**Title**
Essays in Access-Based Beliefs



**Author**
Alladi, Vinayak

**Publication Date**
2019

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

# Essays in Access-Based Beliefs

A dissertation submitted in partial satisfaction of the
requirements for the degree of Doctor of Philosophy

in

Economics

by

Vinayak Alladi

Committee in charge:

Professor Charles Sprenger, Chair
Professor James Andreoni, Co-Chair
Professor Michael Callen
Professor Uri Gneezy
Professor Karthik Muralidharan
Professor Paul Niehaus

2019

The Dissertation of Vinayak Alladi is approved and it is acceptable in quality and form for publication on microfilm and electronically:

Co-Chair

Chair

University of California San Diego
2019

DEDICATION

To my parents and maternal grand parents who supported and loved me unconditionally in the toughest journey of my life. To my teachers who never stopped believing in me. To my idols who give me resilience when it matters most. To my loved ones who fill me with joy everyday.

EPIGRAPH

I find I'm so excited that I can barely sit still or hold a thought in my head. I think it's the excitement only a free man can feel. A free man at a start of a long journey whose conclusion is uncertain. I hope I can make it across the border. I hope to see my friend and shake his hand. I hope the Pacific is as blue as it has been in my dreams. I hope.

*Red, Shawshank Redemption*

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

This dissertation could not have happened without the countless positivity, encouragement, insightful suggestions and ideas, financial and emotional support, small and large acts of kindness, of so many individuals. I will attempt to thank as many as I can, but will no doubt leave out many deserving people: I'd firstly like to thank Charlie Sprenger for believing in me and supporting me relentlessly in this journey and whose commitment and dedication to his work and his students is a constant source of inspiration. I would like to thank James Anredoni, Prashant Bharadwaj, Mike Callen, Simone Galperti, Uri Gneezy, Mark Machina, Karthik Muralidharan, Paul Niehaus and Joel Sobel all of whom challenged me in different ways to become better, were patient and understanding throughout adversity, and provided me with knowledge and guidance that has shaped the course of my academic career. I no doubt am leaving out specific mention of contributions from other incredibly patient, insightful and wonderful faculty at UCSD. I would also like to thank Karthik Muralidharan especially here for his contagious passion for economics, for being a model of work ethic for me, and for helping me find a path to graduate school. Matthew Goldman helped me tremendously to understand difficult concepts throughout and was a constant source of encouragement and a great friend. The daily grind of the PhD requires the emotional support of loved ones and well-wishers, family and friends. In this regard, I thank my mother and father, Vani and Krishnaswamy Alladi, as well as my sister, Meghna Alladi and my brother-in-law, Tejus Menon, who were always there. Miles Berg, Andy Brownback, Paul Feldman, Sieuwerd Gaastra, Kilian Heilmann, Alex Kellogg, Claudio Labanca, Yucheng Liang, Ce Liu, Naglakshmi Nyayapathi, Kalpakam Ramesh, Pablo Ruiz Junco, Shailendra Singh, Ketki Sheth, Yuehui Wang, Wei You, Xu Zhang all deserve special mention for their roles in getting me here. Roy Allen, Richard Brady, Shreyans Bhansali, Quihui Chen, David Coyne, Naveen Bhasavanhally, Erin Giffin, Zheng Huang, Alex Imas, Michael Kaiser, Remy Levin, Shanthi Manian, Abhilash Mudaliar, Alejandro Nakab, Arman Rezaee, Ling Shao, Asha Subrahmanyam, Koji Takahashi were always there. Thank you to my extended family, cousins, aunts and uncles who helped push me forward. Venkatesh Periappa, Krishnamurthy Tata, Shantha Periamma and Siddhu Chitappa for encouraging me, believing in me, and leading by example. Jayashri Atta and Yogi Atta for helping me access funds so I could run my experiments. My family in the bay area who were a constant source of support. Thank you to the members of euro econ et. al, you know who you are, the teammates on Cousins Cousins, the friends on the tennis court, officemates of econ 112 and 122. My research assistants, in particular, Michelle Chau and Dun Lin Fletcher. I thank the National Science Foundation and Russell Sage Foundation for funding for this dissertation, as well as Christ University Bangalore for allowing me to gather data. Seminar participants in and outside of UCSD, as well as visitors to UCSD all listened to me intently and gave me wonderful comments on the work. I would be remiss to not mention the support of the UCSD Economics graduate administration, and the kindness of many individuals, including those outside of the university, cashiers, custodians, doctors, the people at whole foods, uber drivers, valet parkers, who made my life in San Diego more comfortable, adjusted to my quirks, and helped provide me with an environment in which I could work. Last but not least, friends, family and folks before graduate school who helped me get here and made me who I am.

Chapter 1 and 2 in full are being combined and prepared for submission for publication of the material. The dissertation author, Vinayak Alladi was the primary investigator and author of this material.

Chapter 3, in full is currently being prepared for submission for publication of the material. The dissertation author, Vinayak Alladi was the primary investigator and author of this material.

VITA

| | |
|---|---|
| 2004 | Bachelor of Arts, Economics, Northwestern |
| 2007 | Master of Science, Economics, London School of Economics and Political Science |
| 2019 | Doctor of Philosophy, Economics, University of California, San Diego |

ABSTRACT OF THE DISSERTATION

Essays in Access-Based Beliefs

by

Vinayak Alladi

Doctor of Philosophy in Economics

University of California San Diego, 2019

Professor Charles Sprenger, Chair
Professor James Andreoni, Co-Chair

This dissertation studies a notion called access-based beliefs, the idea that access to an alternative has a direct impact on beliefs of its value. As an example, individual's may lower their beliefs of alternatives to which they have low access, a phenomenon colloquially known as sour grapes. Using multiple settings in the lab and field, I demonstrate existence of access-based beliefs consistent with the sour grapes hypothesis and provide evidence for how it may be prevalent in real-world settings with an emphasis on the poor. In chapter 1, in an environment where learning channels are controlled for, I provide a proof-of-concept of access-based beliefs, showing that lower access to a subjective lottery reduces revealed

beliefs of its value. In chapter 2, I study access-based beliefs in a real-world setting. I show that lowering access to a real-world investment (an SAT prep package) lowers beliefs of its expected effectiveness and reduces investment in it amongst a subsample of the population that would be expected to have the highest demand for it. In chapter 3, I explore the notion of relative access, that is, the value of an alternative is a function of whether or not others possess it. I find that a small but significant portion of individuals display a preference for conformity or non-conformity in consumption based on the societal distribution.

# Chapter 1

# Access-Based Beliefs: Theory and Lab Evidence

## 1.1 Introduction

Standard economic theory has traditionally held that beliefs are only a function of information (Aumann, 1976; Harsanyi, 1968) and that beliefs about an alternative's value are independent of the probability of it being in the choice set (Savage, 1954). The latter condition, i.e. separability, is a core assumption of the canonical model of decision-making under uncertainty, expected utility (Von Neumann and Morgenstern, 1944). If one defines access generally and parsimoniously as the probability that an alternative will be in the choice set,[1] then the above assumptions imply that any relationship between access and beliefs can exist only indirectly, through learning.[2]

In contrast, a growing theoretical literature on motivated cognition challenges the traditional view, arguing that beliefs can be self-serving (Bénabou and Tirole, 2016) and that beliefs of a parameter may depend on how likely it is to be payoff relevant (Bénabou and Tirole, 2011). Existing evidence to suggest a relationship between access and beliefs comes in the form of experimental work from outside economics which finds

[1] Access could be defined in other ways, e.g., as the price of an alternative. However, this may lead to additional complications like wealth effects.

[2] Increased access could lead to more experience with an alternative, or could facilitate Bayesian updating about its value.

that the desirability of an alternative may be a function of its accessibility. However, these studies fail to adequately control for learning (Kay, Jimenez, and Jost, 2002) and use un-incentivized desirability ratings to study an alternative's value (Pyszczynski, 1982).

This paper presents the first deliberate exploration in economics of the relationship between access and beliefs. It defines access-based beliefs using decision-theory which helps to formalize phenomena described colloquially as "sour grapes" and "grass is greener". It improves upon prior empirical work from other disciplines by creating a decision environment in the lab that precludes learning, and by adopting a belief elicitation method that is based on incentivized choices. It then shows that the access-based beliefs empirically demonstrated in this paper are predicted by a special class of models in motivated cognition that combine anticipatory utility and reference dependence.

Access-based beliefs is defined as beliefs of an alternative's value that depend on the probability it will be in the choice set. As an example, individuals may avoid disappointment by lowering beliefs of alternatives to which they have low access, a phenomenon commonly known as "sour grapes". Such a mechanism could have significant economic effects. For instance, Anyon (1997) uses in-depth case-studies to argue that poor people have pessimistic outlooks on the value of schooling due to the difficulty of gaining access, a mechanism that ultimately impacts educational outcomes. Kay, Jimenez, and Jost (2002) find that lowering the perceived probability of a candidate winning an election reduces desirability of that outcome. In contrast, studies in marketing show that increasing the perceived inaccessibility of products can increase their desirability and demand (Lynn, 1989; Verhallen, 1982), which can be termed as "the grass-is-greener effect".[3]

In a lab setting that precludes learning, I consider a between-subjects design in which participants are randomly assigned a low (10%) or high (70%) chance of obtaining a subjective lottery with an unknown distribution over two outcomes. Before uncertainty is

[3]Conventionally, sour grapes and grass-is-greener are considered ex-post phenomena. I explore the ex-ante version, before uncertainty resolves, so that it may influence payoff-relevant decisions. Readers can therefore also interpret the theory as *anticipatory* sour grapes (grass-is-greener).

realized, subjects are asked to make a series of contingent choices between the subjective lottery and a collection of objective lotteries (with known distributions) of increasing expected value. The indifference point measures the participant's valuation of the subjective lottery, and under a subjective expected utility model (SEU), I infer the beliefs of the subjective lottery.[4] The null hypothesis, from expected utility and other models of ambiguity, is that there should be no difference in the average implied beliefs between the low and high access conditions. The alternate hypothesis is access-based beliefs, which predicts a difference. The lab experiment was conducted at two locations, UC San Diego and Christ College in Bangalore, India.

The results of the lab experiment are consistent with the existence of access-based beliefs in the direction of sour grapes. Lowering access to the subjective lottery by 60 p.p. reduces mean beliefs of a high payout by 2.6 p.p. and is significant at the 10% level. This represents an approximately 0.2 standard deviation shift in mean beliefs between the high and low conditions, and translates to a 3.3% expected reduction in dollar earnings calculated using the average baseline beliefs of 50% observed in the data. An examination of the distribution of beliefs for high and low access conditions shows that most of the divergence occurs for implied beliefs of the high payout that are coherent under SEU, meaning they are at least 50%.[5] When I restrict the analysis to this relevant subsample, which comprises 75% of the data, I find a slightly larger and more significant treatment effect - a 60 p.p. reduction in access lowers beliefs by 3 p.p. on average and is significant at conventional levels (p-values range from 0.024 to 0.096 depending on specification).[6]

To provide a theory for these findings, I propose an adaptation of the Brunnermeier

[4]More precisely, the indifference point is the subject's probability equivalent of the subjective lottery which allows me, under subjective expected utility, to calculate their implied belief (Andreoni, T. Schmidt, and Sprenger, 2015).

[5]Under subjective expected utility (SEU), being allowed to choose the high-payout state implies that beliefs less than 50% are incoherent, i.e., they do not sum to one. Under non-expected utility models (Klibanoff, Marinacci, and Mukerji, 2005; Schmeidler and Gilboa, 2004), individuals may have coherent beliefs but are averse to subjective uncertainty (ambiguity) and may inherently dislike the alternative.

[6]The Results section shows that there is no differential selection into this subsample by treatment.

and Parker (BP) (2005) optimal-beliefs model to subjective lotteries. In BP, individuals choose beliefs to maximize the tradeoff between anticipation of pleasant future outcomes and the costs of self-delusion, resulting in beliefs of high-utility states that are overly optimistic. I add another source of utility, reference dependence (where the reference point is the expected outcome under the chosen beliefs), which allows individuals to modify their beliefs to avoid disappointment. This addition creates a tradeoff between anticipatory utility and reference dependence which can account for both optimistic and pessimistic beliefs of the subjective lottery as access to it changes.[7] The model predicts a sour-grapes effect under parameter assumptions consistent with past experimental work. It can also predict grass-is-greener effects under more extreme assumptions, which the standard BP model cannot do.

This paper contributes to several literatures. First, a key property of models in motivated cognition, including BP, and other models of anticipatory utility, is that beliefs about a parameter should depend on how likely it is to be payoff-relevant. Coutts (2015) shows in the lab that lowering the payoff gained from an event lowers beliefs of the event's likelihood. I show lab evidence of the converse, i.e., lowering the likelihood of an alternative lowers beliefs of its payout, and demonstrate theoretically that such sour grapes is an implication of these canonical models. As such, access-based beliefs can be viewed as a consequence of optimal belief models which have the above-mentioned property.

By studying access-based beliefs as a two-stage lottery with a subjective lottery as an outcome, this paper serves as one of the first tests between models of subjective uncertainty (ambiguity) and models of motivated cognition. The former maintain that beliefs of the subjective lottery do not change unless new information arises and are agnostic about where priors come from (Savage, 1954; Schmeidler and Gilboa, 2004). They also often predict an aversion to subjective lotteries, which could increase with a higher

[7]This tradeoff is also explored in Sarver (2012) which models optimally chosen reference points. I discuss similarities and differences between the two in the last section of the paper.

likelihood of access (Klibanoff, Marinacci, and Mukerji, 2005). By contrast, my findings suggest instead that beliefs may change without new information, and that implied beliefs are increasing with greater access to the subjective lottery in a potentially self-serving way. Ultimately, they suggest that models of ambiguity may be missing elements of motivated cognition that might matter for decision-making under uncertainty.

Access-based beliefs relates to the endowment effect, the finding that individuals place a higher value on goods they possess (Kahneman, Knetsch, and R. H. Thaler, 1991; R. Thaler, 1980) or goods they are likely to possess (Ericson and Fuster, 2011; Heffetz and List, 2014). Similarly, this paper explores how varying the likelihood that a subjective lottery pays out impacts its valuation. An obvious question is whether or not the lab portion of this paper is simply a special case of past endowment effect experiments. However, this is not the case; I show that the standard model used to explain the endowment effect, i.e. reference dependence, when applied to this study, does not predict an effect.[8] While a higher chance of the subjective lottery (and therefore, of the high payout) raises the reference point, choosing the subjective lottery more in the elicitation will not increase the likelihood of the high payout since the outcomes in the subjective and objective lotteries are identical.[9]

The rest of this paper is organized as follows: In Section 1.2, I discuss the experimental design and in Section 1.3, I present the theoretical predictions from standard models; in Section 1.4, I present the results from the lab experiment and in Section 1.5, I propose a model of optimal beliefs that could rationalize access-based beliefs. In Section 1.6, I discuss the overall results of the paper and conclude with a discussion of the connection to poverty and potential policy implications of access-based beliefs.

[8] The standard explanation of the endowment effect comes from reference dependence where outcomes are treated as gains or losses with respect to a reference point. In this case, losing an object that one is (likely to be) endowed with hurts more than gaining an object that one is not (likely to be) endowed with.

[9] Only when the outcomes in the subjective and objective lotteries are different can reference dependence, specifically, Koszegi-Rabin (KR) (2006) reference dependence, predict the effect found in the lab.

## 1.2 Experimental Design

### 1.2.1 Framework for Laboratory Exploration

I model access-based beliefs as a two-stage compound lottery. The first stage of the lottery is objective and represents access to an alternative while the second stage is subjective, and allows for beliefs over the value of an alternative.[10] This paper focuses on *ex-ante* beliefs before uncertain access to alternatives is resolved. Alternately, one could explore *ex-post* beliefs, after uncertainty is resolved. However, if outcomes have been determined and there are no longer any payoff relevant choices to be made then ex-post beliefs arguably matter less except potentially for their impact on affective or psychological states.

The primitives are a set of outcomes, $\{H, L\}$ where $H > L$, a known probability distribution in the first-stage, $\{\alpha, 1-\alpha\}$, and the agent's beliefs over the subjective probability distribution in the second-stage, $\{\tilde{p}, 1-\tilde{p}\}$. The first stage is an $\alpha$ chance of a subjective lottery and a $(1-\alpha)$ chance of a fixed outside option, $F$ (Figure 1.1). Conditional on realizing the subjective lottery, there is a $\tilde{p}$ chance of a high paying outcome, $H$ and a $(1-\tilde{p})$ chance of low paying option, $L$.

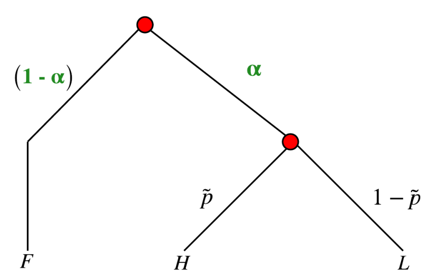


**Figure 1.1.** Access-Based Beliefs as a Compound Lottery

A sour grapes effect is when a decrease in $\alpha$ leads individuals to become less

[10] A subjective lottery is one in which the possible outcomes are known but the distribution of outcomes remains unknown.

optimistic about the subjective lottery and lowers their beliefs of $\tilde{p}$. The grass-is-greener effect would be the opposite response. As $\alpha$ decreases, individuals raise their beliefs of $\tilde{p}$ and the subjective lottery becomes more desirable.

### 1.2.2 Implementing Access-Based Beliefs in the Lab

#### Design Overview

The experimental design closely resembles the framework for access-based beliefs outlined above. In a between-subjects design, the probability of obtaining the subjective lottery is varied and a participant's implied belief about its composition is elicited before uncertainty is resolved but after they know their probability of obtaining it. The elicitation measure used is the probability equivalent of the subjective lottery - a method borrowed from Andreoni, T. Schmidt, and Sprenger (2015). With some probability (high or low) subjects gain access to subjective lottery, with some probability they face an outside option, and with some probability (fixed for both treatment groups) they are paid based on one of the ex-ante choices they made.[11]

The details of the experiment are explained as follows: (1) how the subjective lottery is represented and how the probability it obtains is varied (2) the choices subjects make before the lottery is resolved and how beliefs are elicited from those choices (3) how the choices are incentivized.

#### The Subjective Lottery

The subjective lottery is represented by an opaque JAR Beferred to as "JAR A". "JAR A" contains 20 red or green balls of unknown proportions. Subjects designate which color pays \$H and which pays \$L.[12]

[11] The choices can be thought of as contingent or conditional choices meaning, with some probability, they matter in the decision tree.

[12] Subjects make this choice to avoid any suspicion that the experiment is rigged against them as they might believe the jar is purposefully filled with a low number of winning balls.

### Implementing Uncertain Access to the Lottery

Access to the subjective lottery is determined by a first-stage objective lottery in which subjects select 1 of N sealed envelopes, each containing a single card that determines access, and ultimately, the subject payments. The distribution of the envelope's contents is known to subjects:

1. $\alpha \cdot N$ of the N envelopes contain a card that says **"JAR A' '** $\Longrightarrow$ Drawing such a card means that a participant's payout will be determined by a random draw of a ball from the subjective jar. Depending on the color of the drawn ball and the color they chose to be the winning color, they will be paid.[13]

2. $(1-\alpha) \cdot N$ of the N envelopes contain a card that says **"F' '** $\Longrightarrow$ Drawing such a card means that subjects get a fixed payout, $F$.

### Eliciting Beliefs

Beliefs of the subjective lottery are elicited through incentivized choices *before* the lottery is carried out but *after* the participant knows the chances she faces.[14] Importantly, the subjective lottery (JAR A) over which subjects have their beliefs elicited is the same subjective lottery (JAR A) to which access is varied. Specifically, each participant's probability (uncertainty) equivalent of the subjective lottery is measured. Using a multiple price list (MPL),[15] I find out the probability of $H that would make them indifferent to the subjective lottery as described below.

The left-side is a subjective lottery which was described in the previous section, "JAR A". The right side is a 20-ball jar of known proportions (the objective prospect)

[13] For example, if green is chosen as the winning color and red is drawn, the participant gets the low payment in the subjective lottery.

[14] The belief elicitation is incentivized by assigning a small probability that a participant's payout from the experiment is determined by a randomly chosen row of the MPL, following the convention of random incentive mechanisms.

[15] The multiple price list with real payments in economics was motivated and popularized by Coller and Williams, 1999, and Harrison, Lau, and Williams, 2002.

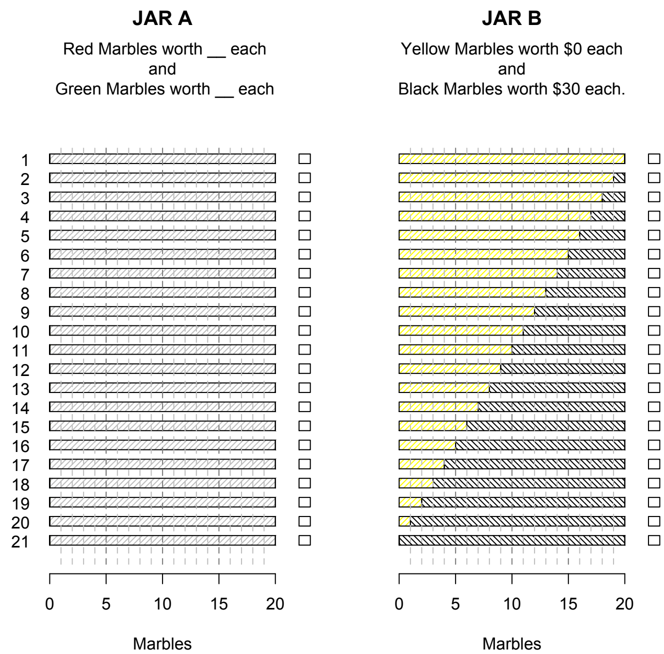


**Figure 1.2.** Multiple Price List over Subjective and Objective Lottery

containing yellow or black balls worth \$H and $\$L_0$ respectively, referred to as "JAR B". Proportions vary from 20 yellow and 0 black to 0 yellow and 20 black and for each row, subjects indicate whether they prefer a draw from the subjective or objective jar by checking the appropriate box. The point subjects switch from the left side(subjective lottery) to the right side (objective lottery) reveals their indifference point between the two options.[16]

To infer beliefs from these choices, I apply Subjective Expected Utility (SEU), the canonical model of decision-making under uncertainty in the literature. SEU states that individuals evaluate a subjective lottery according to the weighted sum of the utility of each outcome in the lottery, where the weights are the subjective probability that the outcome occurs.[17] Under SEU, the indifference point reveals to us the implied beliefs of

[16]Because the MPL is discrete and one of the lotteries must be selected as the preferred lottery, I can only identify their indifference point over an interval.

[17]That is, individuals apply expected utility to subjective lotteries, and are probabilistically sophisticated, meaning they have subjective probabilities of the outcomes that sum to one (see section 3 for details of SEU).

the subjective lottery. Let $\tilde{p}$ be the subjective probability of the number of high paying balls in the subjective lottery, $q$ be the probability of a high paying ball in the objective lottery, and $u$ be any well-behaved utility function. At the indifference point, the following equation holds:

$$\begin{aligned}\text{Expected Utility of the Subjective Lottery} &= \text{Expected Utility of the Objective Lottery}\\ \tilde{p}\cdot u(H)+(1-\tilde{p})\cdot u(L) &= q^*\cdot u(H)+(1-q^*)\cdot u(L_0)\\ \tilde{p} &= q^*\frac{u(H)-u(L_0)}{u(H)-u(L)}+\frac{u(L_0)-u(L)}{u(H)-u(L)} \qquad (1.1)\end{aligned}$$

where $q^*$ is the probability of the objective lottery at the indifference point. Notice that if $L = L_0$, then $\tilde{p} = q^*$, meaning beliefs can be directly calculated without knowing the shape of the function u(), i.e. without assumptions on risk aversion.

**Incentivizing Belief Elicitation Task**

In order to incentivize the choices that subjects make, there are two additional envelopes so that subjects are actually choosing from 1 out of (N + 2) envelopes as opposed to 1 out of N. These 2 envelopes contain a card that says ”Task”. Drawing such a card means that a random row from the belief elicitation task is chosen, and the participant is paid according to a drawn ball from the jar they preferred for that row.[18] Denote “$R$” as the probability a subject’s payout is determined by JAR A or the outside option, and “$(1-R)$” as the probability it is determined by the elicitation task. The complete decision tree subjects faced is illustrated in Figure 1.3.

[18]This is known as a random incentive mechanism - a method commonly used in the experimental literature and in experiments with subjective lotteries (Ahn, Choi, Gale, and Kariv, 2007; Halevy, 2007). Appendix **??** examines the incentive compatibility of the mechanism for this experiment.

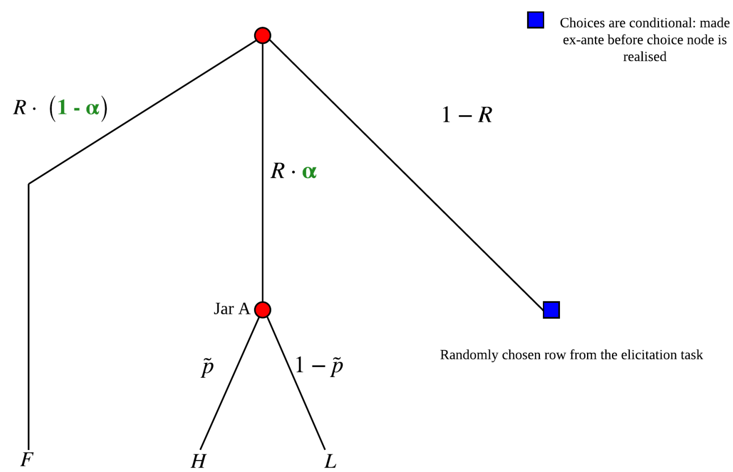


**Figure 1.3.** Complete Decision Tree

The distribution of envelopes is as follows:

- $R \cdot \alpha$ proportion of envelopes are "JAR A" $\implies$ a payment from the subjective jar.
- $R \cdot (1 - \alpha)$ proportion of envelopes are "\$F" $\implies$ a fixed payment of \$F
- $(1 - R)$ proportion of envelopes are "TASK" $\implies$ a payment from a randomly chosen row of the elicitation task according to a drawn ball from the preferred jar for that row.

**Final Design, Timing and Procedures**

The timing of the experiment is as follows: in Step 1, subjects choose, but do not open, one of 10 envelopes in front of them. In Step 2, subjects make their choices on the multiple-price list and in Step 3, envelopes are opened and payments are made based on the envelope contents. Importantly, subjects are aware of the full experimental procedure, including the distribution of envelope contents and how their payments will be determined, before they make their choices.

With this procedural setup, a between-subjects design with two access conditions was implemented as follows: [19]

1. **High Access Condition:** 7 envelopes say "JAR A", 1 envelope says "F" and 2 envelopes say "Task"

2. **Low Access Condition:** 1 envelope says "JAR A" , 7 envelopes say "F' and 2 envelopes say "Task"

A crucial feature was to make the chance that the elicitation pays out constant across both groups so that any differences in beliefs could not be explained by differential incentives to do the task or a differential probability of obtaining the objective lottery.[20]

Two design choices were made with the thought they would increase the psychological conditions under which access-based beliefs are more likely to be observed: (1) Subjects choose, but do not open, one of the envelopes at the start of the experiment so that it sits with them for a while. Subjects are also asked to write the letter "Y" on the envelope they chose, and the letter "N" on the envelopes they did not choose. (2) I gradually increase the attractiveness of the subjective lottery relative to the outside option and to the objective lottery under the prediction that doing so will increase the size of access-based beliefs. As a result, three versions of the experiment are tested in this paper:

1. Version 1 of our experiment has only two outcomes, H and L, where the fixed outside option, F, also equals L. This allows us to calculate beliefs without any assumptions on risk aversion.

[19] A within-subjects design was also considered where beliefs would be elicited twice for the same person at different access levels. This could increase statistical power by controlling for potentially large heterogeneity in beliefs, however, this could generate very little variation in beliefs, as individuals may not change their responses on an identical elicitation task. Furthermore, if responses did change it may also be hard to rule out miscomprehension of the task.

[20] An alternate design considered was varying the probability of access to the subjective lottery without an outside option. While perhaps simpler to explain, this would also change the incentives to do the elicitation and it would change access to the objective lottery in the elicitation as well, thus adding potential confounds.

**Table 1.1.** Versions of Experiment

| | JAR A (subjective lottery) | JAR B (objective lottery) | Fixed Payout |
|---|---|---|---|
| Version 1 | H and L | H and $L_0 = L$ | F = L |
| Version 2 | H and L | H and $L_0 = L$ | F ¡ L |
| Version 3 | H and L | H and $L_0 < L$ | F ¡ L |

2. Version 2 of the experiment maintains only two outcomes in the elicitation task (the MPL) but changes the outside option to a lower value, $L_0 < L$. This increases the attractiveness of JAR A relative to the outside option and still allows you to calculate implied beliefs without any more assumptions.

3. Version 3 of our experiment is the strongest version in terms of creating the feeling of low access to a desirable alternative. As in version 2, the outside option is lowered, but in addition, the low payment of the objective lottery is reduced to $L_0 < L$. This further increases the attractiveness of JAR A relative to the outside option as the outcomes in the objective lottery are now weakly dominated by the subjective one. However, this necessitates some assumptions on risk aversion discussed in the Results section.[21]

## 1.3 Theoretical Predictions from Standard Models

### 1.3.1 Subjective Expected Utility

The canonical model for decision-making under ambiguity (subjective lotteries)[22] is Subjective Expected Utility (SEU) (Anscombe and Aumann, 1963) which has the following primitives: A finite state space $S$, a set of consequences $X$, a set of objective probability

[21] For the analysis, pooled and separate results from all three versions are presented, and the hypothesis that increasing the attractiveness of the subjective lottery from version 1 to 3 generates a larger effect is tested.

[22] When referring to the experiment, I use the term "subjective lottery"; when referring to existing theory, I use either "ambiguity" or subjective lottery, but these terms are interchangeable.

distributions over consequences $\Delta(X)$, and a set of acts $F$, functions which map $S$ into $\Delta(X)$. Individuals have preferences over acts in $F$. If these preferences satisfy the axioms of completeness, transitivity, monotonicity, continuity, independence and non-degeneracy, they can be represented by a Subjective Expected Utility (SEU) functional. For a given act, $a$:

$$U_{SEU}(a) = \int_S p(s)u(a(s))ds \tag{1.2}$$

where $p(s)$ is the subjective probability associated with state $s \in S$, $\int_S p(s)ds = 1$ and $u(a(s))$ is the (von Neumann-Morgenstern) expected utility of the lottery that act $a$ yields in state $s$.

As can be seen in (2), SEU is linear in the probabilities (a consequence of the independence axiom)[23] which implies that a change in the probability of an outcome does not influence how the outcome is valued. In other words, for the compound lottery in Figure 1.1, a change in $\alpha$ does not impact the valuation of the subsequent subjective lottery. This means that access-based beliefs are inconsistent with SEU.

### 1.3.2 Non-SEU Ambiguity Models

There are other prominent models in the ambiguity literature which relax the independence axiom (Klibanoff, Marinacci, and Mukerji, 2005; Schmeidler and Gilboa, 2004). However, these models mostly focus on rationalizing the Ellsberg paradox, the finding that individuals have preferences which reflect an aversion to ambiguous lotteries. The Smooth-Ambiguity model of Klibanoff, Marinacci, and Mukerji (2005) predicts that as exposure to the ambiguous lottery increases the individual becomes more averse to ambiguity, a finding consistent with the grass-is-greener hypothesis. None of the above models would predict sour grapes unless changes to the underlying assumptions are made.

[23]Independence axiom: For any three acts $a, b, c \in F$ and for any $\alpha \in (0,1), a \succsim b \iff \alpha a + (1-\alpha)c \succsim \alpha b + (1-\alpha)c$

### 1.3.3 Reference Dependence

Reference dependence (Markowitz (1952) and Tversky and Kahneman (1991)) is a model in which outcomes are treated as gains or losses with respect to some reference point, with losses being felt more intensely than commensurate gains. The aversion to expected loss has been used to explain why individuals with a low expectation of receiving a consumption good lower their valuation of it (Ericson and Fuster, 2011; Heffetz and List, 2014). In comparison, access-based beliefs examine changes in implied beliefs of an alternative's value caused by changes in the likelihood of possession, where the alternative is a lottery as opposed to a consumption good. As such, the rationalization of loss aversion around a reference point cannot explain access-based beliefs because the identified channel is a change in beliefs (or preferences) exclusively over the subjective lottery and not tradeoffs between different sources of gain-loss utility.[24]

### 1.3.4 Motivated Cognition

Models in motivated cognition suggest that beliefs enter the utility function directly and can be chosen in a self-serving way. The Brunnermeier and Parker (BP) (2005) optimal beliefs model posits that individuals choose beliefs (probabilities) to maximize the tradeoff between anticipation of pleasant future outcomes (anticipatory utility) and the costs of self-delusion. BP predicts higher beliefs of an argument that is more likely to be payoff relevant. Adapted to subjective lotteries, BP would predict that a lower chance of a pleasant, subjective lottery would lead to lower (but still inflated) beliefs, thus going in

[24]Unless the model treats subjective lotteries as a consumption object (a pen or mug, for example), I demonstrate that no reference dependence model can fully explain access-based beliefs in the case of two-outcome lotteries (Appendix 1.7.3) I examine three cases of reference dependence (1) the reference is fixed at some value $r > 0$ (2) the referent is the expected value of the gamble and (3) the referent is stochastic as in the case of Koszegi-Rabin (KR) (2006). I show that in all three cases the optimal report of the agent in the MPL is equal to her true beliefs of the subjective lottery and does not vary with access. The exception is that when there are three outcomes in the lottery, KR reference dependence could explain the behavior. Also, for extreme parameter values, such as unreasonably large loss aversion assumptions, reference dependence predicts that changes in access can move optimal beliefs, not marginally, but to the boundaries.

the direction of sour grapes.[25]

**Table 1.2.** Predictions of Different Models

| **Model** | **Prediction on choices if $\alpha$ changes** | **Reason for prediction** |
|---|---|---|
| Subjective Expected Utility | No Effect | Valuation is linear in probabilities |
| Gilboa-Schmeidler Max-Min | No Effect | Valuation is linear in probabilities |
| KMM Smooth ambiguity | Grass-is-greener | Individuals are ambiguity averse |
| Reference Dependence | No Effect | No incentives to change choice |
| Motivated Cognition | Sour Grapes | Individuals have anticipatory utility. Low access leads to lower, but still inflated beliefs |

## 1.4 Results

I first present results from the lab experiment (Section 1.4.1) , and then proceed to the field experiment (Section 1.4.2), and finally, to a discussion of the heterogeneity of the

### 1.4.1 Data Collection

The experiment was conducted in two locations, the University of California, San Diego (UCSD) and Christ University in Bangalore, India (CUB). Data collection procedures were the same in both places. Sessions were conducted on undergraduates in large universities from a variety of majors and recruiting procedures were identical. Out of 34 lab sessions, 18 sessions were conducted at UCSD and 12 sessions at CUB. Version 1 and 3 were conducted at UCSD while version 2 and 3 was conducted at CUB. I conduct

[25] In section 5, I show how adding reference dependence to the standard BP model allows beliefs to be optimally pessimistic. The model predicts a sour grapes effect under most reasonable parameter assumptions, which is consistent with my experimental findings, but is capable of predicting grass-is-greener under more extreme parameter assumptions, which the standard BP model cannot do.

pooled analysis for the main treatment effects including a fixed effect for location, and then separate analysis by location in the section on heterogeneity.[26]

### 1.4.2 Calculating Implied Beliefs

The outcome of interest for treatment effects analysis is the implied beliefs of the subjective lottery. Under SEU, the implied beliefs can be calculated from the indifference point on the elicitation task (the MPL). From eq (1) in section 2.2.4, when the researcher knows the exact indifference point, $q^* = \tilde{p}$. However, in practice, the MPL is discrete. A 1-row change in the indifference point represents a 5% difference in beliefs, so as an approximation, one can take the midpoint between two rows to measure beliefs.[27] This applies for versions 1 and 2 where both lotteries had the same outcomes.

