-162.
- Hillier, Frederick S., and Gerald J. Lieberman 1980. *Introduction to Operations Research*, 3rd Ed., Holden-Day, Oakland, California.
- Hosseini, Jinoos 1986. "Decision analysis and its application in the choice between two wildcat oil ventures", *Interfaces*, vol. 16, no. 2, March-April, pp. 75-85.
- Huss, William R., and Edward J. Honton 1987. "Scenario Planning—What Style Should You Use?", *Long Range Planning*, Vol. 20, No. 4, pp. 21-29.
- Janis, I. L. 1972. *Victims of Groupthink*, Houghton Mifflin, Boston.
- Janis, I. L., and L. Mann 1977. *Decision Making: A Psychological Analysis*

of Conflict, Choice, and Commitment, The Free Press, New York.

- Jauch, Lawrence R., and Kenneth L. Kraft 1986. "Strategic Management of Uncertainty", *Academy of Management Review*, 1986, Vol. 11, No. 4, pp. 777-90.
- Jeffrey, Richard C. 1968. "Probable Knowledge", in Henry E. Kyburg and Howard E. Smokler (eds.), *Studies in Subjective Probability*, 2nd Ed., Robert E. Krieger, Huntington, NY., 1980, pp. 225-238 (originally published in Imre Lakatos (ed.), *The Problem of Inductive Logic*, North-Holland, Amsterdam, 1968).
- Jungermann, Helmut 1983. "The Two Camps on Rationality", in Roland W. Scholz (ed.), *Decision Making Under Uncertainty*, Elsevier Science Publishers, North-Holland, Amsterdam.
- Kahneman, D., and A. Tversky 1979. "Prospect theory: an analysis of decisions under risk", *Econometrica*, 47(2), March, 263-91.
- and — 1982a. "The simulation heuristic", in Kahneman, Daniel, Paul Slovic, and Amos Tversky (eds.), *Judgment under uncertainty: Heuristics and biases*, Cambridge University Press, Cambridge, UK.
- and — 1982b. "Variants of uncertainty", in Kahneman, Daniel, Paul Slovic, and Amos Tversky (eds.), *Judgment under uncertainty: Heuristics and biases*, Cambridge University Press, Cambridge, UK.
- Kahneman, Daniel, Paul Slovic, and Amos Tversky (eds.) 1982. *Judgment under uncertainty: Heuristics and biases*, Cambridge University Press, Cambridge, UK.
- Kamlah, Andreas 1987. "The Decline of the Laplacian Theory of Probability: A Study of Stumpf, von Kries, and Meinong", in Lorenz Krüger, Lorraine J. Daston and Michael Heidelberger (eds.), *The Probabilistic Revolution*, The MIT Press, Cambridge, Mass., 1987, pp. 91-116.
- Kandel, Abraham 1986. *Fuzzy Mathematical Techniques with Applications*, Addison-Wesley, Reading, Mass.
- Karni, Edi, and Zvi Safra 1988. "Preference Reversals' and the Theory of Decision Making under Risk", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, pp. 163-172.
- Kast, Robert 1988. "A Generalization of Rational Behavior", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, pp. 419-434.

- Keirstad, B. S. 1972. "Decision-Taking and the Theory of Choice", in Carter, C. F., and J. L. Ford, (eds.), *Uncertainty and Expectations in Economics*, [Essays in Honour of G. L. S. Shackle] Augustus M. Kelley, New Jersey.
- Keynes, J. M. 1921. *A Treatise on Probability*, MacMillan and Co., London.
- Klir, George J. 1985. "Complexity: some general observations", *Systems Research*, Vol. 2, No. 2, pp. 131-140.
- 1989. "Is there more to uncertainty than some probability theorists might have us believe?", *Int. J. General Systems*, Vol. 15, pp. 347-78.
- 1990a. "Dynamic aspects in reconstructability analysis: the role of minimum uncertainty principle", *Revue Internationale de Systemique*, Vol. 4, No. 1, pp. 33-43.
- 1990b. "A Principle of uncertainty and information invariance," *Int. J. General Systems*, Vol. 17, pp. 249-75.
- 1992a. "The role of methodological principles of uncertainty in Economics," in J. Lee Auspitz, et al. (eds.), *Praxiologies and the philosophy of economics*, Transaction Publishers, New Brunswick, NJ., pp. 61-92.
- 1992b. "Multimodel Representation and Management of Uncertainty", in Novák, V., J. Ramík, M. Mares, M. Cerny, and J. Nekola (eds.) *Fuzzy Approach to Reasoning and Decision Making*, Academia, Prague, and Kluwer, Dordrecht, pp. 75-86.
- 1993. "Developments in Uncertainty-Based Information", in *Advances in Computers*, vol. 36, pp. 255-332.
- Knight, F. H. 1921. *Risk, Uncertainty and Profit*, Houghton Mifflin, New York.
- Knobloch, Eberhard 1987. "Emile Borel as a Probabilist", in Lorenz Krüger, Lorraine J. Daston and Michael Heidelberger (eds.), *The Probabilistic Revolution*, The MIT Press, Cambridge, Mass., 1987, pp. 215-233.
- Koopman, Bernard O. 1940. "The Bases of Probability", in H. E. Kyburg and H. E. Smokler (eds.) *Studies in Subjective Probability*, John Wiley and Sons, NY, 2nd ed., 1980, pp. 119-131 (originally published in *Bulletin of the American Mathematical Society*, Vol. 46 (1940), pp. 763-774).

