Conditional Prosocial Preferences: A Behavioral Approach to Economic Allocation Decisions

by

Conor M. Artman

(Under the Direction of W. David Bradford)

Abstract

While on one hand many studies have determined that consumers' preferences for fair economic payoffs may be conditional, few have attempted to examine what contexts are sufficient for these conditional behaviors. In three different binary dictator games, I examine preferences for fair outcomes in economic decision-making contexts and what mechanisms drive more self-interested behavior. One condition allows participants to manipulate probabilistic uncertainty, the second allows for a deferment of choice option, and the third allows for deferment and probability manipulation of choices. From data collected on the Amazon Mechanical Turk while at Dan Ariely's lab at Duke University, I examine the statistical significance of the distribution of responses across three conditions and suggest that the mechanism driving more self-interested behavior may act through licensing behavior under uncertainty more than through lack of transparency.

Index Words: conditional preferences, binary dictator game, behavioral economics, uncertainty, transparency

Conditional Prosocial Preferences: A Behavioral Approach to Economic Allocation Decisions

by

Conor M. Artman

A Thesis Submitted to the Graduate Faculty of the University of Georgia in Partial Fulfillment

of the

Requirements for the Degree

Master of Arts

Athens, Georgia

2015

Conditional Prosocial Preferences: A Behavioral Approach to Economic Allocation Decisions

by

Conor M. Artman

Approved:

Major Professor: David Bradford

Committee: Adam S. Goodie

Scott E. Atkinson

Electronic Version Approved:

Julie Coffield

Interim Dean of the Graduate School

The University of Georgia

May 2015

Table of Contents

Chapter 1

Introduction

As the burgeoning field of behavioral economics has attracted greater attention, there have been closer intersections between economic questions and those which have been typically reserved to topics more common to fields like psychology or behavioral neuroscience. Among the primary concerns with consumer behaviors, the concept of fairness in economic contexts provides a first extension for models concerned with consumer preferences, since it introduces the thought that consumers may derive utility from social outcomes as well as self-interested ones. Indeed, if a core tenet of economics is to examine the notion of scarcity and rationally self-interested agents, then a natural step toward relaxing some of the rigidities common to purely rational assumptions is to examine: what a precise notion of fairness may be, if fairness preferences are stable, and what drives preferences for fairness.

With this in mind, there have been a number of studies which have developed precise mathematical notions of economic fairness. Rabin (1993) is one such example, where he works under the hypothesis that people are likely to behave altruistically to those who are altruistic toward them, and are similarly likely to harm others who harm them, all under a gametheoretic framework. Particularly, his theory is developed around people's beliefs of payoffs as well as the material payoff itself, rather than solely the material payoff. Eventually he defines a fairness equilibrium as a Nash equilibrium reflecting this thought: a fairness equilibrium results from either a mutual-min Nash equilibrium (a Nash equilibrium when the goal is to minimize the other's payoff), or a mutual-max Nash equilibrium (a Nash equilibrium when the goal is to maximize the other's payoff). Others, in experimental settings, have examined people's preferences under uncertainty generally, and if people's preferences for fair economic outcomes are stable. Dana et al. (2006) provides instances where people's preferences are shown to fluctuate dramatically from a baseline of strong preferences for fair outcomes to outcomes reflecting predominantly self-interested ones, namely through a series of dictator games obscuring options and outcomes of participants' choices while leaving the payoffs the same. There have been few, if any, studies so far that have examined what factors or contexts drive the stability of fairness preferences.

In this study, I attempt to determine what factors drive what I posit are weak preferences for fair economic allocations in three experimental games. Broadly, my basic hypothesis is that preferences for fair economic outcomes are dependent on the presence of uncertainty (when one's actions can be tied to exact outcomes probabilistically), the lack of transparency in outcomes (when one's actions cannot be directly tied to exact outcomes, intentionally or unintentionally), or a combination of the two (with further elaboration provided in later sections), and that a licensing effect occurs that dominates prosocial (or fair) preferences when uncertainty is present. More specifically, I present participants on the Amazon Mechanical Turk with three conditions of a binary dictator game (henceforth abbreviated to BDG), and collect information on their demographic information. Additionally, I suggest later that there may exist a relationship between underlying psychological characteristics of participants and their preferences for fairness (as revealed by their choices in each game's condition).

From here on, this thesis will take the following development. The second section describes the experimental design I used and why it was selected for my hypotheses. The third section explains the hypotheses and expectations for this study. The fourth section describes the data collection and process. The fifth section describes and justifies my strategy for analyzing the data. Finally, the sixth section will synthesize my explanations and findings in order to draw conclusions about the initial hypothesis, express the limitations of this study, and discuss its relevance for any future work.

Chapter 2

Experimental Design

This study aims to identify an underlying mechanism that exists between the interaction of uncertainty, transparency, and preferences for fair or unfair allocations in economic decision-making. In exploring the existing literature for how other researchers have handled questions involving fairness in economic contexts, I found that economic games provided a fruitful framework to work in as a way to control and reduce the potential complexities that one might run into if trying to proceed in a purely theoretical, econometric, or statistical way. I decided an adjusted one-shot dictator game would be the simplest and most direct way to try to observe economic decisions about fairness.

A dictator game, in its most basic form, involves two people: the “dictator” and the recipient. The person assigned to the position of dictator is given an allocation that he or she must split under some particular rules, and the recipient’s role is simply to receive whatever has been picked by the dictator. I chose a binary dictator game as the basic game for my study, where the dictator game is adjusted by allowing for only two options for the dictator to split a total amount of $10. The first option (or option A in this study) is evidently the best option for the dictator if he or she wants to maximize his or her total payoff, since it offers $9 for the dictator and $1 for the recipient. The second option, option B, is an even $5 split between the dictator and the recipient, and is clearly the more equitable one. This game was constructed with the goal of simulating the simplest possible form of making a fair economic decision. I chose to make the game a one-shot game, since if the game had multiple plays, other factors related to fairness might play a more central role, like compounded sympathy or guilt, which might make disentangling the effects of uncertainty and transparency intractable (this could seem trivial in passing, but even when a concerted effort is made to ensure complete anonymity between participants, the amount given by the dictator is still often greater than 0, so that sequential play could intensify very weak preferences for equitable outcomes; see Hoffman et al. 1994). Additionally, since BDGs have been played and used many times across a variety of studies, I chose to rely on previous work in establishing a baseline for comparison. 75% or higher selection rates toward the fair choice is common in a BDG with one fair and one unfair option (see Camerer 2003, chapter 2 or Kahneman et al. 1986, for example), so the baseline of comparison that I use is that roughly 25-30% of people should choose the unfair option and roughly 70-75% of people should choose the fair option, on average. With all of this in hand, I used three different conditions of the BDG, each with a different tweak to test for specific features I outline in my hypotheses in the next section.