In Version 3 of the experiment assumptions on risk aversion were made to calculate implied beliefs.[28] Without loss of generality, if $L_0 = 0$, as in version 3, and u(0) = 0, the belief to choice mapping is:

$$p = q^* \frac{u(H)}{u(H) - u(L)} + \frac{-u(L)}{u(H) - u(L)}$$

Notice that if one assumes greater risk aversion, i.e. a larger ratio of $\frac{u(H)}{u(H)-u(L)}$, this implies a larger change in the $p$ (beliefs) holding $q$ (the probability equivalent) constant. Therefore, to be conservative in our estimates, I assume risk neutrality throughout, which gives us a lower bound on the treatment effect.[29] In the Results section I examine both pooled and separate treatment effects by version.

[26]See appendix 1.7.4 for further comparisons between UCSD and India data.

[27]For example, switching one's preference from the ambiguous to risky prospect between the 10th and 11th rows of the task sheet is coded as .475, since the implied belief is between .45 and .5. Switching one's preference from ambiguous to risky prospect between the 12th and 13th rows is similarly coded as .575.

[28]Another option could be to do a separate elicitation task, but this may make the entire set of decisions more difficult to model.

[29]Appendix 1.7.5 provides further details of the mapping under different assumptions of risk aversion using a power utility function.

### Dealing with Multiple Switchers

Multiple switchers, defined as subjects who fail to indicate a unique switch point in the elicitation, is a general concern when using multiple price lists. In the US, the rate of multiple switching was 16.6%, a similar rate to other experiments that employ the same methodology.[30] However, in India, the rate was considerably higher at 51.5%, possibly because subjects had not been exposed to lab experiments before. As a result, to ensure comprehension, a comprehension questionnaire was administered before beginning the experiment for the Indian sample. For all treatment effects analysis, I remove multiple switchers, since it is impossible to calculate an implied belief for them.[31]

One unexpected outcome in the Indian data was that the multiple switching differs by treatment.[32] Given that the treatment was randomized and that the effect remains after controlling for observables, it suggests that this was an outcome of the treatment, rather than selection on unobservables. One possible mechanism is that subjects with low access concentrated on the instructions more, and this had an impact of reducing multiple switching in the Indian sample where comprehension was lower and experience with the task was less. Since there is no obvious theoretical link between higher concentration and the sour grapes effect, there seems no obvious reason to interpret the treatment effects any differently.

## 1.4.3 Sample Balance

No statistically significant differences are found between low and high access groups on a set of participant demographics. This is true for the sample of non-multiple switchers shown in the table below along with the full-sample of subjects, multiple switchers included,

[30]Such multiple switching is frequently found in multiple price list experiments and normally occurs for 10-15% of subjects (Holt and Laury, 2002; Meier and Sprenger, 2010).

[31]As a robustness check, I include them for analysis at the individual-choice level.

[32]In a regression of an indicator for multiple switching on treatment, I find a lower rate of multiple switching in the low access treatment of - 15.5 p.p. (p-value of .034) and after controlling for observables, this drops to -13 p.p. and become just marginally significant (p-value .091).

thus showing successful randomization.

**Table 1.3.** Sample Balance

| | (1) Female | (2) Took Uncertainty Class | (3) Age | (4) Year of College | (5) Quantitative Major |
|---|---|---|---|---|---|
| Low Access | 0.554 | 0.181 | 20.131 | 2.494 | 0.281 |
| | (0.040) | (0.031) | (0.162) | (0.089) | (0.036) |
| High Access | 0.596 | 0.160 | 19.994 | 2.292 | 0.337 |
| | (0.039) | (0.029) | (0.151) | (0.093) | (0.037) |
| Observations | 318 | 323 | 323 | 319 | 323 |
| P-value Low = High | 0.449 | 0.605 | 0.535 | 0.118 | 0.276 |

Robust Standard Errors in Parentheses. p<0.1*, p<0.05**, p<0.01***. Each row shows the mean for that demographic variable in the low and high access group respectively. The last row shows a t-test of the difference in means. I find no significant different in means between the two groups.

### 1.4.4 Main Treatment Effects

The table below summarizes the outcome variable, "beliefs", for the full sample of data. One cannot reject that mean beliefs of the high outcome are 50%.

**Table 1.4.** Summary of Dependent Variable: Beliefs

| | Mean | Standard Deviation | Min | Max |
|---|---|---|---|---|
| beliefs | .506 | .152 | .036 | .964 |

The sample consists of all subjects who indicated a unique switch-point in the multiple price list.

To provide a visual depiction of the treatment effects, I plot the distribution of beliefs for low and high access conditions below (Figure 1.4). The figure illustrates a shift in the distribution of the high access distribution to the right for beliefs greater than 50%, while below 50%, the distributions look quite similar. I discuss why 50% is a relevant partition of the data from a theoretical standpoint in the regression analysis below. Although the difference in distributions above 50% is striking, a two-sample KS test and a Wilcoxon rank-sum test fails to reject the equality of the overall distributions with p-values of 0.18 and 0.11 respectively, likely reflecting a lack of power.

I next estimate average treatment effects using the estimating equation below:

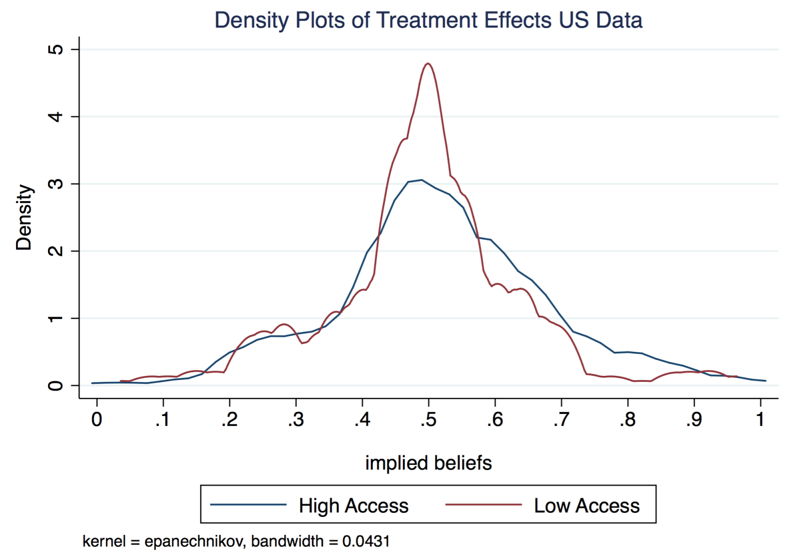


**Figure 1.4.** Kernel Density Estimates of Belief Distributions By Treatment

$$\text{Beliefs}_i = \alpha + \beta \cdot \text{Low Access}_i + \boldsymbol{\gamma} \cdot \boldsymbol{x}_i + \boldsymbol{\nu} \cdot \boldsymbol{z} + \epsilon_i \tag{1.3}$$

Beliefs$_i$ are the implied beliefs for individual i, Low Access$_i$ is a binary indicator for the treatment, the vector $\boldsymbol{x}$ are a set of demographic indicators and the vector $\boldsymbol{z}$ are a set of fixed effects for experiment version (versions 1 to 3) and study location (UCSD vs. India). The coefficient of interest is $\beta$.[33]

Table 1.5 presents the results of OLS estimation of equation (3). Column 1 shows the average treatment effects without any controls. I find that low access reduces beliefs by 2.6 p.p. on average from a baseline of 52 p.p. This represents a .16 standard deviation shift in the belief distribution, and is significant at the 10% level (p-value of .100). When

[33]In all our analysis outliers are dropped. Outliers are defined as subjects who, when offered a choice between the subjective lottery (with payouts of \$30 and \$10) and a certain payout of \$30, choose the subjective lottery. Similarly, I also consider as outliers subjects who, when offered a choice between the subjective lottery and a certain payout of \$10, prefer the certain payout. As a robustness check, I include outliers in the analysis and find that this does not impact our main results.

I control for participant demographics, experiment version, and study location (column 2) I find identical point estimates (2.6 p.p.) and levels of significance (p-value of .102). These results suggest that the probability of obtaining the subjective lottery does impact its valuation, and the effect goes in the direction of sour grapes (as opposed to the grass-is-greener). The size of the effects in terms of expected economic value (in dollars) lost is about $.78 on average, which translates to a 3.3% loss in expected winnings at mean beliefs of 50%.

Columns 3 to 6 look at the treatment effects for a sample of subjects whose beliefs were above and below 50%. This is the region in which the density plots of the high access and low access distributions seem to diverge the most. This is a relevant cutoff to examine because when subjects get to choose the winning color in the subjective lottery, beliefs below 50% represent either: (A) a violation of SEU and coherent beliefs, as subjects could have simply chosen the other color to be the winning color, and therefore, would have increased their expected earnings given those beliefs,[34] or (B) an aversion to ambiguity (subjective lotteries). Either way, incoherent beliefs or an aversion to the subjective lottery would be cases in which one may not expect to find access-based beliefs and, in particular, sour grapes. Importantly, the treatment has no impact on the proportion of subjects who violate SEU, which helps rule out selection on unobservables into this subsample.[35]

When treatment effects are analyzed for this subsample, similar effect sizes with more precision are found. Low access reduces beliefs by 3 p.p. (p-value of .053) representing a .2 standard deviation shift in the belief distribution (column 3) with similar results when I control for observables (column 4). Columns 5 and 6 examine treatment effects for beliefs below 50% without and with controls respectively and finds precise effect sizes of

[34] Another way of seeing the violation is that beliefs under 50% mean that subjects are probabilistically unsophisticated, in that beliefs over the states of the world do not add up to one.

[35] See the rows in the first bottom panel of table 1.5 - the number of observations in high and low access treatments below and above the cutoff shown is nearly equal; regressions in table 1.7.9, appendix 1.7.9 also confirm this result.

0.[36] Columns 7 to 10 examine treatment effects at a lower cutoff point, to ensure that some subjects who may have had beliefs of exactly 50% are not excluded, a possibility when using multiple price lists.[37] At this cutoff, the treatment effect increases to 3.1 p.p. (p-value of .024) (column 6) and is similar in magnitude and significance when controlling for observables (column 7). Below 46%, the treatment effect is precisely 0 (columns 8 and 9).

As a robustness check, differential rates of multiple switching between the treatment groups (17.9% lower multiple-switching in the low access group was driven by the Indian portion of the study) was dealt with in three ways: (1) Calculate treatment effects for the sample of subjects only at UCSD where there was no differential switching. (2) Construct the Lee bounds for the treatment effects taking into account multiple switching.[38] (3) Perform choice-level analysis of the treatment effects using the entire data, including multiple switchers, to look at the impact of treatment on the probability of choosing the subjective lottery. I find that the treatment effects are robust to (1) and (3), but given their size, it is unsurprising that they do not survive both sides of the Lee bounds (details in appendix 1.7.7).

As noted, a hypothesis was that increasing the relative attractiveness of the subjective lottery with respect to the outside option (version 1 to 2), and then with respect to the objective lottery (version 2 to 3), would increase the sour grapes effect. I test this hypothesis in columns 11 and 12 of table 1.5 under the same specification as equation (1.3) by adding interaction terms between treatment and version number. Furthermore, I find that increasing the relative attractiveness of the subjective lottery does increase the sour grapes effect. The interaction terms between treatment and version

[36] In table 1.7.9, appendix 1.7.9 I calculate treatment effects on the actual row number at which subjects switched rather than on beliefs for robustness. Overall, this has no impact on the main findings - the treatment effects either increase or decrease, depending on the specification.

[37] Subjects who preferred the 50-50 objective lottery to the subjective lottery were recorded as having beliefs of 47.5% (the midpoint between 45% and 50%.

[38] The Lee bounds are a trimming procedure to bound treatment effects when there is possible selection (Lee, 2009).

**Table 1.5.** Main Treatment Effects

| | Dependent Variable : Beliefs of a High Payout from Subjective Lottery | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | full distribution | | 50 percent partition | | | | 46 percent partition | | | | By Version | |
| | | | Above 50 | | Below 50 | | Above 46 | | Below 46 | | | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) | (12) |
| Low Access | -0.026 | -0.026 | -0.030* | -0.027* | -0.000 | 0.008 | -0.031** | -0.028** | -0.015 | 0.003 | | |
| | (0.016) | (0.016) | (0.016) | (0.016) | (0.015) | (0.011) | (0.013) | (0.013) | (0.019) | (0.014) | | |
| version 2 = 1 | | 0.072* | | 0.050 | | 0.038 | | 0.020 | | -0.004 | 0.013 | 0.074 |
| | | (0.039) | | (0.040) | | (0.028) | | (0.033) | | (0.029) | (0.042) | (0.055) |
| version 3 = 1 | | 0.045** | | 0.103*** | | -0.103*** | | 0.108*** | | -0.115*** | 0.034 | 0.056** |
| | | (0.019) | | (0.015) | | (0.017) | | (0.014) | | (0.019) | (0.022) | (0.025) |
| India = 1 | | -0.061* | | -0.010 | | -0.031 | | -0.007 | | -0.009 | | -0.060* |
| | | (0.032) | | (0.027) | | (0.025) | | (0.026) | | (0.022) | | (0.032) |
| Low Access × version 1 | | | | | | | | | | | -0.012 | -0.015 |
| | | | | | | | | | | | (0.012) | (0.013) |
| Low Access × version 2 | | | | | | | | | | | -0.024 | -0.021 |
| | | | | | | | | | | | (0.043) | (0.046) |
| Low Access × version 3 | | | | | | | | | | | -0.037 | -0.036 |
| | | | | | | | | | | | (0.028) | (0.028) |
| Constant | 0.519 | 0.525 | 0.632 | 0.558 | 0.395 | 0.437 | 0.587 | 0.528 | 0.335 | 0.370 | 0.500 | 0.519 |
| | (0.012) | (0.018) | (0.011) | (0.018) | (0.011) | (0.013) | (0.010) | (0.014) | (0.013) | (0.018) | (0.009) | (0.018) |
| Observations | 368 | 354 | 183 | 176 | 185 | 178 | 269 | 258 | 99 | 96 | 368 | 354 |
| Controls | No | Yes | No | Yes | Yes | Yes | No | Yes | Yes | Yes | No | Yes |
| R-squared | 0.007 | 0.065 | 0.021 | 0.245 | 0.000 | 0.369 | 0.019 | 0.265 | 0.006 | 0.403 | 0.013 | 0.066 |
| P-value: Low Access = 0 | 0.100 | 0.102 | 0.053 | 0.096 | 0.992 | 0.504 | 0.024 | 0.033 | 0.436 | 0.851 | 0.794 | 0.782 |
| Low Access Obs | 190 | 183 | 90 | 86 | 100 | 97 | 139 | 133 | 51 | 50 | | |
| High Access Obs | 178 | 171 | 93 | 90 | 85 | 81 | 130 | 125 | 48 | 46 | | |
| P-value: v1 = v2 = v3 = 0 | | | | | | | | | | | 0.392 | |
| P-value: v1 = v2 = v3 | | | | | | | | | | | 0.697 | 0.796 |
| P-value: v1 = v2 | | | | | | | | | | | 0.784 | 0.900 |
| P-value: v1 = v3 | | | | | | | | | | | 0.408 | 0.501 |
| P-value: v2 = v3 | | | | | | | | | | | 0.794 | 0.782 |
| Lower Lee bound | -0.058*** | | -0.062*** | | -0.020 | | -0.063*** | | -0.043* | | | |
| Lower Lee Bound p-value | 0.003 | | 0.000 | | 0.288 | | 0.000 | | 0.094 | | | |
| Upper Lee bound | 0.007 | | -0.018 | | 0.043 | | -0.017 | | 0.023 | | | |
| Upper Lee Bound p-value | 0.721 | | 0.330 | | 0.013 | | 0.290 | | 0.349 | | | |

Robust Standard Errors in Parentheses. p<0.1*, p<0.05**, p<0.01***. All regressions estimated using OLS. The dependent variable measures beliefs of the probability of a high payout from the subjective lottery. The treatment effect is indicated by the variable "Low Access" representing participants who had only a 10% chance of the subjective lottery. The "above 50%" partition (columns 3 and 4) represents individuals who have beliefs above 50%. The "above 46%" partition (columns 7 and 8) represents individuals who we cannot rule out have beliefs equal to or above 50%. Refer to text for additional details. The controls used include Gender, dummies for year of college, and whether the participant took decisions-under-uncertainty class. The results do not change much, with or without controls. v1, v2, v3, refer to the interaction terms of treatment with version 1,2 and 3 of the experiment respectively, and p-values report tests of their equality to each other and to zero. The lower and upper Lee bounds are constructed based on the differential rates of multiple-switching.

become increasingly negative, going from -1.5 p.p for version 1, -2.0 p.p. for version 2, and -3.9 p.p. for version 3 but we are underpowered to detect these effects. Notably, the interaction terms are always negative, though not individually significant, demonstrating the consistency of sour grapes evidence across different versions of the experiment.

Overall, the main finding from our treatment effect analysis is a consistent sour grapes effect in the range of 2.6% to 3.1% with p-values between .024 and .102 depending on specification. These findings are robust to (1) many replications of the experiment (over 30 sessions), (2) conducting the experiment in two different locations and (3) changes in the participant payouts (from versions 1 to 3).

## 1.5 Access-Based Beliefs as Optimal Beliefs

In this section, an existing model is adapted to explain access-based beliefs and capture observations from the lab. A candidate model needs to satisfy three criterion: (1) It should predict that access to an alternative impacts beliefs of its value (2) It should allow beliefs to be either optimistic or pessimistic relative to some true benchmark (3) It should predict sour-grapes under some parameter values and grass-is-greener under others since both directions are at least theoretically possible.

One candidate is the canonical model of optimal beliefs, Brunnermeier and Parker (2005) (BP), which is capable of explaining access-based beliefs. In BP, beliefs are optimally chosen to tradeoff two sources of utility: (1) anticipatory utility - the pleasure derived from higher beliefs about future expected consumption (2) expected consumption utility - utility derived from actual consumption. Importantly, optimal beliefs that deviate from the truth can lead to lower real consumption by influencing choices. The model predicts that beliefs of states in which consumption utility is high end up being inflated or optimistic. Stated otherwise, beliefs of a parameter depend on the likelihood that the parameter is payoff relevant. Applied to our setting, when the probability of the subjective lottery increases, an individual has a greater incentive to inflate beliefs to raise anticipatory utility, experiencing first-order gains from doing so, at the expense of second-order losses in expected outcomes. However, a caveat to BP is that it cannot account for optimal beliefs being lower than the truth (criterion 2) since anticipatory utility always pushes beliefs upwards and it also cannot account for a grass-is-greener effect (criterion 3) since the marginal benefit to increasing beliefs is always increasing in the level of access to the subjective lottery.

As a solution, I add reference dependence to standard BP, creating a tradeoff between two psychological sources of utility, anticipation and gain-loss. This allows for a richer set of predictions for how beliefs change with access. The relative strength of anticipatory utility and loss aversion helps determine whether: (A) optimal beliefs are lower or higher than true beliefs (criterion 2) (B) a sour grapes or grass-is-greener effect is observed as the chance to obtain the subjective lottery increases (criterion 3). The tension explored is similar to Sarver (2012) in which the individual optimally chooses an arbitrary reference point to maximize the tradeoff between anticipatory utility and gain-loss utility.[39]

### 1.5.1 Model

Consider a decision-maker (DM) who faces the following 2-stage decision problem. In stage 1, the DM chooses a set of optimal beliefs to maximize long-run wellbeing and in stage 2, the DM chooses actions to maximize expected consumption utility given the optimal stage 1 beliefs. For simplicity, the model assumes that choices are made, uncertainty is resolved, and all utility is experienced in one period.[40] Denote $p$ as the individual's true beliefs of a parameter, $\tilde{p}$ as optimally chosen beliefs of the same, and $c$ as

[39]However, our model differs from Sarver in two crucial ways: (1) Instead of choosing an arbitrary reference point, the reference point is dictated by the expected utility of the lottery under optimal beliefs (2) Unlike in Sarver, the individual in our model considers the costs of self-delusion in terms of foregone consumption.

[40]Access-based beliefs can be captured in a static setting. Adding a time dimension would not change not change our results in any meaningful, qualitative way besides adding extra parameters.

consumption. Let $u(\cdot)$ be an increasing, twice differentiable function over consumption. Then, the agent's long-run utility function is:

$$U = \gamma \cdot \underbrace{\mathrm{E}_{\tilde{p}}(u(c))}_{\text{Anticipatory Utility}} + (1-\gamma)\Big[\,\theta \cdot \underbrace{\mathrm{E}_{p}(u(c))}_{\text{Consumption Utility}} + (1-\theta) \cdot \underbrace{\mathrm{E}_{p}[\mu(u(c)|r)]}_{\text{Gain-loss Utility}}\,\Big] \tag{1.4}$$

,where $\gamma$ is the weight on anticipatory utility, $\theta(1-\gamma)$ is the weight on consumption utility and $(1-\theta)(1-\gamma)$ is the weight on gain-loss utility. The first term, anticipatory utility is denoted as $\mathrm{E}_{\tilde{p}}[u(c)]$, which is the expected utility evaluated under optimal beliefs. In a world where this was the only source of utility, optimal beliefs would be those that maximized the state in which consumption utility was highest. The second term, expected consumption utility is denoted as $\mathrm{E}_{p}[u(c)]$ evaluated using true beliefs, $p$, which disciplines beliefs from being completely optimistic. The third term, $\mu(u(c)|r)$, refers to gain-loss utility which is the evaluation of outcomes as gains or loss with respect to a reference point[41]. Since our model aims to capture the tradeoff between anticipatory utility and gain-loss utility a natural choice for the reference point is the expected value of the lottery, $\mathrm{E}_{\tilde{p}}[u(c)]$, which is directly affected by beliefs.[42]

The DM's stage 1 maximization problem is to choose beliefs $\tilde{p}$ to maximize (1.4) subject to the following constraints on beliefs: (1) $\tilde{p} \in [0,1]$ and (2) $\tilde{p} + (1-\tilde{p}) = 1$.[43] Constraint (1) says that beliefs are probabilities of states and therefore must be between 0 and 1. Constraint (2) says that the sum over all states must be equal to 1.

In stage 2, the DM makes choices $q$, given a set of beliefs, $\tilde{p}$, to maximize expected consumption utility given by $\mathrm{E}_{\tilde{p}}[u(c)]$ where $\mathrm{E}_{\tilde{p}}$ is the expectations operator under optimally chosen beliefs from stage 1, and $c = g(q(\tilde{p}))$, that is, consumption is a function of the

[41]see appendix section 1.7.3 for a full description of reference dependence

[42]As I show in the online appendix - section 3 on reference dependence, when the reference point is choice-acclimating the individual faces the strongest incentive to minimize gain-loss utility.

[43]Since there is no learning in our environment, one does not have to place other restrictions on optimal beliefs that are assumed in BP, for example, that they following the law of iterated expectations.

choices made under optimal beliefs. Trivially, optimal choices are $q^* =$ such that they maximize expected consumption utility.

To derive predictions of the model, I use the context of the experiment. This is a sufficiently general application that is portable to other contexts and allows us to compare model predictions to the data. To refresh, the DM faces a compound lottery with two outcomes, H and L.[44] The compound lottery is an $R \cdot \alpha$ probability of obtaining the subjective lottery, an $R \cdot (1 - \alpha)$ probability of the outside option, L and $(1 - R)$ is the probability that choices made from a multiple price in which the DM chooses between the subjective lottery and a series of objective lotteries are payoff relevant.

The choice variable $q$ is defined as the indifference point on the MPL . An assumption is that the DM chooses beliefs, $\tilde{p}$, over the subjective lottery but not over the objective probabilities, which are $R$ , $\alpha$ and those in the objective prospect, $q$. However, the main predictions of the model would still remain if individuals were allowed to choose all probabilities.[45] The model also assumes that the DM evaluates subjective lotteries using SEU, and thus reduces compound lotteries when calculating probabilities across objective and subjective parts.

The model can be solved using backward induction: In stage 2, the DM chooses an indifference point, $q$, given a set of beliefs, $\tilde{p}$, to maximize expected consumption utility given by $\mathrm{E}_{\tilde{p}}[u(c))]$ where $c = q(\tilde{p})$. It can be shown trivially that $q^* = \tilde{p}$ maximizes expected consumption utility. In stage 1, the DM will choose beliefs, $\tilde{p}$, of the subjective lottery to maximize long-run wellbeing in equation (1.4).

To distill the model's essential prediction, i.e., how a change in $\alpha$ impacts $\tilde{p}$, a few simplifying assumptions are made. One can eliminate consumption utility, which only

[44]Adding more outcomes simply increases the complexity of the expected value calculations, and requires us to make assumptions on risk aversion, but does not change any of the theoretical results.

[45]If individuals could optimally choose both subjective and objective probabilities, this would lead to similar results. For example, if the individual was maximizing anticipatory utility she would want to believe that $R = \alpha = \tilde{p} = 1$.

serves to change the weight on the costs of self-delusion, by assuming $\theta = 0$.[46] Setting $L = 0$ and $\lambda_H = 1$ and $\lambda_L = \lambda$ to normalize loss aversion (both without loss of generality), the objective function becomes:

$$U = \gamma \tilde{p}_H \Delta + (1-\gamma)\Big[(1-p_H)\lambda \cdot -r + (\Delta - r)p_H\Big]$$

,where $\tilde{p}_H$ = probability of high outcome under optimal beliefs, $p_H$ = probability of high outcome under true beliefs (where the optimal choice, $q$, is still determined by optimal beliefs), and $r$ is the reference point. Taking the derivate of the utility function with respect to $\tilde{p}$ yields the following First Order Condition:

$$\underbrace{\frac{d\tilde{p}_H}{d\tilde{p}} \cdot \frac{dU}{d\tilde{p}_H}}_{\text{Psychological Utility Tradeoff}} + \underbrace{\frac{dp_H}{d\tilde{p}} \cdot \frac{dU}{dp_H}}_{\text{Costs of Self-Delusion}} = 0 \tag{1.5}$$

$$\frac{d\tilde{p}_H}{d\tilde{p}} \cdot \Delta(\gamma - (1-\gamma)(\lambda(1-p_H) + p_H)) + \frac{dp_H}{d\tilde{p}} \cdot (1-\gamma)(\Delta + r(\lambda - 1)) = 0 \tag{1.6}$$

,where:

$$\frac{d\tilde{p}_H}{d\tilde{p}} = R \cdot \alpha + (1-R)\tilde{p} \tag{1.7}$$

$$\frac{dp_H}{d\tilde{p}} = (1-R) \cdot (p - \tilde{p}) \tag{1.8}$$

Equation (1.7) is the impact of changing $\tilde{p}$ on the probability of a high outcome under optimal beliefs. Equation (1.7) illustrates two important points: (1) the marginal impact of increasing beliefs $\tilde{p}$ on $\tilde{p}_H$ is always positive. (2) $\frac{d\tilde{p}_H}{d\tilde{p}}$ is increasing in $\alpha$ which shows that the marginal utility of higher beliefs is increasing in $\alpha$. (1.8) represents the

[46] While it is conceptually interesting to view consumption and gain-loss utility separately, consumption utility only changes the weight on the costs of self-delusion and so can be removed to simplify exposition

penalty of moving optimal beliefs, $\tilde{p}$, away from the truth, $p$, in either direction. When $\tilde{p} < p$ ($\tilde{p} > p$), the derivative is positive (negative) and there is an incentive to increase (decrease) $\tilde{p}$. The penalty also increases as the probability of the elicitation, $(1-R)$, goes up.

Equation (1.7) and (1.8) help us understand the first-order-condition in equation (1.6), better. The first term in (1.6) is psychological utility and captures the tradeoff between anticipatory and gain-loss utility. $\frac{d\tilde{p}_H}{d\tilde{p}}$ is always positive, so the sign of the first term depends on $\frac{dU}{\tilde{p}_H}$ which depends on the relative strength between anticipatory utility and gain-loss utility. A higher $\gamma$ increases the weight on anticipatory utility relative to gain-loss utility pushing beliefs higher. A higher $\lambda$ increases loss aversion pushing beliefs lower. If these were the only two forces in the model, optimal beliefs would be pulled to the boundary of 0 or 1, depending on which force was stronger.[47]

The second additive term of equation (1.6) is the cost of self-delusion. $\frac{dp_H}{d\tilde{p}}$ represents the penalty of moving beliefs away from the truth and, as long as individual's are loss averse ($\lambda > 1$), $\frac{dU}{dp_H}$ is always positive. Thus, if the first term in equation (1.6) is positive, i.e. anticipation outweighs gain-loss, then $\tilde{p} > p$ so that beliefs will be optimistic and will increase as $\alpha$ is raised leading to a *sour grapes effect.* However, if the first term is negative, and gain-loss utility outweighs anticipation (which does not occur under reasonable assumptions as shown in simulations) then $\tilde{p} < p$ and beliefs decrease as $\alpha$ increases leading to a *grass-is-greener effect.*[48] This implies that sour grapes and grass is greener are essentially a tradeoff between the relative weights of anticipatory utility and loss aversion in the model. The F.O.C. is difficult to tract analytically but below I show simulation results of optimal beliefs as $\alpha$ changes.

[47]Note that $\alpha$ also enters $\frac{dU}{d\tilde{p}_H}$ through the $p_H$ term and utility is always increasing in $\alpha$ for $\lambda > 1$.

[48]Note that this does not guarantee optimal beliefs will start at a higher level than true beliefs. For that to occur, it would require unrealistic conditions on the parameters.

### 1.5.2 Model Simulations of Sour Grapes and Grass is Greener

In this section I plot the utility function in (1.4) as a function of subjective beliefs, $\tilde{p}$ for different values of $\alpha$ ranging from 0.125 to 0.875 in increments of 0.15. $\alpha = 0.125$ and $\alpha = 0.875$ are the actual values used in the experiment. For each plotted function I label the utility maximizing belief. Through out I assume that the true probability of a high outcome in the subjective lottery is $p = 0.5$.[49]

Figure 1.5 shows conditions under which sour grapes is predicted. The ratio of losses to gains, $\frac{\lambda_L}{\lambda_H} = 2$, which is considered a rule-of-thumb for what has been observed empirically in the lab and the weight on consumption utility, $\theta$, relative to gain-loss utility is 0.4. Increasing $\theta$ further renders the effect of reference dependence on beliefs non-existent. I assume the weight on anticipatory utility is $\gamma = .5$, which puts equal weight on anticipatory utility relative to the combination of consumption and gain-loss utility. When $\alpha$ moves from 0.125 to 0.875, as it does in our experiment, optimal beliefs move from .42 to .746 . Unlike BP, I also reproduce the pattern that beliefs are pessimistic (less than the true belief $p = 0.5$) at lower levels of $\alpha$ and become optimistic (greater than $p = 0.5$) at higher levels of $\alpha$. [50]

Figure 1.6 shows conditions under which a grass-is-greener effect is predicted. Almost all the assumptions required to observe such an effect are unreasonable. For example, the weight on consumption utility, $\theta$ has to be close to 0 (in this case, exactly 0), for reference dependence to have any effect. The ratio of losses to gains has to be considerably smaller than 2 as well, in this case, $\frac{\lambda_L}{\lambda_H} = 1.2$, which does not fit most experimental data. Finally, I maintain the assumption that $\gamma = 0.5$, as the effect is sensitive

[49] This is the expected belief if subjects thought that the jar was filled with green or red balls with equal probability. In our experiment, subjects are told the jar is filled without instruction by someone unaffiliated with the study.

[50] The impact of change in $\alpha$ on beliefs increases with $\alpha$ The intuition is as follows: At very low levels of $\alpha$, an increase or decrease in beliefs has little impact on anticipatory utility or the reference point, so the optimal belief is governed mostly by the costs of self-delusion, thus leading optimal beliefs closer to true beliefs. As $\alpha$ increases, the impact of optimal beliefs on psychological utility grows. As our design only measures two levels $\alpha$, high and low, I do not test for this non-monotonicity.