- Kopp, Sheldon 1972. *If You Meet the Buddha on the Road, Kill Him*, Science and Behavior Books, Ben Lomond, California.
- Kosko, Bart 1993. *Fuzzy Thinking: The New Science of Fuzzy Logic*, Hyperion, New York.
- Kotter, J. P. 1982. *The General Managers*, The Free Press, New York.
- Krüger, Lorenz 1987. "The Slow Rise of Probabilism: Philosophical Arguments in the Nineteenth Century", in Lorenz Krüger, Lorraine J. Daston and Michael Heidelberger (eds.), *The Probabilistic Revolution*, The MIT Press, Cambridge, Mass., 1987, pp. 59-90.
- Kunreuther, Howard, and Joanne Linnerooth 1983. *Risk Analysis and Decision Processes*, Springer-Verlag, Berlin.
- Kyburg, Henry E., and Howard E. Smokler 1980. "Introduction", in Henry E. Kyburg and Howard E. Smokler (eds.), *Studies in Subjective Probability*, 2nd edition, Robert E. Krieger, Huntington, New York, pp. 3-22.
- Laszlo, Ervin 1988. "Footnotes to a History of the Future", *Futures*, October, 1988 pp. 479-92.
- Lawrence, P. R., and J. W. Lorsch 1967. *Organization and Environment*, Irwin, Homewood, Illinois.
- Lee, Kai N. 1983. "The Path Along the Ridge: Regional Planning in the Face of Uncertainty", *Washington Law Review*, Vol. 58, pp. 317-342.
- Lee, Kai N., and Jody Lawrence 1986. "Adaptive Management: Learning from the Columbia River Basin Fish and Wildlife Program", *Environmental Law*, vol. 16:431-460.
- Lindley, D. V. 1987. "The probability approach to the treatment of uncertainty in artificial intelligence and expert systems," *Statistical Science*, 2, No. 1, pp. 17-24.
- Linnerooth, J. 1984. "The political processing of uncertainty", *Acta Psychologica*, 56, pp. 219-231.
- Linstone, Harold A. 1984. *Multiple Perspectives for Decision Making*, North-Holland, New York.
- Linstone, Harold A., with Ian I. Mitroff 1994. *The Challenge of the 21st Century*, State University of New York Press, Albany, NY.

- Loomes, G., and R. Sugden 1982. "Regret theory: an alternative theory of rational choice under uncertainty", *Economic Journal*, 92(368), December, pp. 805-24.
- Loomes, Graham 1988. "Different Experimental Procedures for Obtaining Valuations of Risky Actions: Implications for Utility Theory", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, pp. 37-57.
- Lopes, Lola L. 1988. "Economics as Psychology: A Cognitive Assay of the French and American Schools of Risk Theory", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, pp. 405-16.
- Lorraine, Hilary 1991. "The California 1980 Medfly Eradication Program: An Analysis of Decision Making Under Nonroutine Conditions", *Technological Forecasting and Social Change*, Vol. 40, No. 1, August 1991, pp. 1-32.
- Machina, M. 1982. "Expected Utility Analysis without the Independence Axiom", *Econometrica*, 50, pp. 277-323.
- Mandel, Thomas F. 1983. "Futures Scenarios and Their Uses in Corporate Strategy", in Kenneth J. Albert (ed.), *The Strategic Management Handbook*, McGraw-Hill, New York.
- Mannheim, Karl 1936. *Ideology and Utopia: An Introduction to the Sociology of Knowledge* (translated by Louis Wirth and Edward Shils from the German), Harcourt, Brace and World, New York.
- March, J. G., and J. P. Olsen 1976. *Ambiguity and Choice in Organizations*, Universitetsforlaget, Bergen.
- March, J. G., and H. A. Simon 1958. *Organizations*, Wiley, New York.
- Mason, Richard O., and Ian I. Mitroff 1981. *Challenging Strategic Planning Assumptions: Theory, Cases and Techniques*, Wiley, New York.
- McCall, Morgan W., Jr. 1988. "Developing Executives Through Work Experience", *Human Resource Planning*, Vol. 11, No. 1, pp. 1-11.
- McCord, M., and R. de Neufville 1983. "Fundamental Deficiency of Expected Utility Analysis", in French, S., R. Hartley, L. C. Thomas and D. J. White (eds.), *Multi-objective Decisionmaking*, Academic Press, London, pp. 279-305.
- McCord, Mark R., and Christos Leotsarakos 1988. "Investigating Utility

