5-1-2019

Deception: The Role of Uncertain Consequences

Subhasish Dugar
University of Utah

Arnab Mitra
Portland State University, arnab@pdx.edu

Quazi Shahriar
San Diego State University

Citation Details

Dugar, S., Mitra, A., & Shahriar, Q. (2019). Deception: The Role of Uncertain Consequences. European Economic Review, 114, 1-18.

Deception: The Role of Uncertain Consequences

Subhasish Dugar

Economics Department
University of Utah
260 Central Campus Drive
Gardner Commons, Room 4039
Salt Lake City, UT 84112
Email:subhasish.dugar@economics.utah.edu
Phone: +1 (801) 587-1294

Arnab Mitra

Department of Economics
Portland State University
1721 SW Broadway
Cramer Hall, Suite 241
Portland, OR 97201
Email: amitra@pdx.edu
Phone: +1 (503) 725-3937

Quazi Shahriar[1]

Department of Economics
San Diego State University
5500 Campanile Drive
San Diego, CA 92182
Email: qshahria@sdsu.edu
Phone: +1 (619) 594-1675

ABSTRACT

We study how decisions to lie extend to risky environments. We provide experimental evidence from a sender-receiver game where there is uncertainty over the amount by which a sender's lie reduces its receiver's payoff, which is known only to potential liar. Even though all reduction amounts are equiprobable, ex-post beliefs elicited from senders suggest that, unlike truth-tellers, most liars underestimate the extent of the actual reduction in the receiver's payoff and appear to exploit this self-serving bias, resulting in substantially more lying relative to a baseline treatment without the uncertainty. Subsequent treatments confirm the bias by either providing additional evidence or by removing possible confounds. An intervention treatment nudging senders toward correcting the bias reduces lying.

Keywords: Self-serving bias, Uncertainty, Deception, Experiment

JEL Classification Codes: C72, C91, D81, D83

[1] The authors are grateful to the editor of the journal and two anonymous reviewers for their prescient comments on earlier versions of the manuscript. AM acknowledges financial support from the Faculty Enhancement Award 2016-17, Portland State University.

1. Introduction

The consequences of a lie for the party lied to are often subject to uncertainty. When deciding to tell a "black lie" – lies that financially benefit the liar at the expense of the person lied to (Gneezy, 2005) – potential liars may not always know with certainty the *amount* by which their lie would reduce the material payoff of the person lied to. In many situations, potential liars only know the distribution of possible amounts that ultimately governs the size by which their lie would reduce the payoff of the other person. To organize ideas, consider the vignette below, borrowed from Gneezy (2005) with changes to fit our research:

> *Mr. Johnson is about to close a deal and sell his car for $1,200. The engine's oil-pump does not work well, and if Mr. Johnson doesn't tell the buyer, the engine will overheat on the first hot day, resulting in equally likely damages (the cost of fixing the pump) of either $300 or $500 for the buyer. Being winter, the only way the buyer can learn about this now is if Mr. Johnson were to tell him. In that case, Mr. Johnson will have to reduce the price by at least $300. If Mr. Johnson doesn't tell the buyer, the buyer will learn about it only on the next hot day, resulting in equally likely damages of either $300 or $500.*

The same pattern of uncertain consequences of lying may arise in markets for "green goods", expert financial advice, and medical treatments, to name a few. Common to all of these examples and the vignette above is that an informed party's lie can raise own material payoff and reduce the uninformed party's material payoff, albeit stochastically rather than deterministically. Using Gneezy (2005) type sender-receiver cheap-talk games, we study how the presence of uncertain consequences of a black lie affects people's decision to tell such a lie. Studying this topic is important not only because such situations are pervasive in real life, but also because experiments have shown that people pursue selfish interests by favoring self-serving interpretations of risk and ambiguity (Haisley and Weber, 2010;

Konow, 2000) or adopting self-serving responses to objective risk (Exley, 2015). In essence, these studies show that people often exploit uncertainty to form subjective beliefs vis-à-vis objective risks as these willful self-interested distortions seem to diminish the severity of the material harm that their actions impose on others and thus help people resolve conflict between what is morally right and what is personally profitable in favor of the latter. While above studies and the related literature examine how self-serving readings of objective risks facilitate selfish behavior in situations involving distributional fairness, they do not explore situations allowing for possible untruthful behavior, and therefore, little evidence has been presented about whether uncertainty of the kind we explore here is likewise susceptible to systematic manipulations by people that help them act more dishonestly.[1]

To precisely tease out the effect of aforesaid type of uncertainty on untruthful behavior, we design treatments that may or may not embed uncertain outcomes for the recipient of a black lie. In our *Control* treatment, the sender who is the informed party can send a (costless) message recommending the receiver to choose one of the following two payoff options that is supposed to maximize the receiver's income: $(\pi_{s,} \pi_r) = (\$5, \$19)$ or $(\$12, \$5)$. Upon reading the message, the receiver chooses one of the options that determines each party's final payoff. Since the receiver is uninformed about the payoff consequences

[1] There exists an experimental literature that uses dictator, bargaining, social dilemma and other types of allocation games to show that people use a variety of excuses, before or after their selfish decisions, to maximize own income. While in the studies cited above excuses are driven by self-serving *interpretations* of risk and ambiguity, the following studies document excuses that are *not* driven by subjective manipulations of risk or ambiguity and instead these excuses take many other forms such as people avoid learning how their choices affect others (Bartling et al., 2014; Grossman, 2014; Dana et al., 2007); people develop self-serving biases in bargaining decisions (Babcock and Loewenstein, 1997; Messick and Sentis, 1979); or people rely on the possibility that their decisions did not generate the unfavorable outcome (Falk and Szech, 2013; Linardi and McConnell, 2011; Andreoni and Bernheim, 2009). Grossman and Van der Weele (2017) provide an excellent review of this literature. We are thankful to an anonymous reviewer for bringing these studies to our attention. There are a few theories arguing that people treat moral dispositions as constraints in order to attain selfish goals which they try to circumvent whenever possible (Dahl and Ransom, 1999; Rabin, 1995).

(for both the parties) of each of the two options, the sender has an incentive to lie (i.e. recommend the second option) to maximize own income. If the receiver believes in the message and chooses accordingly, then the lie reduces the receiver's payoff by $14 with certainty. In our *Uncertainty* treatment, with everything else being the same, believing in the same lie fetches the receiver either $8 or $2, realized as per the outcome of a (fair) coin toss. Thus, the lie reduces the receiver's payoff probabilistically, by either $11 or $17, with an expected loss of $14. Given that receivers behave the same way, the expected consequences of the lie are the same for both parties across two treatments. By studying the frequency of lying decisions with and without the uncertainty, we address the question of how preferences for dishonesty respond to uncertain outcomes than to a certain outcome for the affected party.

We discover that people are significantly more likely to tell a lie, 29 percentage-points more lies to be exact, if the negative impact of their lie on the recipient is probabilistic rather than deterministic. For our *Uncertainty* treatment, we elicit senders' belief, in an incentive compatible manner, about the coin toss outcome *after* they have sent their message. The data reveal that untruthful senders' average belief that the coin toss would reduce the receiver's payoff by $11 as opposed to $17 is significantly higher than 50% whereas the same for the truthful senders in that treatment stands exactly at the objective likelihood of 50%. Thus, we gather preliminary evidence that the effect of uncertainty on immoral behavior operates via senders' biased belief about the $8 outcome.

To further explore the extent of the bias in untruthful senders' belief and also to rule out alternative explanations, we conduct additional experiments with a game that includes elements from above two games (called the *Smaller-reduction* treatment) and a modified dictator game that involves uncertain payoff for the recipient for one of the dictator's choices (called the *Dictator* treatment). The data from these additional treatments yield two corroborating insights. First, some of the untruthful senders' belief in the *Uncertainty* treatment may have been so self-servingly biased in favor of the $11 reduction that they virtually ignored the possibility of the $17 reduction in the receiver's payoff. Second, the higher incidence of lying in the *Uncertainty* treatment is not attributable to untruthful senders' risk-seeking preferences over the receiver's money. We design a final treatment, the *Nudging* treatment, to test if a context-relevant subtle cue would mitigate the belief-driven bias and thereby reduce the rate of dishonesty in the presence of the uncertainty. Precisely, we elicit senders' belief about the coin toss outcome *before* they decide to lie (or not). We find that the gentle nudge mends senders' belief toward the objective likelihood of the coin toss outcome and generates a level of lying that is statistically identical to that in the *Control* treatment.

As a whole, we conclude that it is the untruthful senders' biased belief that produces higher incidence of lying in the *Uncertainty* treatment. To arrive at this conclusion, first we rule out a potential confound, namely untruthful senders' risk-seeking preferences over the receiver's money, based on the results of the *Dictator* treatment. Second, results from the *Nudging* treatment extend further support to our conclusion: when senders' belief about the likelihood of the $11 reduction in the receivers' income coincides with the objective risk, rate of lying drops to the level that is statistically identical to that in the *Control* treatment. Finally, we estimate regression models using data from the *Nudging* and *Uncertainty* treatments and the results confirm that a biased sender-belief increases lying significantly.

Liars' biased belief brings to the forefront the question of whether their adoption of a biased view of the objective risk is an ex-ante or an ex-post phenomenon. While the ex-ante interpretation contends that such biased belief is the upshot of a deliberation by the untruthful senders on the likelihood of the $8 outcome *before* rather than after their decision to lie, the ex-post interpretation posits that the biased belief is a manifestation of untruthful senders engaging in mental excuses to maintain a positive self-image *after* having lied. We evaluate both interpretations, but available evidence does not allow us to favor one explanation over the other. Although we cannot precisely pin down the timing of formation of the biased belief in the *Uncertainty* treatment, we are still able to demonstrate that the uncertainty over the receiver's payoff in the *Uncertainty* treatment leads to more lying via the channel of biased belief which our belief elicitation technique captured.

With this, we contribute to three distinct strands of literature. Our first contribution is to a large experimental literature mostly drawing on Gneezy (2005). This body of literature has established that people exhibit substantial "lying aversion", although this psychological trait has been found to be sensitive to a variety of economic factors.[2] Despite these advances, economists so far have not studied how preferences for dishonesty extend to uncertain environments. Our study advances the literature by examining the role of uncertainty in preferences for honesty.