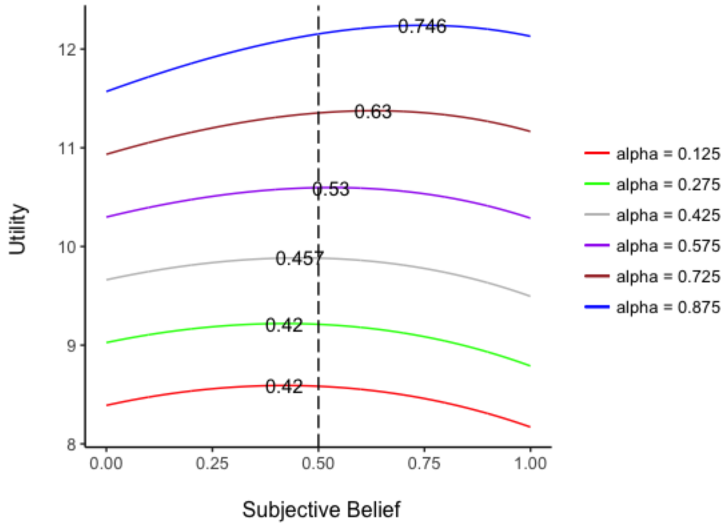


**Figure 1.5.** Impact of Access on Beliefs Under Optimal Beliefs Model and Conditions for Sour Grapes Effect

Above figure plots the utility function in (1.4) on the y-axis as a function of subjective beliefs, $\tilde{p}$, on the x-axis, for different values of $\alpha$ For each plot, I label the utility maximizing beliefs in black. The figure shows an example of parameter values under which the model predicts sour grapes. See text for details.

to changes of $\gamma$ in either direction. Now when $\alpha$ moves from 0.125 to 0.875, optimal beliefs move from .36 to .12. However, I do not find optimal beliefs greater than true beliefs at low levels of $\alpha$ as might be expected for the grass-is-greener effect. In order to generate this pattern, one would need to assume a $\frac{\lambda_L}{\lambda_H} < 1$ , which would also be unreasonable given most evidence. Overall, I find that the model predicts sour grapes for a more reasonable set of parameter values, and this prediction fits with what was observed in the lab and field.

## 1.6 Conclusion

Access-based beliefs are the idea that individuals formulate low beliefs about a subjective prize in order to cope with a low probability of obtaining it. This notion could have significant economic effects. For instance, poor people have pessimistic beliefs on the value of schooling due to the difficulty of gaining access, a mechanism that could ultimately reduce investment in education. A problem with testing such a theory in the real world is

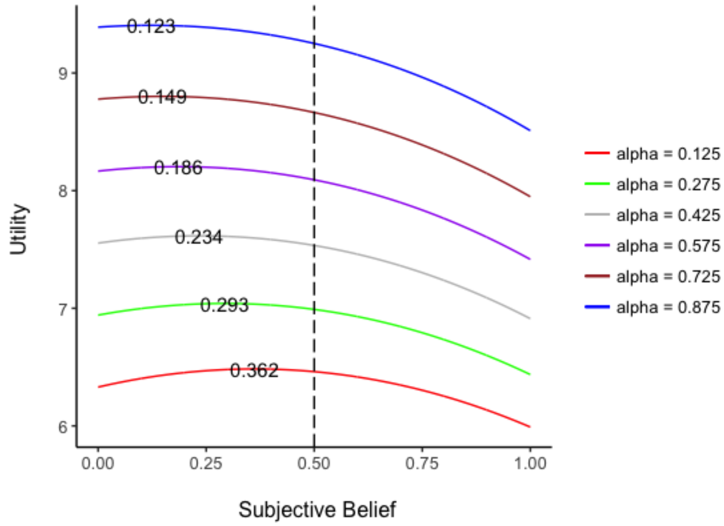


**Figure 1.6.** Impact of Access on Beliefs Under Optimal Beliefs Model and Conditions for Grass-Is-Greener Effect.

Above figure plots the utility function in (1.4) on the y-axis as a function of subjective beliefs, $\tilde{p}$, on the x-axis, for different values of $\alpha$ For each plot, I label the utility maximizing beliefs in black. The figure shows an example of parameter values under which the model predicts grass-is-greener. See text for details.

that there are potential learning confounds. For example, there may be reverse causality where access itself could be a function of value. Furthermore, a change in access may lead to a change in the information that individuals have about an alternative, thus shifting beliefs. In this paper, access-based beliefs are examined in a lab setting where learning channels can be controlled for to isolate the mechanisms of interest.

In a lab experiment, I elicit beliefs of an ambiguous lottery as the objective chance to obtain it varies. I find that lowering access to the subjective lottery, reduces implied beliefs of its value, consistent with the colloquial notion of sour grapes. I further find that this effect is concentrated amongst individuals with high initial beliefs of the lottery. This is consistent with the intuition that individuals who desire an alternative are most likely to engage in sour grapes.

To provide a theoretical explanation of the effect, I adapt an optimal beliefs model from the motivated beliefs literature in which individuals' tradeoff anticipatory utility with gain-loss utility when choosing beliefs of the subjective (ambiguous) lottery. I demonstrate that beliefs increase when the likelihood of obtaining the subjective lottery goes up under most reasonable parameter assumptions. I rule out popular current and past models in ambiguity and reference dependence as unable to explain access-based beliefs. In this regard, the paper can be viewed as one of the first direct tests between motivated beliefs and theories of decision-making under ambiguity.

An interesting implication of the study is that it provides an alternative explanation of the endowment effect. The endowment effect is the finding that initial ownership of a good (or an expectation to own it) increases its valuation. If we think of the initial endowment as providing people with access, then there is clearly a conceptual connection between access-based beliefs and the endowment effect. However, in the paper we show that the prominent model of decision-making used to explain the endowment effect, namely reference dependence, does not make the prediction of an endowment effect when the outcome is a subjective lottery. By contrast, motivated belief models are capable of providing an explanation for both the standard endowment effect and the findings in this paper. Nonetheless, access-based beliefs and reference dependence are two different accounts of the endowment effect, and a promising area of future work might be to develop further tests distinguishing the two theories.

In the next chapter we provide evidence for access-based beliefs in real-world settings and discuss some suggestive evidence of their prevalence in the domains of education and health.

Chapter 1 and 2 in full are being combined and prepared for submission for publication of the material. The dissertation author, Vinayak Alladi was the primary investigator and author of this material.

## 1.7 Appendix

Please contact the author for any additional information

### 1.7.1 Optimal Beliefs Without Reference Dependence

In this section, I show that experiment 1 can be modeled using Brunnermeier and Parker's (2005) optimal beliefs framework and that the sour grapes prediction can be derived. I adapt BP (originally designed for objective uncertainty) to our setting of ambiguity in the following sense. Just as a BP-agent can hold objective beliefs and subjective beliefs, I imagine that an individual facing ambiguity can hold two sets of subjective beliefs; one being what the individual truly believes, and the second being what the individual deceives himself into believing in order to feel happier. I call the first set the "rational subjective beliefs" and the second set the "optimal subjective beliefs" and the BP model constrains the individual to act on the second set.

First, I write down the optimal beliefs utility function in its general form: $P$ is a set of rational subjective beliefs, $q$ is a set of optimal subjective beliefs,$\gamma$ is the weight on consumption utility , $(1-\gamma)$ is the weight on anticipatory utility, $Q$ is the action taken by the individual, and U is a well-behaved utility function. (Subjective) Expected Utility is then:

$$U(q,P,Q) = \gamma U_{t+1}(P,Q) + (1-\gamma)\mathrm{E}_q[U(Q(q))] - \{\mathrm{E}_P[U(Q(P))] - \mathrm{E}_P[U(Q(q))]\} \tag{1.1}$$

The first term is the entire future consumption utility, the second term is the anticipated utility experienced today, and the third term is the economic costs of optimal beliefs, that is, the difference in expected utility when the agent holds rational subjective beliefs and acts accordingly minus the expected utility when the agent holds rational subjective beliefs and acts according to the optimally chosen subjective beliefs.

The agent's decision-problem is to solve for the optimal set of subjective beliefs using a two-stage backwards induction process: (1) At the end, the agent takes an optimal action, given the optimally chosen beliefs from stage 1. (2) Using backward induction, the agent maximizes utility by calculating the optimal set of subjective beliefs to choose given he knows what actions will be taken for any set of chosen beliefs.

Thus, I begin at the final step where agents choose their optimal action, in this case, their report on the multiple price list $Q$ given the optimally chosen beliefs q. The optimal report is the point at which the ambiguous and objective lottery are indifferent. Note I use $q$ for optimally chosen subjective beliefs to distinguish from $p$ which represented any general beliefs of the urn, whether chosen optimally or not.

$$u(L) + q \cdot [u(H) - u(L)] = u(L) + Q \cdot [u(H) - u(L)] \tag{1.2}$$

$$q = Q \tag{1.3}$$

Moving to the second stage, I re-write equation (1) above explicitly for experiment 1 and derive the optimal beliefs $q$. Without loss of generality, I allow the probabilities of the risky urn to vary continuously, even though in the actual MPL they are discrete. Furthermore, I make one small change to notation: In figure 1, I used $\alpha$, $\beta$ and $(1-\alpha-\beta)$ as the probability that either the urn, the outside option, or the elicitation will payout respectively. For notational convenience, here I use $R$ to represent the chance that either the urn or the outside option pays out, and $(1-R)$ as the chance that the payout is based on the elicitation. Without loss of generality, I assume linear utility for notational convenience and $\Delta$, to represent the difference in payoffs $H-L$. Finally, I let consumption utility, $\gamma U_{t+1}(P,\alpha,Q) = C$, since it does not depend on $q$ at all and ultimately drops out of the solution.

$$U(q;P,\alpha,Q) = C + (1-\gamma)\{L + \Delta[R + (1-R)(\int_0^{Q^*} q\,\mathrm{d}v + \int_{Q^*}^1 v\,\mathrm{d}v)]\} + \{L + \Delta[R + (1-R)(\int_0^{Q^*} P\,\mathrm{d}v + \int_{Q^*}^1 v\,\mathrm{d}v)]\} \tag{1.4}$$

The next step is to use backward induction and plug in the solution of the second stage in terms of p into the first stage, yielding:

$$U\,(q;P,\alpha,Q) = C+(1-\gamma)\{L+\Delta[R+(1-R)(\int_0^q q\,\mathrm{d}v+\int_q^1 v\,\mathrm{d}v)]\}+\{L+\Delta[R+(1-R)(\int_0^q P\,\mathrm{d}v+\int_q^1 v\,\mathrm{d}v)]\} \tag{1.5}$$

evaluating the integrals I get :

$$U\,(q;P,\alpha,Q) = \gamma U_{t+1}(P,\alpha,Q)+(1-\gamma)(L+\Delta[R+(1-R)(q^2+\frac{(1-q^2)}{2})]+L+\Delta[R+(1-R)(qP+\frac{(1-q^2)}{2})] \tag{1.6}$$

The derivative of the utility function above with respect to $q$:

$$U'\,(q;P,\alpha,Q) = (1-\gamma)\Delta(R\alpha+(1-R)(2q-q))+\Delta(1-R)(P-q)=0 \implies q^* = \frac{P}{\gamma}+\frac{(1-\gamma)R\alpha}{\gamma(1-R)} \tag{1.7}$$

taking the derivative of $q^*$ with respect to $\alpha$ I get our prediction of sour grapes :

$$\frac{dq^*}{d\alpha} = \frac{(1-\gamma)R}{\gamma(1-R)} \tag{1.8}$$

The derivative will always be positive, thus demonstrating that an increase in access to the ambiguous urn, i.e. an increase in $\alpha$, leads to an unambiguous increase in beliefs about the urn. The expression for the derivative is intuitive. If $\gamma = 1$, meaning that there is no weight on anticipatory utility, then $q = P$. Conversely, as $\gamma$ approaches zero, the weight on anticipatory utility increases, and the higher is $q$ relative to $P$. Furthermore, when $R$, the chance that the participant's payoffs are *not* determined by the incentivized task, increases, the incentive to inflate beliefs also increases. As constructed, $q$ is always greater than or equal to $P$, i.e. optimal subjective beliefs are always greater than or equal to rational subjective beliefs, and adapting the model further to allow for a more explicitly for sour grapes is work in progress.

### 1.7.2 Non-SEU Models of Ambiguity

#### Gilboa-Schmeidler Max-Min

The most popular non-SEU model is the Gilboa-Schmeidler Max-Min Expected Utility (MEU) in which the probability measure is no longer unique, and individuals evaluate the lottery by entertaining the most pessimistic distribution from a set of possible distributions over the outcomes. Once again using notation from Andreoni, T. Schmidt, and Sprenger, 2015, where $a$ is the subjective act, I write the MEU of the compound lottery presented in figure 1:

$$U_{MEU}(\alpha a+(1-\alpha)L)=\min_{p\in P}[\alpha(p\cdot u(H)+(1-p)\cdot u(L))+(1-\alpha)u(L)] \tag{1.9}$$

While the model can explain the Ellsberg paradox it cannot generate the type of violation predicted by sour grapes. This is because the model assumes a form of the independence axiom known as certainty independence which states that when mixing with an objective lottery, individuals treat the subjective part equally. Specifically, under MEU, a single belief governs the subjective lottery for any mixture with an objective one while for sour grapes, the beliefs over the subjective lottery change for different objective mixtures.

#### Kilbanoff, Marinacci, Mukerji Smooth-Ambiguity

Another prominent model in the literature, Klibanoff et al. (2005, KMM for short), can predict sour grapes provided a change to one of its core assumptions is made. Unlike Gilboa-Schmeidler Max-Min, KMM does predict an interaction between changes in probability mixtures and valuations of the subjective lottery but the interaction goes in the opposite direction of sour grapes.

The KMM model features a decision maker who has a proper probability distribution over states but does not treat the expected utilities obtained in each state equally. The utility functional is of the form:

$$U_{KMM}(f) = \int_S p(s)\phi(u(a(s)))ds \tag{1.10}$$

where u(a(s)) is the expected utility in state s and these expected utilities are then aggregated over states by the utility aggregation function $\phi$. When mixing a subjective and objective lottery, the impact of a change in the probability mixture depends upon the shape of the $\phi$ function. If the $\phi$ function is concave, then increasing $\alpha$ will decrease the overall utility of the subjective part of the lottery, consistent with ambiguity aversion. However, if the $\phi$ function is convex, then increasing $\alpha$ will increase the overall utility of the subjective part of the lottery, consistent with ambiguity-seeking behavior. Thus, KMM can explain sour grapes, but it would require changing an assumption that makes it inconsistent with past results.

### 1.7.3 Reference Points

#### Reference Dependence Preliminaries

To fix ideas, Let x be a consumption outcome drawn according to measure F, and r be the reference point. Then, the expected utility of a gamble is outcome is: $U(F|r) = \int u(x|r)\, dF(x)$ where $u(x|r) = m(x) + \mu(m(x) - r)$. The function $m(\cdot)$ represents consumption utility and $\mu(\cdot)$ represents gain-loss utility relative to the referent, $r$.

For simplicity I assume that for small-stakes decisions, consumption utility, $m(\cdot)$, can plausibly be taken as approximately linear, and a piecewise-linear gain-loss utility function is adopted as follows:

$$\mu(y) = \begin{cases} \lambda_H \cdot y & y \geq 0 \\ \lambda_L \cdot y & y < 0 \end{cases}$$

,

where $\lambda_L > \lambda_H$ represents the degree of loss aversion. For a general reference dependence model, I can write the expected utility of our experiment as:

$$U(F|r) = p_H H + (1-p_H)L + p_H \lambda_H (H-r) + (1-p_H)\lambda_L (L-r) \tag{1.11}$$

**Fixed Reference Point**

**Proposition 1.1.** *Under reference-dependent preferences with linear gain-loss utility and a referent fixed at a value* $r$*, conditional on* $[(H-L)+\lambda_H \cdot (H-r)+\lambda_L \cdot (r-L)] > 0$*, which is true under reasonable parameter assumptions, the optimal report of the agent is* $p^* = p$*, her true valuation of the ambiguous lottery. In addition, the optimal report does not depend on the probability of access to the urn,* $\alpha$*.*

The intuition for the result, and of the later results for reference dependence, is essentially the same. Our choice variable $q$ enters our objective function solely via $p_H$, the probability of a high outcome. Therefore, as in the case for a rational agent without reference dependence, our problem reduces to finding the $q$ that maximizes $p_H$ and checking whether this is a maximum of the function. With reference dependence, I need to check second order conditions to ensure that the incentives to increase consumption utility outweigh those to mitigate the loss part of gain-loss utility. In the case that our objective function is linear in $p_H$, I can verify this by checking the sign of the coefficient that multiplies $p_H$. Under all reasonable parameter assumptions, the sign of the coefficient is positive so that the optimal choice is to maximize $p_H$ by setting $q = p$.

The second part of the proposition is finding the impact of access to the ambiguous lottery, $\alpha$, on our optimal choice. Since $\alpha$ does not enter $p_H$ it has no marginal impact on the optimal choice so I can rule out any effect. A change in $\alpha$ could potentially change the function's maximum to a minimum, but this would lead to a discontinuous jump of the optimal choice to its boundaries for small changes in $\alpha$ which I do not observe in the data.

*Proof.* The optimal report is the report that maximizes equation 7 above. I can re-write the equation as:

$$U(F|r) = p_H \cdot (H - L) + L + p_H \cdot \lambda_H \cdot (H - r) + (1 - p_H) \cdot \lambda_L (L - r) \tag{1.12}$$

$$U(F|r) = p_H \cdot [(H - L) + \lambda_H \cdot (H - r) + \lambda_L \cdot (r - L)] + L + \lambda_L (L - r) \tag{1.13}$$

As long as:

$$[(H - L) + \lambda_H \cdot (H - r) + \lambda_L \cdot (r - L)] > 0 \tag{1.14}$$

then equation 9 be maximized by maximizing $p_H$, the only part of the equation where $q$ enters. As was shown in proposition 1, the optimal choice, $q^* = p$ and does not vary with the probability of getting the ambiguous lottery, $\alpha$.

Under most reasonable parameter assumptions the above condition will hold. By assumption , $H > L$ , so the first term is greater than zero. As long as the following condition on the reference point holds, $L \leq r \leq H$ this also guarantees that the remaining terms in brackets are greater than or equal to zero. It would be unreasonable to have a reference point that is lower than the lowest outcome in the gamble , so I do not consider that case. However, even if this were true, the resulting function would have a maximum either at $p^* = 0$ or $p^* = 1$ regardless of $\alpha$. □

**Expectations-Based Reference Points**

Another possibility for the reference point is based on what you expect to get in the lottery. There are two formulations of expectations-based reference points that I consider:

**Case 1:** Reference Point = Expected Utility of the Lottery

This notion of a reference point has been explored in models of disappointment aversion (Bell, 1985; Loomes and Sugden, 1986). Its impact in our model is to introduce an additional source of disutility from maximizing the probability of a high outcome since it mechanically increases the reference point from which gains and losses are calculated. In other words, the reference point becomes endogenous, or choice-acclimating as expectations respond to our choices. This source of disutility was not previously seen under a fixed $r$.

Consider the same objective function as equation 9, except that the reference point, $r$, is now defined as the expected utility of the lottery $p_H H + (1 - p_H)L$. Let $\Delta = H - L$ , then the objective function is:

$$U(F|r) = p_H\Delta + L + p_H\lambda_H(H - (p_H\Delta + L)) + (1 - p_H)\lambda_L(L - (p_H\Delta + L)) \qquad (1.15)$$

**Proposition 1.2.** *Under reference-dependent preferences with linear gain-loss utility, where the referent is the expected value of the lottery, the optimal report of the agent is* $p^* = p$*, her true valuation of the ambiguous lottery, conditional on* $\Delta\lambda_L > (\Delta + H\lambda_H + 2r\alpha p\Delta(\lambda_L - \lambda_H))$*, which is true under reasonable parameter assumptions. In addition, the optimal report does not depend on the probability of access to the urn,* $\alpha$*.*

*Proof.* beginning with equation 11 above, I assume, without loss of generality, that L = 0. This yields:

$$U(F|r) = p_H\Delta + p_H\lambda_H(H - p_H\Delta) - (1 - p_H)\lambda_L p_H\Delta \qquad (1.16)$$

collecting the $p_H$ terms and the $p_H^2$ terms separately I get:

$$U(F|r) = p_H(\Delta + \lambda_H H - \lambda_L \Delta) + p_H^2 \Delta(\lambda_L - \lambda_H) \tag{1.17}$$

I can write equation 13 in terms of $q$ by substituting for $p_H$ as follows:

$$\begin{aligned} U(F|r) = (R\alpha p + (1-R)[qp + \frac{1-q^2}{2}])(\Delta + \lambda_H H - \lambda_L \Delta) + \\ (R\alpha p + (1-R)[qp + \frac{1-q^2}{2}])^2 \Delta(\lambda_L - \lambda_H) \end{aligned} \tag{1.18}$$

I expand the equation by collecting all terms that multiply with $(1-R)[qp + \frac{1-q^2}{2}])$, and all terms that multiply with $(1-R)[qp + \frac{1-q^2}{2}])^2$, separately. The remaining additive terms that do not include $q$ are represented as $K$.

$$\begin{aligned} U(F|r) = K + [qp + \frac{1-q^2}{2}])(\Delta + H\lambda_H - \Delta\lambda_L + 2r\alpha p\Delta(\lambda_L - \lambda_H)) + \\ ((1-R)[qp + \frac{1-q^2}{2}])^2 \Delta(\lambda_L - \lambda_H) \end{aligned} \tag{1.19}$$

It is important for us to sign the term in brackets that multiplies the linear q term. The only term in this expression which is negative is $\Delta\lambda_L$. So, unless $\Delta\lambda_L > (\Delta + H\lambda_H + 2r\alpha p\Delta(\lambda_L - \lambda_H))$ ,which is not the case in our experiment and will generally not be true for reasonable estimates of $\lambda_L/\lambda_H \leq 2$ , the typically observed ratio in lab experiments that measure loss aversion, the sum of the terms will be positive. Thus, assuming a positive value for this expression, the above equation can be maximized by setting $q^* = p$ as shown form the first order condition below:

$$\frac{dU(F|r)}{dq} = (1-R)(p-q)(\Delta + H\lambda_H - \Delta\lambda_L + 2r\alpha p\Delta(\lambda_L - \lambda_H)) +$$
$$(1-R)[qp + \frac{1-q^2}{2}])^2(p-q)\Delta(\lambda_L - \lambda_H) = 0 \quad (1.20)$$

When $q^* = p$ the first order condition is zero. This is a local maximum as the second order condition $\frac{dU(F|r)}{dq} < 0$ under all reasonable assumptions of the parameter values.

Furthermore, the optimal report, $q^*$, does not depend on $\alpha$. Unlike in the case of the fixed reference point, $\alpha$ does enter multiplicatively in our first order condition. However, changing $\alpha$ does not impact the local maximum of the function at $q^* = p$ as shown in the first order condition above.

□

There is now a tradeoff between increasing $p_H$ to maximize consumption utility and reducing $p_H$ to lower the reference point, $r$. The optimal report thus depends on which of the two channels is stronger. Under the assumption of linear gain-loss utility, with reasonable parameter assumptions, the incentives to increase $p_H$ are stronger so the optimal report is $q^* = p$. Additionally, the optimal report continues to be independent of access to the subjective lottery since changes in the parameter $\alpha$ have no impact on the value of $q$ that maximizes $p_H$.

**Case 2:** Koszegi-Rabin Reference Dependence

Under KR preferences , the reference point is the entire distribution of the lottery. This specification comes from the Koszegi-Rabin model of reference dependence (2006) (KR) which provides a well-defined structure to the formation of a reference point. KR reference dependence assumes that " a person reference point is her probabilistic beliefs about the relevant consumption outcome held between the time she first focused on the decision determining the outcome and shortly before consumption occurs." (1141).

Consider again x to be a consumption outcome drawn according to measure F. Let r represent the referent drawn according to measure G. Then the KR utility formulation is: $U(F|G) = \iint u(x|r)\, dG(r)\, dF(x)$ with $u(x|r) = m(x) + \mu(m(x) - m(r))$. The function $m(\cdot)$ represents consumption utility and $\mu(\cdot)$ represents gain-loss utility relative to the referent, $r$. I assume the linear gain-loss utility described in the previous section.

The referent, in this case, is an endowment of an objective-subjective lottery. Given this lottery, an individual makes a series of choices on an MPL. Since these choices impact the final lottery she will face, the individual's referent lottery depends on her choices in the elicitation.

KR define an equilibrium concept called personal equilibrium (PE) in which an individual's plan is optimal given her expectations, her expectations are rational given her plans, and that the plans are credible in that the individual has no incentive to deviate from her plan when choices are made. Equilibrium at the choice level may be too fine a degree of granularity for choice experiments in which individuals make a series of decisions (as argued in Sprenger (2015)). In our experiment, the series of choices (displayed in the uncertainty equivalent task in Figure 1.2) yields a plan of the following form: choose the subjective lottery over the risky prospect up to the point $q^*$ (the individual's uncertainty equivalent), and choose the risky prospect from this point on.

For this plan to be a PE it must be that the participant is indifferent between the subjective lottery and the risky prospect at the value of $q^*$, meaning that there is no incentive to deviate. In other words, the individual will choose a $q^*$ that maximizes her KR utility. [51]

Under these preliminaries, the KR expected utility can be written down as follows:

[51] In the language of KR, the referent lottery can be described more precisely as choices on the elicitation such that her choices maximize her KR utility given those expectations.

$$U(G|G) = p_H \cdot p_H \cdot u(H|H) + (1-p_H)\cdot(1-p_H)\cdot u(L|L) + p_H \cdot (1-p_H)\cdot u(H|L) \\ + (1-p_H)\cdot p_H \cdot u(L|H) \quad (1.21)$$

The first term refers to the chance of expecting H as the referent and obtaining H as the consumption outcome. The second term is similar for L. The third term refers to the chance of expecting H as the referent and obtaining L as the consumption outcome.The fourth term refers to the chance of expecting L as the referent and obtaining H as the consumption outcome. With H > L, the KR model predicts loss aversion to be present in the fourth term. [52]

Analyzing the model yields proposition 1 below. Under reasonable parameter assumptions on the degree of loss aversion, the optimal choice is $q^* = p$, and because $\alpha$ does not enter the expression for $p_H$, it has no marginal impact on the optimal choice that maximizes $p_H$. Thus, KR preferences cannot explain the effect. I prove this below.

**Proposition 1.3.** *Under Kosezgi-Rabin reference-dependent preferences with linear gain-loss utility, where the referent is stochastic, the optimal report of the agent is her true valuation of the ambiguous lottery, , $p^* = p$, if the condition $(\Delta + (1-2R\alpha p)((\lambda_H - \lambda_L H)\Delta) > 0$ is met, which holds under reasonable parameter assumptions. In addition, the optimal report does not depend on the probability of access to the urn, $\alpha$.*

*Proof.* As before, without loss of generality, I assume linear utility and re-write equation

[52] With KR preferences, note that the referent is impacted differently by the optimal $q$ than under models of disappointment aversion, where the referent is the expected utility of the gamble. Specifically, each possible outcome is compared to every other possible outcome when calculating gains and losses so that rather than $(L-r)$ and $(H-r)$, the gain-loss terms are $(L-H)$ and $(H-L)$. This means that the choice variable, $q$, does not enter into the value of the gain or loss, but only in its probability.

17 as:

$$U(G|G) = p_H \cdot H + (1-p_H) \cdot L + p_H(1-p_H)[\lambda_H(H-L) + \lambda_L(L-H)] \tag{1.22}$$

let $\Delta = H - L$:

$$U(G|G) = L + p_H\Delta + p_H(1-p_H)[\Delta(\lambda_H - \lambda_L)] \tag{1.23}$$

Substituting for $p_H$:

$$\begin{aligned} U(G|G) = L + (R\alpha p + (1-R)[qp + \frac{1-q^2}{2}]) \cdot \Delta + \\ (R\alpha p + (1-R)[qp + \frac{1-q^2}{2}]) \cdot (1 - (R\alpha p + (1-R)[qp + \frac{1-q^2}{2}])) \\ \cdot [\lambda_H(H-L) + \lambda_L(L-H)] \end{aligned} \tag{1.24}$$

As I do in proposition 3, I expand and collect terms that multiply with $(1-R)[qp+\frac{1-q^2}{2}])$, and all terms that multiply with $(1-R)[qp+\frac{1-q^2}{2}])^2$, separately. The remaining additive terms that do not include $q$ are represented as $K$.

$$\begin{aligned} U(G|G) = K + ((1-R)[qp + \frac{1-q^2}{2}])(\Delta + (1-2R\alpha p)((\lambda_H - \lambda_L)\Delta) - \\ ((1-R)[qp + \frac{1-q^2}{2}])^2\Delta(\lambda_L - \lambda_H) \end{aligned} \tag{1.25}$$

The first order condition becomes:

$$\frac{dU(G|G)}{dq} = (1-R)(p-q)(\Delta + (1-2R\alpha p)((\lambda_H - \lambda_L H)\Delta) -$$
$$(1-R)[qp + \frac{1-q^2}{2}])^2(p-q)\Delta(\lambda_H - \lambda_L) = 0 \qquad (1.26)$$

Once again, for all reasonable values of the parameters , the expression $(\Delta + (1 - 2R\alpha p)((\lambda_H - \lambda_L H)\Delta) > 0$, unless $\lambda_L/\lambda_H$ is much greater than 2. The squared part of the function is multiplied by a coefficient that is always negative, $\Delta(\lambda_H - \lambda_L)$. However, because this part of the function is subtracted, it increases the function overall. This means that the first order condition is zero when $q^* = p$.This is a local maximum as the second order condition $\frac{dU(F|r)}{dq} < 0$ under all reasonable assumptions of the parameter values.

Furthermore, the optimal report, $q^*$, does not depend on $\alpha$. Unlike in the case of the fixed reference point, $\alpha$ does enter mutliplicatively in our first order condition. However, changing $\alpha$ does not impact the maximum of the function $q^* = p$ as shown by the first order condition above. □

### 1.7.4 Differences in Payouts Between US and India

The payouts in India were calibrated so that the amounts would be equivalent, in purchasing power parity terms, to 30 and 10 US dollars respectively. Using a measure of PPP from the World Bank and from inputs of the economics faculty at Christ University, I chose the high payout amount to be 500 Rupees, which is one of the denominations of Indian currency. The low amount was chosen to be 150 Rupees so that the ratio of high to low approximately equaled \$30 to \$10 to make it comparable with sessions in the US. Although both the high and low payouts are slightly lower in India (after adjusting for PPP and converting to US dollars) I feel, if anything, the excitement around getting money for participation was higher in India where students had not been exposed to this kind of incentivized participation before.

### 1.7.5 Choice to Belief Mapping in Version 3

In figure 1.7 below I show how the mapping varies for a standard power utility function with constant relative risk aversion, $u(x) = x^{\gamma}$ at three different values of $\gamma = (0.3, 0.6, 1)$

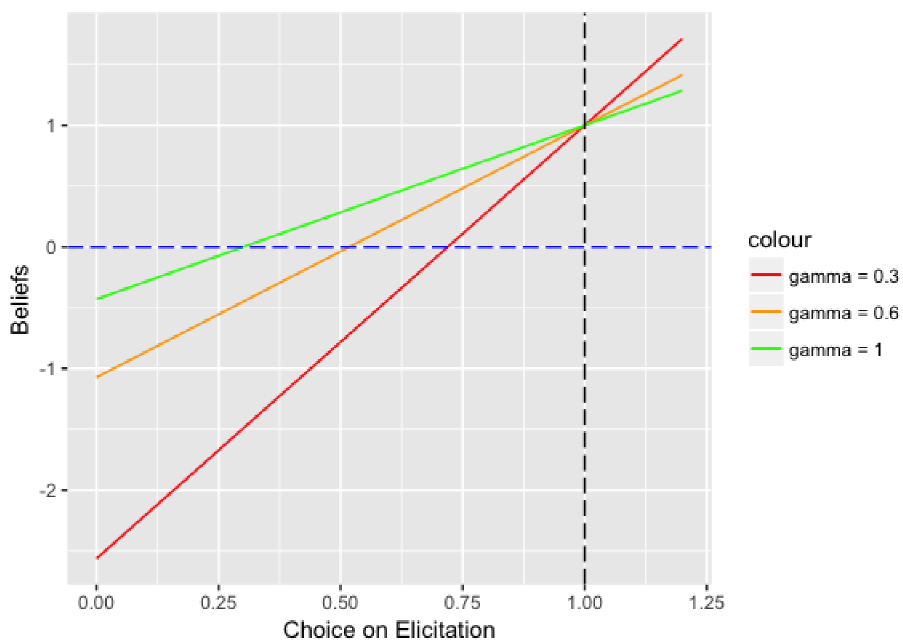


**Figure 1.7.** Choice to Belief Mapping at Different Levels of Risk Aversion

As shown above, as risk aversion decreases from 0.3 to 1, the slopes of the choice-to-belief functions also decrease. Thus, in order to be conservative with our estimates of the treatment effect, I choose risk neutrality as a bench mark for converting choices into beliefs and assume that u(x) = x. This will give us a lower bound on the size of the treatment effect relative to a mapping that assumes risk aversion.

Figure 1.7 also shows that the level of beliefs is always higher for lower levels of assumed risk aversion in the domain of our elicitation choice. This may impact comparison of the level of beliefs across different versions. I discuss this in the results section.

With our assumption that u(x) = x, the mapping from choices to beliefs using the amounts paid out in India becomes:

$$p = q\frac{500}{500-150} - \frac{150}{500-150}$$

This means for every higher row at which a participant switched from preferring JAR A to JAR B, the change in beliefs is %7.14, while in versions 1 and 2, the same change in row will be %5 as shown previously.

It was possible for implied beliefs to be greater than 1 due to the belief intervals being %5, or due to the mapping function, in the case of version 3. So ,in these instances, beliefs were capped at 1. Similarly, beliefs could have been negative in version 3 due to the mapping function so these beliefs were capped at 0.

