People Are More Moral in Uncertain Environments

Yiting Chen Songfa Zhong*

November 2024

Abstract

We conduct a series of experiments and document a robust behavioral pattern whereby people behave more morally in uncertain environments than degenerate deterministic ones. We show that this pattern is weakened when the moral implication of behavior is diminished or when uncertainty pertains to others rather than oneself. These findings are incompatible with standard models that respect dominance. We propose a mechanism based on the anxiety aspect of uncertain environments whereby people act morally as if their moral behavior can help deliver a better outcome. We further delve into the complexity aspect of uncertainty to arrive at a more comprehensive understanding of these findings.

Keyword: uncertainty, morality, lying aversion, social preference, experiment
JEL Classification: C91, D81, D91

---

*Chen: Department of Economics, Lingnan University, yitingchen26@gmail.com. Zhong: Department of Economics, The Hong Kong University of Science and Technology, and Department of Economics, National University of Singapore, zhongsongfa@gmail.com. We thank the helpful comments and discussions of Soo Hong Chew, Syngjoo Choi, Itzhak Gilboa, Lorenz Goette, Yoram Halevy, Jian Li, Yucheng Liang, Wooyoung Lim, George Loewenstein, Bin Miao, Jianying Qiu, Xiangyu Qu, Emanuel Vespa, Peter Wakker, and Chen Zhao, seminar participants at Carnegie Mellon University, CUHK-HKU-HKUST Joint Theory Seminar Series, Erasmus University Rotterdam, Hong Kong Baptist University, Hong Kong University of Science and Technology, Tsinghua University, and Virtual East Asia Experimental and Behavioral Economics Seminar Series, and participants at Risk, Uncertainty and Decision (2023); Decision: Theory, Experiments, and Applications (2022); and ESA Global Around-the-Clock Virtual Conference (2020). First version: November 2020. For financial support, Chen acknowledges National Natural Science Foundation of China (72303195) and Zhong acknowledges Singapore Ministry of Education (Academic Research Fund Tier 1 from the National University of Singapore).

1 Introduction

Models in decision-making under uncertainty commonly assume that people think through each of the possible outcomes and weigh the outcomes according to their probabilities. This assumption is appealing both normatively and descriptively. However, it has been challenged by a growing body of evidence often arising from two complementary perspectives. The first is related to the *anxiety* aspect of uncertain environments: Uncertainty arouses the feelings of fear, anxiety, and stress. Therefore, people may cope with the uncertainty by acting in a certain way to gain a sense of control.[1] The second perspective concerns the *complexity* aspect of uncertain environments: It is cognitively difficult for people to think through contingencies and aggregate probabilistic outcomes. Therefore, they may make suboptimal decisions and violate some fundamental assumptions in standard models.[2] Building on these two lines of inquiry, this paper examines how people make moral decisions in uncertain environments compared with certain ones.[3]

We consider environments in which people make a binary choice: whether to act morally with a cost or not to act. In uncertain environments, they receive a lottery $(h, p; l)$ that yields high outcome h with probability p and low outcome l otherwise. In two corresponding degenerate certain environments, they receive high outcome h in one and low outcome l in the other. In standard models of decision-making under uncertainty, such as expected utility and prospect theory, people ought to respect the principle of dominance: If they choose the same option in both degenerate certain environments, regardless of moral or immoral, they will choose the same way in uncertain environments. By contrast, alternative behavioral patterns that

[1]See Henslin (1967); Langer (1975); Quattrone and Tversky (1984); Shafir and Tversky (1992); Tversky and Shafir (1992); Risen and Gilovich (2008); Stefan and David (2013).

[2]See Shafir and Tversky (1992); Tversky and Shafir (1992); Gneezy, List, and Wu (2006); Charness and Levin (2009); Cason and Plott (2014); Esponda and Vespa (2014); Li (2017); Martínez-Marquina, Niederle, and Vespa (2019); Esponda and Vespa (2023); Enke and Graeber (2023); Oprea (2024).

[3]To clarify our terminology, in this paper, we do not differentiate risk and uncertainty in the sense of known and unknown probabilities. We also do not provide a definition of morality, and instead assume that it is more moral to tell the truth than to lie and to share than not to share.

violate dominance are permissible under the two perspectives on uncertainty. From the anxiety perspective, people may obtain a sense of control by acting in a more moral manner under uncertainty than certainty. This could be attributed to a perceived link between moral behavior and favorable consequence, which is reflected in widespread beliefs such as a just world, moralistic gods, or karma.[4] From the complexity perspective, people who fail in contingent reasoning may behave differently in uncertain environments, either more or less morally, compared with certain environments. Based on this setting, we conduct a series of laboratory experiments to document a set of behavioral observations. These observations challenge standard models and provide insights into the two perspectives on uncertainty.

In the main experiment, we incorporate uncertainty into the dice game paradigm proposed by Fischbacher and Föllmi-Heusi (2013), which is widely adopted to examine truth-telling behavior. Subjects receive a lottery $(h, \frac{n}{6}; l)$ in the form of six boxes numbered from 1 to 6 with n box(es) containing h and $6 - n$ box(es) containing l. They roll a die in their mind—randomly choosing a number between 1 and 6—to select one of the six boxes (Kajackaite and Gneezy, 2017). Subsequently, subjects are informed that one exact box out of the six contains an additional 4 Chinese yuan (referred to as RMB4 hereafter). They are asked to report their initial box selection to receive the corresponding payoff in that box. Reporting the box with the RMB4 (reporting +4) indicates either the truth based on mental die-rolling with a $\frac{1}{6}$ chance or a lie to maximize the payoff. Although lying cannot be observed individually, it can be measured at the aggregate level by the difference between $\frac{1}{6}$ and the observed proportion of reporting +4. After subjects make their decisions, the uncertainty of lottery $(h, \frac{n}{6}; l)$ is resolved in front of them at the end of the experiment.

We include three spreads between h and l—$(40, \frac{n}{6}; 0)$, $(30, \frac{n}{6}; 10)$, and $(22, \frac{n}{6}; 18)$—and seven levels of winning probability, $\frac{n}{6} \in \{0, \frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}, 1\}$, which give rise to 21 decisions in our within-subject experiment. Our design allows us to compare lying

[4]People tend to believe in a just world, in which moral behavior will be rewarded with a desirable fate and immoral conduct will be punished with a negative fate (Lerner, 1980; Bénabou and Tirole, 2006). Religious individuals may believe that moralistic gods would reward the righteous and punish wrongdoers (Purzycki et al., 2016; Enke, 2019). According to the principle of karma, current actions will have consequences in the uncertain future (Converse, Risen, and Carter, 2012).

behavior between the uncertain and the degenerate deterministic environments, and to examine whether the degree of uncertainty, measured by the spread between h and l, matters.

We observe that subjects are less likely to lie in uncertain environments than in certain environments. Specifically, the proportions of reporting +4 are 59.1 percent under uncertainty and 74.8 percent (78.2 percent) under certainty with high (low) outcomes. Moreover, the pattern is more pronounced when the spread between h and l is wider. Namely, it is significant for $(40, \frac{n}{6}; 0)$ and $(30, \frac{n}{6}; 10)$, but insignificant for $(22, \frac{n}{6}; 18)$.

This pattern contradicts standard models with the dominance property. To explore potential mechanisms, we conduct two supplementary experiments by adjusting two key components in the main experiment—morality and uncertainty—to examine whether these two changes affect the behavioral pattern.

The first supplementary experiment reduces subjects' moral concern. The new design is based on the "no dice" condition in Fischbacher and Föllmi-Heusi (2013), in which subjects directly choose their preferred payoff rather than report a die-rolling result to determine the payoff. Fischbacher and Föllmi-Heusi show that some subjects do not choose the highest payoff, and speculate that such behavior may reflect an aversion to greed, which is a weaker moral concern than an aversion to lying. We modify our main experiment to incorporate their design. Here, subjects receive a lottery $(h, \frac{n}{6}; l)$ in the six-box frame and are informed which box contains an additional RMB4. Without rolling a mental die and reporting the number, subjects directly choose their preferred box. We find that subjects are less likely to choose the box with +4 under uncertainty than under certainty, but this difference is substantially smaller than the behavioral pattern in the main experiment.

The second supplementary experiment reduces the uncertainty concern for subjects. Specifically, an anonymously paired partner receives the lottery $(h, \frac{n}{6}; l)$ in the six-box frame, while subjects as decision-makers receive a fixed amount of money. Similar to the main experiment, subjects decide whether to lie for the additional RMB4 through the mental-die-rolling and reporting process. Different from the main experiment, subjects do not face uncertainty themselves, and thus their desire for a good outcome of uncertainty should be weaker for the anonymous partners than for themselves. We find that subjects are less likely to lie under uncertainty than certainty, but this difference is inconsistent across different payoff conditions and is substantially smaller than that observed in the main experiment.

We conduct three additional experiments to further investigate the robustness and generalizability of the main pattern. In the first experiment, we involve one box that contains a loss instead of a gain of RMB4 as in the main experiment. In the second experiment, we employ a dice game in which subjects decide whether to tell the truth after the uncertainty of $(h, \frac{n}{6}; l)$ is resolved (but kept unknown), instead of before the uncertainty is resolved as in the main experiment. In the third experiment, we explore the domain of other-regarding behavior and adopt a modified dictator game, in which subjects receive the lottery $(h, \frac{n}{6}; l)$ and decide whether to share half of the realized payoff with an anonymously paired recipient. Results from these three experiments support a robust and generalizable pattern whereby people are more moral in uncertain environments compared with certain ones.

