6-27-1978

"Environmental Simulation and Long-Term Planning"

Michael A. Arbib

Michael A. Arbib
"Environmental Simulation and Long-Term Planning"
Portland State University
June 27, 1978

PSU Library Special Collections and University Archives
Oregon Public Speakers Collection
http://archives.pdx.edu/ds/psu/11430
Transcribed by Julian Younger, July 11 - July 18, 2020
Audited by Carolee Harrison, September 2020

HOST: The Portland State chapter of Sigma Xi, the Scientific Research Society of North America, organized and planned by Dr. Michael Fiasca in our School of Education, who is here and is an instructor. I think most met you yesterday. From time to time, maybe this afternoon, we'll get you two to stand up again, and let 'em see who helps make things happen here at Portland State, and involved in trying to bring more scientific information to the public. We hope that you will announce this to other people that may be interested, it's—as I say—open to the public; we weren't able to announce this much earlier than last Thursday because of plans for room facilities, and now we hope we can fill it before the end of the week. As you know, there is a lecture tomorrow at 9 and 1 by Harrison Brown, who is well-known to many in his writings on food and energy, as well. Also on Thursday will be the 9 and 1:00 lecture; this will be Christopher Salter, speaking on urbanization.

Some of the disorganization that you seem to… probably thought was occurring, wasn't really disorganization, it was something that was in the schedule, it's only that it was delayed a bit… [laughs] and I'm sure you'll hear a little bit about the reason for it, and it involves O'Hare Airport, since our speaker this morning, Michael Arbib, professor of Systems Neuroscience at the University of Massachusetts in Amherst, had to come through Chicago to get here, and the 6-hour delay caused some difficulty in getting all of these sequences worked together.

In order to just keep on moving, he's the author of *Cybernetic Society*, he's published many other papers and other books in variety of forms, and has been an important lecturer in this general subject recently… well, travels a great deal internationally, also. He's from Australia. And there are a lot of other good things that could be said, but during the day you may get to visit and find out more. I present Dr. Michael Arbib.

MICHAEL ARBIB: Thank you. [pause] Well, hello. There is something dispersing about arriving at an auditorium to find that the view graphs you were waiting for are now slides, and that the slides are in random order, and that 10% are missing. But, other than that everything seems in good order. [laughter] The subject is environmental simulation in long-term planning. You may wonder a bit about… [laughter] about the ability of this group to inform you about planning. The hope is that you will learn from our mistakes.

I'm going to be emphasizing, both this morning and this afternoon, the impact that computers will have. This morning the emphasis will be on planning, and this afternoon the emphasis will be on the way in which citizens may hope to participate in the governmental process and the impact that computers may have on that. It's perhaps appropriate, then, to tell you what happened to me yesterday. I was on a plane, waiting to take off from Chicago for Portland when they discovered that there was some mechanical failure, and so they had to cancel that flight. They found another plane, but unfortunately the plane was smaller, so about 40 of us were left there, watching it take off. And as several of us were there, rather dejected, the man next to me turned and said, "It's those damn computers." I said, "I don't think it's the computers, after all, it was an engine failure, wasn't it?" and he said, "Well yes, I suppose so."

So, what we have, though, from that little anecdote, is a sort of two-fold thing. Firstly, that many people have a view of the computer which is inaccurate and tends to make it a scapegoat. Anything that goes wrong in modern society must be the computer's fault, which I would like to dispel in discussion with you. And, the second aspect is that a lot of people don't really think. I mean, he knew very well that the problem was simply that the plane we were on didn't work, and that the replacement was a smaller plane, but it was much easier to spout out a little pat formula than it was to really think about the facts and their relationships. Now, a lot of you are teachers, and so as I go along talking about these subjects, I will not only try to give you an overview of what's going on with them, but also try to think about what they mean for education, and to what extent the concepts we're discussing here should become part of the everyday knowledge of not just the few of you who have chosen to come here, but of every right-thinking citizen.

“A computer in every home.” Well, that’s not a joke. You may have seen the advertisements for Radio Shack’s not very good home computer, or the somewhat better PET computer. These are now selling for 6 or 700 dollars, and if you remember what happened with pocket calculators where they plunged from about 150 dollars to about $6.95, the odds are that in a few years our home computer will be a lot more powerful than it is today and cost a couple hundred of dollars. So that computing power is going to become less and less something out there and just part of the home. And, a person I’ve corresponded with from Australia made the very interesting observation that in education in the future, the problem will probably be the teachers and not the students, because the teachers will still be scared of computers while the students have been used to playing with the home computer from the age of three and consider them a perfectly natural part of the environment.

Okay, so much for prologue. Let’s… the trouble with slides is that one tends to turn out lights, and the trouble with turning out lights is people tend to go to sleep. So, we’re going to see whether we can keep enough light on… can we turn off the front lights, but leave on some lights at the back? Now, is that visible enough? I think we need more light at the back and less light at the front, if that’s possible. [pause] Isn’t this exciting? All right, now I think that’s bright enough for you to stay awake, and dark enough up here for you to see what’s going on. And the pointer works, isn’t this wonderful? It’s all going to be all right. I needn’t have worried. Okay. We’re saying that we have an increasingly complex world, and a world that we no longer take for granted. If we want to sort of look back in history, it’s not so long ago that people accepted that flood and pestilence and famine were just part of the ordained order of things and there was nothing we could do about it. But we know now that if we are careful about the balance of little creatures and chemicals, we can do something about pestilence and famine, and we know that if we have good satellites feeding large computers, we can make weather predictions that will allow us to take many precautions against flood. So that, we’re moving into an age in which more and more we understand that we can understand. But, as we understand more, we come to see more and more complexity in the world around us, so that we can’t keep straight in our own heads the many interrelationships. So, we now use the computer to try and capture those relationships for us.

Now, once upon a time we used models in the sense of scale models, like a scale model airplane in a wind tunnel: let the wind stream by it in the wind tunnel, see whether it was stable, see whether it would crash. But finally we realized that with the increasing power of computers, we didn’t always need to build that scale model, at least straight away, when we could build our model in the form of mathematical equations, represent those equations in the computer, and look at what would happen in different circumstances. So, this idea of computer simulation is when we have a complex system, whether it’s an airplane or an economic system or an ecological system, we try to set down what we know in a precise form which can be represented as a program in the computer. Now, a program is a sequence of instructions that says, “Take the data, do this to it, test what happens, do that to it, test what happens,” and so on until we come up with some results. So, we can run this program to see how the system will behave for different choices of input and parameters; or we can carry out a control analysis where we’ll find out from the program what sort of inputs will yield some desired behavior. And the point I want to stress is that a complicated system requires a very complicated program. Let’s take a very closed ecosystem, Skylab, the satellite that went up and was meant to be a life support system for three human beings for months at a time. The program for that required the writing of something like 400,000 instructions, and that gives you some idea of the complexity that we’ll have to face if we’re to try and capture, meaningfully, the complexity of the world around us.