Our second contribution is to an emerging literature presenting evidence that people often make excuses before or after their selfish acts to pursue self-regarding motives at

[2] See Gneezy (2005) and Fischbacher and Föllmi-Heusi (2013) for an entry into this literature. The results from these two studies have been attributed to "lying aversion" (Gneezy, 2005, p. 392); "guilt aversion" (Battigalli et al., 2013); as well as individuals' concern for reputation of appearing honest before others (Abeler et al., 2016); and the size of the lie in terms of reputation (Gneezy et al., 2016). Lying aversion has been found to be sensitive to factors such as *certain* monetary consequences of lies for both the liar and the affected party (Gibson et al., 2013; Freeman and Gelber, 2010; Gneezy, 2005), lying via truth-telling (Sutter, 2009), social contagion effects (Innes and Mitra, 2013), bargaining for a better price (Dugar and Bhattacharya, 2018); an appeal for honesty (Pruckner and Sausgruber, 2013), gender (Dreber and Johanneson, 2008), size of the lie (Dugar and Bhattacharya, 2017; Gneezy et al., 2016; Fischbacher and Föllmi-Heusi, 2013; Lundquist et al., 2009), cooperation in prior play (Ellingsen et al., 2009), reputation for honesty (Abeler et al., 2016), and impacts of prior unfair treatments on lying (Gawn and Innes, 2018; Houser et al., 2012).

others' expense (see footnote 1). We add to this literature by documenting that systematic self-serving interpretations of objective risk can also be exploited to act more untruthfully.

Our third contribution is to a thin but growing body of literature demonstrating that in different strategic environments providing subtle moral reminders or cues just before people face the temptation to make selfish choices at others' and sometimes even at one's own expense can significantly reduce self-interested behavior or increase prosocial behavior (Chen and Gesche, 2016; Pruckner and Sausgruber, 2013; Shu et al., 2012; Bateson et al., 2006; Rigdon et al., 2009; Babcock et al., 1995). Rabin (1995) theoretically studies self-serving biases in moral reasoning and recommends "salience injection" as a possible antidote to improve a person's moral conduct. Thaler and Sunstein (2003) advocate "libertarian paternalism" whose aim is to alter individual behavior without removing any existing options or significantly modifying economic incentives or curtailing people's freedom of choice. Researchers in judgement and decision-making also suggest that when morality and self-interest clash, external interventions may mitigate such biases (Moore et al., 2010; Pronin, 2007; Chugh et al., 2005; Pronin et al., 2002). Our intervention treatment adopts the general strategy as laid out in above studies and offers a new candidate that qualifies as a morality-improving nudge.

2. Related Literature

Psychologists have long been interested in the role of self-serving biases in decisions with moral underpinnings. Our primary argument that untruthful senders' biased belief facilitates lying accords well with Festinger's (1957) theory of cognitive dissonance which states that when opposing motivations (e.g. maximize own payoff and also minimize harm to others) clash, individuals may alleviate the resultant psychological discomfort by holding self-centered belief(s) or by adjusting their actions to avoid the discomfort. Messick and Sentis (1979) show in an experiment that when individuals decide on fair wage allocations, they claim more money for themselves than other individuals for identical work. Similar evidence of self-serving actions has been observed in the perceptions of a fair settlement in hypothetical disputes in which perceivers play the role of plaintiff or defendant (Babcock et al., 1995; Loewenstein et al., 1993). Bocian and Wojciszke (2014) allow their subjects to observe other individuals' counter-normative behavior (breaking a rule or cheating for gain), which was judged as immoral. However, this judgment became much more lenient when the observers could make financial gains from the same counter-normative behavior. The impact of self-interest on moral judgments thus can be mediated by people's financial gains from those immoral acts. Trivers (2011) argues that people try to assure themselves that serving self-interest does not contravene their moral principles. Studies on moral hypocrisy (Batson et al., 2002; Batson and Thompson, 2001) show that people behave in a selfish but immoral way when moral behavior would incur even mild costs. Yet, people like to keep a facade of morality and often self-indulgently believe this pretense.

In economics literature, Exley (2015) designs a laboratory experiment to investigate whether prospective donors use risk as an excuse for not giving to charities. In her design, subjects make four types of binary decisions in all four combinations of risky or certain payoffs for themselves and risky or certain payoffs for a charity. When subjects face zero tradeoff in payoffs between themselves and the charity (i.e. when they cannot use risk as an excuse for not giving) they exhibit similar patterns of risk aversion in self lotteries and charity lotteries. By contrast, when subjects face a tradeoff between the two payoffs (i.e. when they could resort to risk as an excuse for not giving) they exhibit more risk-aversion with regard to charity-risk relative to self-risk. The study concludes that subjects overweight the likelihood that charity lotteries return zero-dollar payoffs and use this as an excuse to choose certain amounts for themselves over charity lotteries and thus donate less. Relatedly, Exley (2018) finds that people exploit poor charity performance as an excuse for not giving.

Another study that is also closely related to ours is Haisley and Weber (2010) who show that people who want to increase their own income but also want to keep up a pretense that they are behaving fairly, will self-servingly overestimate the expected value of their matched recipient's income in dictator games involving a choice between a "fair" option and an "unfair" option. The fair option gives both subjects relatively equal allocations, while the unfair option gives more to the dictator and less to the recipient, and also makes the recipient's allocation dependent on a lottery with known probabilities (objective risk) or with unknown probabilities (ambiguity) even though the objective distributions of outcomes are identical in the two lotteries. They find that dictators choose the unfair option more frequently under the ambiguous lottery than under the simple lottery, and dictators' estimates of the expected value of the recipients' allocations are inflated under the ambiguity, indicating that dictators are more likely to form self-serving belief about ambiguity.

Dana et al. (2007) demonstrate that uncertainty over the negative consequences of one's behavior can open 'moral wiggle room'. In their Hidden Information treatment each dictator receives $6 for choosing option A and $5 for choosing option B, but the receiver's payoffs from A and B are determined by a coin flip prior to the game: could be $1 and $5, respectively (as in the Baseline), or flipped ($5 and $1, respectively). The dictators could *learn* the receiver's true payoffs costlessly by clicking a button, although that decision was kept private from the receiver. The treatment thus allows a dictator to remain ignorant to the exact consequences to the recipient. The dictators choose more unfair choice (i.e., A) (63% compared to 26% in the Baseline) and more than 50% choose to remain ignorant. Those who do not reveal all choose option A. The authors' interpretation is that uncertainty about the consequences of one's action is a 'veil' that allows people to make selfish choices while maintaining the illusion that one is a non-selfish type.

Clots-Figueras et al. (2015) study lying behavior in a modified Trust game with cheap-talk wherein receivers always knew the true value of the multiplier, *k*, while senders only had information about the distribution of *k*. They find that a large fraction of the receivers lie about the true value of *k* and close to half of all the senders are gullible and transfer more than the average amount. Thus, their study shows that a specific type of uncertainty can lead to greater lying in situations involving trust.

While the above studies are most closely related to our work, there are other studies in the economics literature that also examine the role of self-serving biases albeit in contexts that are arguably distinct from ours or the studies mentioned above. Rodriguez-Lara and Moreno-Garrido (2012) study how competing justice principles (egalitarian, accountability and libertarian) account for dictators' choices in situations where the dictator and the receiver differ in their relative contributions to the earned surplus available to the dictator to divide. They find that dictators behave according to the "most selfish" of the justice principles and conclude that principles of justice impose constraints on selfish agents which they attempt to bypass by adopting the principle that best serves their material interest.

Rustichini and Villeval (2014) study moral hypocrisy in an experiment where subjects are asked to make hypothetical choices in Dictator, Ultimatum and Trust games and also make judgments regarding the fairness and unfairness of all possible transfers in these games. After a week, the same subjects play the same games for real money and make similar judgements, which allows the authors to measure the divergence, if any, between the actual and hypothetical choices and also examine whether the later fairness judgments conform more to the initial fairness judgments or to the actual choices. The authors find that the actual choices violate the initial fairness judgments, and the later fairness judgments gravitate to the actual choices. They interpret the evidence of self-deceptive tweaking of judgments to actual choices as moral hypocrisy.

Finally, we connect to a body of literature in economics that analyzes how people can use informational asymmetry deceptively to maximize own payoff. Specifically, a number of laboratory experiments have examined the effect of private information on bargaining outcomes using either a standard or a modified version of the ultimatum game with one-sided incomplete information (e.g. Güth and Huck, 1997; Croson, 1996; Güth et al., 1996; Kagel et al., 1996; Pillutla and Murnighan, 1996; Rapoport et al., 1996a; Rapoport and Sundali, 1996b; Straub and Murnighan, 1995; Mitzkewitz and Nagel, 1993). These studies make potential values of the pie and their probabilities commonly known to both proposer and responder, however, only the proposer learns which value of the pie has been realized.[3] These papers broadly demonstrate a result which has come to be known in the literature as "hiding behind the cake effect": even when the pie size is relatively large, proposers use the information asymmetry self-servingly to make significantly lower offers and responders accept such offers. Recently, Ockenfels and Werner (2012) conduct a dictator game experiment on the Internet in collaboration with a German newspaper. Individuals playing in the role of recipients either knew or did not know the size of the cake

[3] Of course, these studies differ along several dimensions each of which raises an important question in the context of the literature. However, we desist from an in-depth discussion of them here.

to be distributed by the dictator. The treatments vary in terms of whether recipients were informed only about the amount dictators gave to them or were also informed about the realized cake size after the experiment. If dictators were only concerned about maximizing own income, they would not care about what recipients knew about the cake size. By contrast, dictators who care about being perceived as fair would tend to hide behind the large cake when recipients are only incompletely informed. They find that some dictators hide behind the large cake, showing that these dictators care about what recipients believe. As implicit in our descriptions of the related literature, there is virtually no study that investigates uncertain consequences of a lie and the interplay between biased beliefs and lying in this environment.[4]

3. The Experiment

The Games and Treatments

Table 1 describes the two one-shot games we employ. As mentioned in the previous section, players play these games in a pair of sender-receiver. Each game begins with the sender recommending one of the two options (payoff distributions), labeled as A and B, to his receiver by sending her one of the two cheap-talk messages given below.

Message-A: "Option A will earn you (my receiver) more money than Option B"

Message-B: "Option B will earn you (my receiver) more money than Option A".

As stated before, the sender is informed while the receiver is uninformed of the payoff distribution under each option. After receiving the message from her sender, which is the only piece of information the receiver ever obtains, the receiver chooses between A

[4] There is a literature on principal-agent games where the degree of uncertainty on whether agent can be blamed for the payoff loss of principal was varied to disentangle simple guilt and guilt from blame (see Charness and Dufwenberg 2006, Vanberg 2008). However, their main focus is not on the nature and extent of lying.

and B, which determines each party's payoff. The receiver only learns about her own earnings but not about any other payoffs to either player in the game. As a result, she remains uninformed about the payoff distributions of the options even after the game is over and never learns whether she has been lied to. The sender knows that the receiver remains in the dark throughout. This is the general structure of all of our games.