To understand these findings, we explore the implications of both the anxiety and complexity perspectives on uncertainty. The anxiety perspective can account for the main findings: People who face uncertainty may obtain a sense of control by behaving morally. Moreover, under some assumptions, it could also help understand the weakened patterns in the two supplementary experiments. Specifically, suppose that subjects obtain a weaker sense of control when their behavior has weaker moral concern, and they have a weaker desire for control when uncertainty is imposed on others, they would then have weaker incentives to act morally under uncertainty in the two supplementary experiments compared with that in the main experiment. We further propose two frameworks for this explanation based on the notions of magical thinking and quasi-magical thinking. Namely, people act morally as (or as if) they perceive a connection between their moral behavior and the resolution of uncertainty.

Moreover, uncertainty is inherently more complex than certainty. In this regard, complexity probably plays some role in our experiments and contributes to some of the observed differences between uncertainty and certainty. While it is worth noting that complexity perspective is silent about whether people would be more or less moral under uncertainty and whether the difference would be strengthened or weakened in the two supplementary experiments, additional assumptions can be imposed to accommodate these observations.

This paper contributes to three lines of research. First, by extending to the domain of moral behavior, our study adds to numerous anomalies in decision-making under uncertainty (Shafir and Tversky, 1992; Tversky and Shafir, 1992; Gneezy, List, and Wu, 2006; Charness and Levin, 2009; Cason and Plott, 2014; Esponda and Vespa, 2014; Li, 2017; Martínez-Marquina, Niederle, and Vespa, 2019; Esponda and Vespa, 2023). While the literature has proposed the complexity-based explanation for these anomalies, we suggest that the anxiety-based explanation may provide a complementary view. More specifically, when people find uncertain environments complex and have difficulty thinking through all possible outcomes, they act morally as (or as if) they believe that their moral behavior can lead to a favorable outcome. Relatedly, in their decomposition of the failure in contingent reasoning, Martínez-Marquina, Niederle, and Vespa (2019) suggest that "one reason for the difficulties in the probabilistic treatment may come from the subject's belief that her actions can influence which state of the world realizes." In their study of the disjunction effect, Tversky and Shafir (1992) show that students choose to have a vacation to celebrate (seek consolation) if they pass (fail) the exam, but choose not to have a vacation if they do not know the outcome of the exam. Similarly, one could argue that students opt not to take vacations as if doing so could potentially jinx their exam results.

Second, our study is closely related to notions including illusion of control, tempting fate, magical thinking, and quasi-magical thinking (Henslin, 1967; Langer, 1975; Quattrone and Tversky, 1984; Shafir and Tversky, 1992; Tversky and Shafir, 1992; Risen and Gilovich, 2008; Stefan and David, 2013). As summarized by John Dewey (1929), "in the absence of actual certainty in the midst of a precarious and hazardous world, men cultivated all sorts of things that would give them the feeling of certainty." To reduce uncertainty about themselves and their surroundings, people tend to identify with social and religious groups (Hogg, 2007) and behave as if their actions can affect the resolution of uncertainty, for example, throwing the dice harder for larger numbers, tolerating cold for a longer period to be "diagnosed" with a longer life expectancy, and voting to "induce" other like-minded persons to vote (Henslin, 1967; Langer, 1975; Quattrone and Tversky, 1984; Stefan and David, 2013). These studies suggest that people may act as (or as if) they perceive a connection between their actions and the outcomes of uncertainty in individual choice settings, or between their actions and other people's actions in strategic settings. Here we show that this anxiety aspect underpins moral decision-making under uncertainty.[5]

Third, this paper contributes to studies on morality under uncertainty. In one strand of the literature, people consider ex ante and ex post fairness when they allocate the winning odds between themselves and others or decide the probabilities of two payoff distributions (Krawczyk and Le Lec, 2010; Brock, Lange, and Ozbay, 2013; Sandroni, Ludwig, and Kircher, 2013; Andreoni et al., 2020).[6] In another strand of the literature, people have some wiggle room or excuses to behave selfishly when the uncertainty concerns others (Dana, Weber, and Kuang, 2007; Haisley and Weber, 2010; Exley, 2016; Gino, Norton, and Weber, 2016; Garcia, Massoni, and Villeval, 2020). In contrast, we observe that when individuals themselves face uncertainty, they are more likely to behave morally. In this regard, our study distinguishes itself from the above studies concerning the role of uncertainty in moral decision-making.

We organize the rest of the paper as follows. Section 2 details the design of experiments and Section 3 reports the results. We discuss the implications and explanations in Section 4 and offer some concluding remarks in Section 5. Supplemental Materials (Chen and Zhong, 2024) contain further details about experiments, results,

[5]Kellner, Reinstein, and Riener (2019) compare donation after winning a lottery with commitment to donate before winning, and observe that subjects are more willing to donate in the latter case. Whereas the authors discuss a similar anxiety-based explanation, their observations are also consistent with other alternatives, including ex ante fairness, loss aversion, and signalling. Chew and Li (2021) suggest that sin stock aversion can in part be due to a belief in karma whereby investing in sin stock may lead to bad outcomes, but social preference may also play a role.

[6]See Machina (1989); Karni and Safra (2002); Trautmann and Wakker (2010); Fudenberg and Levine (2012); Saito (2013) for theoretical discussions. Our paper sheds light on some observations in this literature. For example, in Brock, Lange, and Ozbay (2013), the dictator decides the number of tokens to share with the recipient in different treatments with uncertainty. They observe that the dictator gives more in the treatments when both players face uncertainty, compared with treatments when only the recipient faces uncertainty. While this observation is not consistent with standard models, including ex ante and ex post fairness (Krawczyk and Le Lec, 2016; Brock, Lange, and Ozbay, 2016), it is in line with our observations.

and theoretical discussions.

2 Experimental Design

In this section, we describe the design of the main experiment, two experiments to shed light on the mechanisms, and three experiments to examine the robustness and generalizability of results from the main experiment.

2.1 Main Experiment

Our main experiment builds on the dice game in Fischbacher and Föllmi-Heusi (2013), which provides a paradigm to examine truth-telling behavior. In their experiment, subjects report the outcome of a die that they roll privately and receive a monetary payoff based on their report. If subjects only care about the monetary payoff, they report the outcome with the highest monetary payoff regardless of the actual result of the die roll. If subjects have a strong preference for truth-telling, they report the actual outcome regardless of the resulting monetary payoff. Moreover, subjects can partially lie by reporting an outcome that delivers a falsely higher, but not the highest, payoff. A notable feature of this paradigm is that while lying behavior is undetectable at the individual level, it can be inferred at the aggregate level. In a meta-analysis of 90 experimental studies based on this paradigm, Abeler, Nosenzo, and Raymond (2019) show that individuals exhibit a preference for being honest and for being seen as honest. Our main experiment incorporates exogenous uncertainty in this paradigm, and is denoted as the *Dice Game* experiment.

Design. We endow subjects with a lottery $(h, p; l)$ that pays high outcome h with probability p and low outcome l otherwise. We refer to p as the winning probability hereafter. We examine whether and how the endowed lottery affects lying behavior in four steps: Subjects receive a lottery, randomly choose a number between 1 and 6 in their mind, learn which number carries an additional payoff a, and report the number they randomly chose in their mind. If they report the number that yields the additional payoff a, they receive both the lottery and a—that is, $(h + a, p; l + a)$;

otherwise, they only receive the lottery $(h, p; l)$. If subjects are perfectly honest, approximately $\frac{1}{6}$ of them will report the number with a. The deviation from the expected proportion $\frac{1}{6}$ reveals the prevalence of dishonesty at the aggregate level.

We vary the parameters of the lotteries $(h, p; l)$ in two ways. First, we include seven levels of the winning probability $p \in \{0, \frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}, 1\}$. Subjects make decisions under degenerate deterministic situations when $p = 1$ and $p = 0$, and under uncertain situations when $p \in \{\frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}\}$. This allows us to examine the truth-telling behavior under uncertainty compared with certainty. Moreover, we are also interested in whether the effect of uncertainty varies with the winning probability, from $\frac{1}{6}$ to $\frac{5}{6}$. Because the absolute monetary cost of truth-telling is fixed to be a regardless of the winning probability, we expect the effect of uncertainty to be similar across these five uncertainty conditions. Second, we include three pairs of high payoff h and low payoff l—$(40, p; 0)$, $(30, p; 10)$, and $(22, p; 18)$ and set $a = 4$. This enables us to examine the effect of the spread between high and low payoffs, which reflects a sense of riskiness. We expect a limited effect of uncertainty under the payoff pair $(22, p; 18)$, due to its relatively weak sense of riskiness. The combination of varying winning probabilities and payoff pairs gives rise to 21 lotteries. Correspondingly, in our within-subject design, each subject makes 21 rounds of decisions.

Note that we modify the dice game design in several respects to facilitate implementation. First, for the privately observed outcome that underpins lying decisions, we employ a setting with two states. One state occurs with probability $\frac{1}{6}$ and delivers an additional payoff, and the other state occurs with probability $\frac{5}{6}$ and carries no extra monetary incentive. This is to simplify the choice environment and help avoid vagueness in the moral evaluation of partial lies. Second, we adopt a mental die-rolling process rather than a physical one (Kajackaite and Gneezy, 2017). Specifically, subjects are asked to randomly choose a number between 1 and 6 in their mind before they learn about which number carries the additional payoff. This is to facilitate implementation of the online experiments, as explained below in footnote 7. Kajackaite and Gneezy (2017) show that the tendency to lie is stronger under the mental die-rolling process than in Fischbacher and Föllmi-Heusi's (2013) setting. However, the potential difference between using a mental die and a physical die should not impact our investigation, because our focus is the comparison of truth-telling behavior between uncertainty and certainty. In addition, we include an experiment to examine the robustness with respect to using a physical die, as introduced in Section 2.3.