Participants take a survey online on the Amazon Mechanical Turk. The survey involves an explanation of how the game will be carried out, conveying that they will participate in a study meant to examine economic games. The participants will be compensated through mTurk’s payment setup (which they will know going into the survey) in addition to a random lottery scheme. Participants are informed that 1 in 15 participants will receive an extra $10 added as a bonus on top of what they are normally paid for completing a survey on mTurk (50 cents) and that they should imagine taking the survey with this in mind. When they begin the study, a part of their instructions include treating the potential $10 as the $10 they are splitting, as whatever split they decide will be paid to them and the (hypothetical) recipient. After gathering responses, I go back through and randomly credit $9 or $5 depending on the participant's choice for every 1 in 15 participants.

The first condition is the uncertainty condition. It involves presenting the participant with the two choices, where option B is obviously "fair" (a payoff of $5 each) and option A is obviously not (a payoff of $9 for one participant and $1 for another). The participant views the two options for one full minute, along with their options for selecting between option A and B. In this condition, the participant is given probabilities to assign to each payoff, so that the computer randomly chooses between the two options based on the probabilities assigned, which is why this condition is referred to as the uncertainty condition. The participant decides a probability to be assigned each option rather than directly choosing the option. The probabilities range from a 100% chance of one option and 0% of the other, with incremental options in between (i.e., 70% chance for one option and 30% for the other, 60% chance for one and 40% chance for the other, etc.). This first variant is meant to observe people's preferences in two ways: first to see how strong people's preferences for fair outcomes are when their actions are less certain, and second to see if uncertainty changes people's behavior relative to a baseline of predominantly fair-leaning preferences in the literature. Because there is an implicit understanding that there is factually no uncertainty due to a 100% is selection option, if people have either strong or nonexistent preferences for prosocial outcomes, a redundant set of options should not change their choices.

In the second condition, which I refer to as the transparency condition, participants only have to directly choose either option A ($9 for the participant and $1 for the other person) or option B ($9 for the participant and $1 for the other person). But, the participant is told he or she has ten seconds to respond after the one minute elapses, with the possibility of being cut off randomly during that ten seconds. If the participant is cut off, then the computer will choose for the participant with even probability on both options, and the participant is told this as part of his or her instructions for this condition. The point of the transparency condition is to give the participant an option to hide behind, or a sense of plausible deniability: if they prefer the self-interested option but do not want to select it outright, the computer can pick for them (hence allowing for a lesser degree of transparency in their actions). Additionally, across all conditions, participants are told the recipient will not see how they chose the allocation they did, and will only see the amount given to them, so that anonymity of actions is also assured. A pilot run indicated that the average response time for an unadjusted, straightforward BDG is between 2 and 3 seconds, so that whatever the participant's preferences may be, if their preferences are consistently strong (or nonexistent) then there should be no big difference in response time, as a redundant option has been added to this condition.

The third condition is the first and second condition combined, and I refer to it as the mixed uncertainty condition. The point of this condition is to see how participants choose to assign probabilities (if at all) as compared to the first and second conditions, so that we can see if uncertainty or less transparency in actions causes more self-interested behavior dominantly.

Finally, copies of the questions asked, the conditions I described, the consent form, and the participants' debriefing are listed at the end of this document. Appendix A contains the consent form example, Appendix B the surveys and BDGs across the three conditions, and Appendix C includes the questionnaires answered after the BDGs.

Chapter 3

Hypotheses & Expectations

Owing to the experimental approach that I am using, a formally stated hypothesis is necessary before proceeding in lieu of what would otherwise be a statement of functional forms or assumptions needed to proceed with a theoretical and then econometric analysis, normally. Additionally, before stating my hypothesis, it is essential to define precisely what I mean in this study by “uncertainty” and “transparency”.

First, by uncertainty I mean the usual statistical and economic interpretation of the word, which is to say any choice which is not strictly deterministic—if selected—is defined to be a choice under uncertainty. In particular, however, the individual can perfectly observe what outcomes are available to him or her (so that one of these options is picked with 100% probability), but the range of options presented may be in terms of implicit or explicit probabilities. This leads the consumer to always have a general, but uncertain, degree of knowledge about which option is likely to be picked (unless a 100%-0% option is picked if the consumer is deciding between two choices, for instance).

As for transparency, I mean any option that leaves the individual unable to identify which outcome he or she has been selected into, even with some degree of uncertainty (which in this study will take the form of a deferment option in one condition of the BDGs). This means that while the individual may or may not have an idea of what the chances are of a certain outcome, he or she cannot associate a direct choice on their part with the resulting outcome. In effect, lack of transparency in this study is any purposefully or non-purposefully ignorant choice in the form of deferring to choose. A choice is randomly assigned for the individual when this happens, though the individual does not know if the assigned choice is randomly between all options or automatically set to a default option. (In contradistinction, uncertainty allows for knowledge of chance between choices in all cases at the time of choice, and so in all cases allows for a general understanding of which choice is likely (or unlikely) to be selected by the individual choosing the probabilities.)

I will state and explain my hypothesis in three components, labelled H1, H2, and H3, which correspond to each condition respectively:

H1: Participants choose more self-interestedly in under the uncertainty condition than under the lack of transparency condition, since it allows for people to allocate *some* chance of the fair option. This operates through weak (but consistent) preferences for fairness economically and a licensing effect, psychologically. Economically, this would reflect existing but weak preferences for fairness because the rational, subgame perfect choice still exists as a deterministic choice (namely a 100% chance for the self-interested allocation and a 0% chance for the fair allocation), so that participants should only give some chance of the fair option if they derive some form of benefit (or utility) from it. If people's preferences for fairness are strong and consistent, their responses should be dominantly skewed toward the fair option, as what has been found to be common in typical BDGs with no uncertainty or decreased transparency (Dana et al. 2006).

Psychologically, people will use probability weights as a sort of self-imposed cost in order to choose more self-interestedly, so that people license the self-interested choices to themselves by choosing a lottery with positive probability for both options rather than choosing the optimal lottery that assigns 0% probability to the fair option.

H2: Participants choose more self-interestedly in the lack of transparency condition than in BDGs without uncertainty or decreased transparency, but less so than in uncertainty condition. Economically, if people's preference for fairness are strong and consistent, there should be no difference in unadjusted BDGs, the uncertainty condition, and this one, so that the fair option should be dominantly picked. I claim that since there are weak preferences for fair outcomes, an option for loss of agency introduces the same licensing effect as in the uncertainty condition. But, since participants do not know how much they are paying in terms of probabilities for each option when they defer their choice, there is an element of risk for choosing the fair option if they truly lean more toward the unfair option or vice versa. As a result, there will be greater amounts of people choosing the self-interested option outright than in the standard BDG, and that the number of people who choose to defer will be less than the number of people who choose to act self-interestedly directly. Overall, though, I expect there to be a majority of responses that lean toward the fair allocation since I believe the small risk introduced and the fewer options to self-license will lead participants to choose the fair option predominantly.