Okay, let’s have a look at the next slide. Is there some way I… oh, I press this? Okay. This is actually a very exciting process, because you see random slides are missing, so that I never know which slide will be next, but we’ll do what we can as we go along. All right. Here’s… let’s first look at this curve, which would correspond to a rather simple model of population growth. Let’s say the idea is that every family has, let’s say, four children. That would be the simplest case. Then that would mean that after one generation, there would be twice as many people, because two people on the average in the first generation would yield four people on the average in the second generation. And the four people in the second generation would yield eight people in the third generation, and so each generation there would be a doubling. The point is that many population growth processes can be done in this way. There is some fixed time interval, whether it’s five days or a generation or a hundred years will depend on the particular system, but where we can see a time which doubles and doubles and doubles again. The missing slide showed nicely the doubling, but we can sort of see it there.

Now. So that might be our first model of a population growth situation. We see that each generation there are slightly more offspring than parents, and so there is exponential growth. But, in any real situation, we have to recognize the fact that there are only finite resources, only so many mouths can be fed. So we talk about carrying capacity. Think of a community with just so many fields, and a particular level of agriculture, there’s a limit to how much food they can produce, and even if people go onto a starvation diet and so on, there’s still a fixed upper bound on the number of people who can live. So, one thing we might imagine as a possible effect of limited resources is that we start with exponential growth, as we’re well away from the carrying capacity, but as we begin to get up to the carrying capacity, growth slows down, more people die, for example, if we’re looking at a population model, and finally we balance out here.

Now, in fact, though, it's not clear *a priori* whether that is really what would really happen. Here's another possibility. Suppose that in fact, what happens is that people keep having children, and let's assume that to start with the children don't need so much food, so we get to this stage where, in fact, the population in numbers is much greater than the carrying capacity, because at the first, there are a lot of people who aren't eating very much, the children; but now the children grow up and start having their own children and suddenly, the number is disastrously above carrying capacity. So much so that there are food riots and people start tearing up the crops to eat the roots; and in fact so much so that not only do people start starving at a horrifying rate, and we see a precipitous drop in the population, but in fact they have destroyed the agricultural base in their mad rush to find any source of food, and the population goes to extinction. Now, the point I'm making is that here we have two quite—you know, from what most of us know about the interaction of population and agriculture and so on—two quite plausible views of what could happen if the population grew without limit. Either it will in the end just sort of naturally settle down to an appropriate level matched to the available resources, or a tragedy of this kind, where we go so far above our resources we kill the base on which we depend, and population goes from there.

And, the idea is that we want to start thinking very carefully about how we build models in such a way that we don't just say, "Well, it seems to me that if you have a lot of population eventually there's going to be so much you're going to run out of resources," but instead you say, "All right, here's a model; let's set it up carefully, let's run it on the computer, let's look at its predictions, and now looking at these predictions, let's compare them with some actual studies in the field." Now whether the studies are that we look at some animal populations on an island and see to what extent they might give us an idea for the way in which humans might respond, or whether we look at some historical data about what happened during a famine in some particular region of a country some time in the past. Whichever way it is, we can compare the predictions with the outcomes, and on this basis we say, "All right, the model wasn't quite right." And so on that basis… I guess this arrow should really come up to here. Is this focused well enough for all of you? Doesn't seem quite right from my corner but, it'll do.

So, we come back and we revise the model. So here's the idea, that in a complex situation, we'll start with an initial model which captures what we know at the time, and then as we start using the model, to some extent we'll rely on the predictions but we'll also spend part of our time comparing the predictions, and saying, "All right, well, we thought the model would happen, that seemed to work all right, but this other stuff, that was really wrong. Let's revise the model." And so there's this idea of trying to accumulate a body of scientific knowledge which can be run in the form of a computer program to always give us, if you will, our best bet. But it's going to be a continual process as we try to capture more accurately the way in which the world runs.

Ah yes, the well-known “sceniaro,” known by some as a “scenario.” All right. Now, suppose though, that we are at a stage where we’ve worked hard and we’ve built up a computer program that gives us a fairly good feel for a certain range of issues, let’s say economic investment and its impact on industrial growth, or the way in which we distribute resources between fertilizers and insects which can eat other insects which have been known to eat the crops, and so on. And now we want to start seeing what we should do, and one way we can do this is all right, let’s assume our model works for now, and now, using this famous “sceniaro,” let’s specify some of the things we think might happen, the sort of decisions we might make: how much money do we want to put into fertilizer? What is going to be the level of demand for different crops? What sort of rainfall can we expect to happen? And on the basis of this we can run the model for each scenario and observe the trends. So we can say all right now, if we’re scared about a draught coming up in the next five years, perhaps we should have this crop mix, whereas if we can be confident that the rainfall will be at a different level, then there’s going to be a real problem with this pest and so a lot of our resources had better go in for controlling that. And on that basis, we may have to… well, this is a bit too dramatic. In general, with the things I’ve been talking about, this is just a matter of where we put our dollars in terms of our actions, but in other cases it may suggest a rather dramatic change. We may have found that we had been thinking in terms of a high level of chemical usage in our agriculture and that all our planning had been based on that, and then we find ten years down the line the essential level of poisoning of the soul… soil, or soul as the case may be, that we have to change. But as I say, always remembering that if the model is correct, and so life is like that, ain’t it? That we make decisions and hopefully we make the right ones, but often we don’t and then we try to learn from our mistakes and bury our dead and start all over again.

Okay. So one of the things that we might have is often that rather than sort of make a prediction right into the future, we might say well look, let’s try out a particular crop mix, a particular fertilizer mix, run the simulation for a while, see what happens, and then if on the basis of what we see, we the user will then try out various things for, let’s say, the second year of the agricultural plan, and so on, round and round. So that… [groans] yes. Now, as you see here quite clearly, the user… in this case, we’re thinking about, in this interactive mode, we’re thinking that we’ll provide a short-term scenario. If things look good we keep on going on, we’re saying this is a good way for us to plan the first year, let’s keep going; but if not, we can try again with a different start. We can experiment, in other words, and see, well perhaps have to do some things that in the short run don’t look so good, I mean the case of fertilizers in agriculture might be a good point, where it looks like you’re really dropping your return by not using fertilizers, but you find that if you go more than five years into the future, it may be that the mix of the fertilizers poisons the soil, so that the overall return from the farmland is reduced.