The game *Control* is a deception game introduced by Gneezy (2005). Payoff wise A is the advantageous option for the receiver while B is the advantageous one for the sender. Message-B, therefore, is untruthful.[5] The lie, if told and believed, raises the sender's payoff by $7 (= $12 – $5) while it reduces the receiver's payoff by $14 (= $19 – $5).[6, 7]

The second game, *Uncertainty* – the mainstay of our research question – is designed to examine if the sender's decision to lie is affected by uncertain consequence of the lie. The payoffs under Option A are the same between *Control* and *Uncertainty*. The difference is that the receiver's payoff under Option B in the *Uncertainty* treatment is stochastic – either $8 ($3 more than $5 in *Control*) or $2 ($3 less than $5 in *Control*) with equal probabilities – resulting in an expected payoff of $5. Message-B is still untruthful but if the sender lies and the receiver believes, it reduces the receiver's payoff by either $11 (= $19 – $8) or $17 (= $19 – $2); the mean damage to the receiver is still $14. The sender's potential benefit from the lie also remains the same at $7. On the expectation, the consequence of the lie is

[5] To remain neutral with labeling, we varied the option labels (A or B) across senders in the experiment. For some senders, Message-A was truthful, and for others, Message-B was truthful. For expositional simplicity, however, we will assume throughout the study that Message-B is untruthful.

[6] For the sake of expositional simplicity, we will be assuming throughout this section that the receiver acts upon the message s/he receives. That is, if the sender recommends message A (B), then the receiver chooses option A (B). In the past experiments (for example, Innes and Mitra, 2013; Gneezy, 2005), about 80% of the receivers followed the recommendations of the receivers. This is also true in our experiment.

[7] Gneezy (2005) shows that senders' propensity for untruthfulness increases with the potential benefit from such untruthfulness, and decreases with the potential damage imposed on the receiver. Using a dictator control treatment, he also shows that decision-making in the deception game is not the same as in the dictator game, as there is a psychological cost to lying in the deception game.

exactly the same between these two games. Below we describe our treatments that correspond to these two games.

The Procedures

We conducted two to four experimental sessions for each treatment presented in Table 1. In each session, we recruited an even number of subjects who participated in one and only one of the treatments for a single-shot play of the relevant deception game. Half of them were randomly assigned the role of the sender and the other half the role of the receiver. The two groups were seated in separate rooms and were never at the same place at the same time. Two paid monitors, one in each room, were randomly recruited from the subject pool in each session to assist the experimenter. In total, we recruited 89 pairs of sender-receiver for the two treatments. All experimental sessions were conducted at the San Diego State University, US. Only undergraduate students without any prior knowledge of Gneezy (2005) type deception experiments were recruited.

After a group arrived at the designated room, each subject was randomly assigned a registration number in a double-blind protocol. That is, the registration number remained anonymous not only to other subjects, but also to the experimenter. The registration numbers were used to randomly match a sender to a receiver. Each subject was then given a copy of the instructions and, in addition, the experimenter read aloud the instructions. Then all questions from the subjects were answered in private. The rest of the procedures in a session differed between the two treatments. In the *Control* treatment, the session moved to the sender's decision-making stage at this point: each sender was given a "Message card" which had Message-A and Message-B printed on it; the sender simply had to circle the message that he wished to send and then put the card face down.

The *Uncertainty* treatment involved one additional step *before* the senders made their decision: we (the experimenter together with the monitor) determined each receiver's payoff in the uncertain situation by tossing a coin. The results of the draws were recorded but were not shared with the senders during the experiment. The instructions explained all of these rules of the experiment. Copies of the instructions are available in the Appendix.

After we were done with the randomization, the senders in the *Uncertainty* treatment made their decision using the Message cards the same way as in the *Control* treatment. Once all senders were done making their decision, the cards were collected and the experimenter carried those to the receivers' room. Each sender's Message card was then privately passed on to his/her matched receiver. In all treatments while collecting the Message cards, they were kept face down so that the experimenter cannot map the choices to the senders. Notice that the decision-making process from the perspective of the receivers was identical between the two treatments as they were unaware of the payoffs under each option. After reading the message, each receiver chose between the two options on a piece of paper. Payments for all subjects were determined and were made privately, in cash, and in a sequential manner. To maintain anonymity, we placed subjects' earnings in separate envelopes with registration numbers written on them and laid them on a table. Subjects (both senders and receivers) picked up their envelopes on their way out.

The senders in the *Uncertainty* treatment responded to an incentive-compatible belief elicitation questionnaire *after* they made their decision to lie (or not) and before they received their final payments. Our procedure closely followed the one used in Vanberg (2008). The questionnaire asked them to guess if the coin toss resulted in the larger or the smaller payoff for their respective receiver by choosing from the five options in Figure 1.

The payoff for each option depended on the actual outcome of the coin toss for each sender-receiver pair and was calculated using a quadratic scoring rule (Schlag et al., 2015; Rey-Biel 2009; Costa-Gomes and Weizsäcker 2008; Nyarko and Schotter 2002; McKelvey and Page 1986; Brier 1950).[8, 9] Anonymity of all the subject decisions was strictly maintained; senders and receivers never learned each other's actual identity during or after the experiment. In addition, the receivers never had the opportunity to find out, during or after the experiment, whether the sender was truthful. The monitors were specifically instructed to maintain anonymity and prevent any form of communication among the subjects.

Hypotheses

When a sender lies to his receiver and if the lie is believed, the certain reduction in the receiver's payoff in the *Control* treatment and the expected reduction in the *Uncertainty* treatment are the same ($14), while the sender benefits by the same amount ($7). Based on the findings in the literature (Mitra and Shahriar, 2018; Innes and Mitra, 2013; Gneezy, 2005) one would expect the extent of untruthfulness to be the same in these two treatments unless uncertain consequences of a lie affect sender behavior in the *Uncertainty* treatment in any significant way.

In accordance with our discussions in Sections 1 and 2, we hypothesize that the uncertain consequences in the *Uncertainty* treatment may provide the senders with a scope to form biased belief which can be self-serving in nature. In particular, if a sender believes

---

[8] The payoffs were calculated for the probabilities 85%, 68%, 50%, 32% and 15%. Since the quadratic scoring rule produces very flat payoffs for extreme probabilities, following Vanberg (2008), we avoided those probabilities in this payoff calculation, but the probabilities were scaled to 100%, 75%, 50%, 25% and 0% for the analysis reported in the next section.

[9] Based on the payoff senders received they could have figured out the outcome of the coin toss in all instances expect when their response in the questionnaire was "Uncertain". Since the senders could possibly learn the outcome only after making their decision and that they didn't even know they have to answer this question during the experiment, the belief elicitation process could not have affected their decisions in any way.

that the realization of the receiver payoff under Option B is going to be $8 (that is, if he lies and if it is believed by the receiver, then the lie would reduce the receiver's payoff by the smaller amount) with a probability higher than 0.5, then the expected harm the lie imposes on the receiver would appear to be less than $14 to the sender. The decision-making environment in the *Uncertainty* treatment, compared to the *Control* treatment, thus allows the sender to exploit the uncertainty to lie more. This observation leads to the following hypothesis.

HYPOTHESIS 1**:** *The proportion of untruthful senders is expected to be higher in the Uncertainty treatment than that in the Control treatment.*

Since we argue that the senders may form a biased belief to convince themselves to lie more or to justify lying more in the *Uncertainty* treatment, we expect the elicited sender-belief data to reflect this bias. We formalize this idea in the following hypothesis.

HYPOTHESIS 2: *In the Uncertainty treatment, the belief of the untruthful senders about the realization of the $8 payoff would be greater than the objective likelihood of 50%.*

4. Results

Figure 2 reports the proportion of lies (message-B). In the *Control* treatment, only 57% of the senders (27 out of 47 senders) lied to their respective receiver. Consistent with our expectation, this number increases in the *Uncertainty* treatment to 86% (36 out of 42 senders). Using a *Z*-test for difference in proportions, we find that the increase in the proportion of untruthful senders in the *Uncertainty* treatment is statistically significant (Z = 2.93, p = 0.003).[10] The data, therefore, provide support for Hypothesis 1.

To find out whether senders' belief was biased in the *Uncertainty* treatment, we take a look at the sender responses in the belief elicitation questionnaire. Recall that the questionnaire elicited senders' belief about the likelihood of the two possible receiver payoffs ($8 and $2) under Option B. If the belief was unbiased, then the average belief for either payoff realization would be approximately 50%. In contrast, in the data, shown in Figure 3, we find that the average belief for $8 was 60% (*s.d.* = 17.19%) among the senders who lied, which is statistically significantly different from the objective likelihood of 50% (the one-sample signed-rank test produces Z = 3.15 and p = 0.002).[11, 12] Hypothesis 2, therefore, is also supported by the data. To understand the skewed belief data, we also take a look at the belief distribution. Among the liars, the belief for the likelihood of $8 was always 50% or higher, whereas the belief among the truth-tellers was both higher and lower than 50%, with the average at exactly 50% (*s.d.* = 15.81%).[13, 14]

The above data indicate that the senders exploit the uncertainty over the receiver's payoff to lie more in the *Uncertainty* treatment. The basic explanation is that biased belief makes the harm the lie imposes on the receiver appear *smaller* than they actually are which leads to the higher level of lying in the *Uncertainty* treatment than in the *Control* treatment.

[10] The tests reported in this section are all two-tailed.
[11] The one-sample non-parametric exact test (Schlag, 2008) produces p = 0.012.
[12] The average belief among all the senders in this treatment was 58%, which is also significantly different from 50% (the one-sample signed-rank test produces Z = 2.91 and p = 0.004, and the one-sample non-parametric exact test produces p = 0.015).
[13] Using a binomial test, we can reject the null hypothesis ($p < 0.001$) that the proportion of senders who lied and whose belief for the likelihood of the $8 realization was higher than 50% is significantly different from zero.
[14] Since there were only few truthful senders, we could not perform any meaningful statistical test using the lying or belief data on the truth-tellers.

