Essays in Development and Behavioral Economics

Essays in Development and Behavioral Economics

A dissertation presented
by

Frank N Schilbach

to

The Department of Economics

in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
in the subject of
Economics

Harvard University
Cambridge, Massachusetts
May 2015

Dissertation Advisor: Author:
Professor Sendhil Mullainathan **Frank N Schilbach**

Essays in Development and Behavioral Economics

Abstract

This dissertation consists of three empirical essays in development and behavioral economics.

Chapter 1 considers the impact of heavy alcohol consumption on savings behavior among low-income males in India. High levels of alcohol consumption are more common among the poor. This fact could have economic consequences beyond mere income effects because alcohol impairs mental processes and decision-making. Since alcohol is thought to induce myopia, this paper tests for impacts on self-control and on savings behavior. In a three-week field experiment with low-income workers in India, I provided 229 individuals with a high-return savings opportunity and randomized incentives for sobriety. The incentives significantly reduced daytime drinking as measured by decreased breathalyzer scores. This in turn increased savings by approximately 60 percent. No more than half of this effect is explained by changes in income net of alcohol expenditures. In addition, consistent with enhanced self-control due to lower inebriation levels, incentivizing sobriety *reduced* the impact of a savings commitment device. Finally, alcohol consumption itself is prone to self-control problems: over half of the study participants were willing to sacrifice money to receive incentives to be sober, exhibiting demand for commitment to increase their sobriety. These findings suggest that heavy alcohol consumption is not just a result of self-control problems, but also creates self-control problems in other areas, potentially even exacerbating poverty by reducing savings.

Chapter 2 (with Esther Duflo, Michael Kremer, and Jon Robinson) investigates agricultural technology adoption in Sub-Saharan Africa. Insufficient knowledge of appropriate use can hamper technology adoption. In the agricultural context, if farmers do not observe

each others' inputs, diffusion of both information on the optimal input mix and of the technology itself may be slow. In the context we examine, conditional on using fertilizer, farmers tend to systematically overuse fertilizer (per treated area) on the intensive margin, hence, making it on average unprofitable and possibly curbing usage at the extensive margin. This paper reports results from a large-scale field experiment, which introduced a simple and salient tool, a blue measuring spoon, to help farmers remember how much fertilizer to use. A randomly selected subset of farmers received the technology for free, and the remaining farmers can purchase it at fertilizer stores at a nominal price. Farmers who were randomly assigned to receive a measuring spoon subsequently improved knowledge of how much fertilizer to use, and were more likely to use fertilizer. Spoon purchases among the remaining farmers were higher when these were more likely to use fertilizer due to a randomly assigned fertilizer discount program, and when communication about agriculture was encouraged. Unlike fertilizer adoption itself, purchase and use of measuring spoons diffused rapidly through social networks.

Chapter 3 (with Tom Zimmermann) provides new empirical evidence on trading behavior among individual investors. The main contribution of this essay is to contrast competing explanations of the disposition effect, investors' tendency to hold losing investments too long and to sell winning investments too soon, based on their predictions for realizing different sizes of gains and losses. We find that for all holding periods longer than one month and for both gains and losses, the probability to sell a stock declines monotonically with the size of the absolute return. This fact is not consistent with the model of realization utility, but it is consistent with a version of prospect theory as outlined below. Moreover, we find that investors' propensity to make any trade is largest for small absolute portfolio returns, a fact that is difficult to explain by the existing theories.

Contents

List of Tables

List of Figures

Acknowledgments

I would like first to thank my fantastic team of advisors, who provided encouragement and guidance throughout my time at Harvard. I am deeply grateful to my primary advisor, Sendhil Mullainathan, for his tireless support, motivation, and patience. I could not have asked for a better mentor. Esther Duflo's course sparked my interest in microeconomic development questions. Esther's rigorous approach to research, her drive, and thoughtfulness make her a unique role model that I aspire to be like. Michael Kremer's relentless pursuit of both knowledge and real-world impact through research is truly inspiring. His advice has made me a better researcher. David Laibson has improved my understanding of economic theory. I have enormously benefited from his suggestions and guidance. I thank my friends and classmates for support, instruction, laughter, and commiseration. Especially Benjamin Schoefer has weathered many days and nights in the Littauer basement with me. I also thank Kate Sturla, Luke Ravenscroft, Manasa Reddy, Andrew Locke, Nick Swanson, Louise Paul-Delvaux, and all of the research staff in Chennai for their outstanding field assistance. Finally, I would like to thank my family for their unconditional love and support. My parents and siblings shaped the person I am now. I have always felt their presence throughout my life, despite the geographical distance.