Or, to take another very topical example: whaling. [wails like a baby] Waaah! No, not that. W-H whaling. Where, countries like Japan and Russia that have a large whaling industry are very reluctant to put limits on the catch, because they observe that this will yield very severe dislocations of an important segment of their economy; but of course what you and I would argue is that this is a short-sighted thing, this is a short-term maximum with a long-term loss, because they will end up by driving the whales to extinction, and then their whole industry will crash, and so that what they should be doing now is working out a… for those of you who want to save the whales at all costs, what I am going to say is unacceptable, but for those of you who eat meat, you may accept the idea that some level of culling of animals is appropriate, and that what one has to do therefore is find the appropriate level at which the various animals concerned, in this case the whales, can continue their existence but man can also continue to get a return from it. So these are the sort of examples where it's often very important to try and extend one's projection into the long term, to see whether we should make some short run sacrifices to get these long-term gains.

Now, notice of course, though, the paradox here, why political debate often gets exciting; because you can imagine, we keep emphasizing that our models are not perfect, and we keep emphasizing we're sort of in a trial-and-error way looking at some but not all of the alternatives, and the longer we go into the future, the more whatever inaccuracies there are in the model build up. It's just like compound interest, you know. 5% a year isn't very much but if you make a 5% error the first year and 5% on top of that the second year, within twenty years your error is 100%, which is pretty impressive, since the maximum error you could make is 50%. [laughter] All right.

Now, to get a little more specific, I want to talk with you for a while about some models of the interaction of population policy, agricultural policy, industrial policy, international transfer of funds, and so on, that have been developed under the sponsorship of the Club of Rome. Perhaps some of you will recall seeing some years ago, around '72, a book called "The Limits to Growth," and a sequel a few years later called "Mankind at the Turning Point." And these were attempts to build computer models of world trends in population and resources available to humans and standard of living and so on. And, what we'll learn from these models, I think, is firstly that we don't know as much as we should. Secondly, though, that the process of modelling is a very important and informative one. And I would think that the attempt to bring our knowledge of geography and economics and so on to this explicit form might well be very important in terms of education in the future. I mentioned earlier that the home computer is becoming a reality, and probably it will have roughly the time course that the color television set had, that when it first came in it was very expensive, and only a few people had color television; now, most sets are color. That sort of investment has become part of the way in which people treat things.

So, we can expect that the student of the future will have a home computer and will be using it for his homework, or her information retrieval, or what have you. And that therefore, very often from now on, our educational materials will not be books, but they will be programs to be run on the home computer. In other words, the student, rather than buying a textbook will buy a cassette, or a stack of cassettes, that will be clicked into the side of the computer and then instead of reading the pages one after another, you will have come on the screen something like this, with pictures as well as text, and which will not simply just sit there and let you read it, but will start asking you questions, or perhaps even generate simple movies up there to show you how things will change over time, so that we will have dynamic books. This is going to make an immense impact. Just as in treating books we can either treat them passively, or in other words we just sit there and turn the pages and let the words flow over our retinas, or we can sort of engage in an active process of saying, “Oh the silly so-and-so, of course that’s wrong,” and “Oh, that’s a good point,” and things of that kind. So it will be in the future, that many people will use computers in a very passive way, that playing games of tic-tac-toe and video games will be the thing they use their home computers for. For other people, there will be the opportunity to go much, much deeper in really thinking through mechanism, and really understanding how things work, because that understanding can be captured within a computer program that can then be run on the computer to explore its implications.

It does seem to me that if you give a student the challenge of writing a program, saying, “I’ve taught you something you think you understand; now write a program that performs correctly,” you have a way of understanding whether the students understand principles in a much better way than you do if all you do is give them a problem to go off and work by themselves and come back with an answer that you grade as either right or wrong. Going back to yesterday’s Chicago airport. A lot of people don’t really think through to the causal connections between things. The computer may be very important in raising our general level of awareness as well as providing new tools for the official planners, whether they be an industry or government.

So, let’s go on now and look at these rather preliminary, rather simple models commissioned for the Club of Rome. And let’s start with this question: What do you do with a curve for human population that goes like this? You may have your own personal answer at this stage, but let’s be polite about it. Well, we look at the trends and what we think at first, probably we say, “Oh, it's an exponential curve." Well, we remember that an exponential curve has a doubling time. We start off here and we say, "Gee, it seems to sort of double every 75 years," and then suddenly we see that in fact it's growing faster and faster and faster. Now, in fact, if you sort of did what we call extrapolation, trying to draw a smooth curve that continues this line, you find that it suddenly goes off to infinity around the year 2020, shall we say. And that's wrong, because we already know that things like carrying capacity are going to hold the population at some maximal level, and presumably long before we reach the carrying capacity of the planet Earth, other things will intervene.

So, what we're going to do is say, "Let's start looking more carefully at the underlying processes. Let's not just simply say, "Oh, well if we extrapolate the curve, we're going to have an infinite population on the planet within 100 years," or say, "Oh well, there's some fixed carrying capacity for the planet, so you know very well it will level off." Let's try and get inside—this is the idea of the computer approach—to get inside… or the systems approach, which the computer helps—is to get inside, look at the processes, look at how they interact, and then try to explore those interactions into the future to see what will happen in different circumstances. So, for example, here is part of, just a fragment, of the analysis done in the "Limits to Growth" study. It was a very crude study in that they just looked at world trade, and so there was no attempt to look at interactions between different countries. So when we say population level, that was the whole population of the whole world. Deaths per year: how many people were dying worldwide. Mortality: the rate at which people were dying worldwide. Pollution: the overall average level of pollution. Birth control, fertility, births per year, and so on. And it's very clear to you as you sit here that of course, we can't make very accurate generalizations if we just average over all the differences between the different parts of the world. So, later on we'll talk about attempts to break the models down, and try to get a better feeling for regional differences. But for now, let's just think about this model.