As we mentioned in Section 1, untruthful senders may have formed the biased belief before they lied to perceive lying as less harmful than it actually was (ex-ante interpretation) or senders may have formed the biased belief after they lied to justify the dishonest act and maintain a positive self-image (ex-post interpretation).[15] The ex-ante interpretation is in line with a sprawling literature on sender-receiver deception games that started with Gneezy (2005) which has consistently shown that senders do take into account the consequences of lying – both pecuniary and/or non-pecuniary benefits and costs – *before* they decide to lie. In fact, Gneezy (2005) and others show that, controlling for the sender's financial benefit from lying, senders lie significantly less when their lie can impose a higher material cost on the receiver. Thus, before deciding to lie (or not) people factor in how much the other side would lose from their lie. On the other hand, the ex-post interpretation that biased beliefs are utilized to justify selfish acts accords well with the findings in Andreoni and Sanchez (2014) who show that selfish players may distort their beliefs to maintain a positive image.[16] Our

[15] The senders in the *Uncertainty* treatment, however, could not have lied more in anticipation of the belief elicitation process. The reason is that when the senders decided whether to lie or not, they were unaware that their belief would be elicited.

[16] We thank an anonymous reviewer for bringing this study to our attention. Andreoni and Sanchez (2014) designed an experiment to examine how social-image concerns influence individuals' belief statements. They used a modified trust game and compared subjects' belief about their opponent's probable actions under two contrasting belief elicitation methods. Under the *stated beliefs* method subjects faced no monetary consequences for stating inaccurate belief, however, their social image was at stake. In contrast, under the *revealed beliefs* method subjects' entire game payoff was at stake for misrepresenting their true belief, however, they were unaware that their belief was being measured. Their results show that under the first method sophisticated selfish players lie about the belief they (inaccurately) held about their opponent's behavior that would justify their selfish decision; also, these subjects had the greatest difference between their stated and revealed beliefs. However, the second method reveals selfish players' actual belief about their opponent's behavior which was revealed to be cooperative; their revealed belief, by contrast, were significantly more accurate and indicated that these subjects realized that their selfishness was not justifiable by their opponent's behavior. Thus, when misrepresenting true belief is relatively less costly, people are more likely to distort their true belief. Our belief elicitation method offers some (even though small) incentives for beliefs, and the Andreoni and Sanchez (2014) results indicate that this should at least partially combat socially desirable reporting. In this connection, we also find evidence that the elicited beliefs are also sincere. First, as mentioned above, the average belief about the $8 outcome among liars is 60% whereas the same among the truth tellers is 50%. Second, we measured the degree of association between beliefs and sender-behavior keeping in mind that both are nominal variables. We reject the null hypothesis that the belief and behavior are independent by performing the Pearson Chi Squared test ($\chi 2 = 6.69$, $p = 0.08$ and Cramer's $V = 0.40$).

experimental design does not allow us to distinguish between these competing arguments regarding the timing of the belief formation process.

We find that the receivers' behaved identically in the two treatments, *Control* and *Uncertainty*, discussed above. The aggregated data from the two treatments show that on average 82% receivers followed their respective sender's message.

5. Further Testing for the Bias and Possible Confounds

Since there are alternative belief elicitation methods and researchers disagree on their appropriateness, we wanted to comprehend the extent of the bias without having to rely on any belief elicitation method.[17] To achieve this, we designed a treatment based on the *Smaller-reduction* game in Table 2. This game is closely related to the *Uncertainty* game. Option A is identical between the two games. Compared to *Uncertainty,* the only difference is that the possibility of the smaller payoff of $2 for the receiver under Option B is completely removed. That is, the receiver's payoff under Option B in *Smaller-reduction* game is $8 with certainty, and therefore, if the sender's lie is believed, then the receiver's payoff would be reduced by the smaller amount ($11). Recall that the mean payoff reduction in the *Uncertainty* treatment, calculated using the objective probabilities, was $14. The lie is thus less harmful in the *Smaller-reduction* game than in the *Uncertainty* game. However, the sender's potential gain from the lie remains the same between these two games. The primary aim of the *Smaller-reduction* treatment is to provide an indirect measure of the extent of the bias in favor of the $8 outcome among untruthful senders in the *Uncertainty* treatment without relying on elicited sender belief. Hypothetically speaking, if senders in the

[17] Trautmann and Kuilen (2015) survey the experimental literature to compare the performance of different belief elicitation methods. However, all the techniques compared in their study are of continuous in nature and there is no comparison to discrete measures such as the one employed in this study.

Uncertainty treatment entirely discounted the possibility of the $2 outcome from the coin toss and, thereby believed that their lie would reduce the receiver's payoff by only $11 as opposed to $17, then the proportion of lies would be statistically indistinguishable between these two treatments. The closer the gap between the levels of lying in the *Uncertainty* and the *Smaller-reduction* treatment, the higher may have been the implied bias (toward the $8 outcome) among the untruthful senders in the *Uncertainty* treatment.

Procedurally, the *Smaller-reduction* treatment was identical to the *Control* treatment. 36 pairs of sender-receiver participated in this treatment. As Figure 2 shows, 89% of the senders (32 out of 36 senders) in this treatment lied to their respective receiver. Recall that 86% of the senders lied in the *Uncertainty* treatment. As expected, the proportion of lying in the *Uncertainty* treatment is lower than the same in the *Smaller-reduction* treatment, but the difference is really small and we fail to reject the null hypothesis that the proportions are statistically the same ($Z = 0.42$, $p = 0.68$). We take this as an evidence that the senders in the *Uncertainty* treatment were so biased that, in effect, they virtually ignored the possibility of the larger damage from their lie altogether to lie more.

One can argue that a comparison of the untruthful behavior between the *Uncertainty* and the *Smaller-reduction* treatments only indicates that the level of lying in the *Uncertainty* treatment is really high, but it still does not necessarily imply that the lying resulted due to biased belief. The high level of lying in *Uncertainty* treatment can also be due to the senders' risk preferences over the receiver's payoff. There are studies (Agranov et al., 2014; Kvaløy et al., 2014; Andersson et al. 2013) that show that people can be risk-seeking when the uncertainty is over someone else's payoff. To see how this argument can possibly have a role to play, suppose that some senders in the *Uncertainty* treatment are risk-seeking over the receiver's payoff, then that can lead these senders to recommend Option B more in the *Uncertainty* treatment relative to the *Control* treatment simply because the resulting receiver payoff distribution under option B in the *Uncertainty* treatment is a mean preserving spread of the certain payoff of the receiver under option B in the *Control* treatment. As a result, senders with risk-seeking preferences might lie more in the *Uncertainty* treatment.

To check validity of this alternative explanation, we design a treatment which is based on the *Dictator* game in Table 2. In this game, a dictator-sender chooses between two monetary allocations: B_C and B_U which basically are the Option B in the *Control* and in the *Uncertainty* treatments, respectively, and the receiver has no choice to make.[18] This treatment thus provides a direct test of risk preferences of the senders over the receiver's money. That is, if senders are risk-seeking over the receiver payoff, then we would expect the senders to choose Option B_U more often than Option B_C. The data, however, revealed the opposite phenomenon: 31% of the senders (8 out of 26 senders) chose Option B_U and the other 69% (18 out of 26 senders) chose Option B_C. So, the uncertain option is chosen 38 percentage-points less of the time. This implies that if senders' risk preferences affect senders' decisions in the *Uncertainty* treatment, if at all, then that can explain a lower propensity of the senders to recommend Option B in the *Uncertainty* treatment compared to the *Control* treatment but not the higher one which we actually observed. We therefore reject senders' risk preferences over the receiver's payoff as a possible explanation for higher rate of lying in the *Uncertainty* than in the *Control* treatment.[19]

---

[18] The realized receiver payoff under Option B_U was chosen using a coin toss as in the *Uncertainty* treatment.
[19] For our conclusion to hold, we implicitly assume that these two preferences are independently distributed in our subject population.

Considering all the findings we have discussed thus far, we arrive at two major conclusions: when lies lead to uncertain payoff consequences for the receiver, as in our environment, senders tend to lie significantly more and biased sender-belief is most likely the main channel through which the uncertainty causes the increase in lying.

6. Correction of the Bias & Robustness Check

The primary objective of this section is to test if the senders' biased belief, as in the *Uncertainty* treatment, can be corrected by making external interventions. To accomplish our goal, we designed a treatment, called *Nudging*, based on the same game as in the *Uncertainty* treatment. In the *Nudging* treatment, we test if an external intervention can gently nudge senders in the direction of correctly calculating the expected loss that their lie imposes on the receiver in the presence of the uncertainty. Our intervention mechanism basically requires them to hold beliefs identical to the objective probabilities of the two possible outcomes from the coin toss. Specifically, the *Nudging* treatment asks each sender to guess the actual realization of the receiver payoff under Option B that would result from the coin toss *before* each sender decides whether to lie or not. We use the same discrete belief elicitation method as before. This marks a key departure from the *Uncertainty* treatment which asked each sender to guess the same, however, only *after* each sender had decided to lie or not. If senders can be made to hold correct belief, then the incidence of lying might turn out to be statistically indistinguishable between the *Nudging* and *Control* treatments.

Note that in the *Uncertainty* treatment sender belief was elicited after the decision-making stage, and as a result the belief elicitation mechanism could not possibly have impacted a sender's actual decision. Moreover, the senders were not even aware during the decision-making stage that they would be subjected to the belief elicitation procedure later. On the other hand, we hope that eliciting incentivized belief before a sender's decision, as in the *Nudging* treatment, may force the senders to think harder about the actual probability distribution of the coin toss outcomes and such forced thoughts might help reduce or at best completely eliminate the bias (in favor of the $8 outcome) in the sender belief.

The procedural details of the *Nudging* treatment are identical to those in the *Uncertainty* treatment, except the timing of the belief elicitation process. 38 pairs of a separate group of subjects who had no prior experience with our games made decisions in this treatment. As Figure 2 shows, we find that 61% (23 out of 38 senders) of the senders lied in the *Nudging* treatment. Recall that 57% of all the senders in the *Control* treatment and 86% of all the senders in the *Uncertainty* treatment were untruthful. The proportion of the untruthful senders in the *Nudging* treatment is thus statistically significantly lower than that in the *Uncertainty* treatment ($Z = 2.56$, $p = 0.01$) but not statistically different from the same in the *Control* treatment ($Z = 0.29$, $p = 0.77$). To check whether our intervention corrected the senders' biased belief, we also elicited the sender belief. We find that the average sender belief about the likelihood of $8 realization is 53% (*s.d.* = 20.79%), and we fail to reject the null hypothesis that this belief is different from the objective likelihood of 50% (the one-sample signed-rank test produces $Z = 1.13$ and $p = 0.26$).[20] Thus, by drawing the senders' attention to the coin toss outcome before their decision, we seem to have reduced the biased belief and the lying behavior.[21] Since the risk level stays the same

[20] The one-sample non-parametric exact test produces $p = 0.98$.