Implementation. We implement each round of decisions in four steps (see Supplemental Materials Figure A.1 for the interface). First, subjects receive the endowed lottery $(h, p; l)$ in a frame of six boxes numbered from 1 to 6. Specifically, the process is described as "*There are x box(es) containing h and $6 - x$ box(es) containing l.*" We construct the winning probability $p \in \{0, \frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}, 1\}$ using different values of x from 0 to 6, which is matched with different spreads of h and l. Second, subjects are asked to randomly choose a number between 1 and 6 as the box they choose in their mind and write the number on a piece of paper. This process helps strengthen the sense of dishonesty if subjects decide to report a different number later. However, to address the potential effect of observability, we make it clear to subjects that they will not be asked to display this record at any time. Third, subjects are explicitly informed that one of the six boxes has an additional payoff RMB4—i.e., "*This box is box y.*" Box y, with the additional payoff, varies from round to round. Last, subjects are asked to report the box they chose in their mind in the second step. Subjects make 21 decisions presented in random order without feedback. One of the 21 decisions is randomly selected to pay each subject. To induce segregation of decisions, we include a 10-second blank screen after each round.

It is important to emphasize to subjects that the experimenter does not manipulate the experiment and has little room to do so. We inform subjects that the distribution of h and l will be determined randomly and independent of the additional payoff a. Specifically, the distribution of a is predetermined and revealed in the third step of the decision-making process, while the distribution of h and l among the boxes will be resolved randomly in front of subjects at the end of the experiment.

We recruited subjects to join an online experiment at the scheduled time and date. After subjects entered the online meeting room, the experimenter shared the screen and read the instructions aloud to subjects. At the beginning of the experiment, subjects answered eight comprehension questions with feedback and explanations. This was to familiarize them with the tasks and to reduce potential misunderstanding. Next, subjects started the 21 rounds of decisions. The whole study ended with a short survey (see Supplemental Material Section D for details). After all subjects in the same session finished the experiment, for each subject, the experimenter randomly chose one round to implement to pay the subject, and randomly drew box(es) to contain the high payoff h for the chosen round. The randomization was done using the RANDBETWEEN function of Excel, and the randomization process was displayed to subjects in real time through the shared screen.

We conducted the online experiment between September and October 2022 with 107 university student subjects in China.[7] The experiment consisted of 9 sessions, with 10 to 20 subjects in each session. On average, the experiment took around 45 minutes, which included reading the instructions and real time randomization. Payment for each subject included a show-up fee of RMB20 plus the payoff from one of the 21 decisions. The average payment was RMB44.1 ($\approx$ USD6.4).

2.2 Two Experiments on Mechanisms

The Dice Game experiment offers an approach to document the potential difference in truth-telling behavior between uncertain and certain environments. There are two key components in this design, a moral component and an uncertainty component. Next, we design two supplementary experiments, in which we separately alter one component and observe whether individuals' behavior is affected by this adjustment. This practice provides insights into the mechanisms underlying our main findings.

Direct Choice Experiment. We conduct a *Direct Choice* experiment in which

[7]This experiment and the following three experiments were conducted online using student samples in China, because of the COVID-19 restrictions in China during the period in which we conducted the experiments. Our last two experiments were conducted in person earlier at the National University of Singapore, because there were no COVID-19 restrictions in Singapore at that time. Moreover, in the Ex Ante Resolution experiment conducted earlier in Singapore (as introduced in Section 2.3), the expected value of the lottery was SGD10 and the additional payoff SGD2; in the four experiments in China, we kept this ratio at 5 but adjusted the amount to align with the local subject payment, that is, RMB20 and RMB4, respectively.

the moral implication of the choice behavior is weakened while the uncertainty in the choice environment remains unchanged. This experiment is based on the "no dice" condition in Fischbacher and Föllmi-Heusi (2013), whereby subjects are given several alternative payoffs and they directly choose one payoff to receive without rolling a die. Their results show that although choosing the highest payoff does not involve lying, a proportion of subjects choose to avoid this payoff-maximizing option. This phenomenon is interpreted as an aversion to being greedy or being seen as greedy (see Arad (2014) and Tjøtta (2019) for related evidence). While greed is a common aspect of moral sentiment, it arouses a weaker moral implication than dishonesty (Fischbacher and Föllmi-Heusi, 2013). Therefore, this design is adopted to investigate whether the difference in behavior between uncertainty and certainty (if any) is weaker under greed aversion, compared with that under lying aversion.

Without the die-rolling requirement, the experiment consists of three steps. First, subjects are endowed with a lottery $(h, p; l)$ in a frame of six boxes and told that "*There are x box(es) containing h and $6 - x$ box(es) containing l.*" Second, subjects are informed which box carries an additional payoff of RMB4. Third, subjects choose their preferred box. For this experiment (and the Second Party and Dice Game Loss experiments introduced below), we closely follow implementation of the Dice Game experiment in other respects.

Second Party Experiment. We conduct a *Second Party* experiment in which we reduce the uncertainty concern for subjects. The target of uncertainty is changed from our subjects to anonymously paired partners. Put differently, subjects face uncertainty as a Second Party or bystander. This design involves two players. Player A is given a lottery $(h, p; l)$ but makes no decisions. Player B, as the second party, receives a fixed payoff and makes decisions with honesty concerns. For the lottery received by Player A, we include six payoff pairs: $(40, p; 0)$, $(30, p; 10)$, $(22, p; 18)$, $(20, p; 0)$, $(15, p; 5)$, and $(11, p; 9)$, with seven levels of the winning probability $p \in \{0, \frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}, 1\}$. Player B is endowed with a fixed amount of RMB21 and may need to decide whether to lie for an additional payoff RMB4 for herself. Correspondingly, in our within-subject design, each Player B makes 42 rounds of decision. We adopt this set of parameters with consideration for social preferences (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Andreoni and Miller, 2002; Charness and Rabin, 2002). In the first three payoff pairs, Player B faces disadvantageous (advantageous) inequality when the high (low) payoff of the lottery occurs. In the last three payoff pairs, Player B always faces advantageous inequality regardless of the outcome of the lottery.

Similar to the Dice Game experiment, there are four steps for each round of decision. First, subjects play the role of Player B and are informed that an anonymously paired Player A receives a lottery $(h, p; l)$: "*There are x box(es) containing h and $6 - x$ box(es) containing l.*" Second, subjects randomly choose a number between 1 and 6 in their mind and write this number on a piece of paper. Third, subjects learn that their own payoffs are also in these six boxes; specifically, "*Box y contains RMB25 and the remaining five boxes contain RMB21.*" Last, subjects are asked to select the numbered box they chose in their mind in the second step.[8]

2.3 Three Experiments on Robustness and Generalizability

Dice Game Loss Experiment. We conduct a *Dice Game Loss* experiment to investigate the tendency to lie to avoid losses under uncertainty versus certainty. Subjects face an ex ante probability $\frac{1}{6}$ to lose RMB4 and they can lie to avoid the loss. The design follows the Dice Game experiment with four steps. The main difference is that, in the third step, subjects are informed that one exact box involves a payoff deduction of RMB4. To allow potential deduction, we add RMB4 to the original parameters and have three payoff pairs $(44, p; 4)$, $(34, p; 14)$, and $(26, p; 22)$.

Ex Ante Resolution Experiment. We conduct an *Ex Ante Resolution* experiment, which allows us to examine the robustness of the effect of uncertainty on truth-telling behavior. This experiment differs from the Dice Game experiment in three respects. First, when subjects make decisions, uncertainty has been resolved

[8]In this experiment, all subjects were asked to make decisions as Player B. After all subjects in the same session finished the experiment, the experimenter randomly assigned a role for each subject using the RANDBETWEEN function. Next, subjects were randomly paired and the payoffs of each pair were determined by Player B's decisions and chances. The remaining randomization followed the Dice Game experiment.

but kept unknown.[9] Second, instead of asking subjects to roll a mental die as in the online experiments, we provide subjects with a physical die and a cup in an in-person laboratory environment. Third, there are three payoff pairs for the lottery $(h, p; l)$, $(20, p; 0)$, $(15, p; 5)$, and $(11, p; 9)$, and there is an additional payoff SGD2 in one of the six boxes.

Dictator Game Experiment. We conduct a *Dictator Game* experiment to test the generalizability of the effect of uncertainty in the domain of other-regarding behavior. We extend our experimental framework to the dictator game, which is a classic paradigm to examine sharing behavior. In a standard dictator game, two players are anonymously paired: The dictator is endowed with a fixed sum of money and the recipient is endowed with nothing. The dictator decides any amount between 0 and the fixed sum to share with the recipient, which captures the degree of departure from narrowly defined selfishness. In our experiment, we endow the dictator with lotteries in the form of $(h, p; l)$. Instead of allowing the dictator to share any proportion of the payoff, we use a binary design: The dictator chooses to share evenly or to share nothing with the recipient. If the dictator chooses to share evenly, both the dictator and the recipient receive the same amount, $\frac{h}{2}$ with probability p and $\frac{l}{2}$ otherwise. If the dictator chooses to share nothing, she receives the originally endowed lottery and the recipient receives nothing. Compared with a continuous set of choices, the binary design helps enhance the perception of equality for the even-split option (Bolton, Brandts, and Ockenfels, 2005). In addition, a meta-analysis of dictator game experiments reveals that sharing evenly and sharing nothing are two modal choices (Engel, 2011).

Each round consists of three steps. First, subjects are endowed with the lottery $(h, p; l)$ in the six-box frame. Second, they learn that one box contains the sharing

[9] To reduce suspicion that the experimenter manipulated the uncertainty, we recorded a video of how we predetermined the distribution of h, l, and the additional payoff across the six boxes. Our random device was an urn with six balls numbered from 1 to 6. Subjects received a sealed envelope with information on all outcomes of uncertainty before the experiment, and were supposed to keep it sealed during the experiment. The link for the recorded video was provided at the end of the experiment. Moreover, we used the prior incentive system proposed by Johnson et al. (2021) to predetermine which decision was chosen for payment. The Dictator Game experiment introduced below adopted a similar ex ante resolution scheme.

ratio 5:5 and the remaining five contain the sharing ratio 10:0. Last, they decide whether to share by choosing a box: If they choose to share, they choose the box indicating the sharing ratio 5:5. There are three payoff pairs, $(19, p; 1)$, $(15, p; 5)$, and $(11, p; 9)$, with seven levels of the winning probability $p \in \{0, \frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}, 1\}$.