What we say here is: births per year, we have a positive effect on population level. Now, when I say positive, I don't mean that I have a value judgement, I'm just saying that the numbers are positively linked: the more births per year there are, plus: the more population there is, but also the more population there is, plus: the more births per year there are. And so if we just look at these two circles in their interaction, we have what we call positive feedback. Now, I'm not talking about that in the sort of awareness movement sense of, you know, you stroke my back and I'll stroke yours, or wherever you carry out your massage, but I just mean in the technical sense that there is a feedback interaction: the level here feeds back to the level here, feeds back to the level here, and the overall effect is positive, the more population the more births, the more population, the more births. Whereas of course, with death it's a negative feedback loop, that population level: the more the population there is, the more deaths there are, and deaths tend to stabilize the population. So, here is a positive feedback loop which tends to give a runaway condition, where we get what we call the population explosion, whereas here we can get increasing… we can begin to stabilize. Now, if we look at this loop, then, we see that one of the possible ways in which population can be controlled is by the fact that the more people there are the more pollution there is, the more pollution there is the greater the death rate, and thus the greater the deaths; and so if we look at this loop, here is one way of stopping the population explosion, but I think you'll agree it's not the way we want to go. In other words, if there are too many people, things will get lousy.

Now, we've already talked about other ways that it'll get lousy: there won't be enough food to go around, there won't be enough space for people; but here's another way. Pollution will start increasing the death rate: emphysema, lung cancer, all those good things. And that will keep the population down. Another possible loop is this one, which looks at birth control strategies, which brings down the birth rate, and thus decreases the positive contribution here—positive numerically, again I say, without a value judgement one way or the other. And so, what we have to do, then, is… you see, whereas this sort of thing here and this sort of thing here are things that we might think of as being what we could hope to study scientifically. How will pollution vary with population level? What will different types and levels of pollution do to affect the death rate? When we start talking about birth control, we're not talking about a natural phenomenon, we're talking about a political process, and as you know it's a very sensitive political process. Is abortion an acceptable method of birth control? Is the pill an acceptable method of birth control? Is abstinence an acceptable method of birth control? And different people will say no to all three.

So, it becomes, then… in a sense, there are two different levels here. You might see, there's the level here that says that if you set a certain birth rate, you can start using your computer model to say, "How can I figure out a birth rate here that will yield a stabilized population level and hold pollution at an acceptable level, too?" And, having done that sort of study, you can now say, "All right, now that I've set the birth rate in this model," now a completely different sort of consideration is, "Now, is that scenario politically feasible?" Can we, given the social and political and religious pressures, hope to realize that level?

So, what I'm trying to get at here is something that we'll keep coming back to: the subtle interaction between the mathematical model, which takes the best scientific data that we have and runs it on the computer to explore the impact of different policy decisions; and the fact that those policy decisions often open up a whole mare's nest… a phrase that I have no idea what it means. if anybody knows why a mare's nest is called a mare's nest, please see me afterwards—and then we have to just get into something that we can't model at the moment, in terms of these numbers. But, you know, our political judgement as to what is feasible, what sort of propaganda or public education do we need to bring about that sort of policy?

Okay. Let’s look at some of the computer output that came out of this sort of modelling in this Club of Rome study called “The Limits to Growth.” Is this… all right. Here’s an idea that… suppose this was the sort of gloomy prediction as to what would happen if we just went on the way we were. I say that we would keep using up resources. Meanwhile, as we keep using up resources, population would be growing. And for a while things would look good, there’d still be plenty of resources left, we’d eat more food, industry would go up… [pondering data] here, and this line here. But, fairly soon, around the turn of the century, things would begin to get grim because pollution, as a result of all this industrial growth and population growth, would rise to high levels. Meanwhile, resources would become less and less, and the result of these two effects would be that suddenly, the industrial system, and as a result, the population, would crash. And so somewhere between now and the year 2100, they predicted that there would be worldwide disaster as we depleted all our resources and finally, lost our whole industrial base and lost our resource base.

So, the next model said, well, suppose that in fact, resources didn’t run out. Now, how could we make that sort of assumption? Well, at the moment we have an increasing amount of nuclear power, and that may continue depending on how the crab and clam shell alliances get on with their work, but there is the feeling that even if there is not nuclear power, perhaps if we’re smart about it, we will be able to start using more and more solar energy, and also that ‘round the corner we may have access to fusion energy, where we sample the power of the H… what is it? Tame the power of the H-bomb, and use water as our fuel, and that sounds a pretty good way to go. So, if we really play our cards right, it may be that by the year 2000, we will have enough energy around that we can keep recycling our resources. That we can… we have only, for example, we know that the Earth is fairly thick, I mean there’s a lot of good stuff down there, but at the moment it’s too expensive to go more than a few miles below the surface. Now, if there’s no problem in energy… in having the energy to exploit everything that’s a hundred miles below the surface, let alone thousands of miles below the surface, we’re in great shape as far as resources are concerned. But then they point out that if that happens, you’ve still got the problem that pollution is going to get out of bounds as everything grows, and so now it’s not going to be that we run out of resources that kills the system, it’s going to be that the pollution levels kill the system, and so everything crashes around here. Can we be a bit more optimistic?

Okay, now this is one of… this is a textbook example of how not to make a slide. I can’t read it, can you? This little thing in the corner here? All right. So, this is the idea of… let’s see what happens if we try to stabilize population. Other words, if we do introduce a rational birth control policy that says, you know, somewhere we have to find a way of making it socially acceptable to keep population in check. And then you see that of course population levels out, but now, what's this? Industrial output still crumbles, and so the final model they advocated in this first study, was one in which… well, this is saying let's try and couple investment to industrial growth, so that instead of just letting capital get used up, we keep ensuring that new plant is built where it's needed, and so on, and then here comes this lovely situation here: resources are held in check, food stabilizes, industrial output stabilizes, population stabilizes, pollution stabilizes.

All right, everything's marvelous. We should be happy. Now, unfortunately a lot of people weren't happy for a number of reasons. Now one of those reasons is the one we've mentioned before: that it's an aggregated model. In other words it treats the whole world as one system, and doesn't begin to look at the situation where we have many different regions—developed regions, underdeveloped regions, oil-rich nations, oil-poor nations, and so on—and so that no model which fails to take account of these international relationships can really be a useful indication of the way things will go. There are other problems, and that is that the model, even though it's rather complicated, is still far too simple in that it doesn't involve enough really detailed knowledge of the relationships between agriculture and industry, for example. Also, many economists have attacked it by saying that, despite their inability to control inflation and unemployment, economists do know a lot, and that's not included in these models.

Now, it's going to be beyond the scope of both my detailed knowledge and our time to get into the ways in which we can build more and more complex models. I'm going to look up just one slightly more complicated model, the second report to the Club of Rome, with you in a few minutes. But, going back to our theme that what we should be thinking about here is not simply our own understanding of these processes, but what they mean for education. I think it's fairly clear that we're getting an idea that there are many subtle relationships that not very many of us understand deeply—well, not even that—have just even a superficial acquaintance with. And so I think that just going through this exercise makes very clear to us that it's going to have to be more and more a part of the educational process if we're going to be, in a meaningful way, citizens in a complex, technological society to get access to these sort of things, and the access will be, at one level, just understanding this sort of modelling, understanding more about the relationships between population growth and economic growth.