- and Value Functions with an 'Assessment Cube'", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, 1988, pp. 59-75.
- MECCA 1991. "MECCA: Model Evaluation Consortium for Climate Assessment", Introductory Brochure, MECCA Communications Office, 1-414-785-5946.
- Michael, Donald N. 1973. *On Learning to Plan—And Planning to Learn*, Jossey-Bass, San Francisco.
- Miles, R. E., and C. C. Snow 1978. *Organizational Strategy, Structure and Process*, McGraw-Hill, New York.
- Miller, D. W., and M. K. Starr 1967. *The Structure of Human Decisions*, Prentice-Hall, Englewood Cliffs, New Jersey.
- Mills, J. 1965. "The Effect of Certainty on Exposure to Information Prior to Commitment", in R. P. Abelson et al. (eds.), *Theories of Cognitive Consistency: A Sourcebook*, Rand-McNally, Chicago.
- Mintzberg, Henry, Duru Raisinghani and André Théorêt 1976. "The Structure of 'Unstructured' Decision Processes", *Administrative Science Quarterly*, June 1976, Vol. 21, pp. 246-75.
- Mitroff, I. I., and H. A. Linstone 1993. *The Unbounded Mind*, Oxford University Press, New York.
- Monat, A., J. R. Averill and R. S. Lazarus 1972. "Anticipatory Stress and Coping Reactions Under Various Conditions of Uncertainty", *Journal of Personality and Social Psychology*, 24, pp. 237-53.
- Mongin, Philippe 1988. "A Comment On J. Harsanyi's 'Assessing Other People's Utilities'", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, 1988, pp. 139-44.
- Morgan, M. Granger, and Max Henrion 1990. *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*, Cambridge University Press, Cambridge.
- Munier, Bertrand R. 1988. "A Guide To Decision-Making Under Uncertainty", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, pp. 1-34.
- Nelson, Richard R., Sidney G. Winter, and Herbert L. Schuette 1976. "Technical Change in an Evolutionary Model", *Quarterly Journal of Economics*, 90, pp. 90-118.

- Newendrop, P.S. 1975. *Decision Analysis for Petroleum Exploration*, Petroleum Publishing Company, Tulsa, Oklahoma.
- Ornstein, Robert 1986. *Multimind*, Houghton Mifflin, Boston.
- Osborn, R. N., J. G. Hunt, and L. R. Jauch 1980. *Organization Theory: An Integrated Approach*, Wiley, New York.
- Patterson, Douglas 1993. Review of *Nonlinear Dynamics, Chaos, and Instability* (Brock, W.A., S.A. Hsieh, and B. LeBaron, MIT Press, Cambridge, 1992), *Journal of Finance*, March, pp. 404-07.
- Perrow, Charles 1984. *Normal Accidents: Living With High-Risk Technologies*, Basic Books, New York.
- Pope, Robin Elizabeth 1988. "The Bayesian Approach: Irreconcilable with Expected Utility Theory?", in B. R. Munier (ed.), *Risk, Decision and Rationality*, Reidel, Dordrecht, pp. 221-230.
- Raiffa, Howard 1968. *Decision Analysis; Introductory Lectures on Choices Under Uncertainty*, Random House, New York.
- Ramsey, Frank Plumpton 1926. "Truth and Probability", in Henry E. Kyburg and Howard E. Smokler (eds.), *Studies in Subjective Probability*, 2nd edition, Robert E. Krieger, Huntington, New York, 1980.
- Random House 1980. *The Random House College Dictionary, Revised Edition*, Random House, New York.
- Ravetz, J. R. 1987. "Usable knowledge, usable ignorance", *Knowledge: Creation, Diffusion, Utilization*, 9, pp. 87-116.
- Resconi, Germano, George J. Klir and Ute St. Clair 1992. "Hierarchical uncertainty metatheory based upon modal logic," *Int. J. General Systems*, Vol. 21, pp. 23-50.
- Robertshaw, Joseph E., Stephen J. Mecca, and Mark N. Rerick 1978. *Problem Solving: A Systems Approach*, Petrocelli Books, New York.
- Russell, B. 1958. *Portraits from Memory and Other Essays*, Allen and Unwin, London.
- Russo, J. Edward and Paul J. H. Schoemaker 1992. "Managing Overconfidence", *Sloan Management Review*, Winter, pp. 7-17.
- Savage, L. J. 1954. *The Foundations of Statistics*, Wiley, New York.