From a psychological perspective, the self-imposed cost of licensing the unfair option to oneself varies from person to person: some people might feel a lesser or greater need to choose a probability that leans more toward the fair option in order to allow for any chance of receiving the self-interested allocation. Additionally, I also expect that since people cannot fully manipulate probabilities to license some chance of the unfair option to themselves, people with preferences leaning slightly more than 50-50 toward the fair allocation will choose the fair allocation, and people with preferences leaning more than 50-50 toward the unfair allocation will choose the unfair allocation. If people are actually licensing the unfair option to themselves by choosing probabilities, then the group of people with about 50-50 preferences for fairness should choose the licensing option through deferring, and should naturally be smaller than either groups of people with preferences greater than a 50-50 split for the fair and unfair allocations.

H3: The mixed uncertainty condition will have the most self-interested behavior, since it offers the greatest variety of self-licensing options through either choosing probabilities or through not choosing anything (resulting in a 50-50 default). But, even if the distribution of responses reflects the most self-interested behavior, I still expect small probabilities to be allowed for the fair option.

From an economic standpoint, all participants now have some way of alleviating whatever risk comes with deferring their choice through choosing appropriate probabilities, with the obvious best choice for optimizing utility being the 100% chance pick of the unfair option if there are no preferences for fairness, or the 100% chance pick of the fair option if there is some strong preference for fairness.

From a psychological standpoint, those who value *not* having to choose between allocations but who still want the higher bundle have the choice of deferment ("If I don't choose, I'm not being selfish, but I still have a chance at winning the better choice."). For those who are motivated by rationalizing their decisions when choosing self-interestedly, they can rely on licensing self-interested behavior to themselves through giving the other person a nonzero probability of winning the fair split ("I'm giving the person *some* chance of winning the higher split, even though I don't have to."). Additionally, participants are told that the recipient of the allocation does not see how the $10 was split, and only sees the outcome that they will be paid out, so participants have no real use for deferring their choice. The recipient will get either $5 or $1, and how you arrive at the figure does not change the recipient's value of the $5 or $1, since he or she cannot see how it was picked.

The key here is that, like the previous conditions, this BDG is unchanged from a rational, expected-utility maximizing individual's perspective. In all conditions, there is a deterministic choice for the fair option and for the unfair option. Economically, if participants are strongly motivated by their preferences for fair outcomes, there should be not too much fluctuation between these conditions. If participants are weakly motivated by preferences for fair outcomes, then we should see significant variation across these three conditions as compared to a typical baseline of three-fourths of people choosing the fair option. For uncertainty to dominate transparency for self-interested choices, we would also hope to see significant statistical variation from fair-leaning distributions of choices for conditions involving uncertainty, and less significant statistical variation for conditions addressing less uncertainty or transparency. We would also hope to see small probabilities assigned to the fair option even with a strongly self-interested distribution of responses if people actually do engage in licensing unfair allocations to themselves assigning other participants a small chance at the fair option.

Chapter 4

Data Collection

Initially, I had wanted to collect data with a combination of two methods: in-person lab participants and data collected through the Amazon Mechanical Turk (or mTurk). But, due to financial and time constraints, I was only able to handle data collected from mTurk. While this may limit the data somewhat, overall I believe that it should not dampen my study too seriously. Reips (2000, 2002) discusses and shows that mTurk's strengths lie in subject pool access, subject pool diversity, and low cost. In particular, the subject pool is relatively stable throughout the year while offering a sampling of participants from a wider background than what is typical. And, in stark contrast to common costs associated with typical in-person experiments, it is not difficult to pay under $2 per participant per hour. Paolacci et al. (2010) manage to replicate classic studies from the judgment and decision-making literature at a cost of around $1.71, which paralleled the same results obtained in-lab with undergraduates. Consequently, because of the success others have had with mTurk and for its cost-effectiveness, I chose mTurk as my study's medium for data collection.

124 people participated in my study, 40 in group 1, 42 in group 2, and 42 in group 3, all assigned randomly. The demographics breakdown is in Tables 1, 2, and 3, (respectively for conditions 1, 2, and 3) with gender, age, education, and ethnicity listed. A few of the tables do not fully add up to the full number of participants for two of the conditions, but this is because 2 participants chose not to respond with demographic information.

Table 4.1 – Age for Condition 1

Age Range	Response	Percentage
Under 10 years	0	0%
10 to 14 years	0	0%
15 to 19 years	1	3%
20 to 24 years	8	21%
25 to 29 years	7	18%
30 to 34 years	5	13%
35 to 39 years	4	10%
40 to 44 years	1	3%
45 to 49 years	3	8%
50 to 54 years	5	13%
55 to 59 years	4	10%
60 to 64 years	1	3%
Total	39	
Mean:	30 to 34 years	

Table 4.2 – Gender for Condition 1

Gender	Response	Percentage
Male	18	46%
Female	21	54%
Total	39	

Table 4.3 – Education for Condition 1

Education Range	Response	Percentage
0-11 years	0	0%
12 years	15	38%
13-16 years	21	54%
17-18 years	2	5%
19+ years	1	3%
Total	39	
Mean:	13-16 years	

Table 4.4 – Ethnicity for Condition 1

Ethnicity	Response	Percentage
Caucasian	30	77%
African-American	3	8%
Asian	2	5%
Hispanic	3	8%
Indian	1	3%
Native American/ Middle Eastern/African/ Mixed/Other	0	0%
Total	39	
Average ethnicity:	Caucasian	

Table 4.5 – Age for Condition 2

Age Range	Response	Percentage
Under 10 years	0	0%
10 to 14 years	0	0%
15 to 19 years	4	10%
20 to 24 years	10	24%
25 to 29 years	8	19%
30 to 34 years	10	24%
35 to 39 years	6	14%
40 to 44 years	0	0%
45 to 49 years	0	0%
50 to 54 years	1	2%
55 to 59 years	1	2%
60 to 64 years	2	5%
Total	42	
Mean:	30 to 34 years	

Table 4.6 – Gender for Condition 2

Gender	Response	Percentage
Male	26	62%
Female	16	38%
Total	42	

Table 4.7 – Education for Condition 2

Education Range	Response	Percentage
0-11 years	0	0%
12 years	7	17%
13-16 years	30	71%
17-18 years	4	10%
19+ years	1	2%
Total	42	
Mean:	13-16 years	

Table 4.8 – Ethnicity for Condition 2

Ethnicity	Response	Percentage
Caucasian	30	77%
African-American	3	8%
Asian	4	5%
Hispanic	2	8%
Indian	1	3%
Mixed	1	2%
Native American/ Middle Eastern/African/ Other	0	0%
Total	42	
Average ethnicity:	Caucasian	