To my family.

Chapter 1

Alcohol and Self-Control: A Field Experiment in India

1.1 Introduction

Heavy alcohol consumption is correlated with poverty, yet the nature and consequences of this relationship are not well understood.[1] Poverty could cause demand for alcohol by enhancing its short-term benefits.[2] But alcohol may also be a cause of poverty. In particular, alcohol is thought to affect myopia and self-control. If these effects are large, then heavy alcohol consumption could interfere with a variety of forward-looking decisions. By affecting savings decisions, insurance take-up, human capital investments, and earnings, alcohol could reduce wealth accumulation and deepen poverty. However, though theoretically possible, we do not know whether such effects are present or economically meaningful in practice.

[1] In many countries, low-income individuals are in fact *more* likely to be abstinent from alcohol altogether. At the same time, in many countries including in India, heavy drinking is more common among the poor. This is described in more detail in the next section.

[2] Alcohol is known to be a powerful anesthetic (Woodrow and Eltherington (1988)), it helps individuals fall asleep (Ebrahim *et al.* (2013)), and it can make individuals feel better about themselves ("drunken self inflation," Banaji and Steele (1989)), or relieve stress and anxiety ("drunken relief," Steele and Josephs (1988)). At the same time, physical pain, poor sleep, low self-esteem, and stress are all correlated with poverty (Poleshuk and Green (2008), Patel *et al.* (2010), Haushofer and Fehr (2014), Patel (2007)).

This paper empirically tests for one such effect: the impact of alcohol on savings behavior. To examine this relationship, I conducted a three-week field experiment with 229 cycle-rickshaw peddlers in Chennai, India, in which all subjects were provided with a high-return savings opportunity. To create exogenous variation in alcohol consumption, a randomly selected subset of study participants were offered financial incentives for sobriety. For a cross-randomized subset of study participants, the savings account was a commitment savings account, i.e. individuals could not withdraw their savings until the end of their participation in the study. This feature allowed me to consider the impact of increasing sobriety on self-control problems in savings behavior. In addition, I elicited willingness to pay for incentives for sobriety to assess the extent to which self-control problems themselves contribute to the demand for alcohol.

The incentives for sobriety significantly increased study participants' sobriety during their daily savings decisions, providing a "first stage" to estimate the impact of sobriety on savings behavior. Individuals who were given incentives for sobriety decreased their daytime drinking as measured by a 33 percent increase in the fraction of individuals who visited the study office sober. The intervention also reduced overall alcohol consumption and expenditures by 5 to 10 percent.

Offering incentives for sobriety increased individuals' daily savings at the study office by 60 percent compared to a control group that received similar average study payments independent of their alcohol consumption. This increase in savings is a combination of changes in income net of alcohol expenditures, and changes in savings behavior for given resources. Assessing the contribution of the former requires an estimate of the marginal propensity to save out of available income. Using an estimate of the marginal propensity to save obtained by separately randomizing study payments via a lottery and observing the impact on savings, I find that the combined effects of increased earnings outside of the study and decreased alcohol expenditures explain about half of the observed increase in savings. The remaining share of the increase in savings appears to be due to the effect of alcohol on time preferences. Consistent with this, the estimated marginal propensity to save is almost twice as large for individuals who were offered incentives for sobriety as for individuals in the control group, though this difference is not statistically significant.

The relationship between the effects of sobriety incentives and commitment savings provides further evidence that increasing sobriety directly affects time preferences. In particular, I find that sobriety incentives and the commitment savings feature were substitutes in terms of their effect on savings. While commitment savings and sobriety incentives each individually increased subjects' savings, there was no additional effect of the savings commitment feature on savings by individuals who were offered sobriety incentives, and vice versa. These patterns are consistent with alcohol increasing present bias. An alternative interpretation is that the incentives mitigated the need for commitment savings by reducing the consumption of alcohol, a key temptation good for this population. However, the intervention mainly reduced drinking or shifted it to later times of the day rather than causing abstinence from alcohol altogether. This makes a direct effect of alcohol on time preferences the more likely explanation.