### 1.7.6 Implementing Payment

A random incentive mechanism was employed to pay subjects - a method commonly used in the experimental literature and in experiments with ambiguity (Ahn, Choi, Gale, and Kariv, 2007; Halevy, 2007). In these mechanisms, one of the subjects decisions is chosen at random as the ”decision-that-counts” and the corresponding lotteries are played out. In our case, subjects are informed that if the envelope they open says “task”, then 1 of the 21 decisions they made will be chosen randomly; the corresponding subjective or objective prospect, depending on which was chosen, will be constructed; and a ball will be drawn from the jar for payout. One of the issues with this mechanism is incentive compatibility, i.e., will the choices elicited through a random incentive mechanism coincide with the choices in a single choice problem?

The answer depends on one main concern first raised by Raiffa (1961) - that subjects use the random incentive mechanism to perfectly hedge against ambiguity. In the classic Ellsberg 2-color problem, subjects make a pair of choices between a subjective urn with red and black balls and a risky urn with 50% red and 50% black balls. The first choice is a bet on red and the second choice is a bet on black. If the participant chooses the ambiguous urn in both cases, then the random payment mechanism would give a bet on one color or a bet on the other color with equal probability. Regardless of the color of the drawn marble the participant is holding an objective 50-50 gamble, hedging the ambiguity completely and potentially misreporting his preferences.[53] One advantage of our experiment is that subjects face only one subjective lottery - a bet on the color of their choice. Thus, it is theoretically impossible for subjects in our experiment to hedge against ambiguity.

A second concern is that if subjects reduce compound lotteries, a random payment mechanism is incentive compatible only if subjects are expected utility maximizers (and thus follow the independence axiom). However, experimental results documenting the reduction of compound lotteries are rare and Halevy (2007) finds that those who do perform reduction tend to satisfy independence. Thus, subjects who respect reduction seem to be rare, and seem to be exactly those for whom independence is a reasonable assumption. Furthermore, the optimal beliefs model that I advocate to explain access-based beliefs does follow subjective expected utility in the second-stage when subjects actually make their choices, and would not suffer the concern about non-expected utility preferences. The other non-SEU ambiguity models I explore, including max-min SEU and KMM smooth ambiguity, do not make the predictions of access-based beliefs even when considering a single choice over lotteries, so worrying about the incentive compatibility under those models is second order.

[53]To elaborate, consider an ambiguity averse participant whose underlying preferences are to bet on the two known rather than on the two ambiguous events. However, if he views the two choices as a single decision problem, he may conclude that by choosing bets on the two unknown events he will win with a probability of 0.5, independently of the event that will materialize which is identical to the probability of winning by choosing the two risky bets. Therefore, he is indifferent between truthfully reporting his preferences, and reporting the opposite preferences.

### 1.7.7 Multiple Switchers

#### Summary

As mentioned, I deal with differential rates of multiple switching (8.3 p.p. lower multiple-switching in the low access group driven by the Indian portion of the study) in three ways:

(1) I conduct identical analysis as in table for just the US portion of the data, which makes up 75% of the sample. I present these results in appendix 1.7.9 . Overall , I find similar magnitudes as in the full sample with slightly less precision. For the full distribution of beliefs, the treatment effect has a point estimate of -1.4 but is insignificant (p-value of .415). However, when focusing on the individuals with coherent beliefs, that is, beliefs larger than 50%, I find treatment effects in the same range as I did for the full sample that are also significant. Controlling for observables, the impact on low access is between -2.3 p.p. and -3.0 p.p. depending on specification with significance at the 10% level.

(2) I present the Lee-bounds for the main treatment effects in the last three rows of table 1.5 in the results section. As can be expected for treatment effects of the size we find, the effect strengthens significantly for trimming at the lower bound, but does not survive trimming at the lower bound.

(3) I perform choice-level analysis of the treatment effects which includes the multiple switchers in two ways: (1) all types of multiple switchers regardless of how random the pattern of switching, this represents about 30% of the data (2) those who display a systematic pattern of switching[54], which is about 1/3 of the multiple switchers, representing about 10% of the data. Overall, I find that the treatment effects are more robust to including multiple switchers who display a systematic pattern of switching. Being

[54]For example, some subjects displayed a pattern of only switching one extra time or switch back and forth. Others switched several times in the middle of the MPL when the two lotteries were of more equal value, but not at the ends.

in the low access treatment leads individuals to be less likely to choose the subjective lottery by about 2 p.p with p-values ranging from .06 to .167 depending on specification. However, when I include all types of multiple switchers the treatment effect still has a small negative point estimate -.009, but is not significant (p-value of .354). I present choice-level results with a more detailed discussion in online appendix - section 4.

**Choice Probabilities**

Our analysis so far is conducted with one observation per individual marking their switch point on the MPL. The primary drawback is that this leaves out multiple-switchers, for whom we cannot infer a belief. Analysis at the choice-level allows us to include multiple switchers and provides some gains in statistical power as each subjects make 21 decisions in our experiment.

To better understand data at the choice level, we first examine choice probabilities graphically. Figure 1.8 below depicts the probability of choosing the ambiguous prospect for each row of the elicitation by treatment condition for the sample of single switchers and multiple switchers , separately.

Panel (A) is for the sample of single-point switchers and therefore should mirror what I learned from the analysis of distributional treatment effects which is that in versions 1 and 2 of the experiment (top row, left graph) I see marginal positive effects of high access on switch points at the higher end of the MPL (in this case rows 12 to 14) and for version 3 , I see more pronounced effects of high access at a broader range of the MPL (rows 9 to 16).

Panel (B) of figure 1.8 shows the graphs for the sample of multiple switchers only. Multiple switchers in the high access group do not seem to exhibit a greater likelihood of choosing JAR A compared to low, which I had no prior on since I do not have a theory on multiple-switching. Secondly, the graphs have a downward slope for both treatment conditions. In other words, the probability of choosing JAR A decreases as

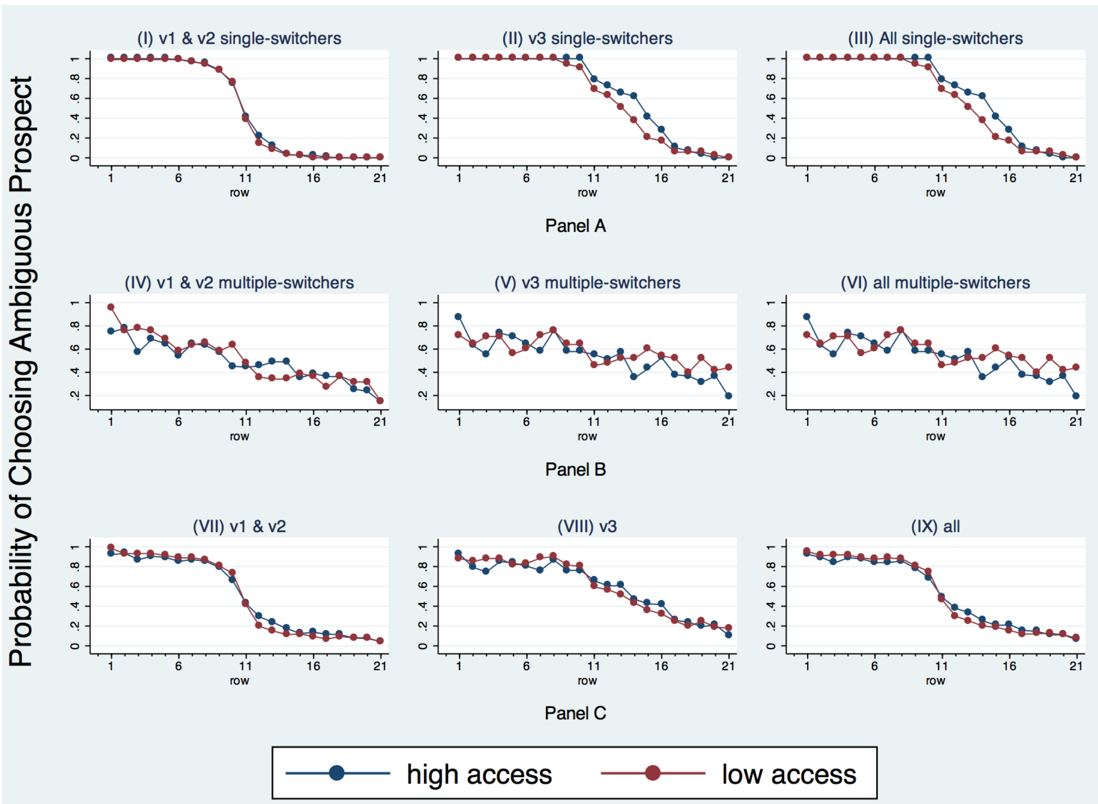


**Figure 1.8.** Choice Probabilities By Row for All participants

the row number increases among the multiple switchers, suggesting that participants are responding to incentives. Furthermore, responses on the post-study questionnaires about how individuals made choices on the task suggest that some types of multiple switching were more systematically thought out than others so that some patterns of multiple switching are not random. In the next section, I categorize types of multiple switchers who were more systematic in their choices and estimate treatment effects for the sample of single-switchers plus these ”systematic” multiple-switchers.

**Systematically Dealing With Multiple Switchers**

Multiple switchers (participants who fail to indicate a unique switch point in the elicitation) represented about 50% of the data collected in India. However, the downward slopes of the choice probabilities in figure 1.8, panel B reflects that some portion of multiple switchers did respond to the incentives of the elicitation and are not choosing between JAR A and JAR B randomly. I classify two kinds of these multiple switchers that were prominent in the data, the ”get-lucky” switchers, who only switched one other time in the MPL, and the ”trembling-hand” switchers, who prefer JAR A initially (for the first at least 5 rows) and JAR B later on (for at least the last 5) , and are undecided in-between.

I calculate treatment effects at the choice-level including these multiple switchers into our sample and estimating the following regression:

$$Choice_{ij} = \beta_0 + \beta_1 LowAccess + \sum_{k=1}^{t} \beta_k x_i + \epsilon_{ij} \tag{1.27}$$

In total, 28 individuals exhibited the ”get-lucky switch” while 19 individuals exhibited ”trembling hands”. These numbers account for 6% and 4% of the data respectively. I present results of this analysis in table 1.6. Overall, it seems that including multiple switchers improves our confidence of the treatment effects in versions 1 and 2 but leads to slightly weaker estimates in version 3 where the number of added observations is low (only

6 individuals in version 3 fall under the category of ”get-lucky” switchers and 6 individuals fall under the category of trembling hands.).

Overall it seems that the analysis of treatment effects at the choice level leads to stronger estimates in the case of version 1 and 2 when I add systematic multiple switchers, but weaker estimates in version 3 where the number of added observations is low (only 6 individuals in version 3 fall under the category of ”get-lucky” switchers and 6 individuals fall under the category of trembling hands.). I present detailed analysis in table 1.6 below.

**Get-lucky Switchers**

Examination of the post-study questionnaires shows that some multiple switchers were trying to ”get-lucky” by switching at some random point of the elicitation in addition to their indifference point. I categorize these multiple switchers as the ”Get-lucky Switchers”, meaning that besides their indifference point, they made one more random switch, but importantly, switched back again.[55]

**Trembling-hands Switchers**

Another pattern found was individuals who exhibited what can be termed ”trembling hands”. These individuals chose JAR A for the first ”n” number of rows, and then seemed to switch between JAR A and JAR B for a a few rows thereafter before finally shifting to JAR B for the remaining ”m” rows. Since there is no exact way to decide how many middle rows is appropraite to distinguish random-switching from trembling hands, I define trembling hands as individuals who chose JAR A for the first ”n” rows, and JAR B for the last ”m” rows and could have any switching pattern in the middle where ”n” and ”m”

[55]To be more precise, the number of switches for these individuals is three since the ”get-lucky” switch consists of them switching from one JAR to the other, but then immediately switching back after that row. This is in addition to a regular switch point. One might ask, what about individuals who switch twice. Those who switched twice were those who began preferring JAR A, than switched to preferring JAR B at some point, and then switched back to JAR A from then on. Or, the opposite could have occurred where individuals began preferring JAR B, than switched to preferring JAR A at some point, and then switched back to JAR B from then on. Either way, there is a portion of their choices that are more irrational than those who switched 3 times

are at least 5 (there are a total of 21 rows). I test the sensitivity of this definition over a range of possible middle row values.

In total, 28 individuals exhibited the "get-lucky switch" while 19 individuals exhibited "trembling hands". These numbers account for 6% and 4% of the data respectively. I show specifically how many fall into each version and how many for each threshold of trembling hands in our regression tables below.

Table 1.6 presents the findings of this analysis for version 1 and 2 where I run the following linear probability model:

$$Choice_{ij} = \beta_0 + \beta_1 LowAccess + \sum_{k=1}^{t} \beta_k x_i + \epsilon_{ij} \tag{1.28}$$

Where $Choice'$ is a "1" or "0" depending on whether individual 'i' chose JAR A in row 'j'. $'Lowaccess'$ is a binary indicator for the treatment and the $'x_i's'$ are a set of demographic questions asked during the post-study questionnaire. Our main test of the theory is whether or not $\beta_1$ is less than 0. I choose a linear probability model as the coefficients are easier to interpret and I am mostly interested in the differences in probability of choosing JAR A by treatment. Indeed, none of our results have boundary concerns.

Column 1 of table 5 shows the results of estimation equation 19 on individuals in versions 1 and 2 with a unique switch-point. I find that the likelihood of choosing JAR A is 1% lower (p-value of .413) in low access groups from a baseline of .505. As predicted, adding all types of multiple switchers to the sample reduces the difference in likelihood to .8%.

In column 3, I restrict our analysis to single-point switchers and the "get-lucky" switchers who switched back and forth between JAR A and JAR B at one other point. Running equation 19 on this sample of individuals increases our point estimate of the treatment effect to 1.9% and the p-value to .123. Although still insignificant, this is a large change in both the treatment effect and our confidence of it.

.

**Table 1.6.** Treatment Effects With Multiple Switchers for Versions 1 and 2

| | Dependent Variable: Probability of Choosing JAR A | | | | | | |
|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| Low Access | -0.010 | -0.008 | -0.019 | -0.019 | -0.020 | -0.018 | -0.018 |
| | (0.012) | (0.010) | (0.012) | (0.011) | (0.011)* | (0.012) | (0.012) |
| Constant | 0.505 | 0.506 | 0.501 | 0.509 | 0.508 | 0.526 | 0.526 |
| | (0.021) | (0.029) | (0.047) | (0.044) | (0.044) | (0.020) | (0.021) |
| Observations | 4762 | 6702 | 5223 | 5495 | 5433 | 5349 | 5265 |
| R-squared | 0.002 | 0.002 | 0.002 | 0.002 | 0.002 | 0.002 | 0.002 |
| Low Access = 0 | 0.413 | 0.436 | 0.123 | 0.100 | 0.086 | 0.120 | 0.128 |
| Controls | Y | Y | Y | Y | Y | Y | Y |
| Single Switch | Y | Y | Y | Y | Y | Y | Y |
| Get lucky Switch | N | Y | Y | Y | Y | Y | Y |
| Trembling Hand | N | Y | N | 5 | 6 | 7 | 8 |
| Other Multiple Switchers | N | Y | N | N | N | N | N |
| Number of extra individuals | | | 22 | 13 | 10 | 6 | 2 |

Robust Standard Errors in Parentheses. p<0.1*, p<0.05**, p<0.01 ***.

In columns 4 to 7, I add to this sample the individuals who exhibited trembling hands. The row of the table labeled ”Trembling Hand” indicates the number of rows at the beginning (”n”) and end of the elicitation (”m”) for which the individual would have had to choose JAR A and JAR B respectively to be considered in this regression. Thus, for column 4, ”n” and ”m” are both 5. For column 5, ”n” and ”m” are both 6 and so on till column 7 where ”n” and ”m” are both 8.

In column 4 I look at the case where I add trembling hand individuals and n = 5. The treatment effect remains at 1.9% as was the case in column 3, and the p-value comes is 0.1. This is the range of treatment effects I see in columns 5 to 8 as well, with p-values ranging from .086 to .128. It seems that most of the impact comes from ”get-lucky” switchers but that there are some marginal gains in confidence of these estimates by adding ”trembling hands” switchers as well.

I conduct the same analysis for version 3 and show results in appendix H (table H.1). The first column shows what I learn from the implied beliefs analysis in table 2. In version 3, individuals in the low access group are 6% less likely (p-value of .058) to choose JAR A from a baseline of 73.9%. Adding multiple switchers in column 2 almost eliminates this entire effect as the point estimate drops to .4% (p-value of .840). In column 3, I add the ”get-lucky” switchers and the point-estimate increases to -3.6% (p-value of .286). In columns 4 to 6 , I add the trembling hand switchers who picked JAR A for the first 10 (9, 8) rows and JAR B for the last 4 (3, 2) rows. The point estimates on the treatment effects are almost identical while the p-values also range from .286 to .249.

### 1.7.8 Structural Analysis of Treatment Effects

Moving from choice probabilities to the actual structural estimation, I jointly estimate beliefs and a noise parameter using a probabilistic choice function. Given the subjective expected utility of the subjective lottery, U(A), and the expected utility of the risky prospect, U(B), the probability that an individual chooses the subjective lottery can be expressed as:

$$Pr(ChoiceA) = F(\frac{U(A) - U(B)}{\nu}), \tag{1.29}$$

where $U(X) = \sum_i p(x_i)u(x_i)$ and $u(x_i) = x^{\alpha}$. The parameter $\nu$ measures stochastic decision error and F is a log normal or logistic CDF transformation that yields the probability of choosing JAR A. The analysis assumes that all individuals within a treatment condition have the same preference parameters and that all heterogeneity comes from decision error. Notice, as $\nu$ tends to zero, the function tends to 1. Similarly as $\nu$ tends to infinity, the function tends to 0. In between, the probabilities moves from 0 and 1, and $\nu$ smooths the choice probability function.

Table 1.7 below shows the results of the structural estimation.We assume a level of risk aversion of $\alpha = 0.7$, the average estimated level of risk aversion in Holt and Laury (2002), and use maximum likelihood techniques to estimate beliefs and a noise parameter. Overall, I find similar results to our reduced-form analysis. Our estimates of belief levels are very similar for versions 1 and 2, where risk aversion is unidentified , and slightly lower for version 3.

In columns 1 to 3 I estimate beliefs separately by treatment condition, and a noise parameter common to both treatments, on the sample of participants who were unique switchers. I find little evidence of the sour grapes effect for versions 1 and 2, a difference in beliefs between high and low access groups of .7% (p-value 0.26), but a larger treatment effect for version 3 alone, a treatment effect of 3.8% (p-value of .16). The combined impact shows a treatment effect of 1.0% (p-value of .075).

**Table 1.7.** Structural Estimation of Treatment Effects

| | No Multiple Switchers | | | Systematic Multiple Switchers | | | All Subjects | | |
|---|---|---|---|---|---|---|---|---|---|
| | (v1 and v2) | (v3) | (all) | (v1 and v2) | (v3) | (all) | (v1 and v2) | (v3) | (all) |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| High Beliefs | 0.498 | 0.495 | 0.497 | 0.501 | 0.468 | 0.494 | 0.496 | 0.459 | 0.480 |
| | (0.005) | (0.020) | (0.005) | (0.004) | (0.020) | (0.005) | (0.004) | (0.014) | (0.005) |
| Low Beliefs | 0.491 | 0.457 | 0.485 | 0.493 | 0.461 | 0.488 | 0.496 | 0.460 | 0.482 |
| | (0.004) | (0.018) | (0.004) | (0.004) | (0.017) | (0.004) | (0.003) | (0.015) | (0.004) |
| High Noise | 1.915 | 3.778 | 2.297 | 2.734 | 5.710 | 3.500 | 5.223 | 12.955 | 7.694 |
| | (0.194) | (0.423) | (0.164) | (0.367) | (0.818) | (0.377) | (0.365) | (1.407) | (0.533) |
| Low Noise | | | | | | | 4.481 | 11.558 | 6.588 |
| | | | | | | | (0.298) | (1.456) | (0.499) |
| Observations | 5327 | 1344 | 6671 | 6019 | 1532 | 7551 | 7393 | 2660 | 10053 |
| R-squared | | | | | | | | | |
| Low=High(Beliefs) | 0.259 | 0.160 | 0.075 | 0.135 | 0.791 | 0.321 | 0.957 | 0.948 | 0.692 |
| Low=High(Noise) | | | | | | | 0.116 | 0.490 | 0.130 |

Robust Standard Errors in Parentheses. p<0.1*, p<0.05**, p<0.01

In columns 4 to 6 I include systematic multiple switchers into the sample and estimate both beliefs and noise separately by treatment. Doing so maintains similar treatment effects in versions 1 and 2 but lowers the treatment with larger statistical significance (p-value of .135), but eliminates much of the effect in version 3.

The estimates of the noise parameter by treatment show that noise seems to be higher in the high access condition than the low. Column 4 for versions 1 and 2 finds a difference of .1 for the noise parameter (p-value of 0.85) and a difference of 1.5 in column 5 (p-value of .12) in version 3 alone. When I pool the results in column 6 from all versions, the difference in noise becomes .6 (p-value of .22).

In columns 7 to 9 I include all participants (unique switchers and multiple-switchers) and estimate both beliefs and noise separately by treatment. Doing so eliminates the treatment effects completely in all specifications as the point estimates are close to 0 with p-values close to 1 in columns 7 and 8 (where I examine treatment effects separately by version) and a p-value of 0.6 for the estimation in column 9 where I pool the treatments

I conducted the same analysis on the full sample of data - single and multiple switchers - as well and estimate both beliefs and noise separately by treatment. Doing so eliminates the treatment effects completely in all specifications as the point estimates are close to 0 with p-values close to 1 where I examine treatment effects separately by version) and a p-value of 0.7 for the estimation in column 6 where I pool the versions.

### 1.7.9 Robustness Checks

**Table 1.8.** Probability of Violating SEU

| | Below 50 perc | | Below 46 perc | |
|---|---|---|---|---|
| | Without Controls | With Controls | Without Controls | With Controls |
| | (1) | (2) | (3) | (4) |
| Low Access | 0.049 | 0.052 | -0.001 | 0.007 |
| | (0.052) | (0.052) | (0.046) | (0.047) |
| version 2 = 1 | | -0.084 | | -0.176 |
| | | (0.128) | | (0.113) |
| version 3 = 1 | | -0.186*** | | 0.049 |
| | | (0.061) | | (0.054) |
| India = 1 | | 0.130 | | 0.142* |
| | | (0.079) | | (0.077) |
| Constant | 0.478 | 0.462 | 0.270 | 0.134 |
| | (0.038) | (0.073) | (0.033) | (0.063) |
| Observations | 368 | 354 | 368 | 354 |
| R-squared | 0.002 | 0.060 | 0.000 | 0.060 |
| P-value TE = 0 | 0.351 | 0.319 | 0.979 | 0.876 |

Robust Standard Errors in Parentheses. p<0.1*, p<0.05**, p<0.01***. This table examines whether treatment had an impact on the probability that beliefs are above and below different cutoffs to see if there might be unobserved selection into these subsamples. Columns 1 and 2 examine the 50% cutoff, while columns 3 and 4 examine the 46% cutoff. The reason for using two cutoffs is that because the belief elicitation is discrete, the 50% cutoff may exclude some individuals who had beliefs of exactly 50%. Therefore, I also examine effects at the 46% cutoff.

**Table 1.9.** Main Treatment Effects - UCSD Only

| | Dependent Variable : Beliefs of a High Payout from Subjective Lottery | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | full distribution | | 50 percent partition | | | | 46 percent partition | | | |
| | | | Above 50 | | Below 50 | | Above 46 | | Below 46 | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
| Low Access | -0.013 | -0.014 | -0.030* | -0.030* | 0.004 | 0.003 | -0.023 | -0.023* | -0.001 | 0.005 |
| | (0.017) | (0.017) | (0.017) | (0.017) | (0.016) | (0.013) | (0.015) | (0.014) | (0.023) | (0.018) |
| Version 3 | | 0.045** | | 0.105*** | | -0.102*** | | 0.109*** | | -0.114*** |
| | | (0.018) | | (0.015) | | (0.017) | | (0.014) | | (0.019) |
| Female (=1) | | -0.032 | | -0.004 | | -0 | | -0.007 | | 0.027 |
| | | (0.017)* | | (0.016) | | (0.014) | | (0.013) | | (0.019) |
| Uncertainty Class (=1) | | 0.033 | | 0.014 | | 0.026 | | 0.011 | | 0.028 |
| | | (0.022) | | (0.022) | | (0.017) | | (0.019) | | (0.029) |
| 2nd year student (=1) | | 0.034 | | 0.025 | | 0.003 | | 0.027 | | 0.012 |
| | | (0.022) | | (0.019) | | (0.018) | | (0.016) | | (0.025) |
| 3rd year student (=1) | | 0.042* | | 0.037* | | -0.008 | | 0.037** | | -0.002 |
| | | (0.024) | | (0.021) | | (0.018) | | (0.018) | | (0.024) |
| 4th year student (=1) | | 0.001 | | 0.015 | | 0.012 | | -0.006 | | 0.005 |
| | | (0.026) | | (0.023) | | (0.022) | | (0.020) | | (0.035) |
| Constant | 0.521 | 0.492 | 0.629 | 0.551 | 0.406 | 0.440 | 0.580 | 0.517 | 0.343 | 0.375 |
| | (0.013) | (0.019) | (0.013) | (0.019) | (0.012) | (0.013) | (0.011) | (0.015) | (0.017) | (0.022) |
| Observations | 274 | 265 | 141 | 136 | 133 | 129 | 208 | 201 | 66 | 64 |
| R-squared | 0.002 | 0.067 | 0.021 | 0.301 | 0.001 | 0.301 | 0.011 | 0.293 | 0 | 0.374 |
| P-value TE = 0 | 0.439 | 0.415 | 0.084 | 0.078 | 0.787 | 0.809 | 0.131 | 0.085 | 0.959 | 0.786 |
| controls | No | Yes | No | Yes | No | Yes | No | Yes | No | Yes |
| Low Access Obs | 138 | 133 | 71 | 68 | 67 | 65 | 106 | 102 | 32 | 31 |
| High Access Obs | 136 | 132 | 70 | 68 | 66 | 64 | 102 | 99 | 34 | 33 |

Robust Standard Errors in Parentheses. p<0.1*, p<0.05**, p<0.01***. All regressions estimated using OLS. The dependent variable measures beliefs of the probability of a $30 payout from the subject lottery. The treatment effect is indicated by the variable "Low Access" representing participants who had only a 10% chance of the subjective lottery. The above 50% partition represents individuals who have beliefs above 50% The 46% partition represents individuals who we cannot rule out have beliefs equal to or above 50%. The controls used in the regression include Gender, dummies for year of college, and whether the participant took the decisions-under-uncertainty class.

**Table 1.10.** Quantile Regressions of Treatment Effects

| | Dependent Variable : Beliefs of a High Payout from Subjective Lottery | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | 10 | 20 | 30 | 40 | 50 | 60 | 70 | 80 | 90 |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| Low Access | 0.000 | -0.000 | -0.000 | 0.000 | -0.000 | -0.040*** | -0.036*** | -0.033* | -0.050** |
| | (0.016) | (0.011) | (0.000) | (0.017) | (0.029) | (0.012) | (0.000) | (0.017) | (0.025) |
| version 2 = 1 | -0.050 | 0.093*** | 0.093*** | 0.107*** | 0.071 | 0.071** | 0.071*** | 0.033 | 0.193*** |
| | (0.036) | (0.026) | (0.000) | (0.039) | (0.069) | (0.030) | (0.000) | (0.040) | (0.064) |
| version 3 = 1 | -0.175*** | -0.082*** | -0.011*** | 0.025 | 0.061* | 0.092*** | 0.118*** | 0.137*** | 0.196*** |
| | (0.019) | (0.014) | (0.000) | (0.020) | (0.034) | (0.014) | (0.000) | (0.020) | (0.028) |
| India = 1 | 0.000 | -0.093*** | -0.093*** | -0.107*** | -0.071 | -0.071*** | -0.071*** | -0.017 | -0.093** |
| | (0.024) | (0.016) | (0.000) | (0.028) | (0.047) | (0.019) | (0.000) | (0.026) | (0.044) |
| Uncertainty Class (=1) | 0.050** | 0.050*** | 0.050*** | 0.036 | 0.000 | 0.040** | 0.036*** | 0.033 | 0.050 |
| | (0.024) | (0.016) | (0.000) | (0.027) | (0.043) | (0.018) | (0.000) | (0.025) | (0.035) |
| 2nd year student (=1) | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | -0.000 | 0.000*** | 0.017 | 0.050* |
| | (0.020) | (0.014) | (0.000) | (0.022) | (0.037) | (0.015) | (0.000) | (0.022) | (0.029) |
| 3rd year student (=1) | 0.000 | -0.000 | 0.000 | -0.000 | 0.000 | 0.010 | -0.000*** | -0.000 | 0.050 |
| | (0.020) | (0.014) | (0.000) | (0.023) | (0.039) | (0.016) | (0.000) | (0.023) | (0.035) |
| 4th year student (=1) | 0.000 | 0.000 | 0.000 | -0.036 | 0.000 | -0.031 | -0.014*** | -0.017 | 0.000 |
| | (0.028) | (0.020) | (0.000) | (0.030) | (0.049) | (0.021) | (0.000) | (0.028) | (0.037) |
| Constant | 0.425 | 0.425 | 0.475 | 0.475 | 0.525 | 0.556 | 0.561 | 0.558 | 0.575 |
| | (0.017) | (0.013) | (0.000) | (0.022) | (0.040) | (0.017) | (0.000) | (0.022) | (0.030) |
| Observations | 354 | 354 | 354 | 354 | 354 | 354 | 354 | 354 | 354 |
| P-value TE = 0 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000 | 0.001 | . | 0.054 | 0.045 |

Robust Standard Errors in Parentheses. $p<0.1$*, $p<0.05$**, $p<0.01$*** . Each column estimates quantile treatment effects in intervals of .1 (deciles). The dependent variable is beliefs of the probability of a high outcome from the subjective lottery and the coefficient of interest is the treatment indicator, ”low access” (row 1), which represents individuals who had a 10% chance of the subjective lottery. As shown in columns 1 to 5, up to the median quantile, the treatment effects are zero. For quantiles past the median, having low access reduces beliefs relative to having high access. (columns 6 to 9).

**Table 1.11.** Main Treatment Effects Using Indifference Row

| | Dependent Variable : Beliefs of a High Payout from Subjective Lottery | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | full distribution | | 50 percent partition | | | | 46 percent partition | | | | By Version | |
| | | | Above 50 | | Below 50 | | Above 46 | | Below 46 | | | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) | (12) |
| Low Access | -0.441 | -0.406* | -0.599* | -0.457* | 0.093 | 0.120 | -0.612* | -0.459** | 0.008 | 0.028 | | |
| | (0.292) | (0.240) | (0.356) | (0.240) | (0.222) | (0.181) | (0.320) | (0.201) | (0.302) | (0.214) | | |
| version 2 = 1 | | 1.061* | | 0.921 | | 0.518 | | 0.352 | | -0.144 | 0.269 | 1.147 |
| | | (0.615)* | | (0.717) | | (0.434) | | (0.545) | | (0.492) | (0.833) | (1.003) |
| India = 1 | | -0.851 | | -0.133 | | -0.434 | | -0.103 | | -0.140 | | -0.843 |
| | | (0.449)* | | (0.377) | | (0.349) | | (0.364) | | (0.309) | | (0.448)* |
| Low Access × version 1 | | | | | | | | | | | -0.233 | -0.275 |
| | | | | | | | | | | | (0.235) | (0.246) |
| Low Access × version 2 | | | | | | | | | | | -0.475 | -0.449 |
| | | | | | | | | | | | (0.852) | (0.912) |
| Low Access × version 3 | | | | | | | | | | | -0.520 | -0.501 |
| | | | | | | | | | | | (0.397) | (0.396) |
| Constant | 12.857 | 11.353 | 14.876 | 12.186 | 10.647 | 9.795 | 13.838 | 11.559 | 10.198 | 8.583 | 11.000 | 11.280 |
| | (0.219) | (0.280) | (0.247) | (0.267) | (0.162) | (0.203) | (0.236) | (0.225) | (0.215) | (0.296) | (0.174) | (0.287) |
| Observations | 368 | 354 | 183 | 176 | 185 | 178 | 269 | 258 | 99 | 96 | 368 | 354 |
| Controls | No | Yes | No | Yes | Yes | Yes | No | Yes | Yes | Yes | No | Yes |
| R-squared | 0.006 | 0.390 | 0.015 | 0.670 | 0.001 | 0.353 | 0.014 | 0.690 | 0.000 | 0.490 | 0.357 | 0.390 |
| P-value TE = 0 | 0.132 | 0.091 | 0.094 | 0.058 | 0.677 | 0.507 | 0.057 | 0.023 | 0.979 | 0.897 | 0.962 | 0.958 |
| Low Access Obs | 190 | 183 | 90 | 86 | 100 | 97 | 139 | 133 | 51 | 50 | | |
| High Access Obs | 178 | 171 | 93 | 90 | 85 | 81 | 130 | 125 | 48 | 46 | | |
| P-value v1 = v2 = v3 = 0 | | | | | | | | | | | 0.392 | |
| P-value v1 = v2 = v3 | | | | | | | | | | | 0.810 | 0.883 |
| P-value v1 = v2 | | | | | | | | | | | 0.784 | 0.853 |
| P-value v1 = v3 | | | | | | | | | | | 0.536 | 0.627 |
| P-value v2 = v3 | | | | | | | | | | | 0.962 | 0.958 |
| Version 1 Obs | | | | | | | | | | | 146 | 142 |
| Version 2 Obs | | | | | | | | | | | 30 | 29 |
| Version 3 Obs | | | | | | | | | | | 192 | 183 |
| Lower Lee bound | -1.071 | | -0.062 | | -0.513 | | -1.300 | | -0.557 | | | |
| Lower Lee Bound p-value | 0.006 | | 0.000 | | 0.113 | | 0.002 | | 0.190 | | | |
| Upper Lee bound | 0.009 | | -0.018 | | 0.593 | | -0.184 | | 0.583 | | | |
| Upper Lee Bound p-value | 0.980 | | 0.330 | | 0.037 | | 0.653 | | 0.099 | | | |

Robust Standard Errors in Parentheses. p<0.1*, p<0.05**, p<0.01***. All regressions estimated using OLS. The dependent variable measures beliefs of the probability of $30 payout from the subjective lottery. The treatment effect is indicated by the variable "Low Access" representing participants who had only a 10% chance of the subjective lottery. The above 50% partition represents individuals who have beliefs above 50% The 46% partition represents individuals who we cannot rule out have beliefs equal to or above 50%. The controls used are Gender, dummies for year of college, and whether the participant took the decisions-under-uncertainty class.