2.4 Summary of the Experimental Design

Our Dice Game experiment combines the standard truth-telling paradigm of Fischbacher and Föllmi-Heusi (2013) with endowed lotteries $(h, p; l)$. The six-box setting provides a natural context for this combination, and it helps intensify the sense of uncertainty when subjects make decisions about honesty. This main experiment provides a direct test for whether subjects are more or less honest under uncertainty than certainty. Building on the Dice Game experiment, we design two supplementary experiments to examine mechanisms and three additional experiments to investigate the robustness and generalizability of the main findings. Table 1 provides a brief summary of these six experiments on how they differ in terms of purpose, design, implementation, and sample size.[10] More details on these experiments are presented in the appendices (see Supplemental Material Table A.1 for summary statistics, Figures A.1-A.3 for interfaces, and Section D for instructions).

Table 1: Overview of Experimental Design

Experiment	Purpose	Uncertainty	Resolution	Decide Whether To	Die-Rolling	Implement	Obs
Dice Game	Pattern	Subjects	Ex Post	Lie to get RMB4	Mental	Online	107
Direct Choice	Mechanism	Subjects	Ex Post	Choose to get RMB4	-	Online	102
Second Party	Mechanism	Partners	Ex Post	Lie to get RMB4	Mental	Online	107
Dice Game Loss	Robustness	Subjects	Ex Post	Lie to retain RMB4	Mental	Online	305
Ex Ante Resolution	Robustness	Subjects	Ex Ante	Lie to get SGD2	Physical	Lab	191
Dictator Game	Generalizability	Subjects	Ex Ante	Share with recipients	-	Lab	148

[10]Note that the Dice Game Loss experiment has a larger sample size than the Dice Game experiment because its statistical power is lower. Specifically, in the Dice Game Loss experiment, around $\frac{1}{6}$ of subjects need to lie to avoid the payment deduction of RMB4, while in the Dice Game experiment, around $\frac{5}{6}$ of subjects need to lie to gain the payment of the additional RMB4.

3 Results

This section reports our main results based on the Dice Game experiment, evidence that sheds light on the underlying mechanisms from the Direct Choice and Second Party experiments, and further support of the main findings from the three additional experiments.

3.1 Main Experiment

Panel (a) in Figure 1 presents the proportion of subjects who report the numbered box with the additional RMB4 (also referred to as the proportion of reporting +4 below). Since individual lying behavior is not observable, we infer the tendency to lie at the aggregate level. The x-axis is the winning probability, and the y-axis is the proportion of reporting +4. For each of the 21 decisions, this proportion is substantially higher than the truth-telling rate of $\frac{1}{6}$ and thus provides a measure of dishonesty. Based on this figure, we observe our main behavioral pattern.

Observation 1. *Subjects are less likely to lie under uncertainty than under certainty.*

The proportion of reporting +4 is 59.1 percent in uncertainty conditions, which is lower than that in certainty conditions with high payoffs (74.8 percent) and low payoffs (78.2 percent). Compared with certainty conditions, the proportion of reporting +4 in uncertainty conditions is reduced by 22.7 percent.

Moreover, the observed difference between certainty and uncertainty conditions is larger when the spread between high and low payoffs is wider. Specifically, under $(40, p; 0)$, the proportions of +4 are 54.0 percent for the five lotteries and 72.9 percent (78.5 percent) for certain payoff 40 (certain payoff 0). The effect of uncertainty diminishes for a small spread: Under $(22, p; 18)$, the proportions are 70.5 percent for the five lotteries and 73.8 percent (81.3 percent) for certain payoff 22 (certain payoff 18). In addition, we observe a stable effect of uncertainty across different winning probabilities. For example, the proportions of reporting +4 are between 47.7 percent and 61.7 percent for the five lotteries under $(40, p; 0)$, each of which is significantly lower than those in the two certainty conditions. For more details, Supplemental Material Table A.2 presents statistical tests of pairwise comparisons between conditions within each payoff pair.

Figure 1: Differences Between Uncertain and Certain Environments

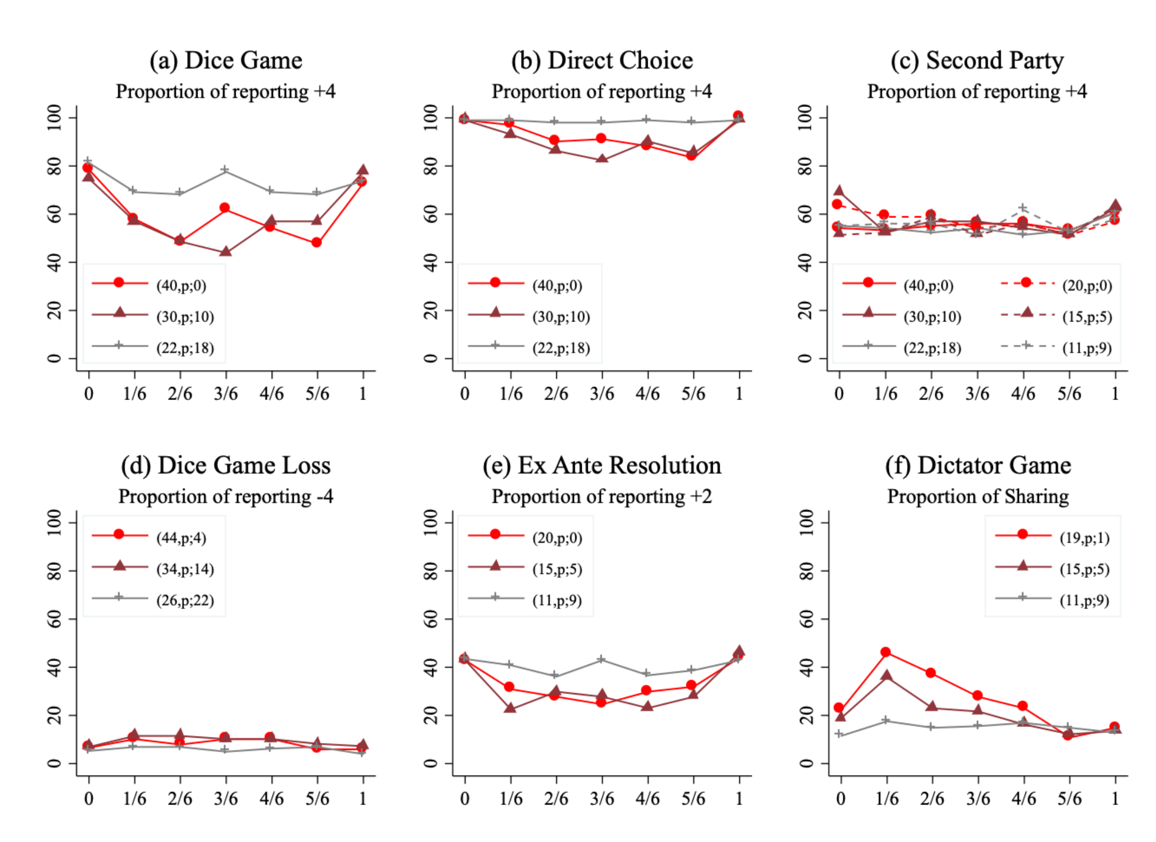

Notes: This figure shows the results of our six experiments, respectively. In each panel, the x-axis is the winning probability and the y-axis is the proportion of subjects who choose the special box as illustrated in the subtitle of each panel.

Regression Analyses. We test these observations using OLS regression analysis. The dependent variable, as a measure of dishonesty, equals 1 if subjects report +4 and 0 otherwise, and the main independent variables are two dummies for the two degenerate certainty conditions, with uncertainty conditions being the reference. We consider a set of control variables and cluster standard errors at individual level.

Table 2: Regression Analyses of the Dice Game Experiment

	OLS: 1_{+4}				
	(1)	(2)	(3)	(4)	(5)
Panel A. Full sample and subsamples by payoff pairs					
	All	All	(40, p; 0)	(30, p; 10)	(22, p; 18)
1_h	0.157***	0.155***	0.188***	0.247***	0.030
	(0.028)	(0.029)	(0.046)	(0.044)	(0.040)
1_l	0.191***	0.190***	0.242***	0.218***	0.108***
	(0.025)	(0.026)	(0.044)	(0.043)	(0.039)
Controls	N	Y	Y	Y	Y
Constant	0.591***	0.299***	0.385***	0.232***	0.381***
	(0.029)	(0.029)	(0.030)	(0.050)	(0.049)
Observations	2,247	2,247	749	749	749
R-squared	0.027	0.383	0.447	0.451	0.424
Panel B. Subsamples by winning probabilities					
	$\frac{1}{6}$	$\frac{2}{6}$	$\frac{3}{6}$	$\frac{4}{6}$	$\frac{5}{6}$
1_h	0.131***	0.195***	0.135***	0.142***	0.171***
	(0.032)	(0.039)	(0.036)	(0.040)	(0.035)
1_l	0.166***	0.230***	0.170***	0.176***	0.206***
	(0.031)	(0.033)	(0.034)	(0.038)	(0.033)
Controls	Y	Y	Y	Y	Y
Constant	0.120***	0.180***	0.357***	0.330***	0.140***
	(0.046)	(0.050)	(0.052)	(0.048)	(0.047)
Observations	963	963	963	963	963
R-squared	0.470	0.465	0.424	0.415	0.470

Notes: 1_{+4} equals 1 if subjects choose the box with the additional RMB4 and 0 otherwise. 1_h (1_l) equals 1 if the condition gives certain high (low) payoff and 0 otherwise. In Panel A, column 1 uses all data without controls. Column 2 further controls for the payoff pair fixed effect, individual fixed effect, duration of the decision (in seconds), and order of the decision (between 1 and 21). Columns 3-5 report results using data on the seven choices under the payoff pair $(40, p; 0)$, $(30, p; 10)$, and $(22, p; 18)$, respectively. In Panel B, each of columns 1-5 uses data on the nine choices, including six choices under certainty and three under uncertainty with the winning probability being $\frac{1}{6}$ to $\frac{5}{6}$, respectively. Standard errors are clustered at individual level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 2 presents our main results. The coefficients of the two dummies are significantly positive without and with controls (Panel A, columns 1-2). On average, subjects under uncertainty show a 15.5 and a 19.0 percentage point decrease in the probability of reporting +4, compared with the certainty of high and low payoffs, respectively (Panel A, column 2). In the regressions using subsamples by payoff pairs, the effect of uncertainty is sizeable and significant under $(40, p; 0)$ and $(30, p; 10)$.