Over 50 percent of subjects exhibited demand for commitment to increase their sobriety, indicating a greater awareness of and willingness to overcome self-control problems than found in other settings, for instance for smoking (Gine *et al.* (2010)), or exercising (Royer *et al.* (2014)). Specifically, in three sets of weekly decisions that each elicited preferences for sobriety incentives in the subsequent week, over half of the study participants chose options that implied weakly dominated study payments. In addition, more than a third preferred incentives for sobriety over unconditional payments, even when the latter were *strictly* higher than the maximum amount subjects could earn with the incentives. These individuals were willing to sacrifice study payments of about ten percent of daily income even in the best case scenario of visiting the study office sober every day. This finding provides clear evidence for a desire for sobriety by making future drinking more costly, in contrast to the predictions of the Becker and Murphy (1988) rational addiction model.[3]

[3]Becker and Murphy (1988) showed that many behaviors of addicted individuals are, at least in theory, consistent with optimization based on stable preferences. Gruber and Kőszegi (2001) subsequently challenged the implicit assumption of time-consistent preferences and replaced it with hyperbolic discounting as formalized

The high demand for commitment does not appear to be the result of misunderstandings on the part of the subjects. Willingness to pay for sobriety incentives did *not* decrease over time among individuals who were asked to choose repeatedly. In fact, past exposure to the incentives *increased* individuals' demand for the incentives. Individuals who had been randomly selected to receive incentives for sobriety for 15 days were more likely to choose incentives for a subsequent week compared to individuals who had received payments independent of their sobriety. Further, individuals whose sobriety increased in response to the incentives were particularly likely to choose the incentives subsequently. Moreover, individuals with lower concurrent inebriation levels were more likely to choose the incentives. Finally, reassuringly, the demand for the incentives decreased in the cost of incentives.

The finding that alcohol *causes* self-control problems builds on psychology research on "alcohol myopia" (Steele and Josephs (1990)). This line of research sought to reconcile the seemingly contradictory effects of alcohol found in a large body of previous research. Depending on circumstances, alcohol can relieve or increase anxiety and tension. It can inflate egos, yet lead to depression. However, according to the "alcohol myopia" theory, a defining feature of alcohol is that it *always* narrows attention, which in turn causes individuals to focus on simple, present, and salient cues. As a result, alcohol has particularly strong effects in situations of "inhibition conflict," i.e. with two competing motivations, one of which is simple, present, or salient, while the other is complicated, in the future, or remote.[4] The behavioral-economics interpretation of this theory is that alcohol causes

by Laibson (1997). Given the similarity of predicted responses of consumption patterns to price changes by the two competing models, Gruber and Kőszegi (2001) were not able to reject Becker and Murphy's (1988) model in favor of their own. The ensuing literature produced suggestive but no conclusive evidence in the smoking domain (Gruber and Mullainathan (2005)). Two recent examples in the context of alcohol consumption found mixed results (Bernheim *et al.* (2012) and Hinnosaar (2012)). Finally, other theories predict demand for commitment as well, including cue-based theories, dual-self models, or temptation and self-control models as in Thaler and Shefrin (1981), Laibson (2001), Gul and Pesendorfer (2001), Bernheim and Rangel (2004), or Fudenberg and Levine (2006). For detailed overviews on the empirical and theoretical literature on commitment devices, see DellaVigna (2009) and Bryan *et al.* (2010).

[4]In a series of studies, Steele and several coauthors aimed to explain a range of social behaviors caused by alcohol, emphasizing the effects of alcohol on aggression and altruism (Steele and Southwick (1985), Steele *et al.* (1985)). These studies and subsequent work on alcohol myopia did *not* study savings decisions or intertemporal

present bias. The findings from my field experiment support this theory in the context of savings decisions. They demonstrate that alcohol-induced myopia can have economically meaningful consequences.

Moreover, this paper adds to the literature on poverty and self-control.[5] With the exception of Banerjee and Mullainathan (2010), this line of research has largely sought to explain choices between overall levels of current and future consumption, rather than to understand how and whether specific goods may cause time-inconsistent preferences. In contrast, this paper argues that focusing on specific temptation goods may not only be an effective way to help individuals overcome their self-control problems regarding the consumption of these goods, but, in the case of alcohol, may also reduce self-control problems in other domains.

This paper also contributes to the growing literature on saving decisions among the poor (Karlan *et al.* (2014b)). The availability and design of savings accounts have recently been found to be important determinants of savings behavior among the poor (Ashraf *et al.* (2006), Dupas and Robinson (2013a), Dupas and Robinson (2013b), Prina (2014), Schaner (2014), Kast *et al.* (2014), Brune *et al.* (2014), Karlan *et al.* (2014a)). Existing studies emphasize the importance of technologies for committing to savings. This paper argues that helping individuals to overcome underlying self-control problems regarding specific goods can be a substitute for commitment devices for overall consumption-saving decisions. More