[21] The results of the *Nudging* treatment lead us to suspect that the incentivized belief elicitation process may have forced some senders in the *Uncertainty* treatment to think objectively about the outcome of the coin toss and thereby may have lessened their bias. This then implies that the level of bias we observed in the *Uncertainty* treatment might be an underestimation of the actual bias in sender belief. The results of the *Smaller-reduction* treatment also seem to support this view. We, however, think that the reason the incentivized belief elicitation process does not correct the bias in the *Uncertainty* treatment as much as it does in the *Nudging* treatment is

between the *Nudging* and *Uncertainty* treatments but the rate of lying differs substantially between them, we claim that it is indeed the biased belief of untruthful senders that led to higher rate of lying in the *Uncertainty* treatment than in the *Control* Treatment. This result also provides further evidence that senders' risk preferences do not play any major role in influencing their choices in the *Uncertainty* treatment, as discussed before.

Despite having presented some evidence that biased beliefs help some senders lie, we recognize that our evidence is subject to some fundamental issues pertaining to measuring individuals' beliefs in strategic contexts. The first issue is that people may misreport their true beliefs to justify their actions despite the presence of financial incentives to truthfully elicit their beliefs. All belief or preference elicitation schemes are prone to this long standing issue, as economists do not have any other measure of preferences except the one they elicit. Although we find the correlation between actions and beliefs to be in line with truthful reporting (footnote 16), this finding is also consistent with the justification motive mentioned above. Therefore, incentivized beliefs or the correlation measures are although fairly useful, but ultimately they provide inadequate information. Second, we conjecture in footnote 21 that the belief elicitation process may have forced some senders in the *Uncertainty* treatment to think more objectively about the coin toss outcome and thereby may have lessened their bias. Thus the bias observed in the *Uncertainty* treatment might be an underestimation of the actual bias. This conjecture, however, runs into a similar issue: it is extremely difficult to know how elicitation changes beliefs, as by definition, we do not observe unsolicited beliefs. Schlag et al. (2015) discuss these measurement problems.

---

that the elicitation in the former takes place not only after the decision to lie has been made but also most likely after beliefs have been formed while in the latter the elicitation takes place before the decision-making stage and possibly before beliefs are formed.

We now move on to a discussion of the robustness of our main findings (discussed in Section 4) to a different belief elicitation protocol.[22] Recall that the belief elicitation procedure we utilized in the *Uncertainty* and *Nudging* treatments was *discrete* in that it asked the subjects to choose from the five options described in Figure 1. These options were then interpreted as five distinct likelihood measures for the realization of $8: 100%, 75%, 50%, 25% and 0%. One can argue that this discrete protocol offers limited number of choices to the senders to choose from and thus may force the distribution to be skewed in a particular direction as a result. To test the robustness of our findings, we re-ran these two treatments (labeled as *Uncertainty_cont_belief* and *Nudging_cont_belief*) in which we utilized a *continuous* belief elicitation protocol. Instead of asking the subjects to choose from only five options, we allowed the subjects to write a number between 0 and 100 up to any decimal points on a piece of paper. The procedure was incentivized using a quadratic scoring rule similar to the one used in the discrete protocol. As discussed below, the data from the two new treatments look qualitatively similar to the corresponding original treatments, and thus provide support to our hypotheses in question.

In *Uncertainty_cont_belief*, 85% (35 out of 41) of the senders lied while 63% (25 out of 40) of the senders in *Nudging_cont_belief* lied; the former is significantly higher than the latter ($Z = 2.35$, $p = 0.02$). The lying proportion in *Uncertainty_cont_belief* is significantly higher than the same in *Control* ($Z = 2.86$, $p = 0.004$) while the proportion in *Nudging_cont_belief* is not significantly different from that in *Control* ($Z = 0.48$, $p = 0.63$). In *Uncertainty_cont_belief*, the average sender belief about the likelihood of $8 realization is 59% (*s.d.* = 16.51%) among the senders who lied, which is statistically significantly

[22] We are thankful to the editor and an anonymous reviewer for this suggestion.

different from the objective likelihood of 50% (the one-sample signed-rank test produces $Z = 3.38$ and $p = 0.001$).[23, 24, 25] In *Nudging_cont_belief*, on the other hand, we again find that the average sender belief about the likelihood of $8 realization is 53% (*s.d.* = 21.88%), and we fail to reject the null hypothesis that the belief is different from the objective likelihood of 50% (the one-sample signed-rank test produces $Z = 1.22$ and $p = 0.22$).[26] Figure 4 presents the histogram of the beliefs in the *Uncertainty_cont_belief* and *Nudging_cont_belief* treatments.

We conclude this section by discussing results from regressions utilizing the above two treatments that extend further statistical support to the fact that the likelihood of lying increases as the belief about the $8 outcome increases.[27] Note that *Nudging_cont_belief* treatment has a nice feature of inducing in senders the objectively correct belief about the $8 outcome for the receiver. As such, if we would like to analyze the role of sender belief in the *Uncertainty_cont_belief* treatment, the *Nudging_cont_belief* treatment could serve as an ideal "base treatment".

In Table 3, we present marginal effects from two Probit models using data on individual senders from the continuous belief sessions. For each model, the dependent variable takes a value of 1 if the sender lies; otherwise it is 0. *Uncertainty Treatment* takes a value of 1 (0) if the observation is drawn from the *Uncertainty_cont_belief* (*Nudging_cont_belief*) treatment. *Belief* is continuous and it takes a value from the set [0%,

---

[23] The one-sample non-parametric exact test produces $p = 0.030$.

[24] The average belief among all the senders in this treatment was 58%, which is also significantly different from 50% (the one-sample signed-rank test produces $Z = 3.40$ and $p = 0.001$, and the one-sample non-parametric exact test produces $p = 0.030$).

[25] A Kolmogorov-Smirnov test rejects the null hypothesis that the distribution of belief among the liars is normally distributed with a mean of 50% ($p < 0.001$).

[26] The one-sample non-parametric exact test produces $p = 0.235$.

[27] We are thankful to an anonymous reviewer for suggesting the regression analysis.

100%]. The results from model (1) basically reproduces the difference in the propensity to lie between the *Uncertainty_cont_belief* and *Nudging_cont_belief* treatments. When we introduce *Belief* as an explanatory variable on the right hand side of model (2), its marginal effect emerges as statistically significant and positive. It shows that a 1% increase in belief increases the probability of untruthfulness by 0.02. (The corresponding elasticity turns out to be 1.19.) The results from model (2) also indicate that, once we control for belief, the marginal effect of *Uncertainty Treatment* goes down from 0.25 to 0.21. This is consistent with the expectation that senders acted on their belief. We would like to indicate that the decrease in the marginal effect of *Uncertainty Treatment* is not substantial. We conjecture that it is due to the likely reduction of the bias in beliefs in the *Uncertainty_cont_belief* treatment caused by the belief elicitation procedure we discussed in footnote 21.

To summarize, we find that untruthful senders' biased belief produces a higher rate of lying. The above claim is supported by three pieces of evidence. First, we eliminate the role of untruthful senders' risk-seeking preferences over the receiver's money as a potential explanation. Second, we show by drawing evidence from the *Nudging* and *Nudging_cont_belief* treatments that when senders' belief about the likelihood of the $11 reduction in the receiver's income coincides with the objective risk, rate of lying plunges to the level that is statistically equal to that in the *Control* treatment. Finally, results from our Probit models statistically establish that biased sender-belief significantly increases lying.

7. Conclusion

We designed a laboratory experiment to study people's propensity to tell a lie when their lie negatively impacts its recipient's material payoff. The negative impact, however, is probabilistic (objective risk is private information to the prospective liar) rather than deterministic even though the expected amounts of loss from the lie are identical under the two situations. We use cheap-talk sender-receiver games where an informed sender can deceive an uninformed receiver by recommending to take an action that is against the receiver's material interest. We find that senders lie significantly more when the loss to the receiver from their lie is probabilistic than when it is deterministic.

Elicited beliefs about the objective risk and the data from additional treatments appear to suggest that untruthful senders severely underestimate the expected amount of loss that their lie imposes on their receiver by adopting self-serving or self-justifying interpretation of the objective risk. Furthermore, the high incidence of lying in presence of the uncertainty is not driven by untruthful senders' risk-seeking behavior over the receiver's money. Finally, we consider two possible interpretations regarding the exact time at which liars develop their biased belief about the objective risk so as to be able to lie more. The first interpretation adopts the view that liars develop the bias before they decide to lie (ex-ante interpretation) whereas the second interpretation embraces the view that they lie first and then they develop the bias (ex-post interpretation). However, our data cannot discriminate between the two alternatives.

The majority of previous research on lying behavior indicates that people display substantial "lying aversion" when lies reduce other party's material payoff deterministically. Our results indicate that people's lying behavior is greatly sensitive to the underlying process that determines the negative impact of a lie. In descriptively more realistic contexts where lies stochastically reduce the other party's payoff, there exists a strong motivation to interpret objective risks in a self-serving or self-justifying manner that leads to more selfish and dishonest behavior. The observed self-regarding interpretation of objective risks is also consistent with a body of literature documenting that people often resort to various types of excuses to make self-interested choices that adversely affect others' welfare.

Drawing cues from several strands of literature that suggest that subtle external interventions aka *nudges* can make ethics salient and thereby correct people's self-serving moral biases, we also made an external intervention. In this new treatment, we revisited the original game with the uncertainty and elicited senders' belief *before* their decision to lie (or not) as opposed to the former occasion where we elicited senders' belief *after* their decision to lie (or not). The data show that varying the timing of the belief elicitation process results in significantly lower amount of lying and the proportion of untruthful senders who believed that their lie would reduce their receiver's payoff by the relatively smaller amount also goes down significantly. Therefore, when people are made conscious of their potential moral bias, many people appear to correct it and are more likely to make an honest choice.

This last result has broad implications for identifying simple but effective strategies that can curtail dishonest practices. In settings allowing for uncertain negative impacts of one's lies on the other party, requiring one to state one's belief about objective risks before (rather than after) making the decision to deceive others appears to make morality salient, which, as we document, can promote honesty. Thus, interventions made right before the "moment of temptation" can help reduce fraudulent acts. By contrast, requiring one to state one's belief immediately after lying may facilitate lying behavior. Therefore, once an individual has acted dishonestly, it is too late to alert them to a standard of morality. Given the vast amount of financial resources devoted to deterring, detecting, and punishing dishonest behavior, a subtle intervention like the one used in this study seems too costly to overlook. One informal practice people frequently adopt especially in situations involving potentials for fraudulent behavior is that they provide a subtle reminder to the informed party (like Mr. Johnson in the vignette in our opening paragraph) about the honesty that is expected from their end. Our nudge treatment is similar in spirit to such gentle but effective reminders.