However, under $(22, p; 18)$, the coefficients become substantially smaller and one becomes insignificant, which suggests a null effect of uncertainty when the gap between high and low payoffs is small (Panel A, columns 3-5). In Panel B, we separately examine the effect of uncertainty by winning probabilities and find that the coefficients of the two dummies are significantly positive under all five probabilities. The regression results confirm our observations from Panel (a) in Figure 1.

Robustness Checks and Additional Analyses. First, we find that the observations are robust to the use of probit regression analysis and to the inclusion of demographic characteristics (Supplemental Material Tables A.3 and A.4).

Second, we include the decision of the previous round in the regression to control for the potential confounding effect of previous decisions. We regress the reporting +4 decision on the +4 decision in the previous round, two dummies for certainty conditions in the current round, and their interaction terms with full controls, and find that the coefficients of the two dummies for certainty conditions are significant and similar in size (Supplemental Material Table A.5).

Third, we examine whether the size of the endowed lottery affects the tendency to lie and how to control for this potential effect in different regression specifications (Supplemental Material Table A.6). We regress the decision to report +4 on the amount of payoff under certainty (column 1), and on the winning probability with controls for payoff pairs under uncertainty (column 2). We find no evidence of the effect of endowment size on lying behavior. Next, we regress the decision to report +4 on a dummy variable that indexes uncertain payoffs and the winning probability (column 3), and on the mean and variance of lotteries (column 4). The results show that subjects are less likely to report +4 under uncertainty and when the variance of the lotteries gets larger.

Last, we conduct an individual-level analysis and classify subjects into different types according to their tendency to report +4 in different conditions. In particular, a subject is classified as the More-Moral type if her proportion of reporting +4 under uncertainty conditions is strictly lower than those under both certainty conditions. Supplemental Material Table A.7 displays the standard of classification, the proportion, and the descriptive characteristics of each type. We show that 54 out of 107 subjects are classified as More-Moral type, in support of the observed pattern at aggregate level.

Taken together, results from these analyses indicate the robustness of the pattern whereby people are more honest in uncertain environments.

3.2 Two Experiments on Mechanisms

Direct Choice Experiment. Panel (b) in Figure 1 shows the proportion of subjects who choose +4 in the Direct Choice experiment. We have two observations.

Observation 2A. *Subjects are less likely to choose the maximum payoff under uncertainty than under certainty.*

The proportion of choosing +4 is 92.0 percent in uncertainty conditions, which is lower than that in certainty conditions with high payoffs (99.3 percent) and low payoffs (99.1 percent). The proportion of choosing +4 in uncertainty conditions is reduced by 7.3 percent compared with that in certainty conditions, which is a smaller percentage change compared with that observed in the Dice Game experiment (22.7 percent).

Observation 2B. *The observed difference between certainty and uncertainty conditions is smaller in the Direct Choice experiment than in the Dice Game experiment.*

We further verify these observations using OLS regression analyses (Table 3). In columns 1-4, the dependent variable equals 1 if subjects choose +4 and 0 otherwise. The main independent variables are the two dummies that index the two degenerate certainty conditions. We observe that subjects are less willing to take the additional payoff under uncertainty than under certainty. On average, uncertainty leads to a 7.0 and a 6.7 percentage point decrease in the probability of choosing +4, compared with high and low payoffs, respectively (column 1). This suggests that subjects exhibit a stronger degree of greed aversion under uncertainty than certainty. Also, this pattern is significant only if the spread between high and low payoffs is large enough (columns 2-4), which is similar to the Dice Game experiment.

Table 3: Regression Analyses of the Direct Choice Experiment

	OLS: 1_{+4}				
	All	(40, p; 0)	(30, p; 10)	(22, p; 18)	+ Dice Game
	(1)	(2)	(3)	(4)	(5)
1_h	0.070***	0.096***	0.107***	0.006	0.070***
	(0.012)	(0.021)	(0.022)	(0.009)	(0.012)
1_l	0.067***	0.081***	0.109***	0.006	0.067***
	(0.013)	(0.022)	(0.022)	(0.013)	(0.013)
1_{DiceGame}					-0.476***
					(0.010)
$1_h \times 1_{\text{DiceGame}}$					0.086***
					(0.031)
$1_l \times 1_{\text{DiceGame}}$					0.124***
					(0.029)
Constant	0.918***	0.989***	0.790***	1.001***	0.775***
	(0.018)	(0.024)	(0.025)	(0.014)	(0.020)
Observations	2,142	714	714	714	4,389
R-squared	0.205	0.336	0.305	0.248	0.435

Notes: 1_{+4} equals 1 if subjects choose the box with the additional RMB4 and 0 otherwise. 1_h (1_l) equals 1 if the condition gives certain high (low) payoff and 0 otherwise. Columns 1-4 use data from the Direct Choice experiment. Column 5 combines data from the Dice Game experiment and the Direct Choice experiment. 1_{DiceGame} equals 1 if subjects are in the Dice Game experiment. We control for the payoff pair fixed effect, individual fixed effect, duration of the decision, and order of the decision. Standard errors are clustered at the individual level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

We find that the magnitudes of the two certainty indexes are less than half of those in the original estimation (Table 3, column 1 vs. Table 2, column 2). Pooling the data from the Dice Game and Direct Choice experiments, we set the dependent variable to be 1 if subjects report or directly choose +4 and 0 otherwise. In the regression, we include a dummy to index the Dice Game experiment and its interactions with the two certainty indexes (column 5). These two interaction terms are both significantly positive, which suggests a stronger difference between uncertainty and certainty in the Dice Game experiment than that in the Direct Choice experiment.

Second Party Experiment. Panel (c) in Figure 1 displays the proportion of reporting +4 in the Second Party experiment. Similarly, we can summarize two observations from this experiment.

Observation 3A. *Subjects are marginally less likely to lie under uncertainty than* *under certainty, when the uncertainty concerns the payoff of paired partners.*

The proportion of reporting +4 is 54.6 percent in uncertainty conditions, which is lower than in certainty conditions with high payoffs (60.4 percent) and low payoffs (58.1 percent). Compared with certainty conditions, the proportion of reporting +4 in uncertainty conditions is reduced by 7.8 percent, which is smaller than the 22.7 percent change observed in the main experiment.

Observation 3B. *The observed difference between certainty and uncertainty conditions in the Second Party experiment is smaller than in the Dice Game experiment.*

Table 4 reports regression results of the Second Party experiment. The dependent variable equals 1 if subjects report +4 and 0 otherwise. The main independent variables are the two dummies, which index the two degenerate certainty conditions of the partner. On average, uncertainty about others leads to a 5.8 and 3.5 percentage point decrease in the probability of reporting +4, compared with the two certainty conditions (column 1). The coefficient of the index for the certain low payoff is significant at 10 percent level. When we examine the six payoff pairs separately in the regression analyses, the signs and significance of the two dummies reveal no systematic pattern (columns 2-7).

Moreover, pooling the data from the Dice Game and Second Party experiments, we include a dummy to index the Dice Game experiment in the regression, as well as the interactions between this dummy and the main independent variables. We find that the two interaction terms are significantly positive, in support of a stronger difference between uncertainty and certainty in the Dice Game experiment than in the Second Party experiment.

It is worth noting that the design of the Second Party experiment shares some features with the literature on moral wiggle room. Supplemental Material Section B reports another experiment, in which we modify our experiment to incorporate the original paradigm of moral wiggle room in Dana, Weber, and Kuang (2007). The new piece of evidence further supports our finding that uncertainty about others does not lead to decision-makers' truth-telling behavior; contrarily, it can lead to dishonesty and information avoidance, which is in line with the notion of moral wiggle room.

Table 4: Regression Analyses of the Second Party Experiment

	OLS: 1_{+4}							
	All	(40, p; 0)	(30, p; 10)	(22, p; 18)	(20, p; 0)	(15, p; 5)	(11, p; 9)	+ Dice Game
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1_h	0.058***	0.061	0.082*	0.077*	0.009	0.096**	0.025	0.058***
	(0.019)	(0.042)	(0.044)	(0.043)	(0.044)	(0.041)	(0.037)	(0.019)
1_l	0.035*	-0.011	0.147***	0.024	0.077*	-0.025	-0.002	0.036**
	(0.018)	(0.040)	(0.044)	(0.039)	(0.046)	(0.037)	(0.036)	(0.018)
1_{DiceGame}								-0.119***
								(0.015)
$1_h \times 1_{\text{DiceGame}}$								0.098***
								(0.034)
$1_l \times 1_{\text{DiceGame}}$								0.155***
								(0.031)
Constant	0.467***	0.011	0.675***	0.069**	0.812***	0.805***	0.576***	0.417***
	(0.021)	(0.027)	(0.043)	(0.030)	(0.084)	(0.073)	(0.093)	(0.024)
Observations	4,494	749	749	749	749	749	749	6,741
R-squared	0.513	0.562	0.544	0.580	0.565	0.613	0.611	0.471

Notes: 1_{+4} equals 1 if subjects choose the box with RMB25 and 0 if subjects choose the box with RMB21. 1_h (1_l) equals 1 if the condition gives certain high (low) payoff for the partner and 0 otherwise. Columns 1-7 use data from the Second Party experiment. Column 8 combines data from the Dice Game experiment and the Second Party experiment. 1_{DiceGame} equals 1 if the subject is in the Dice Game experiment. We control for the fixed effect of payoff pairs, individual fixed effect, duration of the decision, and order of the decision. Standard errors are clustered at individual level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

3.3 Three Experiments on Robustness and Generalizability

Dice Game Loss Experiment. Panel (d) in Figure 1 presents the proportions of reporting -4 across all 21 rounds. The proportion under uncertainty is significantly lower than the completely truth-telling rate of $\frac{1}{6}$, and indicates lying behavior at the aggregate level.[11] Comparing between uncertain and certain environments, we have the following observation.