choice (Giancola *et al.* (2010)). However, many cross-sectional studies, including the ones on alcohol, found a correlation between impulsive “delayed reward discounting" (DRD) and addictive behavior, without establishing existence or direction of causality (MacKillopp *et al.* (2011), Vuchinich and Simpson (1999)). Experimental lab studies consistently found that acute alcohol intoxication reduced inhibitory control in computer tasks (Perry and Carroll (2008)), but the two studies conducted so far did not find effects on impulsive DRD (Richards *et al.* (1999)). In fact, to their own surprise, Ortner *et al.* (2003) found that alcohol intoxication *reduced* impulsivity. My study differs from previous experimental studies in a number of ways. In particular, (i) the duration of the experiment was significantly longer (over three weeks vs. one day), (ii) sample characteristics were markably different (low-income workers vs. college students; higher levels of regular drinking), (iii) stakes were higher (relative to income), and (iv) the main outcome was the amount saved after three weeks (as opposed to impulsive DRD).

[5]This literature goes back to at least Fisher (1930). It was recently revived by several theoretical and empirical contributions. On the theory side, Banerjee and Mullainathan (2010) and Bernheim *et al.* (2014) investigated the possibility of a poverty trap due to the association between poverty and self-control. Recent research on the empirical side includes Mani *et al.* (2013) and Mullainathan and Shafir (2013). For an excellent review, see Haushofer and Fehr (2014).

generally, it argues that time preferences are endogenous, in line with Becker and Mulligan (1997), and, more recently in the context of saving among the poor, Carvalho *et al.* (2014).

The results from this paper have the potential to inform alcohol policy, a much-debated topic in developing countries. In India, states have chosen a wide range of policy options ranging from prohibition (Gujarat) to government provision (Tamil Nadu), and private provision (Delhi) of alcohol.[6] When making such choices, policymakers lack sufficient information on the causes and the impact of alcohol consumption, and the feasibility and effectiveness of policy options. This paper contributes to this knowledge by investigating the relationship between alcohol and self-control, a key aspect in the consideration of policy options such as "sin taxes" or prohibition.

Finally, this paper contributes to our understanding of the effectiveness of incentives to encourage health-related behavior. Financial incentives are among the most successful policies to reduce drug consumption in general (Anderson *et al.* (2009)), and alcohol consumption in particular (Wagenaar *et al.* (2009)).[7] Providing short-run financial or other incentives can have substantial short-term and long-term effects on a number of health-related behaviors (Petry *et al.* (2000), Prendergast *et al.* (2006), Volpp *et al.* (2008), Charness and Gneezy (2009), Higgins *et al.* (2012), Dupas (2014)). In contrast to existing studies, I do not find evidence of effects of short-run incentives on alcohol consumption beyond the incentivized period.

The remainder of this paper is organized as follows. Section 1.2 provides an overview of the study background, including alcohol consumption patterns in Chennai and in developing countries more generally. Section 1.3 describes the experimental design, characterizes the study sample, and discusses randomization checks. Section 1.4 then considers the effect of

[6]See Rahman (2003) for a review of alcohol policy in India. In a major policy shift, Kerala has recently opted to move from government provision of alcohol to prohibition within the next ten years.

[7]This is the case for both incentives in the form of increased prices or taxes, even for heavy drinkers (Chetty *et al.* (2009), Cook and Tauchen (1982)), and in the form of contingency management, i.e. the use of monetary or non-monetary incentives for changing health-related behavior modification, and behavior therapy, especially in the addiction field (Higgins and Petry (1999)). However, the vast majority of these studies were conducted in developed countries such that evidence from developing countries is limited.

increased sobriety on savings, and Section 1.5 investigates the interaction between sobriety and commitment savings. Section 1.6 considers the extent to which self-control problems contribute to the demand for alcohol. Section 1.7 concludes.

1.2 Alcohol in Chennai, India, and Developing Countries

There is scarce information regarding drinking patterns in developing countries, especially among the poor. In this section, I first describe alcohol consumption patterns among low-income individuals in Chennai, India. I then relate the observed patterns to existing data on alcohol consumption in India and in other developing countries.

1.2.1 Alcohol Consumption in Chennai

As a first step toward a systematic understanding of the prevalence of drinking among male manual laborers in developing countries, I conducted a short survey with 1,227 men from ten different low-income professions in Chennai.[8] Surveyors approached individuals from these groups during the day and asked whether they were willing to answer a short questionnaire about their alcohol consumption and take a breathalyzer test.[9] Figures 5.1 through 5.4 show summary statistics of drinking patterns for these professions, based on these surveys.