Table 4.9 – Age for Condition 3

Age Range	Response	Percentage
Under 10 years	0	0%
10 to 14 years	0	0%
15 to 19 years	0	0%
20 to 24 years	8	20%
25 to 29 years	11	28%
30 to 34 years	8	20%
35 to 39 years	2	5%
40 to 44 years	1	3%
45 to 49 years	3	8%
50 to 54 years	1	3%
55 to 59 years	3	8%
60 to 64 years	3	8%
Total	40	
Mean:	30 to 34 years	

Table 4.10 – Gender for Condition 3

Gender	Response	Percentage
Male	27	66%
Female	14	34%
Total	40	

Table 4.11 – Education for Condition 3

Education Range	Response	Percentage
0-11 years	0	0%
12 years	9	22%
13-16 years	29	71%
17-18 years	3	7%
19+ years	0	0%
Total	40	
Mean:	13-16 years	

Table 4.12 – Ethnicity for Condition 3

Ethnicity	Response	Percentage
Caucasian	31	77%
African-American	6	8%
Asian	2	5%
Hispanic	1	8%
Indian	0	3%
Native American/ Middle Eastern/African/ Mixed/Other	0	0%
Total	40	
Average ethnicity:	Caucasian	

For the sake of this study, since the game I used is a simple one, I decided the difference between administering it in person versus online would be a small one, and in fact, there are a number of studies which speak to the merit of mTurk for data collection in situations like mine. For example, as mentioned before, Paolacci et al. gives examples of a number of classic studies which have been successfully and robustly replicated on mTurk, but as for economic games specifically, Amir et al. (2012) report comparable results to those obtained in the laboratory: Amir et al. work under a similar paradigm as my own (an ultimatum game rather than a dictator game) and found that while responders' behavior are mixed, the proposers (or dictators, in my study) exhibited behavior consistent with previously observed behavior in the laboratory, even with a more diverse pool of subjects. Additionally, people have been shown to respond about the same to monetary incentives in person as when galvanized by a lottery system on mTurk, as my study does, as also described in Paolacci et al. And, while my study does deal with smaller amounts of money, Raihani et al. (2013) have shown that US participants tend to be invariant in their responses across amounts of $1, $5, and $10 in dictator games on mTurk. With that given, this works all the better for my study: my pool draws only from United States participants, and if people are to show their preferences for prosocial outcomes and fair allocations, then the lower stakes used in my three BDGs should cause no distortion to their preferences.

As for the caveats of these data, there are a few to name. The first is that the sample for each condition is predominantly Caucasian, 30-34, roughly college-educated, and predominantly male. So, while one of the advantages of mTurk is to be able to get away from completely homogenous subjects, and even though there a decent degree of diversity in my sample, because my sample is modest it is still has a predominantly while, middle-aged, and college-educated profile. Also, while the sample size should be sufficient as a baseline, it is still small enough that the effect sizes observed might be smaller or larger than they should be, and in particular, if there is some measurement error associated with the three BDGs or with the questionnaires, then this sample size would hardly help to alleviate the bias introduced. Finally, one could also argue that there might be an element of a selection bias involved in the data: perhaps those people who are more likely to act self-interestedly are the ones who would spend time extra time online on mTurk to collect extra cash, or perhaps people who are most strapped for cash might take to mTurk for extra money and may further sympathize with the other (fictitious) player on mTurk who is assigned to the the BDG. In either case, if the population of such people is large enough for the sample I have collected, it is possible the data exhibit a bias purely produced by sampling from mTurkers themselves. But, because of the aforementioned studies, I believe this is possible but not likely.

Chapter 5

Data Analysis

If designed well, an experiment should allow for a straight-forward data analysis; in this hope, this section will be straightforward. First I will recall the main questions of the study, and then assign to them the relevant method of analysis that I chose.

This experiment tries to examine if people respond to uncertainty and degrees of transparency with any significance, so my first task is to show, statistically, that this is the case. Chi-square goodness of fit tests would allow for this: my data are categorical in structure, I have a particular theoretical distribution for the probability of each event (the choice of allocation) as a null hypothesis, and I wish to test whether the distribution of observations varies significantly with respect to the null hypothesis that roughly 75% of people prefer the fair option and roughly 25% prefer the unfair one (or similar variations). Specifically, since the uncertainty and mixed uncertainty conditions have 7 options (when one counts all possible probability choices presented to participants, shown in the Appendices), I will compare these two using two initial distributions: one corresponding to a null hypothesis that there are no preferences for fair or unfair outcomes (so that each choice should have equal weight), and one corresponding to a null hypothesis that there are strong preferences for fair outcomes. But, since there is no one correct distribution characterizing strong prosocial outcomes, I will use a variety of null distributions that tend toward strong prosocial outcomes. I will then compare the transparency and uncertainty conditions to see if adding lack of transparency to uncertainty creates a statistically significant difference in the distribution of responses (as compared to a null distribution of even chances for each option, and then as compared to each other).

As with any statistical procedure, there are cautionary limits of the methods outlined to keep in mind. For the goodness of fit test, the expected value of the number of sample observations in each level of the variable in question needs to be at least 5. Given that I have one variable, the choice of the respondents, this should not be an issue. Otherwise, the goodness of the fit test becomes unreliable. Another necessary condition is that the degrees of freedom for the goodness of fit test be at least 2 or greater (given by $N - k - 1$, where N is the number of observations, and k is the number of fitted variables). This also should not be an issue, since the smallest level of degrees of freedom should be 2, and the rest should be at 6.

Chapter 6

Synthesis & Conclusions

For this section, I will discuss what results I came to for the experiment under different goodness of fit tests and then elaborate on their implications. First, I will describe the tests for prosocial preferences under uncertainty that test the existence of prosocial preferences. Then, I will compare the uncertainty and mixed uncertainty conditions under a variety of null distributions that reflect strong prosocial preferences. This should establish whether or not there is evidence in these BDGs for prosocial preferences at all, strong preferences, or weak preferences. From there, I will examine the transparency condition against an even probability null distribution to test the lack of transparency's effect on prosocial preferences. After that, I test the transparency condition against what is typically found in the literature (3/4's of people choosing fairly) to test against strong prosocial preferences. This should show us whether lack of transparency significantly contributes or detracts from prosocial behavior, and should allow us to distinguish between whether lack of transparency is dominant or if uncertainty is dominant. These analyses will conclude with an examination of the distribution of responses across conditions in relation to my findings, and if in fact participants exhibited licensing behavior.

In comparing the uncertainty and mixed uncertainty conditions with a null distribution of even probabilities, or roughly a 14% chance for any given choice since there were 7 choices for probability assignments (see Appendix for exact survey question and options), I find that both conditions differences are statistically significant at the .01 level: the uncertainty condition reported a $\chi(6)^2 = 39$, with p-value = 7.15x10^{-7}, and the mixed uncertainty reported $\chi(6)^2 = 21.47$, p-value = .0015. This means we can reject the null that there are nonexistent or indifferent preferences toward prosocial and/or self-interested preferences. Proceeding, I perform five more χ^2 goodness of fit tests which represent a distribution of choices leaning strongly toward prosocial preferences, as given in Table 6.1 below.