# Chapter 2

# Access-based Beliefs: Field Evidence and Applications

## 2.1 Introduction

The theory of access-based beliefs proposes that beliefs of an alternative's value are based on ones level of access. In particular, the hypothesis is that low access to an alternative leads to lower beliefs of its value, consistent with the notion of sour grapes. The first chapter of the dissertation provided a theoretical framework for studying access-based beliefs, and a proof-of-existence in a lab setting. The key improvement from past psychology work was to completely control for learning channels and measure beliefs through incentivized choices over lotteries. This chapter aims to demonstrate how access-based beliefs may be relevant in real-world settings, and to present evidence from a randomized field-experiment of sour grapes.

To test the importance of access-based beliefs in an applied setting, I repurpose an existing field experiment conducted in Los Angeles public high-schools in which access to a commercially available SAT prep package was randomly assigned to be either 25% (low access) or 75% (high access) for each student within a classroom (Bursztyn, Egorov, and Jensen, 2017).[1] The design is a close, real-world analogue of the lab experiment. After

[1]The paper's original objective was to distinguish between models of social stigma in an environment where effort is stigmatized and another in which ability is rewarded. The critical test was the opposite sign prediction in the two environments of the impact of access on package signup conditional on the decision

privately learning their access level, students were asked whether they would like to signup for a chance to receive the package for free, and immediately after, were surveyed about their expected returns of the package (the number of points they think their score would increase if they got the package). In analyses not done in the original paper, I find that a reduction in access by 50 p.p. lowers mean beliefs of expected returns to the package by 125 points (p = .017) (the SAT was out of 2400 points at the time). For robustness against outliers, I use quantile regression and find that a reduction in access lowers median expected returns by 100 SAT points from a baseline of 200 in the high access condition (p ¡ .001). The effect on the median represents a 0.2 standard deviation shift of the belief distribution.[2]

Overall, the field results confirm the lab findings of the first chapter qualitatively (evidence of sour grapes) and are similar quantitatively (an effect size of $\approx 0.2$ standard deviations in both settings, although the change in access was 10% larger in the lab). The two serve as complementary sets of evidence. The lab experiment controls for learning confounds in the field, and measures belief changes that impact outcomes. The field experiment tests the theory in a setting with real-world relevance and asks individuals explicitly what their expected returns to the package are, thus providing more evidence that belief changes drive the mechanism.

A second hypothesis tested in this chapter is that access-based beliefs are most prevalent for the poor given that they are likely to face the severest forms of low access to critical investments. Given their environments, the poor may be the most likely to use sour grapes as a coping mechanism to deal with lack of accessibility. There is some suggestive evidence that this may be the case. A long-standing development puzzle is that despite

---

to sign up (and a corresponding diagnostic test score conditional on winning the package) being made public. As a control group, half the participant's decisions were kept private and the paper's theoretical model predicts no effect of access on take up in this private condition. The examination of access on beliefs is thus orthogonal to their paper. I thank the authors profusely for sharing their data with me.

[2]It should be noted that while the belief elicitations in this study are un-incentivized, the median expected returns of 100 SAT points is in line with what SAT prep companies advertise and a 1% increase in beliefs is associated with a 3 to 5 p.p. increase in take-up depending on specification.

the high estimate returns to technologies like education, healthcare and migration, poor people appear to underinvest (A. V. Banerjee, A. Banerjee, and Duflo, 2011). A potential explanation that has been proposed is that underinvestment reflects low beliefs of the returns on investment, a claim that has been partially substantiated in the literature on subjective expectations (David McKenzie, 2007; Delavande, 2008; Delavande, Giné, and McKenzie, 2011; Jensen, 2010). The standard explanation for low subjective expectations is that individuals simply have the wrong information or biased sources of information. However, recent field-experiments find that incentives tailored to education increase beliefs of the returns to education, despite no information being provided (Duflo, Dupas, and Benhassine, 2013; S. Sequeira, 2013). While there may be several reasons for this that have to do with learning or signaling, access-based beliefs provides a potential explanation.

To test the hypothesis that access-based beliefs are more likely to be observed amongst the poor, both lab and field experiments were conducted in two distinct socioeconomic settings, low-income areas and high-income areas. The lab experiment was conducted in UC San Diego and a college in Bangalore, India. Compared to UCSD, students at the Indian college come from an environment where the income distribution has a higher variance and lower average. While I find significant evidence of a sour grapes effect in both locations, I find suggestive evidence that the effect is larger in the Indian sample compared to UCSD by 5.4 p.p. and within India, larger amongst the poor (but the results are not significant at conventional levels).[3] Much more strongly, in the field data, I find that low access has a 40% larger negative impact on median beliefs in the low-income school compared to the high-income one (p ¡ .001). Additionally, when decisions are kept private, low access reduces signup for the package lottery by 18 p.p. ($p = .039$) for a sample of high achievers within the low-income school.[4] Taken together, the results suggest that sour grapes may be more prevalent amongst the poor, which may be particularly harmful

[3]For example, a halving of parental income increases the sour grapes effect by 5.6 p.p. ($p = .198$).

[4]Indeed, demand for the package was highest for high achievers in the low-income school compared to other groups, suggesting a potential reason why sour grapes is more likely to be seen amongst this group.

since the poor may benefit the most from important investments.

The findings of the second chapter relate to several emerging psychological theories of poverty. A theoretical literature argues that the poor may suffer from feelings of hopelessness (Lybbert and Wydick, 2018) and sub-optimal aspirations (Genicot and Ray, 2017) that could perpetuate the cycle of poverty. Indeed, there is growing evidence that both aspiration-improving interventions and traditional anti-poverty programs create more optimism about the future, lift people up psychologically, and improve economic outcomes (A. Banerjee, Duflo, Goldberg, Karlan, Osei, Parienté, Shapiro, Thuysbaert, and Udry, 2015; Beaman, Duflo, Pande, and Topalova, 2012; Bernard, Dercon, Orkin, and Taffesse, 2014; Lybbert and Wydick, 2017). This dissertation contributes by uncovering one particular belief channel that might be crucial for understanding the positive and wide-ranging psychological effects observed in these interventions.

The rest of this chapter is organized as follows: In Section 2.2, I discuss the context and experimental design and in Section 2.3 I present the results from the field experiment and the heterogeneity the lab and field results by indicators of poverty; in Section 2.4, I review evidence from literature on education and health behavior that provide suggestive evidence of access-based beliefs and in section 2.5, I conclude with a discussion of the potential policy implications.

## 2.2 Context and Experimental Design

To test our model of access-based beliefs in the field, I repurpose a field experiment in which students in three Los Angeles public high-schools were offered a commercially available SAT prep package for free (Bursztyn, Egorov, and Jensen, 2017). The original purpose of the paper was to test between different models of social stigma. To this end, the authors cross two treatments: a public vs. private treatment in which sign up choices could either be revealed or not revealed to other students in the room and a low vs. high probability treatment in which the probability of receiving the prep package was varied between 25% and 75% conditional on sign up. In addition, for those who win the lottery, the diagnostic SAT score would be revealed for those in the public information treatment. The paper uses the sign of the predicted effect of access on package take up in the public and private condition to distinguish between models of social stigma.

The package included an online app, a diagnostic test, and one-on-one tutoring. Students were given a description of the package and then offered the following lottery:

*If you choose to sign up, your name will be entered into a lottery where you have a 25% (75%) chance of winning the package. Would you like to sign up for a chance to win the SAT prep package? (Please pick one option) Yes / No*

The treatment was varied within classrooms to control for classroom fixed effects, and probabilities of obtaining the package were privately known to each student. Students also answered several survey questions including one about their beliefs about the expected returns of the package. Specifically, after knowing their treatment condition (low or high access), but before uncertainty was resolved, students were asked the un-incentivized question: "How many points do you think this SAT prep package could improve your SAT test scores by?" Below, I utilize responses to this question and the random variation in access to test for access-based beliefs in the field.

### 2.2.1 Connecting the lab and field

While not designed specifically to test access-based beliefs,[5] this field experiment is a close real-world analogue of my lab experiment and presents an opportunity to test the theory in the field. In both experiments, ex-ante beliefs are elicited about the value of

[5]The paper varied the probability of receiving treatment in order to distinguish between two different mechanisms of social stigma, one in which academic effort is frowned upon and another in which low academic ability comes with a social cost.

an alternative for which access varies, and in both cases, the alternative is an investment whose value likely depends on beliefs of its expected returns. There are some notable differences between the two experiments. (1) In the lab, everyone participates in the lottery, while in the field, participation is costless in terms of economic resources, but is a choice variable. (2) In the lab, access varies between 10% and 70%, while the variation is 25% and 75%, slightly less in the field. (3) In the lab, beliefs are elicited indirectly through incentivized choices in the lab experiment while they are elicited directly in the field.

Two additional features of their design are orthogonal to access-based beliefs but end up being important in analyzing the data: (1) The authors crossed the probability treatment (low or high access) with another treatment in which sign up choices could either be revealed or not revealed to other students in the room (public and private).[6] While orthogonal to access-based beliefs, this impacts the analysis I conduct on the effect of access on signing up for the package, but not on the analysis of the effect of treatment on beliefs, since beliefs are not revealed to others. (2) The experiment was conducted in two types of schools, labeled ”smart-to-be-cool” and ”cool-to-be-smart”. In smart-to-be-cool schools, the authors expect greater stigmatization of academic effort and thus chose a lower achieving school with a high share of minority students. In cool-to-be-smart schools, the authors expect that signaling high ability was more likely to be important and thus chose a higher achieving schools with lower minority shares.[7] This impacts the analysis for testing access-based beliefs and the sour grapes hypothesis because it is less likely to be prevalent in schools where academic effort may be stigmatized.

Given this background and experimental design, I proceed to conduct two tests of

[6]In addition, for those who win the lottery, the diagnostic SAT score might would be revealed for those in the public information treatment. The paper uses the sign of the predicted effect of access on package take up in the public and private condition to study a model of social stigma.

[7]The authors also provide subsequent survey evidence confirming that these two types of schools do indeed differ in ways that their model and tests are intended to highlight. Students in cool-to-be smart schools are much more likely to agree that being seen as smart is important for being popular in their school. The difference is large, about 40% of the standard deviation in responses, and statistically significant. Students in cool-to-be-smart schools are also more likely to say that if classmates become more popular because they are studying hard, it is because other students admire hard workers or smart people.

access-based beliefs (1) That access to the prep package influences beliefs of its value. (2) That access to the prep package affects actual decisions and outcomes by examining the impacts on signup for the package.

## 2.3 Results

### Impact of Access on Expected Returns of the SAT Package

To provide a graphical illustration of the treatment effects, Figure 2.1 plots the density estimates of expected package returns for the low and high access conditions. Most of the mass in both distributions is concentrated between 0 to 500 points, with about 75% of the data below 300 SAT points. The median expected return across both treatments is 100 SAT points which is reassuring since most SAT prep package companies advertise that their programs can increase scores by 100 points on average. A visual test of the treatment effect in the range of reasonable expected returns, below 500 points, shows a clear rightward shift of the distribution for the high probability group, indicating a sour grapes effect. A two-sample KS test and a Wilcoxon rank-sum test strongly reject the equality of the overall distributions with p-values of .066 and .021 respectively.

Both distributions have long right tails, in fact, approximately 15% of students believe their score can go up by 500 points or more, while 8% think their scores can go up 1000 points or more, improvements which any prep package may struggle to achieve. Thus, the appropriate regression analysis is likely a quantile regression. In table 2.1, I examine both average and quantile treatment effects of the impact of access on expected returns by estimating equation 2.1 below:[8]

$$\text{Expected Returns}_i = \alpha + \beta \cdot \text{Low Access}_i + \boldsymbol{\gamma} \cdot \boldsymbol{x}_i + \boldsymbol{\nu} \cdot \boldsymbol{z} + \epsilon_i \tag{2.1}$$

[8]This analysis that was not carried out in the original paper, but is mostly orthogonal to the findings in the original paper.

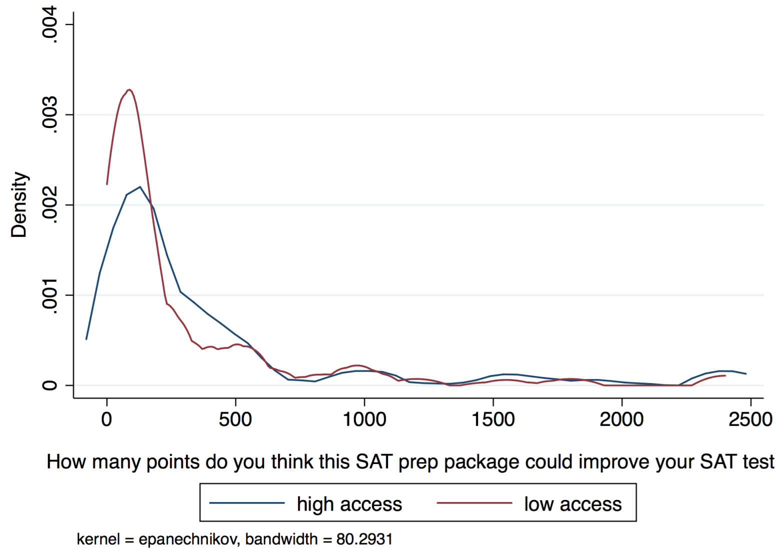


**Figure 2.1.** Kernel Density Estimates of Beliefs of SAT Prep Package Effectiveness By Treatment

,where $\boldsymbol{x}$ indicates student demographic characteristics that include age, ethnicity, gender and $\boldsymbol{z}$ indicates classroom fixed effects.

Column 1 of table 2.1 presents the results for the full sample of individuals excluding two individuals who believed their score would go up by more than the SAT score range.[9] The coefficient on low access is -102 points (p-value of .079) from a baseline of approximately 400 in the high access group. In column 2 , I control for student demographics including age, ethnicity, gender, and classroom fixed effects and I find a slightly larger coefficient (-125 SAT points) and a larger significance level (p-value of .017). These indicate sour grapes effects of about 0.2 standard deviations of the total expected returns distribution, a similar effect size to what I find in the lab and strongly significant.

Given the long right tails of the belief distribution, which reflect some unrealistic expectations of the returns, I examine quantile regressions (columns 3 to 5 of table 2.1) of the treatment effects at the 25th, 50th, and 75th quantiles. I find consistent evidence of a sour grapes effect at all these quantiles. At the 25th quantile, I find a sour grapes effect of

[9]There were only two such outliers, 2800 and 11000

about 25-30 SAT points, significant with and without controls. At the median, they are larger - between 50 and 100 SAT points (significant at the 1% level) - and they remain large but only marginally significant at the 75th quantile (between 100 and 120 points). Overall, the results show fairly convincing impacts of low access on beliefs in a real-world setting in the direction of sour grapes.[10] This concern is unwarranted as a regression of take-up on beliefs find no significant effects.

### Treatment Effects on Take-up

While I find that changing access to the prep package impacts beliefs of its return, an important question is whether or not there is any effect on take-up. This may be unlikely in this context for a few reasons: (1) The package was provided at no charge (2) For students in the condition where signup decisions were kept private, take-up was 80% - an already high number (3) Take-up is likely a more complex function of beliefs and preferences, as demonstrated in the original paper. Therefore, the null hypothesis is that access-based beliefs will not impact take-up in this setting.

The findings are consistent with the null. I show results in appendix **??** where I run OLS on equation (1.3) with an indicator for lottery signup as the dependent variable. Analysis is restricted to the private condition because the public treatment introduces a social stigma channel which affects the way access influences take-up. In column 1, I find that the coefficient on take-up is 1.4 p.p. (p-value of .781) and barely changes when I add controls in column 2. In columns 3 and 4, I restrict analysis to the where there is less social stigma for academic effort and so there might be a higher likelihood of observing a sour grapes effect. However, I find no effects in the restricted sample either. One possible explanation is that in richer schools, students with low access know they will be able to afford the package even if they lose the lottery, and so, do not respond to treatment.

[10] One robustness check is that since beliefs were elicited after the take-up decision (but before students knew if they got the package) there could some effect of take-up on beliefs. This would be hard to identify since take-up is not randomly assigned, is an outcome variable, and beliefs are also outcome variables.

**Table 2.1.** Field Evidence: Impact of Access on Beliefs

| | Dependent Variable : Expected Returns of Package (units - SAT Points) | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Mean Effects | | Quantile Effects | | | | | |
| | | | 25th percentile | | 50th percentile | | 75th percentile | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Low Access | -102* | -125** | -30*** | -25*** | -100*** | -50*** | -120 | -100 |
| | (58) | (52) | (9) | (5) | (15) | (1) | (75) | (80) |
| Hispanic = 1 | | -9 | | -50*** | | -50*** | | -0 |
| | | (61) | | (7) | | (2) | | (87) |
| Male = 1 | | -39 | | -5 | | -0 | | -100 |
| | | (55) | | (5) | | (1) | | (81) |
| Student age | | -4 | | 20*** | | 39*** | | -33 |
| | | (45) | | (5) | | (1) | | (77) |
| Constant | 399 | 509 | 90 | -165 | 200 | -413 | 420 | 1033 |
| | (45) | (743) | (7) | (80) | (11) | (19) | (54) | (1268) |
| Observations | 316 | 316 | 316 | 316 | 316 | 316 | 316 | 316 |
| R-squared | 0.010 | 0.240 | | | | | | |
| Student Demographics | No | Yes | No | Yes | No | Yes | No | Yes |
| Classroom FE | No | Yes | No | Yes | No | Yes | No | Yes |
| P-value: Low Access = 0 | 0.079 | 0.017 | 0.001 | 0.000 | 0.000 | 0.000 | 0.110 | 0.213 |

Robust Standard Errors in Parentheses. p<0.1*, p<0.05**, p<0.01*** . Coefficients in columns (1) and (2) are from OLS regressions, while coefficients from Columns (3) to (8) are quantile regressions. The dependent variable is the answer to the question, "How many points do you think this SAT prep package could improve your SAT test scores by?" and can only range from 0 - 2400. The coefficient of interest is "low access" which is an indicator for individuals who had only a 25% chance of winning the SAT prep package.

Overall it seems that while the expected returns of the SAT prep package were influenced by access, there is no indication that take up was impacted.

### 2.3.1 Access-based Beliefs and Poverty

A deeper understanding of the theory may be gained from examining whether individuals who face severely low access to desirable alternatives have a lower or higher prevalence of access-based beliefs. One could imagine that either (a) greater exposure to low access pushes individuals to find coping mechanisms, or (b) individuals learn to live with low access over time and therefore rely less on sour grapes to feel better.

Some real-world evidence suggests the possibility of the former. Surveys in developing countries find that poor people have low beliefs of the returns to investing in education, health and migration when compared to estimates from data (Attanasio and Kaufmann, 2012; David McKenzie, 2007; Jensen, 2010) and recent field experiments in development contexts find that incentives tailored to education increased beliefs of the returns to education, despite no specific information about the returns being provided (Duflo, Dupas, and Benhassine, 2013; S. Sequeira, 2013), though mechanisms were unexplored. In this section, heterogeneous treatment effects by indicators of poverty are examined in both the lab and field. Overall, I find evidence that the sour grapes effect is larger in poorer environments and greater amongst poorer individuals in both the lab and field, with the field evidence being stronger than the lab.

**Sour Grapes Heterogeneity By Poverty Indicators: Lab Evidence**

The primary reason for conducting the study in a developing country was to test access-based beliefs in an environment where low access to alternatives may be more prevalent. Columns 1 and 2 of table 2.2 show the main difference in treatment effects between the Indian and the UCSD sample. The coefficient on the interaction term between low access and an India indicator is -4.5 p.p. (p-value of .254) without controls (column 1) and -5.2 p.p. (p-value of .204) with controls (column 2). These results are not significant, but are directionally consistent with a larger sour grapes effect in the environment with greater poverty and economic disparity. Further data on self-reported parental income, parental education and student allowance was collected in the Indian sample providing some idea of the socioeconomic background of subjects.[11] Despite being underpowered

[11]Specifically, I asked subjects to report (1) the estimate of their parents total income last year before taxes (2) the education level of the highest earning member of their household (3) their allowance per

due to a small sample, the point estimates on interaction terms between treatment and each of these demographic characteristics are large and negative, showing a consistent pattern of a greater sour grapes effect for poorer individuals (columns 3 to 6, Table 2.2) within India.[12]

**Table 2.2.** Treatment Effects By Poverty Indicators : Lab Results

| | Dependent Variable : Beliefs of a High Payout from Subjective Lottery | | | | | |
|---|---|---|---|---|---|---|
| **Covariate :** | Developing Country | | Log Parental Income | Allowance ¡ 40th percentile | Household Head Graduate | Poverty Index |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Low Access | -0.013 | -0.013 | -0.197* | -0.040 | -0.065 | -0.036 |
| | (0.017) | (0.017) | (0.115) | (0.060) | (0.051) | (0.059) |
| Covariate | -0.007 | -0.033 | -0.027 | -0.009 | -0.015 | -0.032 |
| | (0.030) | (0.038) | (0.032) | (0.064) | (0.099) | (0.041) |
| Low Access×Covariate | -0.045 | -0.052 | 0.056 | -0.060 | -0.076 | -0.067 |
| | (0.040) | (0.041) | (0.043) | (0.092) | (0.120) | (0.061) |
| Constant | 0.521 | 0.519 | 0.629 | 0.611 | 0.588 | 0.644 |
| | (0.013) | (0.019) | (0.129) | (0.116) | (0.085) | (0.054) |
| Observations | 368 | 354 | 72 | 72 | 72 | 72 |
| R-squared | 0.020 | 0.071 | 0.182 | 0.165 | 0.165 | 0.207 |
| Controls | No | Yes | Yes | Yes | Yes | Yes |
| P-value Interaction = 0 | 0.254 | 0.204 | 0.198 | 0.512 | 0.525 | 0.273 |

Robust Standard Errors in Parentheses. p<0.1*, p<0.05**, p<0.01*** . All regressions estimated using OLS and are conducted only on version 3 of the experiment which was done in both India and the US. All poverty indicators are based on subject self-reports. The column headings indicate the covariate for which I analyze heterogeneous treatment effects. In (1) and (2), the covariate is an indicator for the experiment being conducted in India. In (3), it is an estimate of parents total income last year before taxes. In (5) it is the education level of the highest earning member of their household. In (4), it is whether the allowance (aside from boarding and lodging) is below the 40th percentile of the allowance distribution which is $60. In (6) it is a a poverty index ranging from 0 to 3 constructed from columns 3 to 5. Due to some subjects not responding to the parental income question, we lose about 20% of the data in the Indian sample. The coefficient of interest is the interaction term in row 3. The controls used include gender, dummies for year of college, and whether the participant took decisions-under-uncertainty class.

### Sour Grapes Heterogeneity By Poverty Indicators: Field Evidence

Unlike the lab experiment, the field experiment did not collect self-reported measures of income and economic background. However, as mentioned earlier, the authors selected one lower-income and lower-achieving school - the smart-to-be-cool school, and two higher-

month aside from boarding and lodging costs.

[12]In all specifications, I interact treatment with one of these indicators while controlling for the others. The correlation between these measures is at most .38 (between parental income and education level of the household head) and only .16 (between monthly allowance and education level of the household head) suggesting these proxies are independent measures of poverty.

income and higher- achieving schools - the cool-to-be-smart schools.[13]

In table 2.3, column 1, I examine the heterogeneous impacts of low access on median expected returns of the SAT package by an indicator for the poorer school. I find that the interaction term has a coefficient of -83.33 (p-value of .000). In column 3, I conduct the same analysis but with package signup as the outcome variable. The coefficient on the interaction term is -1.0 p.p., directionally consistent with sour grapes, but insignificant. Overall, these results are consistent with the lab: the grapes effect is larger in the poorer environment.

In column 2, I focus only on poor schools and examine the treatment effects by whether a student's grades were below or above the median.[14] The reason being that high achieving students from poor schools are likely to have a higher demand for the SAT prep package since they appear to care about grades and are poorer than their rich-school counterparts, and so may be more likely to respond to the treatment. Indeed, conditional on being in the poor school and in the private treatment condition, the signup rates of high achieving students are 24% higher than low achievers (p-value of .000). I first examine heterogeneity of the treatment effects by self-reported achievement on median expected returns (column 2). The interaction term is -30, showing greater sour grapes for the high achievers, but this effect is not significant.

In columns 4 to 6, I conduct the same analysis, but this time, on the dependent

[13]See Bursztyn, Egorov, and Jensen (2017) for specific ethnic, scholastic and income-based, differences between the two schools that strongly reflect the differences in poverty levels between the schools. As an example, the median income in the poor school ZIP code is about $44,000, and the average SAT score is around 1,200. By contrast, averaging across the second and third schools where it is cool-to-be-smart, the median income is about $66,000 and the average SAT score is around 1,500. They also compare these schools to a sample of138 LAUSD high schools and find that the first school appears close to the median public high school on poverty indicators while the latter two schools are in the bottom 5th percentile of poverty.

[14]The original paper uses this measure to examine heterogeneous treatment effects in the context of peer pressure, and tries to divide the data as close to the median response as possible. I follow the same as this is the only sensible partition. The partition is made at the 70th percentile of responses. Students above the 70th percentile report having Mostly A's or mostly A's and B's, while students below the 70th percentile report having Mostly B's and C's, C's and D's and D's and F's. There are only 5 students who report having mostly A's, and the modal response is mostly B's and C's, so the partition used is the only sensible one.

variable ”take-up”. In all specifications, I control for an indicator of the public treatment and the interaction between low access and public, because each of these variable has its own impact on signup, as theorized in the original paper. As seen in column 4, when I include all students in the poor school, regardless of grades, I find that low access has no impact on signup. I then run the same analysis separately for low and high achievers. I find that for low achievers, low access has no impact on signup (column 5), but for high achievers, who have a higher demand for the package, low access lowers take-up by 18 p.p. (p-value of .039) (column 6). This suggests that for this subsample of high achieving, poorer students, low access has an independent effect on take-up. This effect is orthogonal to the model of negative social stigma being tested, but is as large as some of the negative peer pressure effects found in the original paper. By contrast, access has no impact on signup rates of high achieving students from rich schools who were also in the private treatment.

Thus, it seems that amongst high achieving, poor students, there is evidence of sour grapes in actual take up decisions in a setting that was designed to test a completely different question. This result is consistent with the notion that poor people appear more likely to use sour grapes as a defense mechanism when faced with low access.

## 2.4 Evidence for Access-Based Beliefs in Applied Settings

While no prior study in economics has specifically tested the hypothesis of access-based beliefs, there is suggestive evidence in the form of experimental results and stylized facts that indicate its possible prevalence and implications for real-world problems.

### 2.4.1 Field Experiments in Education

In education, a couple of recent field-experiments have looked at the effects of education programs on perceptions of the returns to education. In particular, S. Sequeira

**Table 2.3.** Treatment Effects By Poverty Indicators: Field Results

| **Dependent Variable :** | Expected Returns Package | | Signup | | | |
|---|---|---|---|---|---|---|
| **Schools Included :** | All Schools | Only Poor Schools | All Schools | Only Poor Schools | | |
| **Covariate :** | Low-Income School | High Grades | Low-Income School | Public Condition | | |
| | | | | High & Low Grades | Low Grades | High Grades |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Low Access | -16.667* | -70.000*** | 0.011 | 0.010 | 0.017 | -0.179** |
| | (9.113) | (21.184) | (0.070) | (0.074) | (0.095) | (0.085) |
| Covariate | 91.667*** | 80.000*** | 0.006 | -0.165** | -0.087 | -0.448*** |
| | (9.767) | (28.567) | (0.082) | (0.081) | (0.101) | (0.124) |
| Low Access × Covariate | -83.333*** | -30.000 | -0.010 | -0.184 | -0.200 | 0.013 |
| | (12.097) | (42.024) | (0.101) | (0.114) | (0.143) | (0.178) |
| Male | -8.333 | -20.000 | -0.008 | -0.022 | 0.010 | -0.062 |
| | (6.029) | (19.083) | (0.050) | (0.057) | (0.072) | (0.101) |
| Student Age | 8.333 | 56.667*** | -0.146** | 0.020 | -0.025 | 0.171* |
| | (5.660) | (16.927) | (0.061) | (0.060) | (0.079) | (0.092) |
| Hispanic | -91.667*** | -20.000 | 0.032 | 0.095 | 0.125 | 0.129 |
| | (8.351) | (44.194) | (0.066) | (0.127) | (0.175) | (0.131) |
| Constant | 66.667 | -746.667 | 3.161 | 0.375 | 1.003 | -1.807 |
| | (92.236) | (276.426) | (0.992) | (0.987) | (1.294) | (1.537) |
| Observations | 316 | 181 | 257 | 257 | 184 | 72 |
| R-squared | | | 0.030 | 0.133 | 0.121 | 0.370 |
| Includes Public Treatment | Yes | Yes | No | Yes | Yes | Yes |
| Classroom FE | No | Yes | No | Yes | Yes | Yes |
| P-value Low Access = 0 | 0.068 | 0.069 | 0.871 | 0.891 | 0.862 | 0.039 |
| P-value Interaction = 0 | 0.000 | 0.476 | 0.921 | 0.108 | 0.165 | 0.943 |

Robust Standard Errors in Parentheses. p<0.1*, p<0.05**, p<0.01***. Columns 1 and 2 are quantile regressions of low access treatment, a covariate, and the interaction between the two on expected returns of the SAT prep package for both. Columns 3 to 6 are OLS regressions of low access treatment, a covariate, and the interaction between the two on signup for the SAT prep package. The coefficients of interest are the "Low Access" treatment indicator (row 1), and the interaction term (row 3). Columns 1 and 3 are for all school types, poor and rich, and the interaction term is treatment × poor school. Columns 2, 4,5 and 6 are restricted to only students in the poor schools and examine the treatment effects by whether or not the student had grades above the median.

(2013) utilize a regression-discontinuity design framework based on receipt of a fellowship to an elite school in India as an instrument to estimate the effect of schooling on an individual's perception of future returns. The paper finds that educational rewards have a positive and statistically significant effect on the perceived returns to education of awardees and a reduction in its variance with evidence that the fellowship changed the individual's overall valuation of education. In other recent work, Duflo, Dupas, and Benhassine (2013) implement both a conditional and an unconditional cash transfer program in Morocco that is labeled and promoted as an education program. The authors find that both programs increased perceptions of the returns to education and improved school enrollment even though no information was provided about the returns to education.

Both papers suggest that perceptions change because the education program provides new information on the returns to education. However, this interpretation is somewhat puzzling given that none of the programs implemented directly provide any information on the returns to education and yet they influence beliefs. The argument made is that the programs serve as a signal about the overall returns to education, for example, by strengthening the perceived link between schooling effort and success (S. Sequeira, 2013) or by signaling that the government must believe that the returns are high (Duflo, Dupas, and Benhassine, 2013). However, the evidence for either of these mechanisms is lacking. Furthermore, S. Sequeira (2013) finds that there are no belief spillovers from the individuals who received the scholarship to their friends or neighbors as one might expect if the learning story was true. Our combined lab and field evidence shows that access-based beliefs exist and can explain the findings in these studies.

### 2.4.2 Empirical Observations in Health Behavior

Access-based beliefs may explain a set of stylized facts in the empirical literature on health behavior. This literature documents that individuals have optimistic beliefs about health risks (Oster, Shoulson, and Dorsey, 2013; Weinstein, 1982) which become more unrealistic in resource-constrained settings (A. V. Banerjee, A. Banerjee, and Duflo, 2011) where an individuals' ability to access healthcare may be low. Furthermore, consumer demand for healthcare in developing countries appears to be highly sensitive to changes in prices (Cohen and Dupas, 2010; Dupas, 2011).

While other explanations are possible to explain each of these facts individually, a recent theoretical paper by Schwardmann (2019) attempts a unifying explanation. Individuals deal with the anxiety of ill-health in one of two ways. They either engage in denial (denialists) of the chance of falling ill, or they take preventative measures (realists) and when an agent has a limited ability to act, the former becomes the more attractive option.[15] Schwardmann notes several interesting implications of this model for market design: (1) Due to greater access, beliefs about the risk of ill-health are higher in low-price environments of perfect competition than under monopoly (2) However, because denialists are more sensitive to price changes, since prices impact their beliefs, monopolists may charge them a lower price (3) Overall, the low demand of denialists may influence the amount of investment in research and development observed in healthcare markets.

## 2.5 Conclusion

Access-based beliefs are the notion that beliefs of an alternative's value are a function of the likelihood it obtains. For example, individuals may have lower beliefs of alternatives to which they have low access, an effect I call sour grapes. I first test this notion in the lab, in chapter 1, and find that a low chance of obtaining a subjective lottery reduces beliefs of the probability that it yields a high payout. I next corroborate this finding in the field, in chapter 2, by repurposing a field experiment in which access to a commercial, free SAT prep package was varied. I find that lower access to the package reduced beliefs of its expected return.

[15] In fact, Schwardmann cites the first two chapters of this dissertation as providing empirical evidence from the lab and field of his proposed model.

The significance of access-based beliefs could be large for groups of people who face very low access to important investments. A critical example might be poor people who often face low access to investments and therefore, may formulate low beliefs about the returns. I find suggestive evidence that access-based beliefs may be more prevalent amongst poor people. Individuals from poorer backgrounds displayed a larger sour grapes effect in both the lab and field contexts. This result is striking and begs the question of whether poor people have learned to use sour grapes as a coping mechanism for low access, or whether some other variable maybe driving the differential prevalence amongst the poor, and whether access-based beliefs could be at the root of a belief-based poverty trap.