The key message of our study is that stochastic consequences of a lie for its recipient entail increased untruthful behavior, which should be taken seriously since many economic contexts involving a possibility of untruthful behavior are characterized by uncertain economic consequences. In this sense, our results have important implications for various real-life economic markets such as markets for financial advice, credence goods (Darby and Karni, 1973) and contractual relations.

In a recent work, Gawn and Innes (2018) show that experience of being lied to significantly erodes trust and trustworthiness and argue that their result has critical implications for smooth functioning of markets and society at large. Our work gains special importance in view of their result. The result from our *Nudging* treatment highlights the link between psychological mechanisms and lying. The lesson that can be gleaned from this result is that strategically similar nudges can be used to lead people toward truthful behavior. If properly designed, such nudges can therefore promote trust and facilitate trade.

Of course, it is better to practice caution when generalizing results with a special population and special set of experimental procedures. Additional research is warranted to examine how lying behavior changes as the underlying stochastic elements of our environment change. Another avenue for future research is to compare the incidence of lying when source of the uncertainty is "ambiguity" (or "Knightian" uncertainty) rather than known objective risk since previous research has demonstrated that people also develop positive and self-serving attitudes toward ambiguity when doing so allows them to behave self-interestedly at another person's expense.

References

Abeler, Johannes, Nosenzo, Daniele and Collin Raymond. (2016). Preferences for truth-telling. *CEDEX Discussion Paper No.* 2016-13.

Agranov, Marina, Alberto Bisin, and Andrew Schotter. (2014). An experimental study of the impact of competition for Other People's Money: the portfolio manager market. *Experimental Economics* 17.4: 564-585.

Andersson, Ola, Håkan Holm, Jean-Robert Tyran, and Erik Wengström. (2013). Risking other people's money: Experimental evidence on bonus schemes, competition, and altruism. Working paper, Research Institute of Industrial Economics.

Andreoni, J., & Bernheim, B. D. (2009). Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects. *Econometrica*, *77*(5), 1607-1636.

Andreoni, James, and Alison Sanchez. (2014). Do Beliefs Justify Actions or Do Actions Justify Beliefs? An Experiment on Stated Beliefs, Revealed Beliefs, and Social-Image Manipulation. *National Bureau of Economic Research Working Paper Series* 20649: 1–38.

Babcock, L., Loewenstein, G., Issacharoff, S., & Camerer, C. (1995). Biased judgments of fairness in bargaining. *American Economic Review*, *85*(5), 1337-1343.

Babcock, Linda, and George Loewenstein. (1997). Explaining bargaining impasse: The role of self-serving biases. *Journal of Economic Perspectives* 11.1: 109-126.

Bartling, B., Engl, F., & Weber, R. A. (2014). Does willful ignorance deflect punishment?–An experimental study. *European Economic Review*, *70*, 512-524.

Bateson, M., Nettle, D., & Roberts, G. (2006). Cues of being watched enhance cooperation in a real-world setting. *Biology Letters*, *2*(3), 412-414.

Batson, C. Daniel, and Elizabeth R. Thompson. (2001). Why don't moral people act morally? Motivational considerations. *Current directions in psychological science* 10.2 (2001): 54-57.

Batson, C. Daniel, Elizabeth R. Thompson, and Hubert Chen. (2002). Moral hypocrisy: Addressing some alternatives. *Journal of Personality and Social Psychology* 83.2: 330.

Battigalli, Pierpaolo, Gary Charness, and Martin Dufwenberg. (2013). Deception: The role of guilt. *Journal of Economic Behavior & Organization* 93: 227-232.

Bocian, K., & Wojciszke, B. (2014). Self-interest bias in moral judgments of others' actions. *Personality and Social Psychology Bulletin*, *40*(7), 898-909.

Brier, Glenn W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78: 1-3.

Charness, Gary, and Martin Dufwenberg. (2006). Promises and partnership. *Econometrica* 74, no. 6: 1579-1601.

Chen, Z. C., & Gesche, T. (2017). Persistent bias in advice-giving.

Chugh, Dolly, Max H. Bazerman, and Mahzarin R. Banaji. (2005). Bounded ethicality as a psychological barrier to recognizing conflicts of interest. *Conflicts of interest: Challenges and solutions in business, law, medicine, and public policy*, 74-95.

Clots-Figueras I, Hernán-González R and Kujal P. (2015). Information asymmetry and deception. *Front. Behav. Neurosci*, 9:109. doi: 10.3389/fnbeh.2015.00109

Costa-Gomes, Miguel, and Georg Weizsäcker. (2008). Stated beliefs and play in normal-form games. *Review of Economic Studies* 75, 729–762.

Croson, R., Boles, T., & Murnighan, J. K. (2003). Cheap talk in bargaining experiments: lying and threats in ultimatum games, *Journal of Economic Behavior & Organization*, 51(2), 143-159.

Dahl, Gordon B., and Michael R. Ransom. (1999). Does where you stand depend on where you sit? Tithing donations and self-serving beliefs." *American Economic Review*: 703-727.

Dana, Jason, Roberto A. Weber, and Jason Xi Kuang. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory* 33.1: 67-80.

Darby, Michael R., and Edi Karni. (1973). Free Competition and the Optimal Amount of Fraud. *Journal of Law and Economics*: 16(1): 67-88.

Dreber, Anna, and Magnus Johannesson. (2008). Gender differences in deception. *Economics Letters* 99.1: 197-199.

Dugar, S., & Bhattacharya, H. (2017). Fishy behavior: A field experiment on (dis) honesty in the marketplace. *Journal of Behavioral and Experimental Economics*, *67*, 41-55.

Dugar, S., & Bhattacharya, H. (2018). The Hidden Costs of Bargaining: Evidence from a Cheating-Prone marketplace. *Working paper*.

Ellingsen, T., Johannesson, M., Lilja, J., & Zetterqvist, H. (2009). Trust and truth. *Economic Journal*, *119*(534), 252-276.

Exley, C. L. (2015). Excusing selfishness in charitable giving: The role of risk. *Review of Economic Studies*, *83*(2), 587-628.

Exley, C. L. (2018). Using Charity Performance Metrics as an Excuse Not To Give. Working Paper.

Falk, A., & Szech, N. (2013). Organizations, diffused pivotality and immoral outcomes.

Festinger, L. (1957). A Theory of cognitive dissonance. Stanford, CA: Stanford University Press.

Fischbacher, Urs, and Franziska Föllmi-Heusi. (2013). Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association* 11.3: 525-547.

Freeman, Richard B., and Alexander M. Gelber. (2010). Prize structure and information in tournaments: Experimental evidence. *American Economic Journal: Applied Economics* 2.1: 149-164.

Gawn, G., & Innes, R. (2018). Do lies erode trust? *International Economic Review*, *59*(1), 137-161.

Gibson, Rajna, Carmen Tanner, and Alexander F. Wagner. (2013). Preferences for truthfulness: Heterogeneity among and within individuals. *American Economic Review* 103.1: 532-548.

Gneezy, Uri. (2005). Deception: The role of consequences. *American Economic Review* 95.1: 384-394.

Gneezy, U., Kajackaite, A., & Sobel, J. (2018). Lying Aversion and the Size of the Lie. *American Economic Review*, *108*(2), 419-53.

Grossman, Z. (2014). Strategic ignorance and the robustness of social preferences. *Management Science*, *60*(11), 2659-2665.

Grossman, Z., & Van Der Weele, J. J. (2017). Self-image and willful ignorance in social decisions. *Journal of the European Economic Association*, *15*(1), 173-217.

Güth, W., Huck, S., & Ockenfels, P. (1996). Two-level ultimatum bargaining with incomplete information: An experimental study. *The Economic Journal*, 593-604.

Güth, W., & Huck, S. (1997). From ultimatum bargaining to dictatorship—An experimental study of four games varying in veto power. *Metroeconomica*, *48*(3), 262-299.

Haisley, E. C., and Weber, R. A. (2010). Self-serving interpretations of ambiguity in other-regarding behavior. *Games and Economic Behavior*, *68*(2), 614-625.

Houser, D., Vetter, S. & Winter, J. (2012). Fairness and cheating. *European Economic Review* 56.8 (2012): 1645-1655.

Innes, Robert, and Arnab Mitra. (2013). Is dishonesty contagious? *Economic Inquiry* 51.1: 722-734.

Kagel, J. H., Kim, C., & Moser, D. (1996). Fairness in ultimatum games with asymmetric information and asymmetric payoffs, *Games and Economic Behavior*, 13(1), 100-110.

Konow, James. (2000). Fair shares: Accountability and cognitive dissonance in allocation decisions. *American Economic Review* 90.4: 1072-1091.

Kvaløy, O., Eriksen, K., & Luzuriaga, M. (2014). *Risk-taking with other people's money* (No. 2014/21). University of Stavanger.

Linardi, S., & McConnell, M. A. (2011). No excuses for good behavior: Volunteering and the social environment. *Journal of Public Economics*, *95*(5-6), 445-454.

Loewenstein, G., Issacharoff, S., Camerer, C., & Babcock, L. (1993). Self-serving assessments of fairness and pretrial bargaining. *The Journal of Legal Studies*, *22*(1), 135-159.

Lundquist, Tobias, et al. (2009). The aversion to lying. *Journal of Economic Behavior & Organization* 70.1: 81-92.

McKelvey, Richard, and Talbot Page. (1986). Common knowledge, consensus, and aggregate information. *Econometrica* 54, 109-127.

Messick, D. M., & Sentis, K. P. (1979). Fairness and preference. *Journal of Experimental Social Psychology*, *15*(4), 418-434.

Mitra, Arnab, and Quazi Shahriar. (2018). Curbing Dishonesty: Interplay between Benefit and Norm. Working paper.

Mitzkewitz, M., & Nagel, R. (1993). Experimental results on ultimatum games with incomplete information, *International Journal of Game Theory*, 22(2), 171-198.

Moore, Don A., Lloyd Tanlu, and Max H. Bazerman. (2010). Conflict of interest and the intrusion of bias. *Judgment and Decision Making* 5.1: 37.

Nyarko, Yaw and Andrew Schotter. (2002). An experimental study of belief learning using elicited beliefs. *Econometrica* 70, 971-1005.

Ockenfels, A., & Werner, P. (2012). 'Hiding behind a small cake' in a newspaper dictator game. *Journal of Economic Behavior & Organization*, *82*(1), 82-85.

Pillutla, M. M., & Murnighan, J. K. (1996). Unfairness, anger, and spite: Emotional rejections of ultimatum offers, *Organizational Behavior and Human Decision Processes*, 68(3), 208-224.

Pronin, Emily, Daniel Y. Lin, and Lee Ross. (2002). The bias blind spot: Perceptions of bias in self versus others. *Personality and Social Psychology Bulletin* 28.3: 369-381.