Table 6.1 – χ^2 goodness of fit tests for Uncertainty & Mixed Uncertainty Treatments (Prosocial Distributions)

Null Distribution	Condition1 $\chi(6)^2$,	$p-val.$	Condition 3 $\chi(6)^2$,	$p-val.$
P1 - (.05, .05, .15, .15, .2, .2, .2)	39.56	5.5x10^{-7}	90.07	2.2x10^{-16}
P2 - (.1, .1, .1, .1, .2, .2, .2)	53.55	9.07x10^{-10}	35.63	3.25x10^{-6}
P3 - (.02, .02, .02, .04, .2, .3, .4)	196.05	2.2x10^{-16}	290.86	2.2x10^{-16}
P4 - (.08, .08, .1, .14, .2, .2, .2)	35.73	3.10x10^{-6}	47.20	1.70x10^{-8}
P5 - (.1, .1, .1, .1, .1, .1, .4)	45.42	3.847x10^{-8}	29.31	5.301x10^{-5}

In all tests with strong prosocial distributions, the uncertainty and mixed uncertainty conditions reject the null distribution, even when allowing for 40% of the distribution of responses to be indifferent or leaning self-interestedly. This lends support to the idea that people neither have nonexistent prosocial preferences nor strongly prosocial preferences (at least in the context of these BDGs). Similarly, when we run these same distributions but in retrograde (reflecting a variety of distributions of strongly self-interested preferences), we arrive at similar results.

Table 6.2 – χ^2 goodness of fit tests for Uncertainty & Mixed Uncertainty Treatments (Self-Interested Distributions)

	Uncertainty Condition		**Mixed Uncertainty Condition**	
Null Distribution	$\chi(6)^2$,	$p-val.$	$\chi(6)^2$,	$p-val.$
P1' - (.2, .2, .2, .15, .15, .05, .05)	102	$2.2\text{x}10^{-16}$	90.07	$6.82\text{x}10^{-11}$
P2' - (.2, .2, .2, .1, .1, .1, .1)	73.55	$6.99\text{x}10^{-14}$	24.14	.0004
P3' - (.4, .3, .2, .04, .02, .02, .02)	387.70	$2.2\text{x}10^{-16}$	183.17	$2.2\text{x}10^{-16}$
P4' - (.2, .2, .2, .14, .1, .08, .08)	66.35	$2.27\text{x}10^{-12}$	30.62	$2.98\text{x}10^{-5}$
P5' - (.4, .1, .1, .1, .1, .1, .1)	73.15	$9.21\text{x}10^{-14}$	21.42	.0015

Since in both cases we reject the null hypothesis, these goodness of fit tests lend support for my hypothesis that people have existing preferences for fairness (since we reject the null distribution corresponding to indifferent preferences), but that they are neither strongly prosocial nor self-interested.

Next we look at the results for the lack of transparency condition where participants are told to decide (deterministically) between the fair and unfair option with a tacit option to defer their choice in lieu of a 50-50 chance at each option. Table 5.3 displays the goodness of fit test for a null distribution with even probabilities assigned to all options (which is the default option for the goodness of fit test in R).

Table 6.3 – χ^2 goodness of fit tests for the Lack of Transparency treatment

Null Distribution	$\boldsymbol{\chi(2)^2}$,	$\boldsymbol{p-val.}$
Q1 – Even probabilities (default)	51.49	6.5x10^{-12}

Since the p-value is sufficiently small (passing at an alpha level of .01) it is safe to say that we can reject the null hypothesis of indifferent preferences in this context. Next, in accordance with the previously mentioned 25%-75% split commonly found in the literature (25% of people choosing the unfair option on average and 75% choosing the fair one on average), Table 6.4 shows goodness of fit tests for three variants of distributions that predominantly favor the prosocial choice while allowing for deferment to the 50-50 choice.

Table 6.4 – χ^2 goodness of fit tests for Lack of Transparency treament

Null Distribution	$\chi(2)^2$,	$p - val.$
Q1' – (.125, .125, .75)	.57	.752
Q2' – (.1, .15, .75)	.48	.784
Q3' – (.15, .1, .75)	.76	.683

Interestingly, I find no significant statistical difference from the null distributions above and the responses by participants in the transparency condition. Since we cannot reject the null hypothesis, it would seem that the lack of transparency in this condition, through providing an extra option for deferment, does not have a significant effect in regard to usual findings in standard BDGs.

Below I have also included the distribution of responses across conditions. In conditions that allowed for a deferment option, I include two tables to reflect the distribution of choices when taking deferment into account and when excluding it. The difference is most notable in the third condition, where around half of all respondents choose to defer to the 50-50 option. Across all conditions we can also see that participants provided a variety of nonzero probabilities for the fair option, though most noticeably in the uncertainty and mixed uncertainty conditions, participants assign smaller weights to the fair option more predominantly whenever the distribution is more self-interestedly leaning, providing some evidence for licensing behavior.

Table 6.5 – Distribution of Choices for the Uncertainty Condition

Option	Response	Percentage
Choice A: $9 for you, $1 for someone else		
Choice B: $5 for you, $5 for someone else		
100% chance of A & 0% chance of B	8	20%
80% chance of A & 20% chance of B	5	13%
60% chance of A & 40% chance of B	3	8%
50% chance of A & 50% chance of B	9	23%
40% chance of A & 60% chance of B	1	3%
20% chance of A & 80% chance of B	5	13%
0% chance of A & 100% chance of B	9	23%

Table 6.6 – Distribution of Choices for the Transparency Condition

(Deferment responses excluded)

Option	Response	Percentage
Choice A: $9 for you, $1 for someone else		
Choice B: $5 for you, $5 for someone else		
Choice A	15	44%
Choice B	19	56%

Table 6.7 – Distribution of Choices for the Transparency Treatment

(Deferment responses included)

Option	Response	Percentage
Choice A: $9 for you, $1 for someone else		
Choice B: $5 for you, $5 for someone else		
Choice A	15	36%
Choice B	19	45%
Defer (50-50 between A and B)	8	19%

Table 6.8 – Distribution of Choices for Mixed Uncertainty Condition

(Deferment responses excluded)

Option	Response	Percentage
Choice A: $9 for you, $1 for someone else		
Choice B: $5 for you, $5 for someone else		
100% chance of A & 0% chance of B	11	52%
80% chance of A & 20% chance of B	1	5%
60% chance of A & 40% chance of B	2	10%
50% chance of A & 50% chance of B	4	19%
40% chance of A & 60% chance of B	1	5%
20% chance of A & 80% chance of B	0	0%
0% chance of A & 100% chance of B	2	10%

Table 6.9 – Distribution of Choices for Condition 3 (Deferment responses included)

Option	Response	Percentage
Choice A: $9 for you, $1 for someone else		
Choice B: $5 for you, $5 for someone else		
100% chance of A & 0% chance of B	11	26%
80% chance of A & 20% chance of B	1	2%
60% chance of A & 40% chance of B	2	5%
50% chance of A & 50% chance of B	25	60%
40% chance of A & 60% chance of B	1	2%
20% chance of A & 80% chance of B	0	0%
0% chance of A & 100% chance of B	2	5%

Now that our analyses and data are complete, there are a few conclusions that can be reached regarding my initial hypotheses. First I will address my hypotheses in the order I named them.