Perhaps the most critical insight of access-based beliefs for policy-makers or firms trying to estimate demand is that belief elicitations may be wrong and influenced by access. Low beliefs may not represent a lack of information, but simply, a lack of access. A potential way to distinguish between whether an individual has wrong information or is engaging in access-based beliefs might be to increase access to an investment and observe the effects on beliefs. While pessimistic individuals would increase their beliefs in both situations, optimistic individuals may react differently. If beliefs are a function of information, then optimistic individuals should lower beliefs following an increase in access as they learn about an alternative's value. However, access-based beliefs suggest that optimistic individuals should increase their beliefs even further when access increases.

The theory of access-based beliefs suggests a novel mechanism for the impact of social programs that expand access, namely that they impact beliefs as well. Further work could look to quantify the impacts different interventions and policies have on beliefs. For instance, if a policy-maker was trying to increase purposefully-held low beliefs about the value of an investment, would providing cash, which would increase access to multiple investments simultaneously, be more effective than providing a subsidy or voucher? Ultimately, the answers to such questions, and to the general impacts of different interventions, may be informed by a greater understanding of access-based beliefs.

Chapter 1 and 2 in full are being combined and prepared for submission for publication of the material. The dissertation author, Vinayak Alladi was the primary investigator and author of this material.

# Chapter 3

# An (Other Person's) Endowment Effect: A Test of Social Reference Dependence

## 3.1 Introduction

The impact of social comparisons and interaction have been documented in a variety of settings, including financial decisions (Bursztyn, Ederer, Ferman, and Yuchtman, 2014), educational outcomes (Zimmerman, 2003), choice of social groups (Sacerdote, 2001) and voting behavior (Harmon, Fisman, and Kamenica, 2017). The theoretical foundations of peer effects fall into two broad categories, as classified in Bursztyn, Ederer, Ferman, and Yuchtman (2014): social learning models, where individuals learn from the actions of others (A. V. Banerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992), and social utility models where utility is derived from making social comparisons (Festinger, 1954). This paper focuses on understanding whether individuals derive utility from making social comparisons in the context of consumption goods. Past work has suggested that conformity in consumption could be driven via the former learning channel, as individuals learn from the consumption choices of others (Moretti, 2011) or via social signaling (Bernheim, 1994), as when individual's engage in conspicuous consumption to signal social status (Veblen, 1899). By contrast, this paper examines how social comparisons may influence consumption choices in a basic exchange setting with standard consumption goods where learning forces and signaling motives are minimized. Given a distribution of consumption, individuals with a preference for conformity may want to emulate the majority, while those with a preference for non-conformity, may choose to deviate from the popular choice and choose the less popular good.

Standard economic theory says that when given a choice between two goods, an individual should choose the good they prefer more. However, an extensive experimental literature on behavior in exchange environments has largely focused on documenting and explaining the endowment effect: the finding that initial ownership of an object increases its valuation and reduces the desire to trade it away (Kahneman, Knetsch, and R. H. Thaler, 1991; R. Thaler, 1980; Tversky and Kahneman, 1991). One unrealistic aspect of these settings is that the distribution of objects is often, though not always initially equal.[1] By contrast, in the real world, endowments and exchanges do not occur in a vacuum and individuals can observe a distribution of an object's possession. Given the strong evidence of peer effects and social comparisons, it seems natural to ask: in addition to one's own endowment, does the distribution of possession of a consumption object in society impact exchange behavior. Consider an initial distribution of endowments where a small portion of people are given a mug while everyone else is given a pen. Standard reference dependence would argue that loss aversion around a reference point would be the correct model to describe exchange behavior, and that regardless of the skewed distribution towards pens, mug (pen) assignees are more likely to want to keep the mug (pen), thus yielding the familiar endowment effect, and lowering exchange behavior overall. However, if individuals make social comparisons, then there may also be a desire to conform, by some mug owners choosing the pen, or to distinguish oneself and feel special, by some pen owners choosing the mug, thus potentially altering the level of trade observed.

[1]For example, in (Heffetz and List, 2014) individuals are randomly assigned mugs and pens, so the distribution is likely to not always be 50-50, leading to more pen or mug owners in a room.

To structure these ideas, I employ the prominent model of reference dependence in the literature, the KR (Koszegi-Rabin) model, in which the reference point is based on an individual's rational expectations, i.e., expectations about future outcomes generated by consumption plans (Kőszegi and Rabin, 2006), and apply it to an exchange setting. I consider a decision-maker who has consumption utility, standard gain-loss utility, and an additional source of social gain-loss utility, which is the utility experienced from social comparisons to the choices of others. In particular, individuals compare their own consumption bundle to the utility they would get from the consumption bundle of others. As an example, an individual that possesses a mug while others possess a pen, feels a utility gain from the mug and a utility loss from not owning the pen scaled by a social gain-loss utility parameter which can differ from the standard gain-loss utility with respect to their own outcomes. The definition of conformity and non-conformity arises from considering values of the social gain-loss utility behavior for which individuals choose to emulate or deviate from the actions of others. Specifically, individuals who are loss-averse with respect to the outcomes of others will increase their consumption of a good as the share of other individuals consuming it increases (conformers), while individuals who are gain-seeking with respect to the outcomes of others will decrease their consumption as the share of other individuals consuming it increases (non-conformers).

In order to make predictions of choice behavior, I first treat the social distribution as an exogenous object and consider the notion of personal equilibrium in KR where an individual's actions and expectations are consistent with each other. This may be relevant in real life situations where my reference group is large, and my own actions have little impact on the social distribution. I show that under this notion of personal equilibrium, an individual's decisions to keep or exchange their endowed object depends on the relative magnitudes of standard and social gain-loss utility. If standard loss aversion is relatively larger, individuals are more likely to exhibit the endowment effect. However, if loss aversion with respect to other people's outcomes is relatively larger, then the distribution of consumption choices will likely impact the total amount of observed exchange.

To test the proposed model of conformity and non-conformity, I conduct a sequential game in the lab where subjects take turns making choices between two standard consumption goods (a mug and a pen) and can observe the choices made by players who moved before them. I use the strategy method to elicit preferences, i.e. subjects are asked to state their preference between a pen and mug at each possible decision node of the game before the game is played. In the end, one of the subject's choices is implemented by assigning a randomly chosen number between 1 and N (where N is the total number of players) to each subject and implementing decisions in ascending order of the numbers assigned. By design, there are no additional material incentives to coordinate on consumption choices besides getting the desired object, and there are limited motives to signal status, as subjects come into the lab not knowing each other. Additionally, social learning is arguably non-existent as subjects get to examine both subjects at the start of the experiment, and are told that there are enough mugs and pens for everyone.

Data was collected on 3 versions of the above design at UC San Diego. A 5-player version of the game, and a 3-player version of the game. The 3-player version was additionally done with and without a randomly assigned initial endowment. A total of 20 sessions of this study were conducted, 10 of them consisted of the 5-player version (data for 50 subjects x 15 decisions each) and 10 of them consisted of the 3-player version (data for 30 subjects x 6 decisions each[2]

Our primary identification strategy is to examine decisions made as the last player to move in the game. The key test of conformity, as defined by our model, is an individual choosing the mug when more prior participants chose the mug and the pen when more prior participants chose the pen. The opposite behavior is identified as non-conformity. Since all uncertainty is resolved by the last row of the game, using these decisions cleanly

[2]The total number of decisions for the with endowment experiment could have been 9 decisions, depending on the initial endowment. This will be fully explained in section 3.1.

identifies an individual's social preferences as no assumptions are required on beliefs about future play. The strength of conformity and non-conformity can be further understood by examining what individuals do for different values of the proportion of mug choosers thus far.

Overall I find that a small portion of individuals exhibit behavior that is consistent with having social preferences as defined by our model. While 60% of subjects choose their preferred object, 13.75% of the subjects exhibit a preference for non-conformity, 7.5% reveal a preference for conformity and the remaining 18.75% choose randomly. To test whether the data observed could have been generated by individuals choosing randomly, I use a binomial distribution to calculate the null distribution of random choice. I strongly reject that our data could have been generated randomly at the 5% level for the 5-player version but do not reject this for the 3-player version, where the probability of observing social preferences randomly is much higher.

I next test for rationalizability of the choices of individuals who I identified as conformers and non-conformers at the last stage of the game, in earlier stages. I use two definitions of rationalizability, one termed weak rationalizability which is the standard notion of rationalizability that admits any feasible beliefs, and the other termed strong rationalizability, which requires some consistency in beliefs across different stages of the game. I find that about half of our subjects identified as having social preferences have choices consistent with even strong rationalizability. Interestingly, I find suggestive evidence that rationality appears stronger amongst non-conformers, as non-conformers commit less violations of rationality, suggesting that non-conformity preferences may be more stable.

A second set of analyses tests the strength of the standard endowment effect in our setting. Our model predicts that the endowment effect could be dampened in our setting where people may respond to the social distribution. Given random assignment into the "with" and "without" endowment treatment groups, we would expect that the presence of the endowment effect should raise the proportion of times individuals endowed with a mug will choose the mug relative to the without-endowment treatment, and vice versa for the pen. While there are differences in the point estimates between these groups, there is no consistent pattern in the signs of the treatment effects (they are sometimes positive and sometimes negative) and I strongly fail to reject the null of no endowment effect (p-values  0.5 on average). However, given the small sample size of the experiment, I am likely underpowered to detect small endowment effects if they exist. Secondly, the theory predicts that we expect to observe less conformity and non-conformity overall in the with-endowment treatment relative to without. While there is some evidence of this, our sample size is too small to reject equality between the two treatments. Taken together, our results do suggest that social considerations appear to crowd out the endowment effect.

The study contributes to a large literature on social comparisons, finding novel evidence that individuals have preferences for conformity and non-conformity even with standard consumption goods. I examine goods that carry no special status, an environment in which there are no material payoffs from coordination and one in which signaling motives are likely low. Simply providing information on the consumption choices of others leads to emulation by about half of the subjects who respond to some groups of subjects and deviance for others. While most studies in this literature find evidence for conformity, or emulating the choices of others, our study finds an equally large number of non-conformers.

Secondly, the study contributes to the literature on expectations-based endowment effects which has found mixed results for the predictions of the KR model. Marzilli Ericson and Fuster (2011) (FE) find evidence in favor of expectations-based reference points, while Heffetz and List (2014) (HL) and Goette, Harms, and Sprenger (2014) (GHS) find results in favor of the status quo. One potential explanation is that social comparisons were very different in FE vs. HL and GS. In FE, the status quo across subjects are identical as no one is given an endowment, but what varies between subjects in the same session is the probability of permission to exchange. In HL and GHS, the initial endowments between subjects within a session vary, but the expected permission to exchange within a session are fixed. I argue that some of these differences may be explained by how subjects compare their own endowments and outcomes to others. In particular, low levels of exchange may be reflective of not only endowment effects, but of social comparisons between groups who were assigned different objects or faced different expectations.

The rest of this paper is organized as follows: In Section 3.2, I present the theoretical model that guides the experimental work and in Section 3.3, I discuss and provide details on the experimental design; in Section 3.4, I present the results from the lab experiment; in Section 3.5, I discuss the results of the paper and suggest further avenues of work and in Section 3.6 I conclude.

## 3.2 Theory

### 3.2.1 Preliminaries

Following Koszegi and Rabin (2006 and 2007), let c represent an outcome vector drawn from the measure F, and r represent a reference point vector drawn from the measure G. Then, the utility function, U(F—G) evaluates the consumption lottery F in comparison to the referent lottery G as follows: $U(F|G) = \iint u(c|r)\, dG(r)\, dF(c)$ where $u(c|r) = \sum_{k=1}^{K} m_k(c_k) + \sum_{k=1}^{K} \mu(m_k(c_k) - m_k(r_k))$, the function $m(\cdot)$ represents consumption utility, and the function $\mu(\cdot)$ represents gain-loss utility relative to the referent, $r$.

Assume that for small-stakes decisions, consumption utility, $m(\cdot)$, can be plausibly regarded as approximately linear, and a piecewise-linear gain-loss utility function is adopted as follows. [3]

$$\mu(x) = \begin{cases} x & x \geq 0 \\ \lambda \cdot x & x < 0 \end{cases}$$

[3] In the case of an exchange setting, this assumption may be even less restrictive as individuals are simply choosing over consumption items.

To incorporate social preferences, I assume that the individual also compares their outcomes to those of others in the population, borrowing a functional form similar to one of the seminal models of social preferences, Fehr and Schmidt (FS). (Fehr and K. M. Schmidt, 1999) Specifically, I assume that individuals experience a utility loss in the form of envy when they are behind others, and a utility gain when they are ahead. Unlike FS, since individuals are not being asked to allocate an object. I believe that guilt is unlikely to play much of a role in an exchange setting.

With this motivation, I augment the KR utility function by adding a social gain-loss component. Specifically, an individual compares their consumption outcome to that of every other individual in society. A question arises as to whether the relevant social comparison is the initial endowment of others (or the object they currently possess) or their expected consumption outcomes. To simplify the analysis, I first assume that the relevant comparison is the initial endowment, and solve for personal equilibria under this assumption.[4] For notation, let $c^1$ indicate the consumption of individual 1 and $c^j$ represent the initial endowments of all individuals, $c^2, c^3, ....c^N$.

$$\begin{aligned} u(c^1|r, c^2, c^3, ....c^N) = & \underbrace{\sum_{k=1}^{K} m_k(c_k^1)}_{\text{Consumption Utility}} + \underbrace{\sum_{k=1}^{K} \mu_s(m_k(c_k^1) - m_k(r_k^1))}_{\text{Individual Gain-Loss Utility}} + \\ & \underbrace{\frac{1}{N-1} \sum_{k=1}^{K} \sum_{j=2}^{N} \mu_o(m_k(c_k^1) - m_k(c_k^j))}_{\text{Social Gain-Loss Utility}} \end{aligned} \tag{3.1}$$

where $\mu_o$, the social gain-loss utility, is assumed to be plausibly linear for small stakes, just as is assumed for individual gain-loss utility, as follows:

[4]In the next section, I consider the case where the comparison point is expected consumption, and following the rational expectations framework laid out in KR, I define a notion of social equilibrium in which the comparison point is based on the expected consumption outcomes of other individuals which have to be choices others would rationally make.

$$\mu(x) = \begin{cases} x & x \geq 0 \\ \lambda_s \cdot x & x < 0 \end{cases}$$

The above utility function is a reference dependent utility function (RDU) with social comparisons, or henceforth referred to as the social RDU. Importantly, I allow the function for social gain-loss utility, $\mu_{social}$ or $\mu_s$ to be different from standard reference dependence, $\mu_{individual}$ or $\mu_I$ and under our assumption of linear gain-loss utility, $\lambda_{social}$ or $\lambda_s$ can be different from $\lambda_{individual}$ or $\lambda_I$. I also preserve anonymity across individuals, in that any interchange of labels would not impact the final utility.[5] Furthermore, I do not place weights on the individual terms in the utility function, but this could matter in some applications if individual gain-loss utility becomes more relevant than social gain-loss utility or vice-versa.

KR define an equilibrium concept called personal equilibrium (PE) in which: (1) an agent's choice maximizes her utility given her expectations (2) an agent's expectations (and hence, her reference point) are rational, in that they are based on the outcomes that would result if the agent follows through on her plan. Formally, consider a choice set $D$ , composed of lotteries, $F$, over consumption outcomes $c$.

**Personal Equilibrium (PE):** A choice $F \in D$, is a personal equilibrium if:

$$U(F|F) \geq U(F'|F) \quad \forall F' \in D$$

### 3.2.2 Exchange Behavior under a Social KR Framework

I now consider exchange behavior in a social KR model. Consider a standard exchange environment where individuals are endowed with either a mug or a pen and

[5]One could imagine augmenting the utility function so that there are weights on each social comparison pair, perhaps because comparisons to individuals with authority or popularity matter more than others but that a status-seeking motive that is not the focus of this model.

choose whether they would rather keep their endowment or exchange it for the other object. Thus, the utility function is over two dimensions, mugs, $m$, and pens, $p$ such that $c^1 = (m,p)$ , $c^j = (m,p)$ and $r^j = (m^j, p^j)$. The value $m \in \{0, M\}$ indicates the consumption utility of having no mug, or one mug, and $p \in (0, P)$ is the utility from having no pen or one pen. Under the linear gain-loss utility specified above, the social RDU in this case is:

$$\begin{aligned} u(c^1|r) = u(m,p|r_m^1, r_p^1, r_m^2, r_p^2, .......r_m^N, r_p^N) = m + p + \mu_s(m - r_m) + \mu_s(p - r_p) + \\ \frac{1}{N-1}\sum_{j=2}^{N}[\mu_o(m - m^j) + \mu_o(p - p^j)] \end{aligned} \tag{3.2}$$

The assumption of linearity of the social comparison term $\mu_o$ allows us to summarize the social gain-loss for a mug owner (respectively, pen owner) by comparing her outcomes to the *proportion* of other individuals, denoted $(1 - \pi_m)$ (respectively, $\pi_m$) who own a pen (mug). Compared to other individuals who own a mug, the mug owner faces no social gain-loss since her consumption is the same in all consumption dimensions.

### 3.2.3 Personal Equilibria in Exchange Environments

Two types of personal equilibria may persist for mug and pen owners; those in which individuals do not expect to exchange (and therefore do not exchange) and those in which they expect to exchange (and therefore exchange). Keep in mind that for mug owners (pen owners) the expectation to not exchange is the one most conducive to generating the standard endowment effect. However, I will show that if the impact of social comparisons in the utility function is large enough, then this could result in a PE in which some individuals who expected to keep their object, decide to exchange. Going through the derivations of equilibrium threshold values below is a useful exercise to later understand how they are impacted by changes in the social distribution of ownership.

**Mug Owner:** Consider a mug owner with a reference point of keeping the mug,

$(r_m,r_p) = (M,0)$. The mug owner has to choose between keeping the mug, yielding consumption outcome $(m,p) = (M,0)$ and giving it up for a pen, yielding consumption outcome $(m,p) = (0,P)$. The social distribution of endowments is characterized by $\pi_m$, the proportion of individuals endowed with the mug, and $(1-\pi_m)$, the proportion of individuals endowed with a pen. Thus, according to the social RDU, the individual can support keeping the mug in personal equilibrium, if:

$$u(M,0|M,0,\pi_m,(1-\pi_m)) > u(0,P|M,0,\pi_m,(1-\pi_m)),$$

or

$$u_m + \underbrace{(1-\pi_m)(u_m - \lambda_s u_p)}_{\text{social gain-loss}} > u_p + \underbrace{(u_p - \lambda_I u_m)}_{\text{indiv. gain-loss}} + \underbrace{\pi_m(u_p - \lambda_s u_m)}_{\text{social gain-loss}} \tag{3.3}$$

In words, keeping the mug is a PE if the utility of keeping the mug is greater than that of exchanging it for the pen, given a mug owner expects to keep the mug.

The highest utility of the pen for which mug owners can support keeping the mug in PE can be found by solving for $u_p$ in eq. (3):

$$u_p{}^{\text{keep mug}} = \frac{u_m(2+\pi_m(\lambda_s-1)+\lambda_I)}{(2+\pi_m(1-\lambda_s)+\lambda_s)} \tag{3.4}$$

The above equation gives the maximum utility of a pen for which the individual can support keeping the mug. Notice that this threshold value is increasing in standard loss aversion , $\lambda_I$, and in the utility of the pen, $u_m$. I later discuss how it varies with $\pi_m$.

A second PE can exist in which mug owners expect to exchange. Given the expectation to exchange, the referent becomes 0 mugs and 1 pen (0, P) and the individual can support exchange in personal equilibrium, if:

$$u(0,P|0,P,\pi_m,(1-\pi_m)) > u(M,0|0,P,\pi_m,(1-\pi_m)),$$

or

$$u_p + \underbrace{\pi_m(u_p - \lambda_s u_m)}_{\text{social gain-loss}} > u_m + \underbrace{(1-\pi_m)(u_m - \lambda_s u_p)}_{\text{social gain-loss}} + \underbrace{(u_m - \lambda_I u_p)}_{\text{indiv. gain-loss}} \tag{3.5}$$

$$\underline{u_p}^{\text{exchange for pen}} = \frac{u_m(3+\pi_m(\lambda_s - 1))}{(1+\pi_m(1-\lambda_s)+\lambda_s+\lambda_I)} \tag{3.6}$$

This equation gives the minimum utility of a pen for which exchanging the mug for a pen could be supported in PE. Any utility lower than this means that exchanging the mug for the pen could not be supported in equilibrium.

**Pen Owners:** Similarly, consider an individual initially endowed with a pen. Again, I consider two types of PE, one in which the individual expects to keep the pen, and the other in which she expects to exchange the pen for a mug. In the first case, the reference point is the initial endowment, $(0,P)$, which is the one most conducive to generating an endowment effect for the pen. The pen owner contemplates both keeping the pen, yielding consumption outcome $(m,p) = (0,P)$, and exchanging the pen for a mug, yielding consumption outcome $(m,p) = (M,0)$.

As before, $\pi_m$ and $(1-\pi_m)$ represents the proportion of individuals endowed with a mug and pen, respectively. Thus, according to the social RDU, the individual can support keeping the pen in PE if:

$$u(0,P|0,P,\pi_m,(1-\pi_m)) > u(M,0|0,P,\pi_m,(1-\pi_m)),$$

or

$$u_p + \underbrace{\pi_m(u_p - \lambda_s u_m)}_{\text{social gain-loss}} > u_m + \underbrace{(1-\pi_m)(u_m - \lambda_s u_p)}_{\text{social gain-loss}} + \underbrace{(u_m - \lambda_I u_p)}_{\text{indiv. gain-loss}} \tag{3.7}$$

$$\underline{u_p}^{\text{keep pen}} = \frac{u_m(3+\pi_m(\lambda_s - 1))}{(1+\pi_m(1-\lambda_s)+\lambda_s+\lambda_I)} \tag{3.8}$$

Note that since the valuations of pens and mugs are closely tied to expected exchange behavior, the minimum value of a pen at which a pen owner can support keeping the pen in PE is identical to the lowest value of a pen at which a mug owner can support trading a mug for a pen, $\underline{u_p}^{\text{keep pen}} = \underline{u_p}^{\text{exchange for pen}}$. In a similar fashion, I can calculate the maximum value of a pen at which a pen owner can support exchanging for a mug in PE: $u_p^{\text{exchange for mug}} = u_p^{\text{keep mug}}$.

### 3.2.4 Social Dynamics of the model

#### Defining Conformers and Non-conformers

To understand how the social distribution may impact equilibrium behavior in this model and how I define conformity and non-conformity, I examine how the utility thresholds for personal equilibria change as I manipulate $\pi_m$. To do this, I differentiate the threshold utilities found in the previous section with respect to $\pi_m$. It is adequate to look at one of the thresholds, since the other threshold varies only up to a constant. Thus, I consider how $u_p^{\text{keep mug}}$ (the maximum utility of a pen at which a mug owner can still support keeping the mug in equilibrium) varies with $\pi_m$.

$$\frac{du_p^{\text{keep mug}}}{d\pi_m} = \frac{u_m(\lambda_s - 1)(\lambda_I + \lambda_s + 4)}{(2+\pi_m(1-\lambda_s)+\lambda_s)^2} \tag{3.9}$$

**Proposition 3.1.** *For $\lambda_s > 1, {u_p}^{keep\ mug}$ is increasing in $\pi_m$, while for $\lambda_s < 1, {u_p}^{keep\ mug}$ is decreasing in $\pi_m$ for all values of $\lambda_I > 0$*

*Proof.* As can be seen, when $\lambda_I$ is greater than 1, the numerator is positive and $\frac{d{u_p}^{\text{keep mug}}}{d\pi_m} > 0$, while when $\lambda_s$ is less than 1, the numerator is negative and $\frac{d{u_p}^{\text{keep mug}}}{d\pi_m} < 0$ This holds true for any value of $\lambda_I$ greater than 0. When $\lambda_I = 1$, the derivative is zero. □

Proposition 3.1 shows that for a mug owner for whom $\lambda_I > 1$, an increase in the proportion of mugs raises the maximum utility of a pen at which they can support keeping the mug in equilibrium. In other words, a higher proportion of mug owners leads this individual to be more likely to keep the mug, and I call such an individual a **conformer**. Conversely, for a mug owner for whom $\lambda_I < 1$, an increase in the proportion of mugs lowers the maximum utility of a pen at which they can support exchanging for the pen in equilibrium. In other words, a higher proportion of mug owners leads this individual to be more likely to exchange for the mug, and I call such an individual a **non-conformer**. When $\lambda_I > 1$, the derivative is zero, and the threshold value is un-impacted by the value of $\pi_m$. The individual has no gain-loss utility with respect to other people's possessions.

Figure 3.1 (figure 3.2) show how the threshold value of a pen for which individuals are willing to keep or exchange their good vary as a function of $\pi_m$ for conformers (non-conformers) relative to individuals who have no kink in social gain-loss utility, i.e. for whom $\lambda_s = 1$. The two dashed lines show the thresholds utility values of a pen at which mug owners can support keeping (in black) and exchanging (in red) their objects in equilibrium for the parameter values $\lambda_s = 1$ and $\lambda_I = 1.2$. The gap between the two lines illustrates the standard endowment effect. For these individuals, who have no social preferences, varying $\pi_m$ has no impact on the willingness-to-pay and willingness-to-accept thresholds.

The black and red solid lines in figure 1 represent the change in these thresholds for exchange when $\lambda_s$ increases to a value of 1.2 and individuals exhibit a preference for

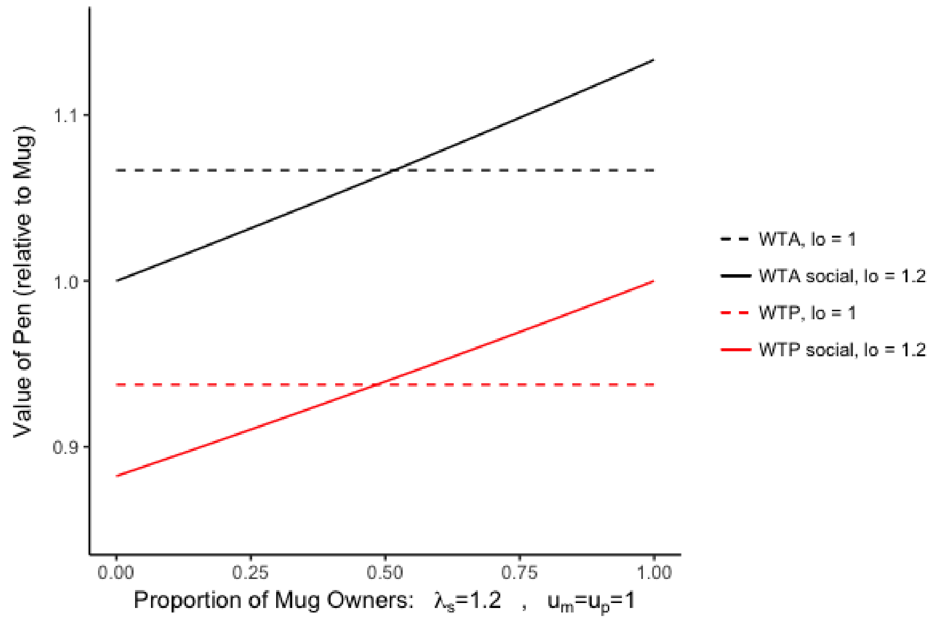


**Figure 3.1.** Personal Equilibria as a function of the endowment distribution for Conformers

conformity. As demonstrated in proposition 3.1, this leads to a lower threshold willingness to keep the mug (exchange for the pen) when $\pi_m < 0.5$ and a higher willingness to keep the mug (exchange for the pen) when $\pi_m > 0.5$. Similarly, black and red solid lines in figure 2 represent the change in these thresholds for exchange when $\lambda_s$ decreases to a value of 0.8 and individuals exhibit a preference for non-conformity. This leads to a higher threshold willingness to keep the mug (exchange for the pen) when $\pi_m < 0.5$ and a lower willingness to keep the mug (exchange for the pen) when $\pi_m > 0.5$.

### Personal Equilibria Thresholds and Individual gain-loss utility $\lambda_I$

A second comparative static of interest is how increasing loss aversion impacts the threshold value of a pen for which the mug owner can still support keeping the mug in equilibrium, i.e., how increasing individual gain-loss utility impacts the likelihood of observing the endowment effect.

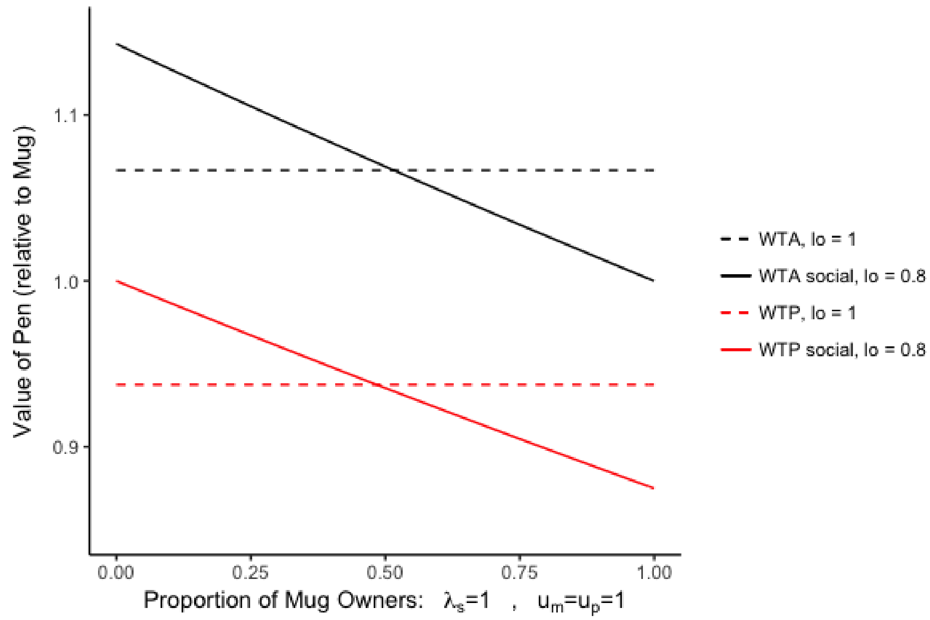


**Figure 3.2.** Personal Equilibria as a function of the endowment distribution for Non-conformers

**Proposition 3.2.** *For a given $\lambda_s$ , ${u_p}^{keep\ mug}$ is increasing in $\lambda_I$*

*Proof.*

$$\frac{d{u_p}^{\text{keep mug}}}{d\lambda_I} = \frac{u_m}{(2+\pi_m(1-\lambda_s)+\lambda_s)} \tag{3.10}$$

This fraction is positive and constant for all values of $\pi_m$ and $\lambda_s > 0$ □

This means that increasing the level of individual gain-loss utility has an unambiguous positive effect on the likelihood of keeping one's object. In figure 3.6 (appendix), I illustrate graphically how the threshold value of a pen at which a mug owner is willing to exchange a mug for the pen in equilibrium is increasing as the level of individual gain-loss utility goes up.

### 3.2.5 Endogenizing The Distribution $\pi_m$

So far, the analysis of behavior has assumed a fixed value of $\pi_m$. However, it is likely that in many settings this is an unrealistic assumption as the number of players may be small and therefore $\pi_m$ is an endogenous object.[6] To analyze the game when $\pi_m$ is endogenous, I propose to define a social equilibrium which combines nash equilibrium with KR's personal equilibrium. This work is ongoing and is not required for establishing the main results of this chapter. Nonetheless, demonstrating the existence (or non-existence) of social equilibria is important for distinguishing this paper from other papers on reference points since the reference point becomes dependent on the actions of others.

To show that a nash equilibrium exists, one must show that there exists at least one fixed point of the best response functions. To do so, one can check that the conditions on these functions under which a fixed fixed point theorem can be applied are satisfied. As we have not reached that step, the remainder of section 3.2.5 assumes that these conditions are satisfied and characterizes one of them in which all individuals are conformers. In particular, one must show that the individuals are playing a best response to other individuals at the nash equilibrium, and to be consistent with KR reference dependence, should also be playing a personal equilibrium. It bears repeating that the remainder of section 3.2.5 is strictly preliminary work, and should not be treated with the same certainty of other results in the paper.

**Simultaneous Games**

Under the assumption of common knowledge and rationality, I consider pure strategy equilibria when individuals choose pens and mugs simultaneously. These equilibria look similar to the types of equilibria found under standard coordination games. The conjecture below attempts to prove that if all individuals in the population are conformists, then there is a conforming equilibrium where everyone chooses the mug or everyone chooses the pen under specific conditions on the parameter values $u_p, \lambda_s, \lambda_I$.

[6] A small literature exists on studying strategic interactions between expectations-based loss averse players (Dato, Grunewald, Müller, and Strack, 2017; Shalev, 2000). Since the focus of this paper is to study a very specific game in order to identify conforming and non-conforming preferences, drawing general comparisons with the theoretical literature on this topic is beyond the scope of this work.