- 1971. "The Elicitation of Personal Probabilities", in Henry E. Kyburg and Howard E. Smokler (eds.), *Studies in Subjective Probability*, 2nd Ed., Robert E. Krieger, Huntington, NY., 1980, pp. 147-191 (originally published in *Journal of the American Statistical Association*, December 1971, Vol. 66, No. 336, Theory and Methods section)).
- Schelling, Thomas C. 1960. *The Strategy of Conflict*, Harvard University Press, Cambridge, Mass.
- Schnaars, Steven P. 1987. "How to Develop and Use Scenarios", *Long Range Planning*, Vol. 20, No. 1, Feb., pp. 105-114.
- Shackle, G. L. S. 1949a. "A Non-Additive Measure of Uncertainty", *Review of Economic Studies*, vol. 17, no. 1, October, pp. 70-74.
- 1949b. "Probability and Uncertainty", *Metroeconomica*, December.
- 1949c. *Expectations in Economics*, The [Cambridge] University Press, Cambridge.
- 1957. *Time in Economics*, North-Holland, Amsterdam.
- 1960. Review of B.S. Keirstead's *Capital, Interest and Profits* (Basil Blackwell, Oxford 1959), *Economic Journal*, LXX, no. 280, Dec., pp. 816-19.
- 1961. *Decisions, Order and Time in Human Affairs*, Cambridge University Press, Cambridge.
- Shafer, G. 1976. *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, NJ.
- Sieber, Sam D. 1981. *Fatal Remedies: The Ironies of Social Intervention*, Plenum Press, New York.
- Simon, Herbert A. 1959. "Theories of Decision-making in Economics and Behavioral Science", *The American Economic Review*, vol. XLIX, No. 3, June, pp. 253-283.
- 1976. *Administrative Behavior*, Macmillan, New York.
- 1978. "Rationality as process and as a product of thought", *American Economic Review*, 68, pp. 1-16.
- 1981. *The Sciences of the Artificial*, 2nd edition, MIT Press, Cambridge, Mass.
- 1983. *Reason in Human Affairs*, Stanford University Press, Stanford,

California.

- 1988/89. "Why Economists Disagree", in *The Journal of Business Administration*, Vol. 18, No. 1, (Winter-Spring 1988/89), pp. 1-19.
- Slovic, Paul 1971. "Limitations of the Mind of Man: Implications for Decision Making in the Nuclear Age", *Oregon Research Institute Bulletin*, 11, pp. 41-49.
- Smith, M. B. 1974. "Probability estimates for petroleum drilling decisions", *Journal of Petroleum Technology*, vol. 26, (June) pp. 687-695.
- Smithson, Michael 1988. *Ignorance and Uncertainty: Emerging Paradigms*, Springer-Verlag, New York.
- Southern California Edison 1988. "Planning for Uncertainty: A Case Study", *Technological Forecasting and Social Change*, Vol. 33, No. 2, April 1988, pp. 119-148.
- Staw, Barry M., Lance E. Sandelands and Jane E. Dutton 1981. "Threat-Rigidity Effects in Organizational Behavior: A Multilevel Analysis", *Administrative Science Quarterly*, December, 1981, pp. 501-524.
- Steinbruner, J. D. 1968. *The Mind and the Milieu of Policy-Makers: A Case History of the MLF*, Ph. D. Dissertation, MIT, Cambridge, Mass.
- 1974. *The Cybernetic Theory of Decision*, Princeton University Press, Princeton, NJ.
- Stevens, S. S. 1968. "Measurement, Statistics, and the Schemapiric View", *Science*, Vo. 161, August 30, No. 3844, pp. 849-856.
- Stigler, Stephen M. 1987. "The Measurement of Uncertainty in Nineteenth-Century Social Science", in Lorenz Krüger, Lorraine J. Daston and Michael Heidelberger (eds.), *The Probabilistic Revolution*, The MIT Press, Cambridge, Mass., 1987, pp. 287-294.
- Stinchcombe, Arthur L. 1990. *Information and Organizations*, University of California Press, Berkeley.
- Sugeno, M. 1977. "Fuzzy measures and fuzzy integrals: a survey", in M. M. Gupta, G. N. Saridis and B. R. Gaines (eds.), *Fuzzy Automata and Decision Processes*, North-Holland, Amsterdam.
- Suppes, Patrick 1975. "Approximate Probability and Expectation of Gambles", *Erkenntnis*, 9, pp. 153-61.

- Swijtink, Zeno, G. 1987. "The Objectification of Observation: Measurement and Statistical Methods in the Nineteenth Century", in Lorenz Krüger, Lorraine J. Daston and Michael Heidelberger (eds.), *The Probabilistic Revolution*, The MIT Press, Cambridge, Mass., 1987, pp. 261-286.
- Thompson, Michael, and Michael Warburton 1985. "Decision Making Under Contradictory Certainties: How to Save the Himalayas When You Can't Find Out What's Wrong With Them", *Journal of Applied Systems Analysis*, Vol. 12, pp. 3-34.
- Townsend, Robert 1970. *Up the Organization: How to stop the corporation from stifling people and strangling profits*, Knopf, New York.
- Turner, B. A. 1978. *Man-made Disasters*. Wykenham, London.
- Tversky, Amos 1972. "Elimination by Aspects: A Theory of Choice", *Psychological Review*, 79, pp. 281-99.
- 1975. "A Critique of Expected Utility Theory: Descriptive and Normative Considerations", *Erkenntnis*, 9, pp. 163-73.
- Tversky, Amos, and Daniel Kahneman 1974. "Judgment Under Uncertainty", *Science*, 185, 1124-1131.
- and —— 1981. "The Framing of Decisions and the Psychology of Choice", *Science*, vol. 211, January 30, pp. 453-8.
- and —— 1982. "Judgments of and by representativeness", in Kahneman, Daniel, Paul Slovic, and Amos Tversky (eds.), *Judgments under uncertainty: Heuristics and biases*, Cambridge University Press, Cambridge.
- Udwadia, F. E., and I. I. Mitroff 1991. "Crisis Management and the Organizational Mind: Multiple Models for Crisis Management from Field Data", *Technological Forecasting and Social Change*, Vol. 40, No. 1, Aug. 1991, pp. 33-52.
- Van Dijk, J.W. Asie 1991. "Foresight Studies: A New Approach in Anticipatory Policy Making in the Netherlands", *Technological Forecasting and Social Change*, Vol. 40, No. 3, Nov. 1991, pp. 223-234.
- Venttsel, Ye. S. 1980. *Elements of Game Theory*, (Translated from the Russian by Vladimir Shokurov), Mir Publishers, Moscow.
- Vickers, J. M. 1975. "Utility and Its Ambiguities", *Erkenntnis*, 9, pp. 287-311