Also, I would think, the ability to play with models of this kind, to get on your home computer and just run some of these models. That when you're interested in something, whereas now you would turn to the encyclopedia and see a fixed article, and you must take whatever is in that article as the available information. Now you can actually start asking questions, and the power of the computer will be available to answer your questions, at least if they're fairly standard questions. I think it's fairly clear that the computer is no magic machine. It doesn't know everything. In a sense, it knows something: it knows what is in its data banks, it knows what is in its programs. So, all right. Here we see the questions we can get answered about different scenarios, about different projections of the future, with the knowledge as summarized in the "Limits to Growth" program. In the future we will summarize more and more knowledge in the form of such programs, and we will be able to ask more complicated questions, explore more complicated alternatives. And so whether we're doing this as expert planners or as citizens, at least we should be able to base our judgements of the political situation on a better understanding of what's going on, of the causes and the effects that are operative.

Okay. [pause] Mesarovic and Pestel, in the book "Mankind at the Turning Point," which was the second report to the Club of Rome, said, let's try and divide the world into regions, and in this way, try to more accurately analyze what's going on. So their regions were sort of North America, Latin America, Arab Africa, African Africa, China, the Soviet Union and Eastern Europe, Western Europe, and so on. Some interesting things. You can't see it in this picture, unfortunately, but for example, one of their regions comprised Israel, South Africa, and Australia and New Zealand. Now, that's strange. The principle was, I guess, everything which sort of approximated a Western European country that wasn't in Western Europe. But, it's fairly clear for those of you who have read the papers in the last five years, that conditions in the three places are somewhat different. And, that again gives you some idea of the problems that we have in modelling. Already, we're saying that to begin to look at some of the important trends in worldwide economic and ecological matters, we've split the world into ten regions, which will give us a much better understanding of the interplay than if we just aggregate into just one big thing of world population, world pollution, and so on. But already we're getting dissatisfied, and so how… there's going to be pressure to break it down further and further. In a while I'll come back to strategies for how we can sort of play this game of levels so that the people in Western Oregon, clamoring for their own regional model, will not force our understanding of the whole world to be at that level of detail, where every group of a million people is explored in great detail.

This is one of these slides I put up to scare people, mainly because it's so simple. [pause] Aha, this is nice, the "casual" model. This was meant to be the causal model, looking at the cause and effect relations, but the illustrators, knowing that we would not have time to study this in detail, [laughs] realized it would be the casual model we would look at together today. What this slide is meant to just show is that, all right, suppose that we've broken the world up into regions, what sort of things do we have to know in each region? There are going to be the general modelling of the environment, technological factors, interaction of population and economic factors, the way in which large special interest groups can influence the situation, and the way in which individual people can affect the groups, and thus affect the demo-economic stratum. A can of worms, in other words. And so, we keep on going with this problem of how do we aggregate at the right level? How much detail should we put in to make sure that the projections we're getting will tell us something, will help us make the right decision, and yet will not clutter us up with so much information that we're just swamped in a sea of figures.

And over here, we're making the distinction that we've already seen before in looking at the difference between the causal model, is the things like, well, the way in which the birth control setting of the fertility rate will affect population level, and the way in which population interacts to cause pollution to affect the death rate. These are things where we can hope that by careful analysis, we can come up with a good scientific understanding of the relationship between these processes. As distinct from the things up here, where we're concerned with the political processes, which set the parameters. In the example we looked at before, the birth control policy. And so, in addition to this causal model, the thing we can run on the computer, there's this stuff up here which is much more, at the moment, in the heads of people making political judgements. You know, what are the norms, what sort of things can the public tolerate? What sort of things do the public want? And how these can interact with these casual models.

Let's just look at a couple of the policy decisions that Mesarovic and Pestel examined in their extended model. Now, it does seem to me there's just too much information on this slide. I don't know how much you can read… is this readable? This stuff on the right hand side? Can you read that, sort of? Sorry, I'll tell you about it. Okay. In fact, most of the models that Mesarovic and Pestel develop in their book are not really ten-region models, but developed versus underdeveloped world models, and try to look at the differences between these two. And they were, in particular, very concerned about the way in which population is growing in the developing world.

In fact, as you know, in the rich nations such as this one, population seems to be levelling off or even declining. There are two theories about this. One is the electric light theory, which has the television theory as a special case, namely that there's something else to do in the evening; though my theory is that it's the cost of college education, and that is that the people who might otherwise not resort to birth control when they figure what it's gonna cost to put four kids through college decide they'll settle for two. Anyway, whatever the explanation, the fact is that in the advanced nations, population does seem to be levelling off. But in the underdeveloped world, it's going like mad, and if current trends were to continue, then we would be blowing through the ten-billion mark around the year 2025, and this is just way beyond supportable levels. It might be possible to feed that many people, it wouldn't be quite clear what that sort of life would be worth that could be supported at that sort of level.

It's interesting, of course, if you think over just what we've been saying, some of it seriously and some of it, apparently, not so seriously, the way in which so many factors interact. That birth control rates is not just a matter of external policy, it depends a lot on living standard and the sort of way in which people perceive their society. If you're a poor peasant and see that you need to have a lot of sons around to till the fields, to keep you in your old age, you're going to have a very different feeling from someone in an industrial society who knows their children won't do anything for them anyway.

So, in these various scenarios, what Mesarovic and Pestel asked just simply was: let's suppose we accept the urgency of bringing the fertility rate down to what we call an equilibrium fertility rate—in other words, one in which the birth rate equals the death rate—so that finally, as we see in all of these scenarios, we finally end up with constant population. What they really want to know is, what is the cost of delay? In other words, can we afford to wait a long time before we introduce good birth control measures in the developing world? And here are the ideas that if we sort of get going fairly quickly, we can level off to a fairly manageable population of six billion. Wait longer, it goes up to eight billion. Even longer, ten billion. The cost of delay of… these are sort of twenty-year differences here in the delay into really getting good population policy going. The costs can be horrendous in terms of human suffering and in the depletion of resources. Unfortunately, as some of you may know, while these simple mathematical predictions… projections seem to be quite compelling in their message that a rational birth control policy is absolutely vital for the developing world, it doesn't seem so to people in the developing world, who seem to believe that our concern is just part of a plot to ensure that there aren't too many other people to steal the resources from us. And so, to many Third World politicians, the feeling is that there is strength in numbers, and that there should be more of them, and that the rich people—the Europeans and the Japanese and the North Americans and so on—should stop telling them how many people they should have, and just start figuring out ways to get more resources out to the Third World, [quieter] which may be true.