**Conjecture 1.** *If everyone is a conformist, i.e.* $\lambda_s > 1 \quad \forall i \in I$*, and expects to keep the mug, then as long as* $u_p < \frac{\lambda_s+\lambda_I+1}{3}$ *and* $u_m = 1$ *, w.l.o.g., and* $\lambda_I > 1$ *then* $\pi^*(c^1, ...., c^N) = 1$ *is an equilibrium. If everyone is a conformist, i.e.* $\lambda_s > 1 \quad \forall i \in I$*, and expects to keep the pen, then as long as* $u_p < \frac{\lambda_s+2}{\lambda_I+2}$ *and* $u_m = 1$ *, w.l.o.g., and* $\lambda_I > 1$ *then* $\pi^*(c^1, ...., c^N) = 1$ *is an equilibrium.*

*Proof.* Suppose everyone expects to keep the mug, then we need to show that:

$$U^i(M, 0|M, 0, \pi^*(c^1, .., c^N) = 1) > U^i(0, P|M, 0, \pi^*(c^1, .., c^N) = 1) \;\; \forall i \in I$$
$$u_m + (1 - \pi_m)(u_m - \lambda_s u_p) > u_p + \pi_m(u_p - \lambda_s u_m) + (u_p - \lambda_I u_s)$$

since $u_m$ is normalized to 1 and $\pi_m = 1$

$$\implies 1 > u_p + \pi_m(u_p - \lambda_s) + (u_p - \lambda_I) \implies u_p < \frac{\lambda_s + \lambda_I + 1}{3}$$

As long as this condition on the $u_p, \lambda_s, \lambda_I$ are satisfied, then everyone choosing mugs will be an equilibrium. Since $\lambda_s > 1$ and under the assumption of loss aversion, $\lambda_I > 1$ this condition says that utility of the pen should not be larger than the average of 1 , $\lambda_s$ and $\lambda_I$. Suppose everyone expects to keep the pen, then we need to show that:

$$U^i(M, 0|0, P\pi^*(c^1, .., c^N) = 1) > U^i(0, P|P, 0, \pi^*(c^1, .., c^N) = 1) \;\; \forall i \in I$$
$$u_m + (1 - \pi_m)(u_m - \lambda_s u_p) + (u_m - \lambda_I u_p) > u_p + \pi_m(u_p - \lambda_s u_m)$$

since $u_m$ is normalized to 1 and $\pi_m = 1$

$$\implies 1 + 1 - \lambda_I u_p > u_p + u_p - \lambda_s \implies u_p < \frac{\lambda_s + 2}{\lambda_I + 2}$$

The conditions on the maximum value of $u_p$ under which everyone choosing the mug is an equilibrium, depend on whether individuals expect to keep the mug or the pen. If they expect to keep the pen, the maximum threshold of $u_p$ at which the mug can be an equilibrium comes down. □

An analogous social equilibrium exists when $\pi^*(c^1, ...., c^N) = 0$ with different cutoff values for $u_p, \lambda_s, \lambda_I$ under which everyone is consuming a pen. In this case, $u_p$ can be lower than the normalized value of $u_m = 1$ and it still may be the case that individuals consume pens in equilibrium if $\lambda_I$ is large enough.

On going work is being done on finding and proving the other pure strategy equilibria in this simultaneous game when different assumptions are made about the distribution of social preferences. The main takeaway looking at simultaneous games is that identification of social preferences is nearly impossible because it relies greatly on assumptions of common knowledge. As a result, I look towards sequential games.

**Sequential Games**

I study a sequential game in which a group of subjects take turns choosing between the pen and the mug, an environment that exactly mirrors the final experimental design I chose. Consider a sequential game with a finite number of players, strategies, and outcomes. The strategies are simply a choice of pen or mug from the set $M, P$ given the proportion of mugs $\pi_m$ chosen so far at every node in the game. The payoffs come from the following utility function of each player, exactly as before:

$$u(c^1|r, c^2, c^3, ....c^N) = \underbrace{\sum_{k=1}^{K} m_k(c_k^1)}_{\text{Consumption Utility}} + \underbrace{\frac{1}{N-1}\sum_{k=1}^{K}\sum_{j=2}^{N}\mu_o(m_k(c_k^1) - m_k(c_k^j))}_{\text{Social Gain-Loss Utility}} \tag{3.11}$$

where I assume a linear social gain-loss utility function as before. Then , we have the following proposition:

**Proposition 3.3.** *Assume the following: (1) $\lambda_I = 1$ (2) Individuals know their own conformity type, $\lambda_s$, but not the conformity types of others. (3) $u_m = 1$ and $u_p$ is distributed* *symmetrically around 1 with finite variance. The optimal payoff maximizing strategy is that conformers, for whom $\lambda_s > 1$, should choose the object that has been chosen most frequently so far regardless of the types of future players. Non-conformers, for whom $\lambda_s < 1$, should choose the object that has been chosen less frequently so far. If both objects have been chosen equally in the past, one is indifferent between the two.*

*Proof.* Proof by Induction.

1. Show that this holds for N=1:

(a) Starting at the last decision node, the average Nth player is assumed to be indifferent between mugs and pens in expectation. She will choose the mug or pen based on the following rule: Conformers, for whom $\lambda_s > 1$, will choose the mug if $\pi_m > 0.5$ and the pen if $\pi_m < 0.5$. Correspondingly, non-conformers , for whom $\lambda_s < 1$ and $\pi_m > 0.5$ will choose the pen if $\pi_m < 0.5$.

$$u_m + \underbrace{(1-\pi_m)(u_m - \lambda_s u_p)}_{\text{social gain-loss}} > u_p + \underbrace{\pi_m(u_p - \lambda_s u_m)}_{\text{social gain-loss}} \tag{3.12}$$

If $\pi_m = 0.5$ then choose the object you like more. On average, assume that mugs and pens are liked equally.

(b) Show that the optimal strategy is true for the $N-1^{th}$ player. The $N-1^{th}$ player cares about his choice with regards to the final distribution of objects chosen. The conjecture says that his optimal strategy is the same as the $N^{th}$, i.e. the $N-1^{th}$ player only responds to past choices. Consider two cases:

Case 1: If the $N^{th}$ player's decision will not change which object is in the majority, then the $N-1^{th}$ player should always follow the optimal strategy of the $N^{th}$ player in the game, described in (1) above, as this is a dominant strategy. Stated differently, regardless what happens in the remainder of the game, the object that will finally be in the majority has already been determined, so it is a dominant strategy to then go with or against the majority so far depending on one's type (conformer or non-conformer).

Case 2: Suppose the $N^{th}$ player's choice will have an impact on which object is in the majority. This can only happen if there is a tie in the proportion of objects chosen so far, i.e. the current distribution is $\pi_m = 0.5$ which also implies that the $N-1^{th}$ player's choice is the one that lead to $\pi_m = 0.5$. This necessarily implies that the $N-1^{th}$ player must have been a non-conformer.

2. Assume that the optimal strategy is true for the $K+1^{th}$ player and show that it must be true for the $K^{th}$ player.

The $K+1^{th}$ last player knows that if he were the $K^{th}$ last player, the strategy of ignoring the future, and only responding to players who have already moved, is optimal. Therefore, the $K+1^{th}$ last player now only needs to calculate if the decisions of the $K^{th}$ last player affects his decision. By the logic of (1) above, the $K+1^{th}$ player only affects the $K^{th}$ last player's strategy if her decision is going to affect which object is in the majority. Again, this can only happen if the distribution thus far is $\pi_m = 0.5$. By the logic in (2) we know that the average $K^{th}$ last player's decision does not affect what the $K+1^{th}$ last individual would do.

$\square$

### 3.2.6 Strategy Method and Rationalizability

As a preview to the experimental design we end up using, we decided to have subjects play a sequential game and used the strategy method to elicit choices at all possible nodes of the game. This provides us with a rich enough data set to understand preferences for conformity and non-conformity at the individual level, especially by analyzing choices made at the end of the game one $\pi_m$ is fully resolved and is fixed. A second set of analyses we do to test how consistent these choices are is to look at the rationalizability of choices made at every node in the game.

We use two criterion for rationalizability: (1) Weak rationalizability, which is a term we use to indicate the standard definition of rationalizability in a game in that there is some admissible belief on future play that is consistent with the choice made at that node (Bernheim, 1984) (2) Strong rationalizability based on an assumption of consistency in beliefs about future play. We define these in turn.

**1. Weak Rationalizability**

Define $\pi_m^s$ to be the proportion of mugs chosen from row 1 to row s and $\pi_m^f$ to be the proportion of mugs chosen at the final choice node of the game. At any node $a$, if there exists a feasible final node such that the choice in that node is the same as the choice in $a$ then the choice is rationalizable. By 'feasible" I mean that $\pi_m^f$ subsumes $\pi_m^s$.

As an intuitive example, let "$M$" be "Mug" and "$P$" be "Pen". If a person has chosen $M$ after seeing $MM$, then they must choose mug at one of the following final distributions in the five player game, $MMMM, MMMP, MMPP$.

**2. Strong Rationalizability:**

Strong rationalizability puts some additional structure on beliefs about future game play. Consider the case that an individual is a conformer based on the decisions made in the last row and chooses $M$ for a particular value of $\pi_m^s$ in an earlier row. If the individual has forecasted $\mathbb{E}(\pi_m^f|\pi_m^s)$, then in the next round, if they see another $M$, they must necessarily believe that $\mathbb{E}(\pi_m^f|\pi_m^{s+1}) \geq \mathbb{E}(\pi_m^f|\pi_m^s)$ for $s+1 > s$, which means they must choose a mug in this situation as well.

This definition of rationalizability is stronger than the first because it places more consistency on the beliefs an individual can have after more uncertainty has resolved. An example of a choice that does not violate weak rationalizability but violates strong rationalizability is the following:

As an intuitive example, suppose an individual chose $M$ after seeing $M$ and $P$ after seeing $P$. Then, under consistent beliefs they must choose $M$ after $MM$ and $P$ after $PP$. However, under weak rationalizability they can choose $M$ after $MM$ and $P$ after $PP$ as long as it is consistent with some choice in the last round of the game.

## 3.3 Experimental Design

Our experimental design is a sequential game used to test the predictions of the model. Subjects in the lab take turns choosing between a UC San Diego mug and a UC San Diego pen and can observe the decisions chosen before them, but not those that will be made subsequently. [7,8] In order to separately identify whether an individual likes the pen or the mug from their preferences for conformity or non-conformity, individuals are asked to make decisions using the strategy method and state their choices ex-ante, before any choices are actually implemented, at every possible node in the game. This allows the researcher to observe what the individual would do if everyone chose a mug, as well as what the individual would do if everyone chose a pen, providing a way to identify conformity and non-conformity.[9,10] After choices are elicited, subjects are randomly assigned a number between 1 and N, where N is the number of players in the game, and decisions are implemented in the order of the assigned number and the experiment concludes.

Figure 3.3 shows an example of this game when N = 3. Choices are elicited at every decision node indicated by the rectangles.[11]

[7]the pens and mugs were procured at the UC San Diego bookstore and are roughly the same dollar value.

[8]The experiment was designed and programmed in Otree.

[9]We considered alternate designs which increased the social pressure, e.g., assigning 18 individuals a mug and 2 a pen asking individuals if they would like to trade. Making just one decision might have reduced concerns of experimenter demand effects. However, the drawback to this would be not identifying preferences for conformity and non-conformity at an individual level, and having much less to say about outcomes from a given session, as discussed in section 2.6 on pure strategy equilibria in a simultaneous move game. The sample size requirements to test our central prediction would also go up drastically.

[10]In the empirical strategy I discuss how I identify conformers and non-conformers using the definition implied by our model.

[11]Note that the order in which decisions are made is irrelevant, i.e., if the previous choices have been Mug and Pen, I do not separately ask what a participant prefers if the mug was chosen first and then the pen, vs the pen was chosen first and then the mug. This reduces the number of decisions considered in the 5-player version of the game that I ran but only by one for the 3-player version. Furthermore, the order in which decisions are made is not something our theory predicts would matter.

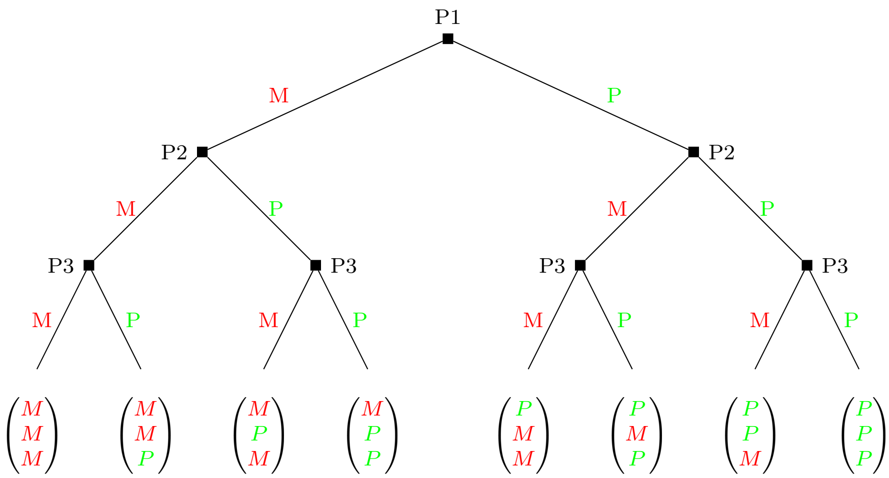


**Figure 3.3.** Game Tree - 3 players

Additional instructions included the following: (1) subjects are told that there are enough mugs and pens for everyone so they should choose their preferred object for each decision. This was to ensure that subjects were fully aware that there are no supply constraint on the goods. A follow-up comprehension question was asked to verify that they understood this. (2) subjects were given a chance to physically examine both objects at the start of the experiment so that they had as much information about them as possible. This minimizes any information asymmetries between subjects about the goods, thus minimizing the channel of social learning. Even if social learning does exist, it would go against us finding evidence of non-conformity. These design elements allowed us to isolate the mechanisms for behavior that our model explores. An example of the decision screen is shown in figure 3.4.

I test the above design in two versions of the game, a 3 player version and a 5 player version where participants make 6 and 15 choices respectively. I chose small groups to reduce the number of choices participants made since I was adopting the strategy method.

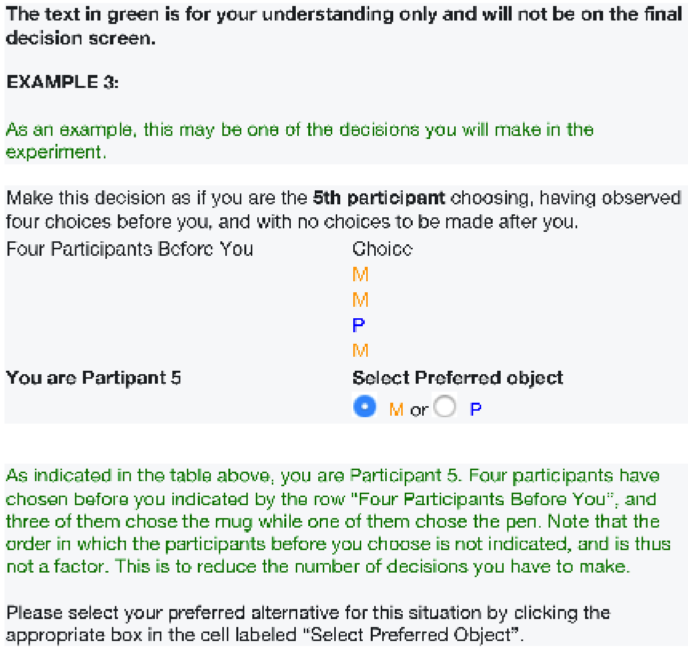

**The text in green is for your understanding only and will not be on the final decision screen.**

**EXAMPLE 3:**

As an example, this may be one of the decisions you will make in the experiment.

Make this decision as if you are the **5th participant** choosing, having observed four choices before you, and with no choices to be made after you.

| Four Participants Before You | Choice |
|---|---|
| | M |
| | M |
| | P |
| | M |
| **You are Partipant 5** | **Select Preferred object** |
| | M or P |

As indicated in the table above, you are Participant 5. Four participants have chosen before you indicated by the row "Four Participants Before You", and three of them chose the mug while one of them chose the pen. Note that the order in which the participants before you choose is not indicated, and is thus not a factor. This is to reduce the number of decisions you have to make.

Please select your preferred alternative for this situation by clicking the appropriate box in the cell labeled "Select Preferred Object".

**Figure 3.4.** Example of subject's decision screen without endowment

### 3.3.1 With Endowment Experiment

In order to study whether the endowment effect persists in our social setting, a version of the 3 player game is conducted with the additional feature that players are randomly assigned one of the objects (either the mug or the pen) to start. I wanted to give the endowment effect the best chance of being observed so I assigned an object at the start of the experiment (right after subjects have had a chance to examine both objects). This ensures that the assigned object stays with subjects for the duration of the experiment till decisions are implemented. For this version of the experiment, subjects have either 6 or 9 choices to make since they make different choices depending on whether a mug assignee or a pen assignee chose pen before them. Below is an example of one of the decisions a subject could make if they were assigned the mug and the other two subjects were assigned the pen.

As shown in figure 3.5, subjects are informed of the distribution of endowments in

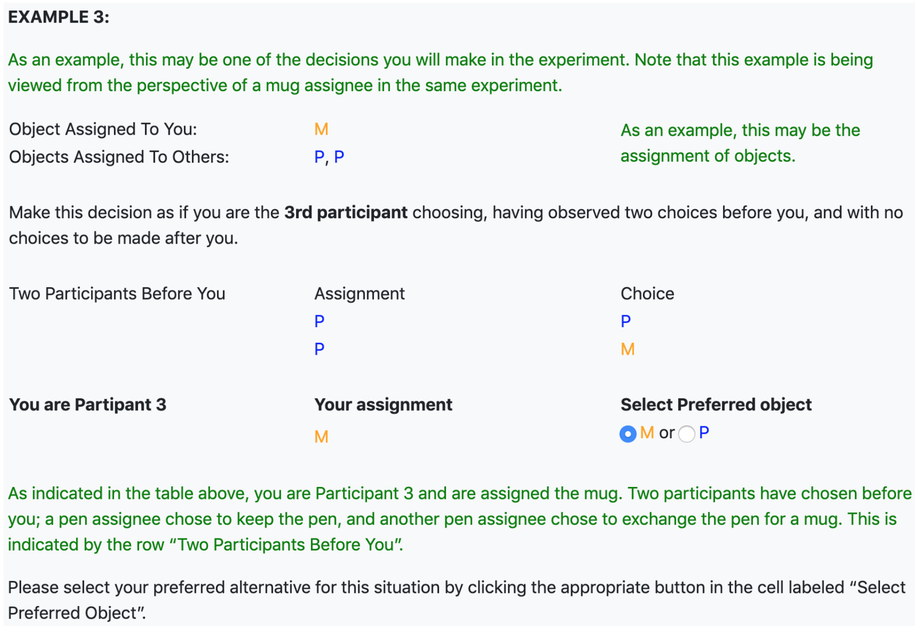
EXAMPLE 3:

As an example, this may be one of the decisions you will make in the experiment. Note that this example is being viewed from the perspective of a mug assignee in the same experiment.

| | | |
|---|---|---|
| Object Assigned To You: | M | As an example, this may be the assignment of objects. |
| Objects Assigned To Others: | P, P | |

Make this decision as if you are the **3rd participant** choosing, having observed two choices before you, and with no choices to be made after you.

| Two Participants Before You | Assignment | Choice |
|---|---|---|
| | P | P |
| | P | M |
| **You are Partipant 3** | **Your assignment** | **Select Preferred object** |
| | M | M or P |

As indicated in the table above, you are Participant 3 and are assigned the mug. Two participants have chosen before you; a pen assignee chose to keep the pen, and another pen assignee chose to exchange the pen for a mug. This is indicated by the row "Two Participants Before You".

Please select your preferred alternative for this situation by clicking the appropriate button in the cell labeled "Select Preferred Object".

**Figure 3.5.** Example of subject's decision screen with endowment

the room. They are also told what choices participants who have moved before them made (and their initial endowments).[12] The language used to discuss a subject's endowment is "ownership neutral". Individuals endowed with the mug are referred to as "mug assignees" and vice-versa for pens. This mirrors literature that has tried to use minimal ownership language in studying the endowment effect in order to avoid priming or language confounds.

## Theoretical Predictions and Hypotheses

### Null Hypothesis

The null hypothesis is that in the stripped down exchange environment we've created, the distribution of consumption goods should have no impact on a person's preferences over them. This is predicted by the following classes of models:

1. **Classical Utility Theory:** under standard utility theory individuals should choose the good that they prefer

[12]This could potentially lead to an additional way in which people conform or not conform which is that they emulate or deviate from the behavior of other individuals who started with the same object as they did. This is beyond the scope of our model and requires a much larger data set to test.

2. **Social learning:** since individuals get to examine both objects, I assume there is not much room for social learning in our experiment. Even if there was, it would go against finding evidence of non-conformity. Examples of social learning models include (A. V. Banerjee, 1992).

3. **Signaling Models and Status Goods:** individuals buy certain consumption goods because they signal status as in the models of conspicuous consumption (Veblen, 1899) or models of conformity generated by signaling (Bernheim, 1994). However, given the nature of goods in our experiment, mugs and pens, and that individuals in the room do not know each other or care about belonging to a particular group[13] I would not expect an effect of our experiment under these models.

4. **Reference Dependence:** Standard reference dependence models would only predict an effect in our setting if the social distribution of objects had an impact on a person's reference point, but this fact has not been considered important in past literature (Kahneman, Knetsch, and R. H. Thaler, 1991; R. Thaler, 1980; Tversky and Kahneman, 1991).

**Alternate Hypothesis**

The alternate hypothesis is that individuals make social comparisons and may respond to the social distribution, consistent with the predictions of our augmented reference dependence model.

**Without vs. With Endowment**

There are two predictions, based on the proposed model, for how exchange behavior might be affected by initially endowing subjects with one of the goods. Assuming that individuals are loss averse with respect to departures from their endowment ($\lambda_I > 1$), as has been typically found in past experiments, then I expect to find that:

[13] Of all the subjects who participated in the experiment, 2 pairs appeared to know each other.

1. Individuals endowed with one of the goods will choose that good more, on average, than individuals who were not endowed with it.

2. As an direct consequence of (1), I expect to see a lower proportion of individuals displaying a social preference in the with endowment compared to the without endowment treatments.

## 3.4 Results

### 3.4.1 Data Collection

In total, data from 20 lab sessions was collected at UC San Diego broken down as follows:

**Table 3.1.** Data Collection

| | # of sessions | # of subjects | # decisions per person |
|---|---|---|---|
| 5-player | 10 | 50 | 15 |
| 3-player without endowment | 5 | 15 | 6 |
| 3-player with endowment | 5 | 15 | 6 or 9 |

### 3.4.2 Sample Balance

Across a set of demographic characteristics, I find no statistically significant differences between the subjects in the 3-player vs. 5-player version, as well as subjects in the with and without endowment, showing that the randomization was successful (Table 3.2).

### 3.4.3 Empirical Strategy

Our empirical strategy has three main components:

1. **Identifying social preferences using decisions as the last player:** First, I provide evidence for the key prediction of our model: that the social distribution

**Table 3.2.** Sample Balance

| | (1) Female | (2) Age | (3) Year of College | (4) Intermediate Micro | (5) Non-chinese |
|---|---|---|---|---|---|
| 5-player | -0.218 | 0.453 | -0.113 | 0.147 | -0.060 |
| | (0.148) | (0.631) | (0.396) | (0.129) | (0.117) |
| 3-player w/ endowment | -0.067 | 0.467 | -0.062 | 0.133 | 0.333 |
| | (0.183) | (0.783) | (0.500) | (0.160) | (0.145)** |
| Constant | 0.667 | 20.267 | 3.133 | 0.133 | 0.200 |
| | (0.129) | (0.554) | (0.347) | (0.113) | (0.103) |
| Observations | 79 | 80 | 79 | 80 | 80 |
| P-value: 5 - player = 0 | 0.145 | 0.475 | 0.776 | 0.258 | 0.609 |
| P-value: 3 - player w/ endowment = 0 | 0.716 | 0.553 | 0.902 | 0.407 | 0.024 |

Robust Standard Errors in Parentheses. p<0.1*, p<0.05**, p<0.01***. The constant represents the mean in the 3-player without endowment version.

of an object has an impact on preferences for the object. To test this hypothesis I calculate the proportion of individuals that can be identified as conformers or non-conformers using only the decisions made when choosing as the last player in the game. Decisions made as the last player provide the strongest test of identification of social type because they require no assumptions on beliefs about future play since $\pi_m$ and $(1-\pi_m)$ can no longer change.

2. **Identifying social preferences across all decisions:** Second, I examine the rationalizability of decisions made by an individual at all rows of the game. I use two notions of rationalizability as defined in the theory section 3.2.6. The first is the standard game-theoretic notion of rationalizability, which I term "weak rationalizability", i.e., that there is some belief about future game play that could justify a particular choice. The second is strong rationalizability based on assuming some consistency of beliefs across rows of the game.

3. **Testing presence of the endowment effect:** Finally, I test whether the endowment effect survives in our setting where individuals could potentially have a desire to emulate or deviate from the behavior of others. Since subjects are randomized into with and without endowment treatments, I would expect that ex-ante preferences over the two goods are identical. Any difference in choice probabilities between the two groups can only result from being endowed with a good.

To carry out our empirical strategy for (1) and (2), I define below what decisions are consistent with our model. For convenience, I will sometimes refer to the ”Nth player choosing” as ”the Nth row”.

### 3.4.4 Identifying social preferences using decisions made as the last player

To identify conformity and non-conformity as defined by our model, the following definition holds true. I define it for the last row only in this section, but more generally in the next section. Consider decisions made in the last row when $\pi_m$ is fixed. For individuals who choose the same object for all decisions they faced, they are trivially consistent with our model. For individuals who switch at least once between choosing a mug and a pen in the last row, the following are necessary and sufficient conditions for consistency with our model:

1. Consider an individual's choices across a pair of symmetric nodes where symmetric is defined as $\pi_m^i$ and $(1-\pi_m^i)$ for $\pi_m^i \neq 0.5$. If an individual chooses the mug (pen) in the first case and the pen (mug) in the other, the individual may be a conformer (non-conformer) only if they choose the mug for all $\pi_m^j \geq \pi_m^i$ and the pen for all $(1-\pi_m^j) \leq (1-\pi_m^i)$.

2. If an individual prefers a mug and two values of $\pi_m^i$ and $(1-\pi_m^i)$, then they must also prefer the mug for any value of $\pi_m^j \leq \pi_m^i$ and $(1-\pi_m^j) \geq (1-\pi_m^i)$ . The same is true if they prefer pen in both situations.

3. At $\pi_m^i = 0.5$ the individual may choose the pen or mug.

**5-player game**

Table 3.3 shows a raw count of the possible combinations of choices I observe for all individuals in the experiment when choosing as the the last player, player 5, in the 5-player version of the experiment. The columns indicate the possible choices that previous subjects could have made, including the value of $\pi_m$, while the rows indicate the possible combinations of choices for the 5th player. For example, in column 2, I see that the previous choices were MMMP (not in any particular order) and the value of $\pi_m = .75$. The left most column indicates which combinations of choices can be classified into "conformers", "non-conformers", "random-choosers", and "always chose pen or mug". "Conformers" and "non-conformers" based on the definition above, and our model. Anyone inconsistent with our model is a "random chooser". And for the purposes of our statistical tests of random choice I distinguish these individuals from those who only chose the pen or only chose the mug for all decisions.

In total there are $2^5 = 32$ possible combinations for the last row. Two of those combinations are for individuals who always choose the mug or always choose the pen. Of the remaining possible ways in which a person can make choices, 4 of them are consistent with the definition of conformity as outlined, and 4 of them are consistent with the definition of non-conformity. The remaining 22 possible combinations are the "random choosers".

The data shows that 10 of the individuals made decisions consistent with the definition of conformity or non-conformity used in the model. 5 of them appear to be conformers, while 5 of them appear to be non-conformers. Of these individuals, only 3 of the 10 demonstrated social preferences when $\pi_m = 0.75$, so it does seem that social preferences are more likely under extreme values of $\pi_m$.

The number of individuals that chose randomly using the strict definition that they did not always choose the same object for the 5th row is only 6. However, this is not a fully accurate representation of the number of random choosers, because 5 of the 34 individuals who chose only one object for the last row made different decisions in previous rows. If we consider these individuals to be random choosers as well, it raises the number

**Table 3.3.** Identifying Conformity and Non-conformity from last row decisions in 5 player game

| | Possible Combinations of decisions by previous 4 players | | | | | |
|---|---|---|---|---|---|---|
| | MMMM $\pi_m = 1$ | MMMP $\pi_m = .75$ | MMPP $\pi_m = .5$ | MPPP $\pi_m = .25$ | PPPP $\pi_m = 0$ | Count |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| **Always chose mug or pen** | M | M | M | M | M | 20 |
| | P | P | P | P | P | 14 |
| | | | | | Total | 34 |
| **Choices consistent with conformity** | M | M | M | M | P | 2 |
| | M | M | M | P | P | 1 |
| | M | M | P | P | P | |
| | M | P | P | P | P | 2 |
| | | | | | Total | 5 |
| **Choices consistent with non-conformity** | P | M | M | M | M | 3 |
| | P | P | M | M | M | 1 |
| | P | P | P | M | M | 1 |
| | P | P | P | P | M | |
| | | | | | Total | 5 |
| **Choices classified as random** | M | M | M | P | M | |
| | M | M | P | M | M | |
| | M | P | M | M | M | 1 |
| | M | P | M | P | M | |
| | M | M | P | P | M | |
| | M | P | P | M | M | |
| | M | P | P | P | M | 1 |
| | M | M | P | M | P | |
| | M | P | M | P | P | |
| | M | P | M | M | P | 2 |
| | M | P | P | M | P | |
| | P | P | P | M | P | |
| | P | P | M | P | P | |
| | P | M | P | P | P | |
| | P | M | P | M | P | |
| | P | M | M | M | P | |
| | P | M | M | P | P | |
| | P | P | M | M | P | 1 |
| | P | M | P | P | M | 1 |
| | P | M | M | P | M | |
| | P | P | M | P | M | |
| | P | M | P | M | M | |
| | | | | | Total | 6 |

**Table Notes:** The columns indicate the possible choices that previous subjects could have made for the 5 player game, including the value of $\pi_m$, while the rows indicate the possible combinations of choices for the 5th player. The left most column indicates which combinations of choices can be classified into "conformers", "non-conformers", "random-choosers", and "always chose pen or mug". "Conformers" and "non-conformers" based on the definition above, and our model. Column 6 refers to the total number of observations in the 5-player version that made that combination of choices. M indicates "Mugs" and P indicates "Pens".

of random choosers to 11. This is an important detail in doing our hypothesis tests.

In table 3.4, column 1, I show the main statistical test, which is whether or not I can reject that our data may have been generated by individuals randomly choosing between mugs and pens. To conduct such a test, I construct the null distribution by assuming that the possible combinations of choices in the last row are the result of random choices. If mugs and pens are equally likely to be chosen (a fact that cannot be rejected from our data) then the likelihood of finding a combination of choices consistent with conformity is $8/30 = .267$ (I remove the individuals who always chose mug or pen from the set of possible combinations otherwise it would be 8/32.) In total, 10 individuals made choices consistent with a social preference from a sample of 21 individuals who did not always choose the same object. Thus, using the probability of observing 10 or more successes in 21 independent trials with probability of success of .267 is .032. I can therefore reject the hypothesis that the data observed is random at the 5% level.

Overall, our findings in the 5-player game suggest that the distribution of objects impacts ones preferences, consistent with the central prediction of our model.

**3-player game**

Turning our attention to the 3-player experiment without any initial endowment, I conduct the same exercise in classifying conformity and non-conformity using data from the last row of the game tree when all decisions by other players have been made (Table 3.5). Table 3.5 is an analogous version of table 3.4 but done for the 3-player experiment, with columns indicating the possible previous choices subject could have observed and the rows indicating possible combinations of the 3rd player. The possible previous choices are "Mug, Mug", "Mug, Pen" and "Pen, Pen" with corresponding values $\pi_m = 1, 0.5, 0$. In total there are $2^3$ possible combinations of choices the 3rd player can make of which 4 of them are consistent with our definition of conformity and non-conformity.

**Table 3.4.** Hypothesis Tests for Random Choice

| | Number of observations in data under each type of choice | | | |
|---|---|---|---|---|
| **Types of Choices** | **5 player** | **3 player without endowment** | | **3 player with endowment** |
| | | using 3rd player decisions | using 2nd & 3rd player decisions | |
| | (1) | (2) | (3) | (4) |
| **Chose Mug or pen for all decisions** | 29 | 10 | 10 | 9 |
| **Choices consistent with conformity (C)** | 5 | 1 | 1 | 0 |
| **Choices consistent with non-conformity NC)** | 5 | 3 | 2 | 3 |
| **Choices classified as random (R)** | 11 | 1 | 2 | 3 |
| N = of trials = C + NC + R) | 21 | 5 | 5 | 6 |
| k = # of success C + NC | 10 | 4 | 3 | 3 |
| Proportion of individuals identified as C or NC (k/N) | 10/21 | 4/5 | 3/5 | 3/6 |
| Probability of observing data under random choice: | 8/30 | 4/6 | 8/30 | *** |
| **Probability of obtaining k or more successes** | 0.032 | 0.46 | 0.12 | 0.435 |

**Table Notes:** C refers to "conformers" , NC refers to "non-conformers and "R" refers to random choosers. The columns indicate the version of the experiment. The first 4 rows of the left most column indicates the classification of people into "types" based on their choices. The last row shows results of a hypothesis test (the p-value) that the proportion of individuals classified as conformers and non-conformers could have been generated by random choice. *** indicates the calculation done for 3 player with endowment which combines the probability of conformity and non-conformity when making 3 random choices and when making 4 random choices since some individuals made 3 choices and some made 4 choices as the last player for the "with endowment" version.

Out of 15 subjects, 10 of them always chose the mug or always chose the pen (and did so in prior rows of the game as well). Of the remaining subjects, 4 of them made choices consistent with the definition of conformity or non-conformity used in the model. 3 of them appear to be non-conformers, while 1 of them appears to be a conformer. 1 of them appears to be making random choices.

The chance of seeing the combinations of choices consistent with conformity and non-conformity randomly is 4/6 which is significantly higher than in the 5-player game where it was 8/30. I conduct the same hypothesis test for the 3-player game as I did for 5-players. I use the binomial distribution to construct the probability of observing

4 successes or greater in 5 trials when the probability of success is 4/6 or .6666. The probability of such en event is 0.46 (as shown at the end of column 2 in table 4) and therefore, I fail to reject the null that the social preferences observed in the 3-player game could have been generated by random choice.

For the 3-player game with endowment, I find that 9 subjects chose either the mug or the pen always. Of the remaining 6 individuals, 3 of them satisfied the definition of non-conformity and the rest were random choosers (table 3.5). However, I fail to reject that this data could be observed through random choice as shown in table 3.4, column 4.