Pronin, Emily. (2007). Perception and misperception of bias in human judgment. *Trends in cognitive sciences* 11.1: 37-43.

Pruckner, Gerald J., and Rupert Sausgruber. (2013). Honesty on the streets: A field study on newspaper purchasing. *Journal of the European Economic Association* 11.3: 661-679.

Rabin, Matthew. (1995). Moral preferences, moral constraints, and self-serving biases. *Department of Economics, UCB.*

Rapoport, A., Sundali, J. A., & Seale, D. A. (1996a). Ultimatums in two-person bargaining with one-sided uncertainty: Demand games. *Journal of Economic Behavior & Organization*, 30(2), 173-196.

Rapoport, A., & Sundali, J. A. (1996b). Ultimatums in two-person bargaining with one-sided uncertainty: Offer games, *International Journal of Game Theory*, 25(4), 475-494.

Rey-Biel, Pedro. (2009). Equilibrium play and best response to (stated) beliefs in normal form games. *Games and Economic Behavior* 65 (2009), 572–585.

Rigdon, M., Ishii, K., Watabe, M., & Kitayama, S. (2009). Minimal social cues in the dictator game. *Journal of Economic Psychology*, *30*(3), 358-367.

Rodriguez-Lara, Ismael, and Luis Moreno-Garrido. (2012). Self-interest and fairness: self-serving choices of justice principles. *Experimental Economics*15, no. 1: 158-175.

Rustichini, Aldo, and Marie Claire Villeval. (2014). Moral hypocrisy, power and social preferences. *Journal of Economic Behavior & Organization* 107: 10-24.

Schlag, Karl H. (2008). A New Method for Constructing Exact Tests Without Making Any Assumptions. Department of Economics and Business Working Paper 1109, Universitat Pompeu Fabra.

Schlag, Karl H., James Tremewan, and Joel J. Van der Weele. (2015). A penny for your thoughts: a survey of methods for eliciting beliefs. *Experimental Economics* 18.3: 457-490.

Shu, L. L., Mazar, N., Gino, F., Ariely, D., & Bazerman, M. H. (2012). Signing at the beginning makes ethics salient and decreases dishonest self-reports in comparison to signing at the end. *Proceedings of the National Academy of Sciences*, *109*(38), 15197-15200.

Straub, P. G., & Murnighan, J. K. (1995). An experimental investigation of ultimatum games: Information, fairness, expectations, and lowest acceptable offers, *Journal of Economic Behavior & Organization*, 27(3), 345-364.

Sutter, Matthias. (2009). Deception through telling the truth?! Experimental evidence from individuals and teams. *Economic Journal* 119.534: 47-60.

Thaler, R. H., & Sunstein, C. R. (2003). Libertarian paternalism. *American Economic Review*, *93*(2), 175-179.

Trautmann, Stefan T., and Gijs Kuilen. (2015). Belief elicitation: A horse race among truth serums. *Economic Journal* 125.589: 2116-2135.

Trivers, R. Deceit and Self-deception. (2011). Fooling Yourself the Better to Fool Others. Allen Lane, Penguin Books, New York, NY.

Vanberg, Christoph. (2008). Why do people keep their promises? An experimental test of two explanations. *Econometrica* 76.6: 1467-1480.

Tables

Table 1. Deception games

Game	Option	Sender's Payoff	Receiver's Payoff
Control	A	$5	$19
	B	$12	$5
Uncertainty	A	$5	$19
	B	$12	$8 or $2, equiprobable

Table 2. Two additional games

Game	Option	Sender's Payoff	Receiver's Payoff	Goal
Smaller-reduction	A	$5	$19	To measure the extent of the bias
	B	$12	$8	
Dictator	B_C	$12	$5	To check sender's risk preferences over the receiver's payoff
	B_U	$12	$8 or $2, equiprobable	

Table 3. Regression results on lying behavior

	Marginal Effects from Probit Models	
	(1)	**(2)**
Uncertainty Treatment	0.25	0.21
	(0.08)***	(0.08)***
Belief		0.02
		(0.004)***
Obs.	81	81
R^2	0.08	0.19
Log likelihood	–41.67	–36.72

Note: *Uncertainty Treatment* is a dummy for *Uncertainty_cont_belief* with respect to *Nudging_cont_belief*. Standard errors are in parentheses. ***, ** and * denote statistical significance at 1%, 5% and 10% level, respectively.

Figures

Figure 1. Sender-belief elicitation as per the discrete method in the *Uncertainty* treatment

Please select one:	*If it was actually $8, you'll get*	*If it was actually $2, you'll get*
☐ Certainly $8	$2.00	$0.50
☐ Probably $8	$1.80	$1.00
☐ Unsure	$1.50	$1.50
☐ Probably $2	$1.00	$1.80
☐ Certainly $2	$0.50	$2.00

Figure 2. Lying proportions

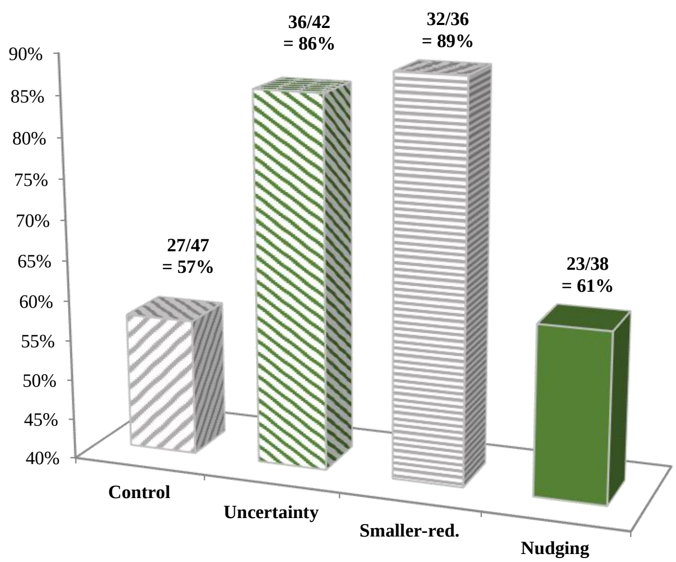

Figure 3. Average sender-belief about $8 realization as per the discrete method in the *Uncertainty* treatment

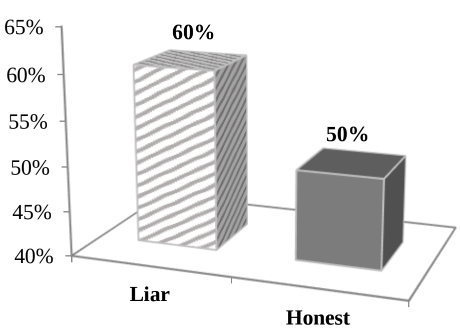

Figure 4. Histogram of sender-belief about $8 realization

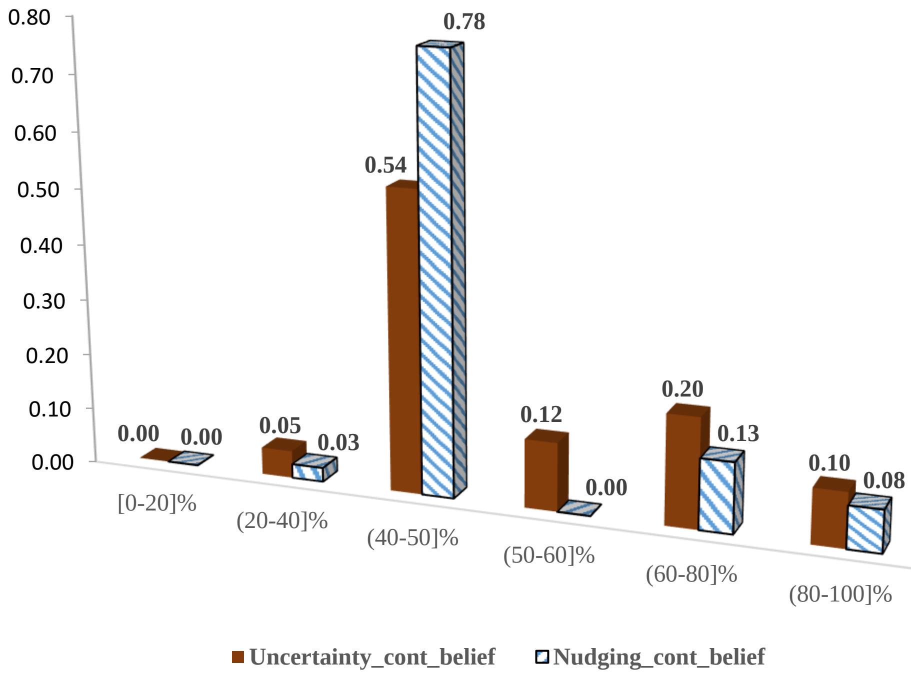

Appendix for Experimental Instructions (not intended for publication)

Control Treatment

Instructions for Senders **Sender's Registration Number: ______________**

Introduction: Welcome to an experiment in decision-making! Your registration number is written at the top of this page. This number will be used to identify you for payment. You must not communicate with the other participants until the end of the session. If you have any question at any time, please raise your hand and one of us will help you.
One of you will be randomly selected to be the "monitor" for today's experiment. The monitor will not participate in the experiment but will help the experimenter in conducting the experiment.

Anonymity: No one (directly or indirectly) associated with this experiment will ever know the real identity of any participant. Put differently, each participant's decisions will remain strictly anonymous, and therefore we will never be able to connect your real identity with your decisions.

The Experiment: This is a short experiment in decision making. In this experiment:

• Each of you who are located in this room will be called a "**Sender**" for this experiment.
• Each of you will be randomly matched with another student in a different room, who will be your "**Receiver**".
• Neither of you will ever know the real identity of the other, as discussed before.
• The money that you earn will be paid to you at the end of this experiment, in cash and in confidence.
• Two possible monetary payment options are available to you and your Receiver. The two payment Options are:

Option A: $5 to you (the Sender) **AND** $19 to the other student (your Receiver).

Option B: $12 to you (the Sender) **AND** $5 to the other student (your Receiver).

• You know what the Options are, but your Receiver **WILL NOT**.

• You will choose **ONE** of two possible messages, which will be sent to your Receiver. The two messages are:

Message A: "Option A will earn my Receiver more money than Option B."

Message B: "Option B will earn my Receiver more money than Option A."

• This is the <u>only task</u> you will need to perform in today's experiment.