Observation 4. *Subjects are less likely to lie under uncertainty than under certainty, when the lying decision is made to avoid losses.*

The proportion of reporting -4 is 8.50 percent under uncertainty, compared with

[11]There is an alternative explanation for Observation 1: Subjects in our main experiment regard the extra gain of RMB4 as a signal of the low outcome of the lottery and thus avoid choosing it under uncertainty. This explanation suggests that, conversely, a loss could be regarded as a signal of the high outcome and thus subjects would tend to over-report the losing state under uncertainty than certainty. By contrast, we observe that people tend to dishonestly under-report the losing state which does not support this alternative hypothesis.

5.68 percent and 6.23 percent under certainty with high and low payoffs, respectively. That is, the proportion in uncertainty conditions is increased by 42.6 percent compared with certainty conditions. OLS regression analyses show that this pattern is statistically significant with the whole sample, and the signs remain robust with a drop in the significance for some of the subsample analyses (Supplemental Material Table A.8).[12] Overall, this finding suggests that subjects are less likely to lie to avoid losses under uncertainty than certainty, which is consistent with Observation 1 in the main experiment.

Ex Ante Resolution Experiment. Panel (e) in Figure 1 displays the proportions of reporting +2 in the Ex Ante Resolution experiment in which subjects make the decision after the uncertainty has been resolved but kept unknown. We have the following observation.

Observation 5. *Subjects are less likely to lie under uncertainty than under certainty, when the uncertainty is resolved ex ante.*

The proportion of reporting +2 is 31.4 percent under uncertainty, compared with 44.3 percent and 43.1 percent under high and low payoffs, respectively. Compared with certainty conditions, the proportion of reporting +2 in uncertainty conditions is reduced by 28.1 percent.[13] This pattern whereby subjects are more honest in uncertain environments is consistent with the observation in the main experiment. Moreover, we observe that the effect of uncertainty is significant when the spread between high and low payoffs is large enough and is invariant with the winning probability. These findings are supported by OLS regression analyses (Supplemental Material Table A.9).

[12]The observed 42.6 percentage change is also higher than the 22.7 percent in the Dice Game experiment. Apart from statistical power, these two experiments differ in other respects such as gain versus loss framing, reference point, the size of the lie (Gneezy, Kajackaite, and Sobel, 2018), and so on. Therefore, we do not want to over-interpret the differences and instead use the Dice Game Loss experiment to examine the robustness of our main pattern.

[13]This percentage change is higher than that in the Dice Game experiment. Apart from whether decisions are made before or after the uncertainty resolution, these two experiments differ in several ways including online versus in-person environment and samples from China versus Singapore. Therefore, we do not want to over-interpret the comparison between the two experiments. Instead, we use the Ex Ante Resolution experiment to examine the robustness of the main findings.

Dictator Game Experiment. Panel (f) in Figure 1 reports the proportions of choosing the box to share payoffs with others in the Dictator Game experiment. We have the following observation.

Observation 6. *Subjects are more likely to share with others under uncertainty than under certainty.*

The proportion of sharing is 22.2 percent in uncertainty conditions, which is higher than those in certainty conditions with high payoff (13.7 percent) and low payoff (17.6 percent). Similar to the Dice Game experiment, the observed difference between certainty and uncertainty is larger when the spread between high and low payoffs is wider. For example, the proportion of sharing is 45.9 percent for $(19, \frac{1}{6}; 1)$ and 14.9 percent and 22.3 percent for certain payoff 19 and certain payoff 1, respectively. By contrast, under the payoff pair $(11, p; 9)$, all seven conditions are statistically indifferent. Moreover, we observe a distinct pattern: The difference between uncertainty and certainty decreases with the increase in the winning probability in this experiment, but is invariant to the winning probability in the Dice Game and Ex Ante Resolution experiments. For instance, the proportion is 45.9 percent for the lottery $(19, \frac{1}{6}; 1)$ and 23.0 percent for the lottery $(19, \frac{4}{6}; 1)$.[14] We further verify these findings using OLS regression analysis (Supplemental Material Table A.10).

4 Discussion

In this section, we discuss our observations in terms of theoretical implications and possible explanations. We also propose theoretical frameworks for the anxiety-based mechanism that is related to notions of magical thinking and quasi-magical thinking.

[14]This distinct pattern can be due to the expected monetary cost of sharing $p\frac{h}{2} + (1-p)\frac{l}{2}$, which increases with the winning probability p in the Dictator Game experiment. By contrast, the cost of telling the truth is fixed at RMB4 in the Dice Game experiment and SGD2 in the Ex Ante Resolution experiment. We explain this formally in Supplemental Material Section C.

4.1 Theoretical Implications

Risk Preferences. In most models of decision-making under risk and uncertainty, decision-makers should respect the principle of dominance. Specifically, if an act is chosen in all deterministic situations, it should remain the favorable option in uncertain situations in which each deterministic situation occurs with the given probability. However, our finding of a stronger tendency to behave morally under uncertainty than certainty violates this principle.

More formally, in our setting, there are two states—s_h yields the high monetary payoff h and s_l yields the low monetary payoff l—and two acts, moral act m and immoral act i. Decision-makers evaluate each act under each state, which generates four possible consequences: $\{h_m, l_m, h_i, l_i\}$. For example, h_m denotes the consequence of the moral act when state s_h happens, and so on. Here the consequence can capture not only monetary payoffs but also moral considerations such as lying aversion, and other-regarding preferences. In our main experiment, we show that a substantial proportion of subjects prefer the immoral act to the moral act in the two deterministic situations—$h_i \succ h_m$ and $l_i \succ l_m$—but choose the moral act in the uncertain situations, $(h_m, p; l_m) \succ (h_i, p; l_i)$. If individuals are expected utility maximizers,

$$U(\text{immoral}) = pu(h_i) + (1-p)u(l_i),$$
$$U(\text{moral}) = pu(h_m) + (1-p)u(l_m).$$

Given that $u(h_i) > u(h_m)$ and $u(l_i) > u(l_m)$, we will have $pu(h_i) + (1-p)u(l_i) > pu(h_m) + (1-p)u(l_m)$.

Similarly, if individuals are rank-dependent utility maximizers with probability weighting function $w(p)$,

$$U(\text{immoral}) = w(p)u(h_i) + (1-w(p))u(l_i),$$
$$U(\text{moral}) = w(p)u(h_m) + (1-w(p))u(l_m).$$

We will also have $w(p)u(h_i) + (1-w(p))u(l_i) > w(p)u(h_m) + (1-w(p))u(l_m)$. In this regard, standard models have difficulty accommodating the documented pattern that people are more moral in uncertain environments.[15]

Preferences for Truth-telling. In models with preferences for truth-telling, individuals make trade-offs among the monetary payment, the direct cost of lying, and the cost of being perceived as a liar or the value of being perceived as honest (Gneezy, Kajackaite, and Sobel, 2018; Abeler, Nosenzo, and Raymond, 2019). These models propose that features of the random process, which generates privately observed information about monetary payment, affect the propensity to lie. First, the probability of observing high payment affects reputational concerns (Abeler, Nosenzo, and Raymond, 2019; Gneezy, Kajackaite, and Sobel, 2018). Put differently, when this objective probability is lower, reporting the high payment looks more suspicious to the audience, and thus decision-makers with image concerns may be less likely to make such a report. Second, the probability distribution of this random process shapes the reference points of decision-makers, and changes lying behavior. Garbarino, Slonim, and Villeval (2019) show that when the probability of observing low payment decreases, the reference point measured by the ex ante expected payment increases. Consequently, decision-makers would suffer a greater loss if reporting the low payment and thus have a stronger tendency to lie.

Our experiment keeps this random process constant: The probability of the high payoff is fixed to be $\frac{1}{6}$ and the monetary gain from lying is fixed to be RMB4 in the Dice Game experiment. Preferences for truth-telling or reference dependence is likely to be similar for $(h, p; l)$ with varying p. To integrate $(h, p; l)$ into these models, we can embed the preference for truth-telling, with reputational concerns or reference dependence, in the consequences $\{h_m, l_m, h_i, l_i\}$. However, as noted above, this specification cannot account for our observations due to dominance violation.

[15]It is also important to consider models that permit dominance violation, including the disappointment theory (Bell, 1985; Loomes and Sugden, 1986); regret theory (Bell, 1982; Loomes and Sugden, 1982); reference-dependence theory (Kőszegi and Rabin, 2007); and models with preference for gambling (Fishburn, 1980; Diecidue, Schmidt, and Wakker, 2004). For example, under models with preference for gambling, when the utility of gambling is sufficiently high, utility for the lottery can be higher than that for the best realization of the lottery. Namely, it is possible that $(h_m, p; l_m) \succeq h_m$ and $(h_i, p; l_i) \succeq h_i$. However, given that $h_i \succ h_m$, in order to have $(h_m, p; l_m) \succ (h_i, p; l_i)$ to account for our findings, we need to have a strong assumption that preference for gambling is substantially stronger under a moral act than under an immoral act.