**Table 3.5.** Identifying Conformity and Non-conformity from Last Row Decisions in 3 player game without endowment

| **Type of Choices Made By 3rd player** | Possible Combinations of decisions by previous 2 players | | | |
|---|---|---|---|---|
| | **MM** | **MP** | **PP** | **Count** |
| | $\pi_m = 1$ | $\pi_m = .5$ | $\pi_m = 0$ | |
| | (1) | (2) | (3) | (4) |
| **Always chose mug or pen** | M | M | M | 9 |
| | P | P | P | 1 |
| | | | Total | 10 |
| **Choices consistent with conformity** | M | M | P | |
| | M | P | P | 1 |
| | | | Total | 1 |
| **Choices consistent with non-conformity** | P | M | M | 3 |
| | P | P | M | |
| | | | Total | 3 |
| **Choices classified as random** | P | M | P | 1 |
| | M | P | M | |
| | | | Total | 1 |

**Table Notes:** The columns indicate the possible choices that previous subjects could have made for the 3 player game, including the value of $\pi_m$, while the rows indicate the possible combinations of choices for the 3rd player. The left most column indicates which combinations of choices can be classified into "conformers", "non-conformers", "random-choosers", and "always chose pen or mug". "Conformers" and "non-conformers" based on the definition above, and our model. Column 6 refers to the total number of observations in the 5-player version that made that combination of choices. M indicates "Mugs" and P indicates "Pens".

### 3.4.5 Identifying social preferences using all decisions

In this section, I check if individuals who displayed social preferences in the last row made decisions in prior rows that are rationalizable. I use the following criterion as defined in section 3.2.6: (1) Weak rationalizability, which accords with the standard definition of rationalizability in a game in that there is some admissible belief on future play that is consistent with the choice made (Bernheim, 1984). (2) a stricter definition of rationalizability I call strong rationalizability, based on a stronger assumption of consistency in beliefs about future play.

#### Rationalizability Results

Table 3.6, columns 1 and 2 report the results of the two rationalizability tests for individuals who made consistent choices in the last row. These are the individuals who were either conformers or non-conformers discussed in the last section. The top panel reports the results for the 5-player game and the bottom panel for the 3-player game without endowment.

Out of 10 individuals in the 5-player game, only 2 individuals violated weak rationalizability. However, when considering strong rationalizability, based on consistent beliefs, 5 of the 10 individuals violate this notion, of which 4 were conformers. This suggests that perhaps the preference for non-conformity is more stable.

Out of 4 individuals in the 3-player without endowment game, only 1 out of 4 individuals who was a non-conformer violated weak rationalizability. That same individual is the only one who violated strong rationalizability as well by always selecting $P$ when making a choice as the 2nd player at all nodes, but choosing $M$ after seeing a mug and a pen as the third player.

One way to make our analysis of rationality more comprehensive is to count the minimum number of choices one would need to change for each individual to such that they would remain consistent with the model.

**Conformity and Non-conformity Across Rows of the game**

Columns 3 to 10 of table 3.6 indicate whether an individual who was classified as a conformer or non-conformer in the last row, displayed preferences for conformity and non-conformity in earlier rows as well. This provides some indication of their beliefs about future play.

Overall, I see that non-conformers are more likely to display preferences for non-conformity at earlier stages of the game. In the 5-player versions, out of the 5 non-conformers, 2 of them were non-conformers in every row, 1 of them was a non-conformer for the last 3 rows, and 2 of them were only non-conformers in the last row. This gives us an indication of how strong their preference for non-conformity was. Compared to the non-conformers, out of the 5 conformers, only 1 was a conformer in every row, 1 was a conformer in the last 3 rows, 1 of them was a conformer for the last 2 rows, and 2 of them were a conformer for the last row. [14]

Turning to the 3-player without endowment version, of the 4 individuals who were identified as conformers or non-conformers when choosing as the last player I see that only one of the non-conformers appeared to be a non-conformer choosing as both the second and the third player.

What I learn from the analysis of rationalizability is that non-conformity preferences appear more stable in the data relative to conformity preferences as non-conformers made less violations. I also find suggestive evidence that non-conformers are more confident in the sequences of choices they will observe at the end of the game. However, there is not much evidence of greater non-conformity than conformity when $\pi_m = 0.75$ in the data (a difference of one observation).

[14] In terms of individuals switching between choices that reveal conformity and non-conformity across the game, I find only 2 instances of it across the 17 individuals who identified as a conformer or non-conformer in all the experiments run.

**Table 3.6.** Violations of Rationality Across All Decisions - 3 and 5 player

| ID | Type in last row | Violation of Rationalizability (1 = Violation, 0 = No Violation) | Violation of Belief Consistency (1 = Violation, 0 = No Violation) | Conformer = 1 | | | | Non-conformer = 0 | | | | Switch C/NC |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) |
| | | | | **Decision Row** | | | | | | | | |
| | | | | **2** | **3** | **4** | **5** | **2** | **3** | **4** | **5** | |
| | | | | 5 player | | | | | | | | |
| 1 | NC | 1 | 1 | 0 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | 1 |
| 2 | NC | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 |
| 3 | NC | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 |
| 4 | NC | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 0 |
| 5 | NC | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 |
| 6 | C | 0 | 1 | 1 | 0 | 1 | 1 | 0 | 0 | 0 | 0 | 0 |
| 7 | C | 0 | 1 | 1 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 |
| 8 | C | 0 | 1 | 0 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 |
| 9 | C | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 |
| 10 | C | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 |
| Total Violations | | 2 | 5 | | | | | | | | | 1 |
| | | | | 3 player without endowment | | | | | | | | |
| 11 | NC | 0 | 0 | 0 | 0 | - | - | 0 | 1 | - | - | 0 |
| 12 | NC | 0 | 0 | 0 | 0 | - | - | 1 | 1 | - | - | 0 |
| 13 | NC | 1 | 1 | 0 | 0 | - | - | 0 | 1 | - | - | 0 |
| 14 | C | 0 | 0 | 0 | 1 | - | - | 0 | 0 | - | - | 0 |
| Total Violations | | 1 | 1 | | | | | | | | | 0 |

**Table notes:** "NC" refers to Non-conformers, "C" refers to Conformers. The top panel calculates rationality for the 5-player version and the bottom panel for the 3-player version. Column 1 indicates whether an individual violated the general notion of rationalizability, and column 2 indicates violation of the stricter notion of rationality where beliefs across rows must be consistent. Columns 3-10 indicate whether the individual indicated choices that would classify them as C or NC when choosing as the 2nd, 3rd, 4th, etc. player, as indicated by the heading "decision-row". Switch C/NC refers to whether the individuals switched between conformer and non-conformer between rows.

### 3.4.6 Endowment Effect Under Social Comparisons

One of the motivating questions of the paper is to test whether the endowment effect survives in an environment with social pressures to exchange. Under our model, as the proportion of individuals who own a particular object increases, social preferences could lead to a dampening of the endowment effect. I examine this by comparing the proportion of times individuals choose mug (pen) when they are randomly endowed with one, to the proportion of times it is chosen in the treatments without any endowments. Since individuals are randomly assigned to treatments, I assume that any difference in these proportions must come from being endowed with the object. I present the results in table 3.7.

As can be seen in column 1 which looks at only the 3-player version, the proportion of mug choices for individuals who were ”endowed with mug” is approximately -11.5 p.p. , and insignificant (p = .51) showing that subjects who were endowed with the mug were no more likely to choose the mug than subjects who were not given any endowment. In column 2, I conduct the same analysis, but also add individuals from the 5-player version. We see that it now appears that individuals endowed with the mug are more likely to choose it compared to the no endowment treatment. However, this is primarily because individuals in the 5-player treatment happened to choose the mugs less than the 3-player treatment overall. I account for this in columns (3) and (4), first by explicitly controlling for fixed effects, and then by weighting the observations by the inverse probability that they are included based on the number of observations for that version of the experiment. I find that while the point estimate in column 4 is positive, individuals endowed with the mug are about 6.3 p.p. more likely to keep the mug, but this is insignificant (p = .69). Thus, this could reflect a lack of power to detect the endowment effect, but combined with the other evidence, it is more suggestive that the endowment effect is crowded out.

The identical analysis is conducted in columns 4-8 for the likelihood of choosing pen conditional on being endowed with the pen. In column 4, restricted to the 3-player version, I find that individuals do appear to be more likely to choose the pen given they are endowed with it. However, given the result on mugs being -11.5 p.p. for this same sample, taken together, this appears to be just a difference in preferences in consumption utility between pens and mugs as opposed to evidence of an endowment effect. Columns 5-8 where I incorporate data from version 5 illustrate that being endowed with a pen leads to no greater likelihood of choosing the pen, and I would consider this a null result.

Finally, the proportion of conformers and non-conformers found in the 3-player version without endowment is .8 ( 4 out of 5) while in the with endowment it is .5 ( 3 out of 6). While this is a large difference in magnitude, the sample is too small to demonstrate whether endowing subjects with a good lowers the amount of conformity and non-conformity observed.

**Table 3.7.** Evidence of an Endowment Effect

| **Dependent Variable :** | Chose Mug = 1 | | | | Chose Pen = 1 | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Endowed with mug | -0.115 | 0.075 | -0.115 | 0.063 | | | | |
| | (0.172) | (0.155) | (0.169) | (0.155) | | | | |
| Endowed with pen | | | | | 0.204 | -0.014 | 0.014 | 0.026 |
| | | | | | (0.201) | (0.185) | (0.197) | (0.185) |
| Version 5 = 1 | | | -0.213 | | | | 0.213 | |
| | | | (0.103)** | | | | (0.103)** | |
| Constant | 0.767 | 0.576 | 0.767 | 0.589 | 0.233 | 0.424 | 0.233 | 0.411 |
| | (0.086) | (0.054) | (0.085) | (0.052) | (0.087) | (0.054) | (0.085) | (0.052) |
| Observations | 156 | 906 | 906 | 906 | 138 | 888 | 888 | 888 |
| Versions | only 3 | all | all | all | only 3 | all | all | all |
| Probability Weights For Version | No | No | No | Yes | No | No | No | Yes |
| R-squared | 0.016 | 0.002 | 0.018 | 0.002 | 0.045 | 0.000 | 0.017 | 0.000 |
| P-value Starting with object = 0 | 0.509 | 0.629 | 0.497 | 0.687 | 0.321 | 0.941 | 0.303 | 0.887 |

Robust Standard Errors in Parentheses. p<0.1*, p<0.05**, p<0.01*** Unit of observation in all regressions is at the individual-choice level. Regressions are all clustered at the participant level.

The explanation I consider for not finding an endowment effect is the one I provide in this paper, which is that preferences to conform and not conform seem to dominate. However, an alternate explanation is the one put forth by (Goette, Graeber, Kellogg, and Sprenger, 2018) which is that if there is heterogeneity in loss aversion , that is, if there are some individuals for whom individual gain-loss utility, $\lambda_I < 1$, then this would lead to a desire to exchange their endowed object. These subjects can be thought of as "gain-seeking". Thus, sample sizes in typical studies of the endowment effect are an order of magnitude too small to detect it, even when the proportion of loss averse individuals is larger than those who are "gain-seeking".

## 3.5 Discussion and Further Work

Overall, the paper presents two pieces of evidence on how the social distribution of consumption impacts preferences: (1) A small but statistically significant portion of individuals respond to other people's choices in a way that is consistent with a utility function in which individuals compare their consumption bundle to others around them. (2) When compared to a treatment where individuals are initially endowed with an object, individuals are no less likely to prefer the object than in the no-endowment case, although the study is underpowered to detect small endowment effects.

While I find a small overall treatment effect as only 22% of individuals demonstrate a social preference in our sample, one of the questions for future work is what would happen if the social pressure we implemented in our experiment was increased. Instead of using the strategy method, where individuals state their preference at all possible nodes of the game, I could implement a design where individuals take turns making a single decision in real time so that only one of the games potential outcomes is realized. Imagine perhaps that individuals physically go up to the front of the room and choose between the pen and mug showing others what they have done. This may increase the amount of social preference observed overall, but it could potentially also change the ratio of conformers to non-conformers, as individuals may feel a stronger urge to conform.

One concern that has been brought up about our design is the possibility of an experimenter demand effect as under the strategy method, subjects are asked choose between the mug and the pen several times. The concern is that this may lead them to think that the experimenter wants them to make different decisions each time. However, I argue that even if this were true, it is not clear that subjects have any idea what model the experimenter has in mind or what a“correct” decision is. As such, I believe that conditional on there being variation in subject’s answers, we are likely to be learning something valuable about preferences for conformity and non-conformity at the individual level. If anything, I believe the results present a lower bound of the true extent of social preferences in an exchange setting since the strategy method itself may be the most subtle treatment one could do, but one is best suited for identification.

## 3.6 Conclusion

There is substantial evidence that economic choices are greatly influenced by the actions of our peers in a variety of domains, including education (Lazear, 2001), firm decisions (Kaustia and Rantala, 2015), paying taxes (Perez-Truglia and Troiano, 2015), and energy consumption (Allcott, 2011), to name a few. In this paper, I focus on an environment which has largely ignored these social factors, behavior in exchange settings. Reference dependence models have typically been used to study behavior in these settings and have focused on explaining the endowment effect, the finding that individuals’ initial ownership of a good increases its valuation. I provide evidence that in a standard exchange setting, individual preferences over consumption goods are a function of the observed distribution of consumption in society. I find that a small, but significant portion of individuals are conformers, displaying a preference for the object that has been chosen the most, and an equally large number are non-conformers, showing a preference for the object chosen least.

To explain the findings, I present a model of reference dependence with an additional reference point to explain why individuals may respond to the social distribution. Our model abstracts from social learning and social signaling, arguing that the desire to conform or not conform is driven simply by the gain-loss utility derived from making social comparisons. Our experiment controls for social learning, by allowing subjects to examine both objects at the start, and for social signaling, by having subjects play the game against anonymous peers. This suggests that even in a basic exchange environment, individuals display social preferences. Additionally, contrary to most models and empirical results that find evidence mostly in favor of conformity, I find that I cannot reject equality of the ratio of non-conformers to conformers in our sample. This constitutes further evidence that perhaps the mechanisms that drive emulation and deviation in our setting are indeed different from the mechanisms that drive social preference in other settings where social learning and social signaling are likely to play bigger roles.

A second result concerns the existence of the endowment effect in a social setting. In an additional treatment where individuals are first endowed with an object and then provided information on the social distribution, the endowment effect appears to dampen. Individuals endowed with a given object are no more likely to choose it than the comparison group who were not assigned anything. However, given recent findings in the literature that not all individuals are loss averse and some may have a preference to trade away their assigned object, our study may be underpowered to detect small endowment effects (Goette, Graeber, Kellogg, and Sprenger, 2018).

Overall, our results suggest that a preference to emulate or deviate from the consumption of choices of others could mitigate exchange asymmetries in settings where the endowment effect is observed. This may be applicable to recent studies that show endowment effects in applied settings. For example, Carney, Lin, Kremer, and Rao (2018) provide evidence from a field-experiment that the endowment effect causes borrowers to dislike taking loans in which the collateral is an asset they already own and is therefore already a part of their reference point. While the prescribed solution is to collateralize loans using the new asset, our paper suggests that simple information on the the choices others have made may reduce the endowment effect as individuals compare themselves to the group. However, our data suggests that information must be more subtle and specific than that. In order to induce action "a" by a conformer, one should show information that more people are taking the action, while inducing action "a" by a non-conformer requires providing information that fewer people are taking the action. In general, a large portion of non-conformers in a sample could explain the muted effects of providing information on other people's behavior, and it may be important to know what the composition of social preferences is in a given population.

Chapter 3, in full is currently being prepared for submission for publication of the material. The dissertation author, Vinayak Alladi was the primary investigator and author of this material.

# 3.7 Appendix

## 3.7.1 Comparative Statics of Social KR Model

### Changing the value of individual loss aversion

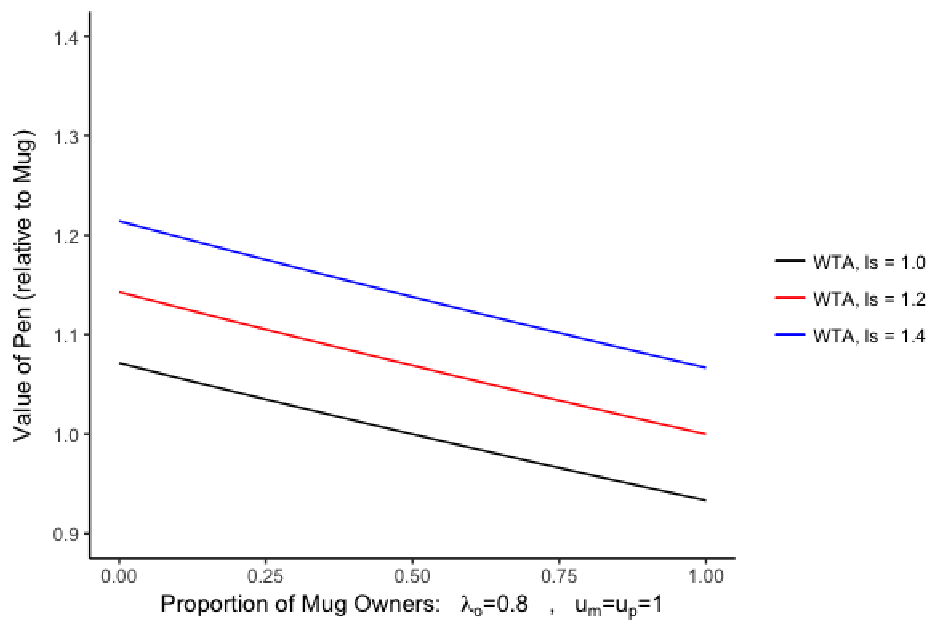


**Figure 3.6.** Impact of Changes in Individual Gain-Loss Utility on Personal Equilibrium Thresholds

### Comparative Statics of Equilibrium Threshold for Pen Owners

$$\frac{d\underline{u_p}^{\text{keep pen}}}{d\pi_m} = \frac{u_m(\lambda_s - 1)(\lambda_I + \lambda_s + 4)}{(1 + \pi_m(1 - \lambda_s) + \lambda_s + \lambda_I)^2} \tag{3.13}$$

## 3.7.2 Preference Consistency

One question of interest amongst individuals who did not always choose the same object is whether or not they are answering completely randomly or something can still be learned about their underlying preferences, as this speaks to how subjects thought about the questions they were being asked.

To do this, I focus on the decisions made when the proportion of mugs chosen thus far by previous subjects is 50%.[15] The columns of table 9 show the proportion of times subjects choose mug in these 50-50 situations while the rows show the percentage of times the mug is chosen overall across all situations. This is done for the 5-player and 3-player games combined. As can be seen, there is a strong correlation between choosing mug more often in the 50-50 cases and choosing mugs overall. This serves as a check for the consistency of our interpretation of the preferences being expressed in the data.

**Table 3.8.** Correlation between choices at a 50-50 distribution and other choices

| Proportion of Times Mug Chosen | Chooses Mugs at least 50% of the time when proportion of mugs chosen thus far is 50% | | | | |
|---|---|---|---|---|---|
| | **0** | **0.333** | **0.5** | **0.666** | **1** |
| 6.7% | 1 | 0 | 0 | 0 | 0 |
| 13.3% | 1 | 0 | 0 | 0 | 0 |
| 16.7% | 1 | 0 | 0 | 0 | 0 |
| 20.0% | 1 | 0 | 0 | 0 | 0 |
| 26.7% | 1 | 0 | 0 | 0 | 0 |
| 33.3% | 1 | 0 | 0 | 0 | 0 |
| 40.0% | 0 | 1 | 0 | 0 | 0 |
| 46.7% | 0 | 0 | 0 | 2 | 0 |
| 50.0% | 0 | 0 | 1 | 0 | 1 |
| 53.3% | 0 | 0 | 0 | 2 | 0 |
| 60.0% | 0 | 0 | 0 | 2 | 2 |
| 66.7% | 0 | 0 | 0 | 0 | 2 |
| 73.3% | 0 | 0 | 0 | 0 | 2 |
| 80.0% | 0 | 0 | 0 | 1 | 1 |
| 93.3% | 0 | 0 | 0 | 0 | 3 |
| | 6 | 1 | 1 | 7 | 11 |

### 3.7.3 With Endowment Last Row

[15]This occurs twice in the 3-player game, when no decisions have been made and when one person has chosen a mug and one person has chosen a pen, and three times in the 5-player, with the addition of two people have chosen mugs and two have chosen pens.

**Table 3.9.** Identifying Conformity and Non-conformity from Last Row Decisions in 3 player game with endowment

| **Endowment (1 = Mugs,0 = Pens) Exchanged (Y = Yes, N = No)** | **1N1N** | **1N1Y** | **1N0Y** | **1N0N** | **1Y1Y** | **0Y0Y** | **1Y0Y** | **1Y0N** | **0Y0N** | **0N0N** | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **Possible Combinations of decisions by previous 2 players (columns)** | **MM** | **MP** | **MM** | **MP** | **PP** | **MM** | **PM** | **PP** | **MP** | **PP** | **Count** |
| | $\pi_m = 1$ | $\pi_m = .5$ | $\pi_m = 1$ | $\pi_m = .5$ | $\pi_m = 0$ | $\pi_m = 1$ | $\pi_m = .5$ | $\pi_m = 0$ | $\pi_m = .5$ | $\pi_m = 0$ | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | |
| **Choose Mug or pen for all decisions** | M | M | . | . | M | . | . | . | . | . | |
| | P | P | . | . | P | . | . | . | . | . | |
| | . | . | M | M | . | . | M | M | . | . | |
| | . | . | P | P | . | . | P | P | . | . | |
| | . | . | . | . | . | M | . | . | M | M | |
| | . | . | . | . | . | P | . | . | P | P | |
| | | | | | | | | | | Total | 11 |
| **Choices consistent with conformity** | M | M | . | . | P | . | . | . | . | . | |
| | M | P | . | . | P | . | . | . | . | . | |
| | . | . | M | P | . | . | P | P | . | . | |
| | . | . | M | M | . | . | M | P | . | . | |
| | . | . | . | . | . | M | . | . | M | P | |
| | . | . | . | . | . | M | . | . | P | P | |
| | | | | | | | | | | Total | 0 |
| **Choices consistent with non-conformity** | P | M | . | . | M | . | . | . | . | . | |
| | P | P | . | . | M | . | . | . | . | . | |
| | . | . | P | P | . | . | P | M | . | . | |
| | . | . | P | M | . | . | M | M | . | . | |
| | . | . | . | . | . | P | . | . | M | M | |
| | . | . | . | . | . | P | . | . | P | M | |
| | | | | | | | | | | Total | 3 |
| **Choices classified as random** | others | | | | | | | | | | |
| | | | | | | | | | | Total | 1 |

"M" stands for mugs, and "P" stands for pens

# Bibliography


Ahn, David, Syngjoo Choi, Douglas Gale, and Shachar Kariv (2007). "Estimating ambiguity aversion in a portfolio choice experiment". In:

Allcott, Hunt (2011). "Social norms and energy conservation". In: *Journal of public Economics* 95.9-10, pp. 1082–1095.

Andreoni, James, Tobias Schmidt, and Charles Sprenger (2015). "Measuring Ambiguity Aversion: Experimental Tests of Subjective Expected Utility". In:

Anscombe, Francis J and Robert J Aumann (1963). "A definition of subjective probability". In: *The annals of mathematical statistics* 34.1, pp. 199–205.

Anyon, Jean (1997). *Ghetto schooling: A political economy of urban educational reform.* Teachers College Press.

Attanasio, O. and K. Kaufmann (2012). "Education choices and returns on the labor and marriage markets: Evidence from data on subjective expectations". In: *Working Paper.*

Aumann, Robert J (1976). "Agreeing to disagree". In: *The annals of statistics*, pp. 1236–1239.

Banerjee, Abhijit V (1992). "A simple model of herd behavior". In: *The quarterly journal of economics* 107.3, pp. 797–817.

Banerjee, Abhijit V, Abhijit Banerjee, and Esther Duflo (2011). *Poor economics: A radical rethinking of the way to fight global poverty.* Public Affairs.

Banerjee, Abhijit, Esther Duflo, Nathanael Goldberg, Dean Karlan, Robert Osei, William Parienté, Jeremy Shapiro, Bram Thuysbaert, and Christopher Udry (2015). "A multifaceted program causes lasting progress for the very poor: Evidence from six countries". In: *Science* 348.6236, p. 1260799.

Beaman, Lori, Esther Duflo, Rohini Pande, and Petia Topalova (2012). "Female leadership raises aspirations and educational attainment for girls: A policy experiment in India". In: *science*, p. 1212382.

Bell, David E (1985). “Disappointment in decision making under uncertainty”. In: *Operations research* 33.1, pp. 1–27.

Bénabou, Roland and Jean Tirole (2011). “Identity, morals, and taboos: Beliefs as assets”. In: *The Quarterly Journal of Economics* 126.2, pp. 805–855.

— (2016). “Mindful economics: The production, consumption, and value of beliefs”. In: *The Journal of Economic Perspectives* 30.3, pp. 141–164.

Bernard, Tanguy, Stefan Dercon, Kate Orkin, and Alemayehu Taffesse (2014). *The Future in Mind: Aspirations and Forward-Looking Behaviour in Rural Ethiopia*. Tech. rep. CEPR Discussion Papers.

Bernheim, B Douglas (1984). “Rationalizable strategic behavior”. In: *Econometrica: Journal of the Econometric Society*, pp. 1007–1028.

— (1994). “A theory of conformity”. In: *Journal of political Economy* 102.5, pp. 841–877.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992). “A theory of fads, fashion, custom, and cultural change as informational cascades”. In: *Journal of political Economy* 100.5, pp. 992–1026.

Brunnermeier, Markus and Jonathan Parker (2005). “Optimal Expectations”. In: *American Economic Review* 95.4, pp. 1092–1118.

Bursztyn, Leonardo, Florian Ederer, Bruno Ferman, and Noam Yuchtman (2014). “Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions”. In: *Econometrica* 82.4, pp. 1273–1301.

Bursztyn, Leonardo, Georgy Egorov, and Robert Jensen (2017). “Cool to be smart or smart to be cool? Understanding peer pressure in education”. In: *The Review of Economic Studies*.

Carney, Kevin, Xinyue Lin, Michael Kremer, and Gautam Rao (2018). *The Endowment Effect and Collateralized Loans*. Tech. rep. Working Paper.

Cohen, Jessica, Pascaline Dupas, et al. (2010). “Free distribution or cost-sharing? Evidence from a randomized malaria prevention experiment”. In: *Quarterly journal of Economics* 125.1, p. 1.

Coller, Maribeth and Melonie B Williams (1999). “Eliciting individual discount rates”. In: *Experimental Economics* 2.2, pp. 107–127.

Coutts, Alexander (2015). “Testing Models of Belief Bias: An Experiment”. In: *Browser Download This Paper*.

Dato, Simon, Andreas Grunewald, Daniel Müller, and Philipp Strack (2017). "Expectation-based loss aversion and strategic interaction". In: *Games and Economic Behavior* 104, pp. 681–705.

David McKenzie John Gibson, Steven Stillman (2007). "A land of milk and honey with streets paved with gold: do emigrants have over-optimistic expectations about incomes abroad?" In: *World Bank Policy Research Working Paper*.

Delavande, Adeline (2008). "Pill, patch, or shot? Subjective expectations and birth control choice". In: *International Economic Review* 49.3, pp. 999–1042.

Delavande, Adeline, Xavier Giné, and David McKenzie (2011). "Measuring subjective expectations in developing countries: A critical review and new evidence". In: *Journal of development economics* 94.2, pp. 151–163.

Duflo, Esther, Pascaline Dupas, and N. Benhassine (2013). In: *Working Pape NBER*.

Dupas, Pascaline (2011). "Health behavior in developing countries". In: *Annu. Rev. Econ.* 3.1, pp. 425–449.

Ericson, Keith M Marzilli and Andreas Fuster (2011). "Expectations as endowments: Evidence on reference-dependent preferences from exchange and valuation experiments". In: *The Quarterly Journal of Economics*, qjr034.

Fehr, Ernst and Klaus M Schmidt (1999). "A theory of fairness, competition, and cooperation". In: *The quarterly journal of economics* 114.3, pp. 817–868.

Festinger, Leon (1954). "A theory of social comparison processes". In: *Human relations* 7.2, pp. 117–140.

Genicot, Garance and Debraj Ray (2017). "Aspirations and inequality". In: *Econometrica* 85.2, pp. 489–519.

Goette, Lorenz, Thomas Graeber, Alexandre Kellogg, and Charles Sprenger (2018). "Heterogeneity of Loss Aversion and Expectations-Based Reference Points". In:

Goette, Lorenz, Annette Harms, and Charles Sprenger (2014). "Randomizing Endowments: an experimental study of rational expectations and reference-dependent preferences". In:

Halevy, Yoram (2007). "Ellsberg revisited: An experimental study". In: *Econometrica* 75.2, pp. 503–536.

Harmon, Nikolaj, Raymond Fisman, Emir Kamenica, et al. (2017). "Peer Effects in Legislative Voting". In: *American Economic Journal: Applied Economics*.

Harrison, Glenn W, Morten I Lau, and Melonie B Williams (2002). “Estimating individual discount rates in Denmark: A field experiment”. In: *American economic review*, pp. 1606–1617.

Harsanyi, John C (1968). “Games with incomplete information played by Bayesian players part II. Bayesian equilibrium points”. In: *Management Science* 14.5, pp. 320–334.

Heffetz, Ori and John A List (2014). “Is the endowment effect an expectations effect?” In: *Journal of the European Economic Association* 12.5, pp. 1396–1422.

Jensen, Robert (2010). “The (perceived) returns to education and the demand for schooling”. In: *The Quarterly Journal of Economics* 125.2, pp. 515–548.

Kahneman, Daniel, Jack L Knetsch, and Richard H Thaler (1991). “Anomalies: The endowment effect, loss aversion, and status quo bias”. In: *Journal of Economic perspectives* 5.1, pp. 193–206.

Kaustia, Markku and Ville Rantala (2015). “Social learning and corporate peer effects”. In: *Journal of Financial Economics* 117.3, pp. 653–669.

Kay, Aaron C, Maria C Jimenez, and Joim T Jost (2002). “Sour Grapes, Sweet Lemons, and the Anticipatory Rationalization ofthe Status Quo”. In:

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2005). “A smooth model of decision making under ambiguity”. In: *Econometrica* 73.6, pp. 1849–1892.

Kőszegi, Botond and Matthew Rabin (2006). “A model of reference-dependent preferences”. In: *The Quarterly Journal of Economics* 121.4, pp. 1133–1165.

Lazear, Edward P (2001). “Educational production”. In: *The Quarterly Journal of Economics* 116.3, pp. 777–803.

Lee, David S (2009). “Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects”. In: *The Review of economic studies* 76.3, pp. 1071–1102.

Loomes, Graham and Robert Sugden (1986). “Disappointment and dynamic consistency in choice under uncertainty”. In: *The Review of Economic Studies* 53.2, pp. 271–282.

Lybbert, Travis J and Bruce Wydick (2017). “Hope as aspirations, agency, and pathways: poverty dynamics and microfinance in Oaxaca, Mexico”. In: *The Economics of Poverty Traps*. University of Chicago Press.

— (2018). “Poverty, aspirations, and the economics of hope”. In: *Economic Development and Cultural Change* 66.4, pp. 000–000.

Lynn, M (1989). “Scarcity effects on desirability: mediated by assumed expensiveness?” In: *Journal of economic psychology* 10.2, pp. 257–274.

Markowitz, Harry (1952). “The utility of wealth”. In: *Journal of political Economy* 60.2, pp. 151–158.

Marzilli Ericson, Keith M and Andreas Fuster (2011). “Expectations as endowments: Evidence on reference-dependent preferences from exchange and valuation experiments”. In: *The Quarterly Journal of Economics* 126.4, pp. 1879–1907.

Moretti, Enrico (2011). “Social learning and peer effects in consumption: Evidence from movie sales”. In: *The Review of Economic Studies* 78.1, pp. 356–393.

Oster, Emily, Ira Shoulson, and E Dorsey (2013). “Optimal expectations and limited medical testing: evidence from Huntington disease”. In: *American Economic Review* 103.2, pp. 804–30.

Perez-Truglia, Ricardo and Ugo Troiano (2015). *Shaming tax delinquents: Theory and evidence from a field experiment in the United States*. Citeseer.

Pyszczynski, Tom (1982). “Cognitive strategies for coping with uncertain outcomes”. In: *Journal of Research in Personality* 16.3, pp. 386–399.

S. Sequeira J. Spinnewijn, G Xu (2013). “Rewarding Schooling Success and Perceived Returns to Education: Evidence from India”. In: *Working Paper*.

Sacerdote, Bruce (2001). “Peer effects with random assignment: Results for Dartmouth roommates”. In: *The Quarterly journal of economics* 116.2, pp. 681–704.

Savage, Leonard J (1954). “The Foundations of Statistics, New York”. In: *SavageThe Foundations of Statistics 1954*.

Schmeidler, David and Itzhak Gilboa (2004). “Maxmin expected utility with non-unique prior”. In: *Uncertainty in Economic Theory*. Routledge, pp. 141–151.

Schwardmann, Peter (2019). “Motivated health risk denial and preventative health care investments”. In: *Journal of Health Economics*.

Shalev, Jonathan (2000). “Loss aversion equilibrium”. In: *International Journal of Game Theory* 29.2, pp. 269–287.

Sprenger, Charles (2015). “An endowment effect for risk: Experimental tests of stochastic reference points”. In: *Journal of Political Economy* 123.6, pp. 1456–1499.

Thaler, Richard (1980). “Toward a positive theory of consumer choice”. In: *Journal of Economic Behavior & Organization* 1.1, pp. 39–60.

Tversky, Amos and Daniel Kahneman (1991). “Loss aversion in riskless choice: A reference-dependent model”. In: *The quarterly journal of economics*, pp. 1039–1061.

Veblen, Thorstein (1899). *The theory of the leisure class: An economic study in the evolution of institutions*. Macmillan.

Verhallen, Theo MM (1982). “Scarcity and consumer choice behavior”. In: *Journal of Economic Psychology* 2.4, pp. 299–322.

Von Neumann, John and Oskar Morgenstern (1944). *Theory of games and economic behavior*. Princeton university press.

Weinstein, Neil D (1982). “Unrealistic optimism about susceptibility to health problems”. In: *Journal of behavioral medicine* 5.4, pp. 441–460.

Zimmerman, David J (2003). “Peer effects in academic outcomes: Evidence from a natural experiment”. In: *Review of Economics and statistics* 85.1, pp. 9–23.