• After seeing your message, your Receiver will CHOOSE ONE of the two Options, either Option A or Option B. Actual payments will be made according to the option your Receiver will choose. For instance, if your Receiver chooses Option A, then you will receive $5 and your Receiver will receive $19. On the other hand, if your Receiver chooses Option B, then you will receive $12 and your Receiver will receive $5.

• No Receiver will ever know the true Options, or the sums to be paid to you under the different Options. To show you how much information is shared with the Receivers we have attached the **Instructions for Receivers** here. (The Receivers in the other room are currently going through these instructions.) As you can see in this document, Receivers are unaware of the actual payments associated with the two options.

A Recap of the Steps:

• First, the experimenter, together with the monitor, will pair each of you with a separate Receiver.
• Second, you will choose **any one** of the two possible messages, **A or B**, which you want to send to your Receiver in a "Message Card" which will be given to you in few minutes. The experimenter will then collect these Message Cards back, carry them to the other room and pass on your message to your Receiver.
• Third, after seeing your message the Receiver will choose **any one** of the two possible options, **A or B**.
Then your earnings will be determined according to your Receiver's choice, and paid to you in private and in an anonymous manner. (If your Receiver chooses Option A, then you will receive $5 and your Receiver will receive $19. On the other hand, if your Receiver chooses Option B, then you will receive $12 and your Receiver will receive $5.)

Now, do you have any questions?

Please wait.

Instructions for Receivers **Receiver's Registration Number:____________**

Introduction: Welcome to an experiment in decision-making! Your registration number is written at the top of this page. This number will be used to identify you for payment. From now on, you should not communicate in any way with the other participants until the end of the session. If you have any question at any time, please raise your hand and one of us will help you.

Anonymity: No one (directly or indirectly) associated with this experiment will ever know the real identity of any participant. Put differently, each participant's decisions will remain strictly anonymous, and therefore we will never be able to connect your real identity with your decisions.

The Experiment: This is a short experiment in decision making. In this experiment,

• Each of you who are located in this room will be called a "**Receiver**" for this experiment.
• You have been randomly matched with another student in another room, who we will call your "**Sender**".
• Neither of you will ever know the real identity of the other, as discussed before.
• The money that you earn will be paid to you at the end of this experiment, in cash and in confidence.
• Two possible monetary payment options (A and B) are available to you (the Receiver) and your Sender in the experiment. We showed the two payment options to your Sender.

• **YOU** (the **Receiver**) will choose **ONE** of the two options, which will determine the payments to you and to your Sender. The only information you will have is a message your Sender has sent you.

• Your Sender was given a choice to send any **ONE** of the two messages, A or B. The two messages are described below.

Message A: "Option A will earn you (the Receiver) more money than Option B."
Message B: "Option B will earn you (the Receiver) more money than Option A."

• You have been given an envelope which contains the messages that your Sender decided to send you. Please open the envelope and take a look at the message your Sender decided to send you.

• We now ask you to choose either Option A or Option B. Your choice will determine the payments to you and to your Sender. You will never be told what sums were actually offered in the option not chosen (that is, if the message sent by your Sender was true or not). Moreover, you will never be told the sum your Sender actually receives.

We now ask you to choose any ONE of the two options.

I (the Receiver) decide to choose (please circle ONLY ONE):

Option A **Option B**

When you are done circling one of the above options, please hand over this page to the experimenter. The experimenter will make the payments according to the option you just chose.

Uncertainty Treatment

<u>**Instructions for Senders**</u> <u>**Sender's Registration Number: ____________**</u>

<u>**Introduction**</u>**:** Welcome to an experiment in decision-making! Your registration number is written at the top of this page. This number will be used to identify you for payment. You must not communicate with the other participants until the end of the session. If you have any question at any time, please raise your hand and one of us will help you.
One of you will be randomly selected to be the "monitor" for today's experiment. The monitor will not participate in the experiment but will help the experimenter in conducting the experiment.

<u>**Anonymity**</u>**:** No one (directly or indirectly) associated with this experiment will ever know the real identity of any participant. Put differently, each participant's decisions will remain strictly anonymous, and therefore we will never be able to connect your real identity with your decisions.

<u>**The Experiment**</u>**:** This is a short experiment in decision making. In this experiment:

• Each of you who are located in this room will be called a "**Sender**" for this experiment.
• Each of you will be randomly matched with another student in a different room, who will be your "**Receiver**".
• Neither of you will ever know the real identity of the other, as discussed before.
• The money that you earn will be paid to you at the end of this experiment, in cash and in confidence.
• Two possible monetary payment options are available to you and your Receiver. The two payment Options are:

Option A: $5 to you (the Sender) **AND** $19 to the other student (your Receiver).

Option B: $12 to you (the Sender) **AND**:
Either $8 or $2 to the other student (your Receiver), with equal chances for $8 and $2.

• You know what the Options are, but your Receiver **WILL NOT**. (Notice that you do not know and will not know which of the two amounts, $8 or $2, your receiver actually gets under Option B. The experimenter, with the help of the monitor, will determine this in a random manner which will be discussed later.)

• You will choose **ONE** of two possible messages, which will be sent to your Receiver. The two messages are:

Message A: "Option A will earn my Receiver more money than Option B."

Message B: "Option B will earn my Receiver more money than Option A."

• This is the <u>only task</u> you will need to perform in today's experiment.

• After seeing your message, your Receiver will CHOOSE ONE of the two Options, either Option A or Option B. Actual payments will be made according to the option your Receiver will choose. For instance, if your Receiver chooses Option A, then you will receive $5 and your Receiver will receive $19. On the other hand, if your Receiver chooses Option B, then you will receive $12 and your Receiver will receive either $8 or $2 (the two possibilities are equally likely).

• No Receiver will ever know the true Options, or the sums to be paid to you under the different Options. To show you how much information is shared with the Receivers we have attached the **Instructions for Receivers** here. (The Receivers in the other room are currently going through these instructions.) As you can see in this document, Receivers are unaware of the actual payments associated with the two options.

Randomization Under Option B:

• Recall that your Receiver's earnings under Option B is either $8 or $2 where the two amounts are equally likely. We will use a simple randomization method to determine which one of these two amounts will actually be paid to your Receiver if he/she chooses Option B.
• In the randomization, the monitor will toss a coin; if the coin toss results in a Head your Receiver will get $8 when you choose Option B, and if a coin toss results in a Tail your Receiver will get $2 when you choose Option B. We will do this for each Sender separately and note down the result of each draw against each Sender's registration number. (You will **not** be told about the results during or after the experiment.)

A Recap of the Steps:

• First, the experimenter, together with the monitor, will pair each of you with a separate Receiver.
• Second, the experimenter, together with the monitor, will determine your Receiver's earnings under Option B using the randomization method described above.
• Third, you will choose **any one** of the two possible messages, **A or B**, which you want to send to your Receiver in a "Message Card" which will be given to you in few minutes. The experimenter will then collect these Message Cards back, carry them to the other room and pass on your message to your Receiver.
• Fourth, after seeing your message the Receiver will choose **any one** of the two possible options, **A or B**.
Then your earnings will be determined according to your Receiver's choice, and paid to you in private and in an anonymous manner. (If your Receiver chooses Option A, then you will receive $5 and your Receiver will receive $19. On the other hand, if your Receiver chooses Option B, then you will receive $12 and your Receiver will receive either $8 or $2 based on the randomization described above.)

Now, do you have any questions?

Please wait.

Dictator Treatment

<u>Instructions for Senders</u> **<u>Sender's Registration Number: _____________</u>**

<u>Introduction</u>: Welcome to an experiment in decision-making! Your registration number is written at the top of this page. This number will be used to identify you for payment. You must not communicate with the other participants until the end of the session. If you have any question at any time, please raise your hand and one of us will help you.
One of you will be randomly selected to be the "monitor" for today's experiment. The monitor will not participate in the experiment but will help the experimenter in conducting the experiment.

<u>Anonymity</u>: No one (directly or indirectly) associated with this experiment will ever know the real identity of any participant. Put differently, each participant's decisions will remain strictly anonymous, and therefore we will never be able to connect your real identity with your decisions.

<u>The Experiment</u>: This is a short experiment in decision making. In this experiment:

• Each of you who are located in this room will be called a "**Sender**" for this experiment.
• Each of you will be randomly matched with another student in a different room, who will be your "**Receiver**".
• Neither of you will ever know the real identity of the other, as discussed before.
• The money that you earn will be paid to you at the end of this experiment, in cash and in confidence.
• Two possible monetary payment options are available to you and your Receiver. The two payment Options are:

Option A: $12 to you (the Sender) **AND** $5 to the other student (your Receiver).

Option B: $12 to you (the Sender) **AND**:
Either $8 or $2 to the other student (your Receiver), with equal chances for $8 and $2.

• You know what the Options are, but your Receiver **WILL NOT**. (Notice that you do not know and will not know which of the two amounts, $8 or $2, your receiver actually gets under Option B. The experimenter, with the help of the monitor, will determine this in a random manner which will be discussed later.)

• You **will CHOOSE ONE of the two Options, either Option A or Option B.** Actual payments will be made according to the option you will choose. For instance, if you choose Option A, then you will receive $12 and your Receiver will receive $5. On the other hand, if you choose Option B, then you will receive $12 and your Receiver will receive either $8 or $2 (the two possibilities are equally likely).

• No Receiver will ever know the true Options, or the sums to be paid to you under the different Options. To show you how much information is shared with the Receivers we have attached the **<u>Instructions for Receivers</u>** here. (The Receivers in the other room are currently going through these instructions.) As you can see in this document, Receivers are unaware of the actual payments associated with the two options.

Randomization Under Option B:

• Recall that your Receiver's earnings under Option B is either $8 or $2 where the two amounts are equally likely. We will use a simple randomization method to determine which one of these two amounts will actually be paid to your Receiver if he/she chooses Option B.
• In the randomization, the monitor will toss a coin; if the coin toss results in a Head your Receiver will get $8 when you choose Option B, and if a coin toss results in a Tail your Receiver will get $2 when you choose Option B. We will do this for each Sender separately and note down the result of each draw against each Sender's registration number. (You will **not** be told about the results during or after the experiment.)

A Recap of the Steps:

• First, the experimenter, together with the monitor, will pair each of you with a separate Receiver.
• Second, the experimenter, together with the monitor, will determine your Receiver's earnings under Option B using the randomization method described above.
• Third, you will choose **any one** of the two possible options, **A or B**.
• Fourth, your earnings will be determined according to your choice, and paid to you in private and in an anonymous manner. (If you choose Option A, then you will receive $12 and your Receiver will receive $5. On the other hand, if you choose Option B, then you will receive $12 and your Receiver will receive either $8 or $2 based on the randomization described above.)

Now, do you have any questions?

Please wait.