Notice, again, a point I keep returning to, but it's very important: that the computer is really a rather small part of this whole story. The computer, as we've been seeing it here, is a tool to help us keep track of complexity as we need to juggle more issues, and to the extent that we can be precise about certain processes that impinge upon our decision making, then the computer can run scenarios for us or carry out control analyses for us to find out what are optimal decisions within certain parameters. But what is clear is that so many of these factors are political, and so many of the factors go beyond what we have precise knowledge of, so that we're using the computer as part of a complex decision making process. There is no possibility, I think, in the short run, anyway—and in the discussion period those of you who come may want to examine some of the science fiction possibilities—but there is certainly no possibility in the short run that we would want to just turn things over to the computer, and say "What should we do?" and follow the computer's judgements.

Another line of exploration was addressing just this concern, the concern of the underdeveloped nations, that part of their problem was that so much of the resources are tied up in the developed world, and that some form of capital transfer should be brought about to start giving more resources to the Third World. And, I guess what happens is that one finds that it's very hard to come up with a policy that could lead to anything like equality of the capital income between nations. Perhaps the best one can do is to head for five to one. Latin America, for example, at the moment is running about eight to one, in standard of living, on the average. In other words, the average person in North America has eight times as much capital goods as a person in Latin America, and for South Asia it's something like twenty times as much. So, if we took even five to one as a modest target, on the basis that not everybody needs two homes… did I say that? I meant two cars, but never mind. They found that that massive transfer of capital, of money and resources and technology, were necessary.

Unfortunately, the next slide is the one that isn't here, but what they found essentially was that if we sort of aim for the year 2025 as our target date, it would cost about seven thousand billion dollars to just ensure a five to one ratio between the developed and the underdeveloped nations. In other words, that much capital would have to be transferred to the underdeveloped nations. Now, here's the interesting thing: remember that number, since you can't stare at it on the screen, seven thousand billion. If, on the other hand, we said "Oh well, there's no hurry," and waited longer, and didn't start transferring these funds until the year 2000, their estimate is that it would cost then eleven billion. In other words, about half as much again: eleven thousand billion, about half as much again. And now here's the interesting point: if we decide to actually be generous and say, "Well, we're sorry for those poor people down there; let's make a big effort and let's try and get the necessary capital down to them," not by the year 2025 as we've just been discussing, but by the year 2000, they've found that it would cost much less, only 2.5 billion. And the point, of course, is that we get one of those positive feedback loops going, those regenerative loops. The earlier we get the money to help the underdeveloped nations develop, the sooner they have sufficient capital to have their own factories, and to start being able to produce more and more of the goods they need themselves, reducing their dependency on resources from elsewhere. So that, the message we really get from Mesarovic and Pestel, I think, is that there is a lot of serious disequilibrium between different nations in the world that major steps, both in population control and in transfer of funds amongst others, are required to try and correct the situation, and that the costs of delay are really unbearable, and that therefore we have to start moving now. And that isn't going to be easy, is it? When the people of Cleveland are reluctant to pay taxes to send children to school, it may be somewhat hard to persuade them to use their taxes to help people on the other side of the world attain a tolerable standard of living. So again: politics.

Okay. Just a few more question… comments about this. We've said that the Mesarovic-Pestel model worked on looking at ten different regions in the world and the interactions between them, although we've emphasized the interactions between just two broad regions: developed and underdeveloped nations. What has been happening since is that various people have been trying to build local models, and saying: well, we want to answer questions, sure, about world trends, but in particular we're interested about what it means for us, what sort of decisions can we as a nation, or as a region, make? And in fact, there are many, many regional models now, with more or less attention to what is going on outside the region. Two I just mentioned here are that in Iran, they have a model which sort of has a very detailed analysis of the Iranian economy, but also has a less detailed model of the Middle East region of which it's a part, and an even less detailed model of sort of the Mesarovic-Pestel level of detail of the whole world. So, they can ask [audio skips] …will be, and we can think of some obvious examples. In terms of agriculture—now this is a question that every country has to ask—how much of your agriculture should be for the home market, and how much of it should be for the export market? Well, it's how much… to what extent do you want to look at your farmlands as a resource for getting foreign capital, and contributing to your balance of trade; to what extent are you concerned primarily with feeding your own people?

Something which makes even clearer than these for a good understanding of international economic relationships, is the question of oil price rises. Remember what we've talked about before is the balance between short-term decisions and long-term decisions. I mean, it's clear to any oil… OPEC minister that oil price rises are a good thing in the short run, because it means more cash is coming in this year. But of course, if what happens is that you lead to the collapse of the economies of the developed nations, then you have, in the old phrase, killed the goose that lays the golden egg. Or, for those of you into energy, the golden erg. But, actually, that is so bad because in fact, of course it's… well, never mind. You can tell it's bad. And then, a certain amount of questions as to what regional cooperation should be.

The West German model, again: a detailed model of West Germany, a less detailed model of the European economic community of which it's a part, and then a very crude sub-regional model of the rest of the world; and there, where the Iranian model gave a lot of attention to agriculture and oil, here, the German economy depends very much on high technology production. And so, things like actually looking at the educational mix to see to what extent resources should be going into producing people who can be good technicians versus people who can be good engineers, versus people doing basic research, versus, I guess, people who can be salespersons, is all part of the sort of more detailed analysis that would distinguish the West German model from models that other nations might want to develop.

And, on the next missing slide, we have… [laughs] the Latin American model, where a group in Venezuela is looking at the relationships between themselves, as an oil-rich nation, and their surrounding countries in Latin America that are oil-poor nations. And, also looking at the fact that Latin America is perhaps the region of highest population growth, and where they're getting into things that… their concern in their model is not simply with population, per se, and the sort of fact as we've discussed before, but the whole question of how do you provide employment for a rapidly growing population? And so there's a whole question about where do you capitalize in terms of industry versus agriculture, and what are your policies to be in terms of incentives for going out and living in the Amazonian jungle versus getting an apartment in the cities to take care of these huge populations.