- Volkema, Roger J., Jaime L. Manzano and James Gaglione 1991.
 "Developing a Cost-Benefit Analysis Process for Small-to-Moderate
 Expenditure Decisions" *Technological Forecasting and Social Change*,
 Vol. 40, No. 1, Aug. 1991, pp. 53-71
- Von Neumann, J., and O. Morgenstern 1944. *Theory of Games and
 Economic Behavior*, Princeton University Press, Princeton
- and — 1947. *Theory of Games and Economic Behavior*, 2nd edition,
 Princeton University Press, Princeton [Expected Utility Hypothesis
 included for the first time in 2nd edition].
- and — 1953. *Theory of Games and Economic Behavior*, 3rd edition,
 Princeton University Press, Princeton.
- Wack, Pierre 1985a. "Scenarios: uncharted waters ahead", *Harvard
 Business Review*, Sep.-Oct., pp. 73-89.
- 1985b. "Scenarios: shooting the rapids", *Harvard Business Review*,
 Nov.-Dec., pp. 139-150.
- Weaver, Warren 1968. "Science and complexity", *American Sci.*, 36, pp.
 536-544.
- Weinberg, G. 1975. *Introduction to General Systems Thinking*, John Wiley,
 New York.
- Wildavsky, A. 1985. "Trial without error: Anticipation vs. resilience as
 strategies for risk reduction." In M. Maxey and R. Kuhn (eds.)
Regulatory Reform: New Vision or Old Curse, Praeger, New York.
- Wilson, D., and A. Wilson 1970. *Toward the Institutionalization of Change*,
 Working Paper 11, Institute for the Future, Middletown, CT.
- Winkler, R. L. 1982. "Information and modeling in risk assessment", in H.
 Kunreuther (ed.) *Risk: A Seminar Serie*, IIASA, Laxenburg, Austria.
- Winter, Sidney G. 1964. "Economic 'Natural Selection' and the theory of
 the Firm", *Yale Economic Essays* 4, pp. 225-72.
- 1971. "Satisficing, Selection, and the Innovating Remnant", *Quarterly
 Journal of Economics*, 85, pp. 237-61.
- 1974. "Neoclassical vs. Evolutionary Theories of Economic Growth:
 Critique and Prospectus", *The Economic Journal*, vol. 84, No. 336
 (Dec. 1974), pp. 886-905.

- Yaari, M. E. 1987. "The Dual Theory of Choice Under Risk", *Econometrica*, 55, 1, pp. 95-115.
- Zadeh, Lofti A. 1973. "Outline of a new approach to the analysis of complex systems and decision processes", *IEEE Trans. Syst. Man, Cybernetics*, SMC-1, pp. 213-228.
- Zadeh, Lofti A. 1978. "Fuzzy sets as a basis for a theory of possibility", *Fuzzy Sets and Systems*, 1, pp. 3-28.
- Zeeman, E. C. 1976. "Catastrophe Theory", *Scientific American*, April, pp. 65-83.
- Zeleny, Milan 1982. *Multiple Criterion Decision Making*, McGraw-Hill, New York.
- 1984. "On the (ir)relevancy of fuzzy sets theories", *Human Systems Management*, 4, pp. 301-306.
- Zentner, René D. 1982. "Scenarios, Past, Present and Future", *Long Range Planning*, Vol. 15, June, pp. 12-20.