Social Preferences. Models of social preferences under uncertainty focus on how individuals take ex ante and ex post fairness into account (Brock, Lange, and Ozbay, 2013; Saito, 2013). In Supplemental Material Section C, we provide discussions of these models and derive their predictions in our Dictator Game experiment. In general, under regular conditions, these models fall short in addressing the increase in sharing behavior in uncertain environments (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Andreoni and Miller, 2002; Charness and Rabin, 2002). To obtain clearer theoretical implications, it is crucial to further investigate possible explanations.

4.2 Possible Explanations

Our behavioral pattern aligns with numerous documented anomalies in decision-making under uncertainty, which often arise from the anxiety and complexity perspectives on uncertainty. Below we discuss the extent to which these two perspectives on uncertainty may contribute to our findings.[16]

The anxiety perspective on uncertainty offers a simple explanation for the main behavioral pattern whereby people are more moral under uncertainty (Observations 1, 4, 5, and 6): They behave morally under uncertainty to cope with the feeling of anxiety and to obtain a sense of control. This explanation is related to beliefs in a just world, moralistic gods, and karma, which link moral behavior with favorable outcomes of uncertainty.

Apart from the main behavioral pattern, we also observe weaker differences between uncertainty and certainty in the Direct Choice and Second Party experiments. Anxiety perspective can also help explain these observations when we impose further

[16] Our experimental design strives to reduce potential confusion and to enhance understanding. All experiments include understanding tests and explanations to enhance subjects' comprehension (see Supplemental Material Section D for the understanding tests in the Dice Game experiment). For example, we test subjects' understanding of probability and the independence between their behavior and uncertainty resolution. Subjects on average scored 7.8 out of eight understanding tests in the Dice Game experiment, and performed similarly well in the remaining five experiments. While confusion is hard to completely eliminate and might play some role in our experiments, it is unlikely to be a key factor in our main findings.

assumptions. Specifically, when behavior has a weaker moral implication, that is, greed in the Direct Choice experiment compared with dishonesty in the Dice Game experiment, people are less likely to gain control through their behavior, which results in the smaller differences in behavior between uncertainty and certainty (Observations 2A and 2B). Similarly, when uncertainty affects others in the Second Party experiment rather than oneself in the Dice Game experiment, the drive for a favorable outcome is diminished, which leads to the less pronounced differences between uncertainty and certainty (Observations 3A and 3B).

From the complexity perspective, people may have difficulty thinking contingently and aggregating across consequences, and thus they may behave distinctively under uncertainty and certainty. On the one hand, this perspective allows the possibility that people are more moral in complex uncertain environments, which aligns with Observations 1, 4, 5, and 6. On the other hand, it is silent about whether people would be more or less moral under uncertainty.[17] For the complexity perspective to account for the patterns in the two supplementary experiments, we may need to further assume that these two environments involve a lower degree of complexity compared with that of the Dice Game experiment. While the Direct Choice experiment is probably less complex since it does not include the mental-die-rolling process, the Second Party experiment appears to be more complex because another subject with uncertainty is involved.

To sum, both anxiety and complexity, as inherent aspects of uncertainty in our experiments, can account for our observations when some assumptions are imposed. It is important to point out that our goal is to elucidate these observations through these perspectives, not to distinguish between them. As we noted in the Introduction, the two perspectives may be intertwined: When people find uncertain environments complex, they may rely on some widespread beliefs to arrive at a decision, and thus behave morally as (or as if) they believe that their moral behavior can lead to a

[17]One may think about various heuristic rules that people adopt to make complex uncertain decisions. For example, they may tend to choose the salient option. Arguably, the salient option can be the box that distinguishes itself from others: carrying an extra payoff, causing a payoff deduction, or bringing a fair allocation of payoffs. However, choosing the salient box indexes a less moral manner in the first case, but a more moral manner in the last two cases.

favorable outcome. Below we present a theoretical framework to capture such link between uncertainty and morality by the notions of magical thinking and quasi-magical thinking.

4.3 Magical Thinking and Quasi-magical Thinking

Magical thinking refers to the belief that one can influence the outcome of uncertainty through some specific acts, even though the acts have no causal link to the uncertainty. For example, people may believe that throwing the dice harder results in higher numbers, tolerating cold for a longer time extends life expectancy, and voting induces others to vote. One way to model magical thinking is to directly assume that decision-makers explicitly hold such a belief that their acts can change the outcome of the uncertainty. More specifically, subjects believe that with probability α, the world is karmic and thus a moral act leads to high payoff h_m, and an immoral act leads to low payoff l_i. With probability $1-\alpha$, the world is objective, and thus high and low payoffs occur with probability p and $1-p$, respectively. Namely, if individuals are expected utility maximizers with karmic belief α, we have

$$U(\text{immoral}) = \alpha u(l_i) + (1-\alpha)(pu(h_i) + (1-p)u(l_i)),$$
$$U(\text{moral}) = \alpha u(h_m) + (1-\alpha)(pu(h_m) + (1-p)u(l_m)).$$

Therefore, the increases in honesty and altruism under uncertainty no longer violate dominance, and they occur when karmic belief α is strong and the gap between $u(h_m)$ and $u(l_i)$ is large.

An alternative approach is quasi-magical thinking, whereby people act as if they believe that their action influences the outcome of uncertainty, even though they do not really hold such a belief when asked (Shafir and Tversky, 1992; Risen, 2016). More specifically, as Shafir and Tversky (1992) observe, it is unlikely that subjects truly believe they can control the outcome of the die by throwing harder, live longer by tolerating cold for a longer time, or induce others to vote by voting. Nevertheless, they behave as if they hold such beliefs. This notion can be captured by an act-dependent probability weighting function, based on the source method (Chew and Sagi, 2008; Abdellaoui et al., 2011). More specifically, moral acts and immoral acts yield distinct probability weighting functions $w_m(p)$ and $w_i(p)$, respectively. If individuals are rank-dependent utility maximizers with an act-dependent probability weighting function, we will have

$$U(\text{immoral}) = w_i(p)u(h_i) + (1 - w_i(p))u(l_i),$$
$$U(\text{moral}) = w_m(p)u(h_m) + (1 - w_m(p))u(l_m).$$

Similarly, the documented pattern no longer violates dominance with $w_m(p) > w_i(p)$ and is more likely to emerge when the gap in decision weights $w_m(p) - w_i(p)$ is larger. Here the act-dependent probability weighting function captures a positive feeling, such as hope or optimism when people behave morally, and a negative feeling such as anxiety or pessimism when they behave immorally.[18]

Taken together, both frameworks can help explain the overall behavioral patterns, but they differ in whether there is an explicit belief. Our tentative evidence is more consistent with the notion of quasi-magical thinking, because we provide explicit winning probabilities and emphasize the independence between choices and uncertainty resolution in our experiments (see Supplemental Material Section C for details). Nevertheless, we acknowledge that it is challenging to measure beliefs in our experiments. As Shafir and Tversky (1992) put it, "Whereas magical thinking involves indefensible beliefs, quasi-magical thinking yields inexplicable actions." Here we focus on observable actions and leave the joint investigation of beliefs and actions for future research.

5 Concluding Remarks

We document a behavioral pattern whereby people are more moral in uncertain environments. The main experiment demonstrates this pattern using a dice game to

[18] Supplemental Material Section C provides details of these two approaches. A more general approach is to reformulate the state space, which can be found in Chapter 11 of Gilboa (2009) on Newcomb's paradox. More discussion of Newcomb's paradox and related theoretical works can be found in Nozick (1969); Jeffrey (1965); Gilboa and Schmeidler (1995); Karni and Vierø (2013); Schipper (2016); Karni (2017); Gilboa, Minardi, and Samuelson (2020).

measure honesty under uncertainty and certainty. Additionally, our two supplementary experiments underscore the significance of morality concerns and uncertainty that people face. To further validate and extend our findings, we conduct three additional experiments that explore the robustness of this pattern and the generalizability to the domain of other-regarding behavior.

Our study offers an explanation for some observations in various applied settings. For example, it has been widely observed that investors exhibit an aversion to sin stocks and a preference for socially responsible investment. Moreover, using a longitudinal dataset of actual donations and a dictator game, a recent study shows that individuals' generosity increases under the COVID-19 threat (Fridman, Gershon, and Gneezy, 2022). While various factors may have contributed to these observations, our study provides some insights. For example, when investors face complex decisions in portfolio choice and feel anxious about their returns, they tend to choose stocks with moral considerations. Also, the increased generosity is in line with our explanation that altruistic giving yields a sense of control during the pandemic, which is also consistent with alternative explanations such as increased feelings of sympathy and the desire to experience positive emotions.

Our study also sheds some light on models of decision-making under uncertainty. An implicit prerequisite of standard models is that individuals "properly" perceive and "fully" attend to the choice situations—that is, the states, acts, and consequences. However, many departures from the expected utility can be interpreted as manifestations of basic human perception. For example, probability weighting is often linked to the perception of likelihoods with a sense of optimism and pessimism (i.e., Diecidue and Wakker, 2001), and to the complexity aspect of uncertainty (Enke and Graeber, 2023; Oprea, 2024; Esponda and Vespa, 2023). Here we propose that the connection between moral behavior and outcomes from uncertainty can be in part viewed as a manifestation of anxiety and complexity, and can be modeled through an explicit belief or a probability-weighting function to capture a sense of control or optimism. This approach can help explain the patterns in our experiments, some phenomena in the aforementioned applied settings, and beyond.

References

Abdellaoui, M., Baillon, A., Placido, L., and Wakker, P. P. (2011). The rich domain of uncertainty: Source functions and their experimental implementation. *American Economic Review*, 101(2):695–723.

Abeler, J., Nosenzo, D., and Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, 87(4):1115–1153.