Well. We're running out of time, of course, so let me skip the next three thousand slides. Oh, what a lot of material we have here. Heh heh heh. Well. Okay. Just briefly, it does seem to me that the sort of technology we've been talking about here, I think this message has been coming through fairly clearly, is going to become part of the way we plan. Using computer simulation and database management, and sort of trying to keep track of the knowledge we have in the form of computer models and updating it, this will become available to planners, and I think it will also become available to citizens. You might sort of think of the idea of mission control, not simply for a space shot, but also for managing the economy and the ecology and so on of a region. And then we can imagine these being linked into networks that people working on planning in one state or one country will, via the computer network, interact with other people. There are various seeds for this sort of collaboration, the food and ag… the various UN agencies. For example, meteorology is now clearly a worldwide thing, where we have satellites tracking the weather worldwide to give us more and more accurate weather predictions.

But, a little bit of reality. It's not going to be all that easy, because we live in a world of intense competition, there are still many wars. As resources dwindle there will be a lot of competition for those resources. So, the idea of a computer network where everybody happily shares information back and forth is, of course, not a realistic one. What we can hope, though, is that there will be networks which can handle a lot of routine type planning operations, while other computer planning centers will be more concerned with what we might call national security, to handle crises and competition and so on.

Now, in the last part of this talk, I want to just indicate another way in which computer technology is going to affect our decision-making. I've been stressing until now the idea of looking at ecological, economic, physical systems; trying to capture in numerical form the relationships between the various properties of these systems; and then using the network, or the computer or whatever, to explore the effect of different decisions. And then as I say: once we've got what looks like the right way to go, the sort of scenario, the choice of strategies that looks as if it will yield the best results, I've emphasized that we have to leave the computer for a while and start thinking about the political consequences and what the electorate will think and so on. And then we can go back to the computer and say well, we've now established different bounds, let's try and do that.

But, the computer is also being used to just store a lot of information and help us get access to the information we need. And, more and more we're trying to work on what we call artificial intelligence. Now what does that mean? Well, at one level it means what it says. We want to build machines which are intelligent. Now, by that, though, what I would like to think that means is that it's not that we're building sort of a race of robots that will take over, but rather that we're concerned with building, if you will, partners—intellectual partners that are machines. Well, what does that mean? Well, suppose you think of a library as a partner, an information partner. It has a lot of information that you want, and it's there for you, and now you want to go and get it out. Now, the way you get it out now is you probably go to a catalog and that gives you a list of a hundred entries, and then you go off and find those hundred things, and you page through them, and finally you might find what you want, but probably either in more detail or less detail than you need. So, one of the ways in which artificial intelligence is going is to say, how can we take a body of information and systematize it in such a way that you could sort of talk to the computer, and it will help you find the answers to your questions?

Now, as I discovered earlier this morning, many of you in Portland have a preliminary feeling for this, when you go down to Portland Mall and get bus route information. Here, of course, it's very simple: you dial in where you want to go to, and the computer responds with where you want to get there from, and you respond with the right keyed-in information, and it says what day of the week and you respond, and then finally it flashes up the information you need to make your decision. So, the idea is how do we sort of go beyond that, so that we can get information out of the computer in much more complex situations? Where it's not just a matter of giving you access to one of a few tables of information, but in some sense helping you explore what is a unique situation by bringing to bear the necessary information to make the right decision.

And so I want to tell you about a purely, well, a new development, something that hasn't reached the sort of mass production stage, but a research project at Stanford University. It's called MYCIN, and as you can… this time the spelling is correct, it's not an exotic French perfume, but it's based on the fact that there are lots of antibacterial agents like puromycin that end with the phrase "mycin." And this is a project… a program that is designed to help you—you being the physician—find the right therapy to fight a bacterial infection. And let's just quickly look at the pieces. There are two basic pieces to the program. There's a consultation system, whose job it is to take data on the patient, clinical information entered by the physician, you know—what's the blood count? What sort of gunk did you find in the sputum? What were the results of various tests?—and take a whole lot of data that has been built up in the past on clinical knowledge of various physicians, and then bring the two together to interact with the user, to help the user find out, through an ongoing consultation, what the right therapy is in the given situation. And an important feature of this is that not only is there a system that will just keep asking you questions and finally tell you, you know, "Inject the following gunk into the patient," but there is also an explanation system that lets you ask questions, and find out whether you believe what the computer is doing. So in other words, we've built up within the computer a set of rules about, you know, if you see this and you see that then it's probably this, but if you see that, you should test for something else before you can make up your mind, that sort of thing. And so you can ask… if you get a question fired at you by the consultation system, coming out of the program, and you don't believe that's the right sort of question to ask you, you can ask it to tell you what rule it's using from its databank. And then you can either say, "Oh yeah, fine," and then go back to here and continue with the consultation, and of course that's one of the things a computer is good for, keeping track of the details. So that you forgot to check something, and the computer has a database and pulls out the right rules to make sure you keep track of these things. But in other cases, you'll say, "Silly computer," and then there's another program which can respond to you in storing new rules on the basis of your clinical experience. And so the idea would be that it would probably be a bad program design to have these rules immediately go into the system, because it could be that you're wrong, but on the other hand it would be a bad idea to just throw away your complaint or your ideas. So, probably the best thing is to actually have an accumulation of new rules, and then a panel of human experts convene at regular intervals to look over those rules and see to what extent they should be added to the database of the system for use in future consultation on the system.

Let's just look at some of these things in a little more detail. One of the most important features of the system, I think you'll agree, is the idea that it's said right from the start that the way we're using the computer is not something that we have to accept, it's not something that, you know… like the man at the airport yesterday, when something goes wrong, it's the computer's fault and there's nothing you can do about it, or like the bureaucrat who won't help you, and in the old days they'd say, "I'm sorry I can't help you but it's the rules." Now they say, "I'm sorry I can't help you, it's the computer." So, what we want is to design our computer systems so that people understand that the computer is not just some box that is fixed and malevolent, and you must do what it says. The computer is something that humans have programmed to do what they want it to do, and that if there's anything wrong with the computer, that means the program needs changing. And so we have to have a structure which allows us to keep changing the computer program as we learn more about what people need and about what the facts are. And so it's very important to be able to acquire new knowledge from experts.

Of course, in other types of systems, one not like this which is a system for experts to use, but a system available to the general public, then the sort of way that it would be used would be somewhat different, is that rather than ability to acquire new knowledge from experts, it would rather be the question of how do you factor in the citizens' complaints to make the system more responsive over time, so it would be sort of a three-way interaction. Here, it's an interaction between just experts, because the experts are the users of the system, and the system. In the future we can expect more and more of this sort of three-way interaction between the computer experts, the users in the general public who know enough to see that the system isn't working properly but perhaps aren't sufficiently expert in computers to see how to set it right, and the people who can help set it right.