VII.2. Quotations without specific citations

- Bacon, Francis "If a man will begin with certainties..." — Columbia
- Barnfield, Richard "Nothing is more certain..." — Oxford
- Chomsky, Noam "As soon as questions..." — Oxford
- Cicero, Marcus Tullius "Where others say..." — Classical
- Clarke, Arthur C. "Any sufficiently..." — Columbia
- Confucius "To know that we know ..." [Russo and Schoemaker 1992]
- Gall, John "Generalized Uncertainty Principle: Complex..." (Gall 1975:40)
- Getty, J. Paul "Without the element of uncertainty..." — unknown
- Greene, Graham "When we are not sure..." — Columbia
- Freund, J. D. "[Statistics includes] the entire science..." [Freund 1962]
- Horace "Carpe diem..." — Bartlett's
- Hoshi, Kensho "If life were eternal..." — Thoughts

- Howe, Edgar Watson "A reasonable probability ..." (cited in Colin Dexter 1992. *The Way Through the Woods*, Macmillan, London)
- Hoyle, Edmond "When in doubt, win the trick." — Humorous
- Hume, David "Reason is and ought to be the slave..." — Columbia
- James, William "There is no more miserable..." — Penguin
- Jung, Carl Gustav "The great decisions..." — Bartlett's
- Keats, John "At once it struck..." — Bartlett's
- Knobloch, Eberhard "The future is never concretely determined..." (paraphrasing E. Borel, in Knobloch 1987:229)
- Kopp, Sheldon "All important decisions..." (Kopp 1972:166-67)
- Laplace "Given for one instant ..." — Bartlett's
- March and Simon "In the case of uncertainty..." [Smithson 1988:264]
- Mill, John Stuart "There is no such thing..." — Bartlett's
- Moore, George "The difficulty in life is the choice." — Bartlett's
- Russell, Bertrand "The fundamental cause..." — Columbia
- "Dogmatism and skepticism ..." "Philosophy for Laymen", *Unpopular Essays*, 1950, p. 361
- Shackle, G. L. S. "In a world..." (Shackle 1961:43)
- Voltaire, François-Marie Arouet "Uncertainty is an uncomfortable..." — [P.R. Stokke et al. 1991]
- Wilde, Oscar "I can stand brute force..." — Humorous
- Zeleny, Milan "There is no single..." [Zeleny 1982:xviii]

American: Bohle, Bruce (ed.) 1967. *The Home Book of American Quotations*, Dodd, Mead and Company, New York

Bartlett's: Beck, Emily Morison (ed.) 1980. *John Bartlett's Familiar Quotations*, Little, Brown and Company, Boston

- Classical: Harbottle, Thomas Benfield 1958. *Dictionary of Quotations (Classical)*, 2nd Ed., Frederick Ungar Publishing, New York
- Columbia: Andrews, Robert (ed.) 1990. *The Concise Columbia Dictionary of Quotations*, Columbia University Press, New York
- Humorous: Adams, A. K. 1969. *The Home Book of Humorous Quotations*, Dodd, Mead and Company, New York
- Oxford: Oxford University 1980. *The Oxford Dictionary of Quotations*, 3rd Ed., Oxford University Press, Oxford
- Penguin: Cohen, J. M. and M. J. Cohen (eds.) 1962. *The Penguin Dictionary of Quotations*, Atheneum, New York
- Thoughts: Seldes, George (ed.) 1985. *The Great Thoughts*, Ballantine Books, New York

VIII. GLOSSARY

actor: a party able to make (usually conscious) decisions, such as a human being, a committee, a firm, etc.

agent: an actor

ampliative: enlarging or extending; ampliative reasoning: “problems involving reasoning in which conclusions are not entailed in the given premises” (Klir 1992a:70).

aMW: average megawatt — a quantity of energy equivalent to the generation (or consumption) of 1 MW (1,000,000 watts) constantly throughout a year. The shape or pattern, however, is not assumed to be constant; it could be variable, or a rate of 2,000,000 watts for half a year, etc.

black box: a model in which an observer can see the inputs to the box, and the outputs from it, but cannot see “inside” the box to discern specifically how the box produces the outputs from the inputs. The behavior is known, but the reasons for it, the inner workings, are not.

BPA: the Bonneville Power Administration, a federal agency created to market the power produced by federal dams on the Columbia River. It sells wholesale power to other utilities and a small number of industrial customers.

cogeneration: adding electric generating capability to a system that is or would be already producing heat. Many industrial processes, such as pulp and paper plants, or thermal enhancement of oil recovery, produce “waste heat”, which can be used to drive a steam turbine to co-generate electricity.

dependent variable: in functions or causal explanations, a variable that is determined by (depends on) values of other variables; an output variable. E.g. z in “ $z = 2x - 3y + 3$ ”. Contrasts with *independent variable*.

discount rate: a percentage rate, usually annual, that indicates how much lower the value (or importance) today is a unit of some commodity would occur (or be received or paid, etc.) one year from today. For example, the right to be paid \$1 a year from now is generally worth somewhat less than \$1 today; this is represented as a positive

discount rate, often equated to the cost of capital for profit-seeking firms.

dominate: one alternative dominates another if it is at least as good as the second on all dimensions of comparison and better than the second on at least one dimension.

DSI: direct service industry — an industrial customer purchasing energy directly from BPA.