H1: *Participants choose more self-interestedly in the uncertainty condition than in the lack of transparency condition, since it allows for people to allocate some chance of the fair option. This operates through weak (but consistent) preferences for fairness economically and a licensing effect, psychologically.*

Comparing the uncertainty condition to the transparency condition, we see that participants do indeed act more self-interestedly in the uncertainty condition. Now from our goodness of fit tests across conditions, we see that while uncertainty creates a significant fluctuation in participants' choices (as compared to a 25%-75% split), lack of transparency does not significantly create fluctuations in participants' choices. Additionally, we have seen that the null distributions corresponding to indifferent preferences as well as strong prosocial preferences were rejected. This consequentially leaves us with the conclusion that prosocial preferences are weak and conditional on uncertainty rather than lack of transparency. If self-interested preferences were robust under *any* degree of uncertainty or transparency, we should have had similar rejection rates in the uncertainty and transparency conditions. But, since we only had statistically significant distributions with the uncertainty condition and 3, we are guided to the conclusion that uncertainty is the dominant driving force. Further, we can assert that there are weak preferences for fairness because under any and all conditions we still observe choices that give some small probability to the recipient for receiving the fair outcome, so that in any condition *some* preference for fairness is expressed, but with its intensity depending on the presence of probabilistic uncertainty. Indeed, under the deepest uncertainty (when prosocial/fair preferences diminish the most in the mixed uncertainty treatment) we can see licensing behavior at work through a heavy progressive shift in probabilities across conditions. The transparency condition shows statistically similar responses to typical BDGs, the uncertainty condition shows a shift toward more self-interested options, and the mixed uncertainty condition shows the heaviest shift toward self-interested options, but even still, when the most self-interested distribution appears in the mixed uncertainty condition, the bulk of the distribution allows for some significant chance of the fair option. In fact, the split in probabilities shifts to a nearly bimodal distribution between the deterministic unfair choice and the 50-50 split in the mixed uncertainty condition, with nearly no choices present for probabilities favoring the fair option. Economically, this would make sense if there were any true risk present, since higher risk would cause deeper convexities in utility functions, and in turn participants should demand higher risk premiums, which would explain the shift to more self-interested behavior. But since there is no true risk (a deterministic choice is always present for whichever choice the subject prefers) this cannot be explained by convexity and risk premiums, so it must be the case that some degree of utility is gained from the potential for prosocial outcomes. This is what marks the difference between a risk premium and licensing: the risk is illusory, yet the participant selects a lottery that should be seen as riskier anyway.

H2: *Participants choose more self-interestedly in the transparency condition than in BDGs without uncertainty or decreased transparency, but less so than in uncertainty condition. But, since participants do not know how much they are paying in terms of probabilities for each option when they choose to defer their choice, there is an element of risk for choosing the fair option if they truly lean more toward the unfair option or for choosing the unfair option if they truly lean more toward the fair option.*

Unlike in the uncertainty and mixed uncertainty conditions, there is some element of risk in the transparency condition since there is a chance the subject will get a choice he or she does not truly want if they pick to defer. This would seem to provide an incentive to choose exactly the preference one would want. As predicted, we do see more self-interested behavior in the transparency condition than in standard BDGs, and we do see less self-interested behavior in the transparency condition than in uncertainty condition, but without statistical significance. This would suggest that either the effect of the treatment was not strong enough in this condition, or that transparency alone is a weaker form of uncertainty, and causes lesser shifts toward self-interested behavior. Taking the latter train of thought, since transparency and how it was presented in this study is a weaker form of uncertainty, no subjects can adequately license self-interested behavior, and they behave largely as if no uncertainty were introduced. This null result then lends some credibility to uncertainty (rather than transparency) being the driving factor in licensing behavior and higher levels of self-interested behavior.

H3: *Mixed uncertainty will have the most self-interested behavior, since it offers the greatest variety of self-licensing options through either choosing probabilities or through not choosing anything (resulting in a 50-50 default). All participants now have some way of alleviating whatever risk comes with deferring their choice through choosing appropriate probabilities, with the obvious best choice for optimizing utility being the 100% chance pick of the unfair option if there are no preferences for fairness, or the 100% chance pick of the fair option if there is some strong preference for fairness. Those who value not having to choose between fair and unfair allocations (through wanting the unfair allocation without wanting to acknowledge the fact) have the choice of loss of agency, as measured by response time. Average response time is 2.4 seconds after viewing choices, and since respondents have five times as much time to respond before being cut off, nonresponse greater than 4-5 seconds will be interpreted as deferring choice.*

As predicted, mixed uncertainty yields the highest rates of self-interested behavior. Additionally, the mixed uncertainty condition's distribution of choices was found to be statistically different from the transparency condition's distribution, as well as the indifferent and strongly prosocial null distributions. This suggests two things: adding more or different degrees of uncertainty in an uncertain environment augments self-interested behavior, and different forms of uncertainty contribute to licensing behavior in environments where probabilities for outcomes can be manipulated. Since another variety of uncertainty through deferment of choice is present, participants have more options in serving their preferences. Since all possible options for fair and unfair allocations are accounted for, allowing for a choice for each level of preference for the unfair/fair option, those who are motivated by rationalizing their decisions when choosing self-interestedly can rely on licensing self-interested behavior to themselves through nonzero probabilities of the recipient winning the fair split. On the other hand, for those who may not be motivated by rationalizing their decisions (but who are interested in the higher payoff), the potential cost of losing the unfair allocation acts as the cost of licensing for the chance of winning the unfair split, causing a predominant mass in the 50-50 split for the mixed uncertainty condition. This would imply consistent but weak preferences for fair outcomes because as uncertainty increased, so did participants' options for licensing behavior (though they were assured the recipient would never see how the choice was made), which should make it easier to choose self-interestedly. But, instead, a bimodal distribution between the unfair option and the 50-50 split formed, showing that when given a variety of ways to anonymously choose self-interestedly people still want to choose equitably in *some* capacity. Additionally, all participants who took longer than average to decide either chose self-interestedly (a lottery favoring 60% or greater toward the self-interested option) or waited until they were cut off (assigning them the 50-50 lottery).