Oh, who needs to read that? We have just passed another missing slide. But, this just gives you an example of the sort of things we can use. I mentioned that another important feature of the system is that—so we've looked at one—I think perhaps the most important feature of the system is the one we've just discussed: the fact that it's understood that the computer program is not something fixed that we must conform to, but rather it was built to help us do things and that therefore it can be reprogrammed to do those things better for us. The other factor that makes the program a useful program and a responsive program is it must have an explanation facility. We can ask it questions about why it's doing what it's doing. And here's an example of using it in terms of general knowledge. Here we're asked… the user types in, "Do you ever prescribe carbenicillin for pseudomonas infections?" A question which often springs to mind during the quiet hours of the night. [laughter] And, what the computer does is it looks at its database and finds that, in fact, this combination of this particular drug and this particular infection are part of these rules: 24, 68, 137, and 138. It says, "I have in fact four rules, which are related to your question. Which one do you wish to see?" And you type in "Rule 024," and then it types out rule 024 for you. And, this is perhaps interesting because it just gives us, phrased in English, the sort of process that's going on in the computer, the way it will decide what to do next. So here's a rule. It says that "If the therapy under consideration is gentamicin," okay, in other words, because of what has happened so far in the consultation, you've decided this is a possible agent to use as the antibacterial agent, "and if you've got to a stage where the tests have established that the organism is pseudomonas, then what you should do is say that well, it's probably not a good idea to only use gentamicin, but you should try combining gentamicin with carbenicillin." So, you see that this is an answer… one of the possible answers to the question about carbenicillin in relation to pseudomonas. And, this is also the sort of rule that you can get. If you can ask… so suppose, somewhere along the line the computer types out "Use gentamicin and carbenicillin," then you as the user might type in, "Why?" and then its answer would be to type out this rule, you see, and then you would say, "Oh yes, I remember we were talking about gentamicin before," you see, and you'd been confused about it suddenly mentioning carbenicillin when you had thought gentamicin was the right therapy, but now it reminds you of this rule. And then, there are other questions, the kind just sort of, "Where are we in the consultation?" you know, "What have we decided are the possible infections? What is our level of confidence in the different type of infection?" and so on.

So, our final slide for the morning… [pause] isn't our final slide, all right… is that what I've suggested here as we've looked at this particular program, which I should tell you is able to do its job of diagnosing bacterial infections better than all but perhaps the top ten percent of specialists. What I want to do in conclusion is look at some of its design features and discuss with you to what extent they imply important features that we should have in computer systems in the future, which are helping people get access to information. And, here they are. [laughs] Some of these seem a little trite: the program should be useful. Though, of course, I think the market for home computer programs will probably be much more for programs that are entertaining, just as PBS doesn't do as well as ABC, so I expect that the bulk of cassette sales for home computer programs will be for playing games and things like this, and relatively few for doing anything madly useful, but that's all right, people are allowed to have fun. But, anyway, in this case, we're talking about these consultation programs to help people make decisions about important situations.

Programs should be educational, when appropriate. The idea is that we don't want the computer to reduce us to the level where we never think, but instead we just say, "Well, the computer said it so it must be right." And let's take a very simple example of this. People are worried about kids using hand calculators, and say that perhaps this is bad because people will lose their numerical skills. Well, my feeling is that it isn't particularly bad, but we have to teach different sorts of skills. In other words, you have to teach kids that if they type in two times two and the answer comes out six million, that something's wrong. But on the other hand, if you type out 2.000369 times 1.764, you recognize that that's roughly two times two, so the answer is roughly four, and if the answer comes out as—and I guarantee this is not the correct one; in the typed-up version we'll make sure this is correct—3.96425, then you say okay, I mean that's in the right ballpark, so I'm going to assume the circuitry is working and accept that answer. So, in other words, what you want is that you as the human have to keep track of the big picture, and you must learn the skills to understand roughly the way things are going, to know when things seem to be going wrong, and you had better check whether the computer is using the right rules or not, and things of this kind, but then you allow… you rely on the computer to carry out the details. Just as when you drive a car, you rely on the steering mechanism, you don't keep getting out and moving the tires in a different direction before you allow yourself to go a little further forward; but nonetheless, you keep your eyes open and retain responsibility for the general direction in which the thing is progressing. And, continuing the analogy, you can crash on the computer as well as in the car.

But, so. The program should be educational, when appropriate. It should be structured in such a way that people who use the program actually come to understand the situation better, rather than turn off and don't understand it. The program should be able to explain its advice. We've seen that working. If the computer tells you to do something, you can ask it why. It should be able to understand questions. If you want to know more about the situation. The program should be able to acquire new knowledge; and we've seen that's very important, and here's the next… the next line is really the same thing: the knowledge base should be easily modified. That the computer programs that are designed to solve special problems should not be things that we're stuck with, but things that can improve over time, can be more useful and more flexible.

This brings me to the end of this morning's sermon. Let me just summarize by saying that the title of this series is "Earth Resources Limited," and we talk a lot about resources in terms of energy, in terms of food, in terms of the ecosystem, but what we keep coming back to is that underlying all this is the question of management. That we live in a world of complex technology, of limited resources, and we have to use our brains, we have to understand the causes and the effects, we have to make decisions. Some of those decisions can be made on the basis of scientific analysis of biological and economic processes, others are much more vague sorts of decisions made on our understanding or feel for political processes. The computer is going to be a very important part of keeping track of that complexity, but it's not going to be the only part, and so it's going to be humans and computers together, building the brave new world of tomorrow.

HOST: Thank you, Michael. For the group that's here from Earth Resources, we'll go through the same pattern, pretty much, as we did yesterday. There's a little more information over at the cafeteria that isn't operating as a cafeteria this summer. But while you're… when you're meeting in groups, might reform with different people, think perhaps of some of the goals that we've suggested that, in group one particularly, the teachers that might be working toward one project together, what you've had suggested in this system, utilizing the computer and simulation might be a good direction to go, and some of the questions could be: where do we obtain the information to be incorporated into such a computer system? Does it require a great deal of computer equipment, the hardware part of it, or can some of it be done in rather inexpensive units, or can some of it be done without a machine at all? How do we develop additional information that can be expanded in a class, and incorporate the utilization of student projects and activities to develop that kind of a computer program or a simulation in a classroom, or in a situation for community leaders, or for council members, to help bring about an understanding of the information? And you can just keep going on questions after questions that may help you formulate the ideas of what you want to do with the next few weeks. So we'll officially have Michael back for questions at 11:00.

ARBIB: [off mic] I think maybe we can get the carousel back, and finish loading these…

HOST: [off mic] Yes, yes. We can go do it in a more comfortable place, too.

[program ends]