EUH: expected utility hypothesis — the set of axioms and proof originally formulated by von Neumann and Morgenstern that provided a justification for adopting maximization of expected utility (MEU) as the leading decision rule for decisions under uncertainty (technically, decisions under *risk*).

ex ante: the *ex ante* perspective for evaluating the quality of a decision considers only what a decision-maker could have known before (“ante”) making a risky decision, most particularly before observing which state of the world will eventuate.

ex post: the *ex post* perspective is from a point after the results of a risky choice are known, i.e. hind sight. Elster, Dror and Lovell argue that from the *ex post* perspective you can judge whether a decision was *lucky* but not whether it was *well made*.

false certainty (false uncertainty): when we feel certain (uncertain) but actually do not (do) have enough information to know; contrasted

Table 19. Certainty vs. uncertainty

	We have adequate information	We have inadequate information
We think that we know	<i>True certainty</i>	<i>False certainty (Overconfidence)</i>
We think that we don't know	<i>False uncertainty (Underconfidence)</i>	<i>True uncertainty</i>

with *true certainty* and *true uncertainty* as in table 19.

fixed cost: the portion of a cost, such as that of building and operating a generating plant, that does not vary with the level of operation (i.e. the costs are *fixed*). See also *variable costs* below.

frequentist: believer in probabilities based on long-run frequency; see *objectivist*. Contrasts with *subjectivist*.

fuzzy: imprecise — applies to both qualitative descriptions, e.g. “the man is

tall”, and quantitative, e.g. “the man is seven feet tall”. In neither example is the man’s height specified exactly.

fuzzy gamble: a gamble (choice with more than one possible outcome) where the payoffs, odds, options, or other aspects are uncertain or “fuzzy”. Not only is the outcome not known, but even the specifics of the gamble are not, even cannot be, known well (from Y. Dror).

gamble: an option open to a decision-maker which has more than one possible outcome, usually described probabilistically. For example, $(A, .75; B, .25)$ indicates a gamble where outcome A has a 75% chance of occurrence, and outcome B has a 25% chance. A gamble denoted $(A, .25)$ indicates a 25% chance of a payoff of A and a 75% chance of no payoff (and no penalty either).

GDM: a Generalized Decision Model, including more decision factors about which a decision-maker might be uncertain than the SDM (Standard Decision Model). See section III.2.10

independent variable: in functions or causal explanations, a factor that is allowed to vary, not determined by the other variables; an input variable. E.g. x or y in “ $z = 2x - 3y + 3$ ”. Contrasts with *dependent variable*.

KWh: kilowatt-hour — 1,000 watts generated (or consumed) for a period of one hour.

max(): a function whose result is equal to the largest (maximum) of its arguments, e.g. $\max(-2, 1, 3) = 3$.

maximax: a decision rule commonly invoked in game theoretic analyses: for each option, identify the best (maximum) possible outcome; select the option with the best best-case outcome (the *maxi* mimum of the *maximum* outcomes).

maximin: a decision rule commonly invoked in game theoretic analyses: for each option, identify the worst (minimum) possible outcome; select the option with the best worst-case outcome (the *maxi* mimum of the *minimum* outcomes).

MEU: maximization of expected utility — this is the decision principle that von Neumann and Morgenstern’s Expected Utility Hypothesis (EUH) justifies. Each decision option may have more than one outcome, depending on the state of the world that eventuates. The utility of each outcome is weighted by the probability of occurrence of the state of the world associated with it — this produces, for each option, an *expected utility*. The claim of EUH is that the best choice is the one

with the *maximum expected utility*.

mill: monetary unit equal to 1/1000 of a U.S. Dollar (in contrast, “mil” is 1/1000 of an inch).

mill/KWh: mills per kilowatt-hour, equivalent to dollars per megawatt-hour — common unit for describing rates charged for energy sales.

min(): a function whose result is equal to the smallest (minimum) of its arguments, e.g. $\min(-2, 1, 3) = -2$.

minimax regret: a variant of the maximin decision rule (minimax is the inverse of maximin: if options are scored in *costs* instead of *benefits*, using minimax produces the same result that maximin on benefits would). For minimax regret: for each state of the world, identify the best possible result over all options. Then for each option, identify the largest possible *regret* across the states of the world, which is determined by the difference between the outcome of that option under a particular state of the world and the best possible result for that state of the world. Then select the option with the smallest worst-case regret (the *minimum* of the *maximum* regret).

MW: megawatt — 1,000,000 watts.

MWh: megawatt-hour — energy equivalent to 1,000,000 watts generated (or consumed) for a period of one hour, or 1,000 watts for a period of 1,000 hours, etc.

nominal dollars: contrasted with *real dollars*. Real dollars from more than one year have been translated (by discounting) to a common year for purposes of comparison to eliminate the effect of inflation. Nominal dollars have not been translated; the value of a nominal dollar depends on the year in which it is exchanged, and comparisons of nominal dollars from different years are misleading unless the effects of inflation are accounted for. For example, if an item cost \$1 in both 1970 and 1971, and the inflation rate from 1970 to 1971 was 10%, the price in nominal dollars was constant, but the price in real dollars declined: in 1970, the price in real 1970 dollars was 1.00, but in 1971, after adjusting for the decline in the value of a dollar, was only 0.90 in real 1970 dollars.