In closing, this experiment's goal was to explore whether a tenable argument exists for weak preferences for fair preferences in economic games, and if so, ask how people may be making these decisions. I argued that people attempt to license self-interested behavior to themselves as a means of rationalization in environments with uncertainty, and that as uncertainty increased, people still express some weak preference for fair outcomes. This informs the literature by showing that people's preferences are neither nonexistent nor strong for prosocial outcomes, and that when preferences flip (or seem inconsistent) it can in large part be attributed to the presence of uncertainty. Additionally, we can say that uncertainty may dominates lack of transparency in participants' actions as driving self-interested behavior because, as predicted, the degree of self-interested behavior increased from the transparency condition, to the uncertainty condition, and increased the most in the mixed uncertainty condition. But, as always, there are limitations to consider. My sample size for this study is not particularly large, which could have a significant effect on some of the goodness of fit tests. Indeed, to counteract this, in part, I attempted to apply the goodness of fit test to a number of distributions in order to get a better sense for its accuracy, since at very low probabilities it breaks down. Even given null distributions involving low probabilities, in comparison to more modest null distributions the results remain the same. It is also difficult to extrapolate the results in this study out generally, as no attempt was made to examine the external validity for this study and since the sample was not overly diverse, truncation issues would likely crop up as a critique. Though as a pilot into exploring what conditionally weak prosocial preferences could mean for economic games and decision making, the main attempt has been to try to explore licensing behavior better. Finding some evidence for this, I would make some suggestions for future work in the following ways. There could be conditions involving the fake recipient being able to see how the dictator chooses to see if licensing behavior dissipates. If so, it could be the case that licensing behavior is driven by deceptive motives under the guise of decision-making under uncertainty, and so would be other-regarding. If this were the case, we could test this hypothesis by using a number of scales known to correlate well with creativity, numeracy, or other traits, and see what relationship exists between these traits. Previous work by Gino & Ariely (2012) suggest higher creativity could mean better ways of deceiving oneself and others, so one could examine the relationship between creativity, for example, and licensing behavior in economic games under uncertainty. It could be possible that people reporting higher degrees of creativity may exhibit more self-interested behavior on average than other participants. Additionally, one could try explore how robustly licensing behavior is: do people exhibit the same behavior in allocation decisions when they are a part of a larger coalition acting as the dictator? If higher creativity means higher self-interested behavior, should higher degrees of numeracy lead to lower ability to self-deceive, and in turn, to a tendency toward heavier licensing or a greater preference for prosocial outcomes? These and many other variants could be proposed to see to what point uncertainty causes people to act more self-interestedly through licensing.

Bibliography

Amir, Ofra, and David G. Rand. "Economic games on the internet: The effect of $1 stakes." *PloS one* 7.2 (2012): e31461.

Antinyan, Armenak. "Loss and Other-Regarding Preferences: Evidence From Dictator Game." *Department of Management, Università Ca'Foscari Venezia Working Paper* 3 (2014).

Camerer, Colin F. "Behavioural studies of strategic thinking in games." *Trends in cognitive sciences* 7.5 (2003): 225-231.

Camerer, Colin. *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press, 2003.

Dana, Jason, Roberto A. Weber, and Jason Xi Kuang. "Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness."*Economic Theory* 33.1 (2006): 67-80.

Gino, Francesca, and Dan Ariely. "The dark side of creativity: original thinkers can be more dishonest." *Journal of personality and social psychology* 102.3 (2012): 445.

Hoffman, Elizabeth, et al. "Preferences, property rights, and anonymity in bargaining games." *Games and Economic Behavior* 7.3 (1994): 346-380.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. "Fairness and the assumptions of economics." *Journal of business* (1986): S285-S300.

Mason, Winter, and Siddharth Suri. "Conducting behavioral research on Amazon's Mechanical Turk." *Behavior research methods* 44.1 (2012): 1-23.

Paolacci, Gabriele, Jesse Chandler, and Panagiotis G. Ipeirotis. "Running experiments on amazon mechanical turk." *Judgment and Decision making* 5.5 (2010): 411-419.

Rabin, Matthew. "Incorporating fairness into game theory and economics." *The American economic review* (1993): 1281-1302.

Raihani, Nichola J., Ruth Mace, and Shakti Lamba. "The effect of $1, $5 and $10 stakes in an online dictator game." *PloS one* 8.8 (2013): e73131.

Reips, Ulf-Dietrich. "The Web experiment method: Advantages, disadvantages, and solutions." *Psychological experiments on the Internet* (2000): 89-117.

Reips, Ulf-Dietrich. "Standards for Internet-based experimenting." *Experimental psychology* 49.4 (2002): 243.

Appendix A

Example Consent Form

Participants first see this screen online before starting the survey:

This survey is a part of Duke research. Participation is voluntary and we expect it will take about three to five minutes to complete. You can stop at any time. You will be asked to play a decision game, and after, there will be questions about yourself. Confidentiality and anonymity are assured. If you have any questions about this research, please click here.

If you understand and wish to proceed, click "Proceed".

After clicking proceed, the survey will begin.

Appendix B

Surveys

The following are the three Qualtrics surveys that the Qualtrics randomizer will select from when participants start the survey:

Study 1:

After seeing the consent screen from Appendix A, participants will come to these instructions:

In this survey, you will be randomly paired with another person who chose to take this survey. There are two roles in this survey: someone who decides how a payoff will be made for both people and someone who only receives whatever the other person chooses. There are two choices that you can make for your payoff, and for the person paired with you. But, you will decide by assigning probabilities. You will see a tool which lets you assign a probability to each choice on the next page.

Additionally, every 1 in 15 participants will receive $10 as a bonus in addition to payment for this survey, so play as though you are dealing with this $10.

So, if you choose 50% for one choice and 50% for the other, then there will be a 50% chance the computer picks the first choice, and a 50% the computer picks the second. Similarly, if you choose 40% for the first choice and 60% for the second choice, there will be a 40% chance the computer picks the first choice, and a 60% chance the computer picks the second. You will have the choice to pick values ranging from 0% (no chance of one choice being picked by the computer) to 100% (the computer will only pick that choice).

You will be brought to the two choices, and after one minute, you will be able to assign probabilities to each choice (either A or B) based on the available options. When you are ready, click "Next".

❍ Next

After clicking "Next", participants will wait one minute at this screen:

Remember: 1 in 15 participants in this survey will get this payoff, so that for whichever option you choose, you and the another person will get the specified amount in these two options. After one minute, you will be able to choose how you would like to distribute the probabilities for each option.