Andreoni, J., Aydin, D., Barton, B., Bernheim, B. D., and Naecker, J. (2020). When fair isn't fair: Understanding choice reversals involving social preferences. *Journal of Political Economy*, 128(5):1673–1711.

Andreoni, J. and Miller, J. (2002). Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70(2):737–753.

Arad, A. (2014). Avoiding greedy behavior in situations of uncertainty: The role of-magical thinking. *Journal of Behavioral and Experimental Economics*, 53(1):17–23.

Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, 30(5):961–981.

Bell, D. E. (1985). Disappointment in decision making under uncertainty. *Operations Research*, 33(1):1–27.

Bénabou, R. and Tirole, J. (2006). Belief in a just world and redistributive politics. *Quarterly Journal of Economics*, 121(2):699–746.

Bolton, G. E., Brandts, J., and Ockenfels, A. (2005). Fair procedures: Evidence from games involving lotteries. *The Economic Journal*, 115(506):1054–1076.

Bolton, G. E. and Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, 90(1):166–193.

Brock, J. M., Lange, A., and Ozbay, E. Y. (2013). Dictating the risk: Experimental evidence on giving in risky environments. *American Economic Review*, 103(1):415–437.

Brock, J. M., Lange, A., and Ozbay, E. Y. (2016). Dictating the risk: Experimental evidence on giving in risky environments: reply. *American Economic Review*, 106(3):840–842.

Cason, T. N. and Plott, C. R. (2014). Misconceptions and game form recognition: Challenges to theories of revealed preference and framing. *Journal of Political Economy*, 122(6):1235–1270.

Charness, G. and Levin, D. (2009). The origin of the winner's curse: A laboratory study. *American Economic Journal: Microeconomics*, 1(1):207–236.

Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics*, 117(3):817–869.

Chen, Y. and Zhong, S. (2024). *Supplement to "People Are More Moral in Uncertain Environments"*. Econometrica Supplemental Material.

Chew, S. H. and Li, K. K. (2021). The moral investor: Sin stock aversion and virtue stock affinity. Working paper.

Chew, S. H. and Sagi, J. S. (2008). Small worlds: Modeling attitudes toward sources of uncertainty. *Journal of Economic Theory*, 139(1):1–24.

Converse, B. A., Risen, J. L., and Carter, T. J. (2012). Investing in karma: When wanting promotes helping. *Psychological Science*, 23(8):923–930.

Dana, J., Weber, R. A., and Kuang, J. X. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33(1):67–80.

Dewey, J. (1929). *The quest for certainty : a study of the relation of knowledge and action*. Minton, Balch & Co.

Diecidue, E., Schmidt, U., and Wakker, P. P. (2004). The utility of gambling reconsidered. *Journal of Risk and Uncertainty*, 29(3):241–259.

Diecidue, E. and Wakker, P. P. (2001). On the intuition of rank-dependent utility. *Journal of Risk and Uncertainty*, 23(3):281–298.

Engel, C. (2011). Dictator games: a meta study. *Experimental Economics*, 14(1):583–610.

Enke, B. (2019). Kinship, cooperation, and the evolution of moral systems. *Quarterly Journal of Economics*, 134(2):953–1019.

Enke, B. and Graeber, T. (2023). Cognitive uncertainty. *The Quarterly Journal of Economics*, 138(4):2021–2067.

Esponda, I. and Vespa, E. (2014). Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 6(4):180–202.

Esponda, I. and Vespa, E. (2023). Contingent thinking and the sure-thing principle: Revisiting classic anomalies in the laboratory. *Review of Economic Studies*, page rdad102.

Exley, C. L. (2016). Excusing selfishness in charitable giving: The role of risk. *Review of Economic Studies*, 83(2):587–628.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114(3):817–868.

Fischbacher, U. and Föllmi-Heusi, F. (2013). Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association*, 11(3):525–547.

Fishburn, P. C. (1980). A simple model for the utility of gambling. *Psychometrika*, 45(4):435–448.

Fridman, A., Gershon, R., and Gneezy, A. (2022). Increased generosity under covid-19 threat. *Scientific Reports*, 12(1):4886.

Fudenberg, D. and Levine, D. K. (2012). Fairness, risk preferences and independence: Impossibility theorems. *Journal of Economic Behavior & Organization*, 81(2):606–612.

Garbarino, E., Slonim, R., and Villeval, M. C. (2019). Loss aversion and lying behavior. *Journal of Economic Behavior & Organization*, 158:379–393.

Garcia, T., Massoni, S., and Villeval, M. C. (2020). Ambiguity and excuse-driven behavior in charitable giving. *European Economic Review*, 124(1):103412.

Gilboa, I. (2009). *Theory of Decision under Uncertainty*. Cambridge university press.

Gilboa, I., Minardi, S., and Samuelson, L. (2020). Theories and cases in decisions under uncertainty. *Games and Economic Behavior*, 123(1):22–40.

Gilboa, I. and Schmeidler, D. (1995). Case-based decision theory. *Quarterly Journal of Economics*, 110(3):605–639.

Gino, F., Norton, M. I., and Weber, R. A. (2016). Motivated bayesians: Feeling moral while acting egoistically. *Journal of Economic Perspectives*, 30(3):189–212.

Gneezy, U., Kajackaite, A., and Sobel, J. (2018). Lying aversion and the size of the lie. *American Economic Review*, 108(2):419–453.

Gneezy, U., List, J. A., and Wu, G. (2006). The uncertainty effect: When a risky prospect is valued less than its worst possible outcome. *Quarterly Journal of Economics*, 121(4):1283–1309.

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.

Henslin, J. M. (1967). Craps and magic. *American Journal of Sociology*, 73(3):316–330.

Hogg, M. A. (2007). Uncertainty–identity theory. *Advances in Experimental Social Psychology*, 39(1):69–126.

Jeffrey, R. C. (1965). *The Logic of Decision*. University of Chicago Press.

Johnson, C., Baillon, A., Bleichrodt, H., Li, Z., Van Dolder, D., and Wakker, P. P. (2021). Prince: An improved method for measuring incentivized preferences. *Journal of Risk and Uncertainty*, 62(1):1–28.

Kajackaite, A. and Gneezy, U. (2017). Incentives and cheating. *Games and Economic Behavior*, 102:433–444.

Karni, E. (2017). States of nature and the nature of states. *Economics and Philosophy*, 33(1):73–90.

Karni, E. and Safra, Z. (2002). Individual sense of justice: A utility representation. *Econometrica*, 70(1):263–284.

Karni, E. and Vierø, M.-L. (2013). "Reverse bayesianism": A choice-based theory of growing awareness. *American Economic Review*, 103(7):2790–2810.

Kellner, C., Reinstein, D., and Riener, G. (2019). Ex-ante commitments to "give if you win" exceed donations after a win. *Journal of Public Economics*, 169(1):109–127.

Köszegi, B. and Rabin, M. (2007). Reference-dependent risk attitudes. *American Economic Review*, 97(4):1047–1073.

Krawczyk, M. and Le Lec, F. (2010). 'give me a chance!' an experiment in social decision under risk. *Experimental Economics*, 13(1):500–511.

Krawczyk, M. and Le Lec, F. (2016). Dictating the risk: Experimental evidence on giving in risky environments: Comment. *American Economic Review*, 106(3):836–839.

Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology*, 32(2):311–328.

Lerner, M. (1980). *The Belief in a Just World*. Springer US.

Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–3287.

Loomes, G. and Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 92(368):805–824.

Loomes, G. and Sugden, R. (1986). Disappointment and dynamic consistency in choice under uncertainty. *Review of Economic Studies*, 53(2):271–282.

Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27(4):1622–1668.

Martínez-Marquina, A., Niederle, M., and Vespa, E. (2019). Failures in contingent reasoning: The role of uncertainty. *American Economic Review*, 109(10):3437–74.

Nozick, R. (1969). Newcomb's problem and two principles of choice. In *Essays in Honor of Carl G. Hempel*. Springer Netherlands.

Oprea, R. (2024). Simplicity equivalents. Working paper.

Purzycki, B. G., Apicella, C., Atkinson, Q. D., Cohen, E., McNamara, R. A., Willard, A. K., Xygalatas, D., Norenzayan, A., and Henrich, J. (2016). Moralistic gods, supernatural punishment and the expansion of human sociality. *Nature*, 530(7590):327–330.

Quattrone, G. A. and Tversky, A. (1984). Causal versus diagnostic contingencies: On self-deception and on the voter's illusion. *Journal of Personality and Social Psychology*, 46(2):237–248.

Risen, J. L. (2016). Believing what we do not believe: Acquiescence to superstitious beliefs and other powerful intuitions. *Psychological Review*, 123(2):182–207.

Risen, J. L. and Gilovich, T. (2008). Why people are reluctant to tempt fate. *Journal of Personality and Social Psychology*, 95(2):293–307.

Saito, K. (2013). Social preferences under risk: Equality of opportunity versus equality of outcome. *American Economic Review*, 103(7):3084–3101.

Sandroni, A., Ludwig, S., and Kircher, P. (2013). On the difference between social and private goods. *The B.E. Journal of Theoretical Economics*, 13(1):151–177.

Schipper, B. C. (2016). A note on states and acts in theories of decision under uncertainty. Working paper.

Shafir, E. and Tversky, A. (1992). Thinking through uncertainty: Nonconsequential reasoning and choice. *Cognitive Psychology*, 24(4):449–474.

Stefan, S. and David, D. (2013). Recent developments in the experimental investigation of the illusion of control. a meta-analytic review. *Journal of Applied Social Psychology*, 43(2):377–386.

Tjøtta, S. (2019). More or less money? an experimental study on receiving money. *Journal of Behavioral and Experimental Economics*, 80(1):67–79.

Trautmann, S. T. and Wakker, P. P. (2010). Process fairness and dynamic consistency. *Economics Letters*, 109(3):187–189.

Tversky, A. and Shafir, E. (1992). The disjunction effect in choice under uncertainty. *Psychological Science*, 3(5):305–309.