N.P.V.: net present value: the sum of costs (and revenues) over a period of time, with each year's cost discounted to reflect the belief that a dollar gained or lost in a later time is not quite as important as a dollar gained or lost now.

objectivist: a person who maintains that valid probabilities must be based

on long-run frequency of occurrence, and that “mere” subjective estimates of the likelihood of occurrence of an event are inappropriate for use in expected utility calculations. Also called *frequentist*. Contrasts with *subjectivist*.

optimal: a solution that is known to be *the best* according to some criterion. Contrasts with *robust*.

optimize: to seek an optimal answer, the best possible (according to specific standards). Contrasts with *satisfice*.

PNW: the Pacific Northwest region — Washington, Oregon, Idaho, and western Montana.

preferences over: a person is said to have preferences over e.g. consequences if the person prefers some possible consequences over others. If a decision-maker has no preferences over one dimension of possible consequence, then differences along that dimension between options should not affect the decision-maker's choice.

P.U.C.: Public Utility Commission, a body charged with regulating aspects of utility operations within a state.

real dollars: see definition of nominal dollars.

reservation price: the lowest (highest) price a seller (buyer) would be willing to accept; important to keep secret during negotiations.

risk averse: aversion to risk is typically assumed when a decision-maker prefers a sure payoff to a gamble that has an equal or even a larger expected value. For example, preferring a sure \$1.00 to a fifty-fifty gamble on \$2.50, where the latter option has an expected value of \$1.25. Contrasts with *risk neutral* or *risk prone*. Associated with a concave *utility curve*.

risk neutral: neither *risk averse* nor *risk prone*. Considers gambles to be as valuable as sure payoffs equal to the expected value of the gamble.

risk prone, risk seeking: prefers a gamble to a sure payoff equal to or even greater than the expected value of the gamble. For instance, willing to pay \$1.00 for a 1% chance of winning \$50, a gamble with an expected value of \$0.50. Associated with a convex utility curve.

risky: having the possibility of harm. As used here, riskiness is a joint function of the likelihood and the magnitude of the harm; for example, compared to an action with a one in a hundred chance of breaking a limb, an action with a one in fifty chance of breaking a limb would be

riskier, and so would an action with a one in a hundred chance of death.

risky choice, risky decision: a gamble.

robust: a solution that is *good* under a wide variety of circumstances, less sensitive to deviations from predicted conditions. Contrasts with *optimal*.

satisfice: to choose an option that is good enough, perhaps before all options have been evaluated. Contrasts with *optimize*.

scenario analysis: The essence of Scenario Analysis is the consideration of at least two significantly different scenarios, i.e. *possible* future environments or worlds. Options (or strategies, etc.) are then assessed on the basis of their general level of anticipated success across the range of scenarios, not on the basis of their estimated optimality in the context of a single, “most likely” state of the world.

SDM: the Standard Model of Decisions that underlies traditional decisional analysis in which a decision-maker chooses the option that maximizes his utility. See section II.3.7.

SEU: subjective expected utility — this is the principle of maximization of expected utility, modified by the subjectivists (e.g. Savage) to permit the use of subjective probabilities (of the occurrence of different future states of the world) in lieu of von Neumann and Morgenstern’s insistence on long-run frequency probabilities in calculating the expected utility of each option.

stakeholder: a person or organization who has some kind of interest or *stake* in the consequences of actions under consideration. For example, the set of stakeholders in a possible closing of a nuclear power plant include the owners of the plant, their employees, people in the community affected by levels of employment, people concerned about the risk from nuclear operations, consumers of electricity, alternative suppliers of electricity, and many others.

subjectivist: a person who believes that when probabilities based on long-run frequency are not available, subjective estimates of the likelihood of occurrence of a particular event is the best information we can have on that likelihood, and it is therefore legitimate to use such an estimate in expected utility calculations. Subjectivists hold that all probabilities are essentially opinions, and that subjective estimates have full validity. Contrasts with *objectivist*, *frequentist*.

true certainty (true uncertainty): see *false certainty (false uncertainty)*.

utile: an imaginary or postulated unit of the utility or desirability of an outcome. An outcome measured at 2 utiles is twice as attractive as one measured at 1 utile.

utility: the quantity that people are assumed in the SDM to optimize when they make rational choices (see also *utile* above).

utility curve: a curve representing the desirability of varying quantities of some other variable, often money. In figure 47, a concave utility curve is shown. The curved line shows how the utility or attractiveness of monetary outcome changes with the quantity of money. In a concave utility curve, the marginal utility of money declines: the millionth additional dollar is not as desirable as the first. A convex utility curve is curved the opposite way.

Figure 47. An example utility curve

variable cost: the portion of a item's cost that varies with the way the item is used. For resources that generate electricity, the *variable cost* is typically stated as the additional expense incurred in generating an additional KWh of energy; contrasted with *fixed costs* (see above), which are constant and not affected by the actual pattern of operation. Fuel cost is usually the major part of the variable cost; construction costs comprise the bulk of the fixed cost. There are both fixed and variable components of the operation and maintenance cost.