Choice A: You receive $9 and someone else receives $1 Choice B: You receive $5 and someone else receives $5

100% chance Choice A is picked:0% chance Choice B is picked
80% chance Choice A is picked:20% chance Choice B is picked
60% chance Choice A is picked:40% chance Choice B is picked
50% chance Choice A is picked:50% chance Choice B is picked
40% chance Choice A is picked:60% chance Choice B is picked
20% chance Choice A is picked:80% chance Choice B is picked
0% chance Choice A is picked:100% chance Choice B is picked

After the minute elapses, participants then can decide:

Remember: 1 in 15 participants in this survey will get this payoff, so that for whichever option you choose, you and the another person will get the specified amount in these two options.

Choice A: You receive $9 and someone else receives $1

Choice B: You receive $5 and someone else receives $5

Please choose one of the options listed below:

- 100% chance option A is picked & 0% chance option B is picked
- 80% chance option A is picked & 20% chance option B is picked
- 60% chance option A is picked & 40% chance option B is picked
- 50% chance option A is picked & 50% chance option B is picked
- 40% chance option A is picked & 60% chance option B is picked
- 20% chance option A is picked & 80% chance option B is picked
- 0% chance option A is picked & 100% chance option B is picked

After answering, participants will be brought to two different questionnaires (see Appendix C).

<u>Study 2:</u>

The only changes in Study 1 from Study 2 are slight changes in the instructions, and a slight change in the choice participants make. The instructions in Study 2 are:

In this survey, you will be randomly paired with another person who chose to take this survey. There are two roles in this survey: someone who decides how a payoff will be made for both people and someone who only receives whatever the other person chooses. There are two choices that you can make for your payoff, and for the person paired with you.

You will be brought to the two choices, and after one minute, you will be able to decide which choice (either A or B) you choose. From there, you will have 10 seconds to make your choice, and at some point during the 10 seconds the computer will randomly choose between the two options if you have not already done so. If this happens, the computer will choose evenly and randomly between Choice A and B. When you are ready, click "Proceed".

❍ Proceed

The question changes from being probabilistic to being deterministic:

Remember: 1 in 15 participants in this survey will get this payoff, so that for whichever option you choose, you and the another person will get the specified amount in these two options.

Choice A: You receive $9 and someone else receives $1

Choice B: You receive $5 and someone else receives $5

Please choose one of the options listed below:

- ❍ Choice A
- ❍ Choice B

After they are finished, participants then take the two short questionnaires (in Appendix C) as in Study 1.

<u>Study 3:</u>

The only changes in Study 3 are the in the directions. The directions in Study 3 are a hybrid of Study 2 and Study 1. Aside from that, the only other difference from Study 2 is that Study 3's choices look the same as Study 1's.

In this survey, you will be randomly paired with another person who chose to take this survey. There are two roles in this survey: someone who decides how a payoff will be made for both people and someone who only receives whatever the other person chooses. There are two choices that you can make for your payoff, and for the person paired with you. But, you will decide by assigning probabilities. You will see a tool which lets you assign a probability to each choice on the next page.

Additionally, every 1 in 15 participants will receive $10 as a bonus in addition to payment for this survey, so play as though you are dealing with this $10.

So, if you choose 50% for one choice and 50% for the other, then there will be a 50% chance the computer picks the first choice, and a 50% the computer picks the second. Similarly, if you choose 40% for the first choice and 60% for the second choice, there will be a 40% chance the computer picks the first choice, and a 60% chance the computer picks the second. You will have the choice to pick values ranging from 0% (no chance of one choice being picked by the computer) to 100% (the computer will only pick that choice).

You will be brought to two choices, and after one minute, you will be able to decide what probabilities you want to assign to each choice (either A or B). From there, you will have 10 seconds to make your choice, and at some point during the 10 seconds the computer will randomly choose between the two options if you have not already done so. If this happens, the computer will choose evenly and randomly between Choice A and B. When you are ready, click "Next".

❍ Next

This is what Study 3's question will look like (which they will be able to answer after one minute):

Remember: 1 in 15 participants in this survey will get this payoff, so that for whichever option you choose, you and the another person will get the specified amount in these two options.

Choice A: You receive $6 and someone else receives $1 Choice B: You receive $5 and someone else receives $5

Please choose one of the options listed below:

❍ 100% chance option A is picked & 0% chance option B is picked
❍ 80% chance option A is picked & 20% chance option B is picked
❍ 60% chance option A is picked & 40% chance option B is picked
❍ 50% chance option A is picked & 50% chance option B is picked
❍ 40% chance option A is picked & 60% chance option B is picked
❍ 20% chance option A is picked & 80% chance option B is picked
❍ 0% chance option A is picked & 100% chance option B is picked

As in Studies 1 and 2, there will be two questionnaires afterward.

Appendix C

Questionnaires

After the payoff decisions in Studies 1, 2, and 3, the participants will see these questionnaires. The first 8 questions aim to measure the participant's numeracy, and the list of adjectives with check-boxes aim to get a sense of the participant's creativity. Either before or after these questions, demographic information is collected, and the survey is through.

For each of the following questions, please check the box that best reflects how good you are at doing the following things:

	1	2	3	4	5	6
How good are you at working with fractions?	❍	❍	❍	❍	❍	❍
How good are you at working with percentages?	❍	❍	❍	❍	❍	❍
How good are you at calculating a 15% tip?	❍	❍	❍	❍	❍	❍
How good are you at figuring out how much a shirt will cost if it is 25%?	❍	❍	❍	❍	❍	❍

For each of the following questions, please check the box that best reflects your answer:

	1	2	3	4	5	6
When reading the newspaper, how helpful do you find tables and graphs that are parts of a story?	❍	❍	❍	❍	❍	❍
When people tell you the chance of something happening, do you prefer that they use words “it rarely happens”) or numbers “there’s a 1% chance”)?	❍	❍	❍	❍	❍	❍
When you hear a weather forecast, do you prefer predictions using percentages (e.g. “there will be a 20% chance of rain today”) or predictions using only words (e.g., “there is a small chance of rain today”)?	❍	❍	❍	❍	❍	❍
How often do you find numerical information to be useful?	❍	❍	❍	❍	❍	❍

Please indicate which of the following adjective best describe yourself.   Check all that apply.

- ❑ Capable
- ❑ Honest
- ❑ Artificial
- ❑ Clever
- ❑ Intelligent
- ❑ Well-mannered
- ❑ Cautious
- ❑ Wide interests
- ❑ Confident
- ❑ Inventive
- ❑ Egotistical
- ❑ Original
- ❑ Commonplace
- ❑ Narrow Interests
- ❑ Humorous
- ❑ Reflective
- ❑ Conservative
- ❑ Sincere
- ❑ Individualistic
- ❑ Resourceful
- ❑ Conventional
- ❑ Self-confident
- ❑ Informal
- ❑ Sexy
- ❑ Dissatisfied
- ❑ Submissive
- ❑ Insightful
- ❑ Snobbish
- ❑ Suspicious
- ❑ Unconventional