The overall prevalence of alcohol consumption among low-income men is high (Figure 5.1). 76.1 percent of individuals reported drinking alcohol on the previous day, ranging across professions from 37 percent (porters) to as high as 98 percent (sewage workers).[10] In

[8]The prevalence of alcohol consumption among women in Chennai and in India overall is substantially lower. It has been consistently estimated to be below five percent in India, with higher estimates for North-Eastern states and lower estimates for Tamil Nadu (where Chennai is located) and other South Indian states (Benegal (2005)). In the most recent National Family Health Survey (Round 3, 2005/6), the prevalence of reported female alcohol consumption was 2.2 percent (IIPS and International (2008)). It is highest in the lowest wealth (6.2 percent) and education (4.3 percent) quintiles.

[9]To ensure a high participation rate, individuals were given Rs. 20 ($0.33) for their participation in this short survey. As result, only five out of 1,232 individuals approached declined to participate.

[10]Porters are individuals who help carry luggage or other items at train stations. Sewage workers spend their days working, and sometimes swimming, in waist-deep human sewage. These individuals report drinking

addition, on days when individuals consume alcohol, they drink considerable quantities of alcohol (Figure 5.2). Conditional on drinking alcohol on the previous day, men of the different professions reported drinking average amounts ranging from 3.8 to 6.5 standard drinks on this day.[11] Since alcohol is an expensive good, the resulting income shares spent on alcohol are enormous (Figure 5.3). On average, individuals reported spending between 9.2 and 43.0 percent of their daily income of Rs. 300 ($5) to Rs. 500 ($8) on alcohol. These numbers are particularly remarkable because many low-income men in Chennai are the sole income earners of their families.[12] Finally, 25.2 percent of individuals were inebriated or drunk during these surveys, which all took place during the day (Figure 5.4).[13]

1.2.2 Alcohol Consumption in India and in Developing Countries

The substantial level of alcohol consumption among low-income groups in Chennai shown in Figures 5.1 through 5.4 raises the question of how these numbers compare to other estimates for Chennai, for India, and for developing countries overall. Limited data availability and data inconsistencies make answering this question difficult. In particular, data on breathalyzer scores are rare. However, there is reason to believe that the estimates for Chennai are not unusual compared to other parts of India or other developing countries.

The daily average quantity of alcohol consumed by male drinkers in India, about a quarter of the male population, is only slightly higher than the average of the physical

heavily before and during work to numb themselves, in particular to the smell.

[11]I follow the US definition of a standard drink as described in WHO (2001). According to this definition, a standard drink contains 14 grams of pure ethanol. A small bottle of beer (330 ml at 5% alcohol), a glass of wine (140 ml at 12% alcohol), or a shot of hard liquor (40 ml at 40% alcohol) each contain about one standard drink.

[12]The surveys reported here do not include questions about other family members and their incomes. However, female labor market participation is relatively low in Chennai. In my sample, less than a third of married men report that their wives earned income during the past month.

[13]Compared to other professions, the fraction of inebriated sewage workers is low given their reported expenditures and consumption. Anecdotally, this is explained by the fact that about a month before the surveys took place, one of the workers drowned in the sewage and his family was not given any severance payment because he was found to have been drunk at the time of the accident in an autopsy. After this incident, sewage workers stopped drinking at work, at least temporarily. Most individuals continued drinking alcohol regularly, but they did not drink during work hours.

quantities shown in Figure 5.2 (WHO (2014)). The average male Indian drinker consumes about five standard drinks per day, exceeding the estimates for German, American, and even Russian drinkers in the same WHO (2014) report.[14] In comparison, individuals who drank alcohol on the previous day in Chennai report on average drinking about 5.3 standard drinks per day. Looking beyond India, male drinkers in Uganda (56 percent of the male population) consume about 4 standard drinks per day. The prevalence of male alcohol consumption is somewhat lower in other Sub-Saharan countries, but the physical quantities consumed by drinkers are similarly high.[15] Alcohol consumption has also been steeply on the rise in China in recent years. According to the most recent WHO estimates, male Chinese drinkers (58.4 percent of the male population) consume 2.9 standard drinks per day.

There is also evidence that heavy alcohol consumption is more prevalent among the poor in developing countries. In India, both the prevalence of drinking and heavy alcohol consumption are more common among low-income and low-education individuals (Neufeld *et al.* (2005), Subramanian *et al.* (2005), IIPS and International (2007)).

Moreover, surveys among low-income groups show a commonly held belief that the positive correlation between excessive alcohol consumption and poverty reflects a causal relationship. For instance, in village surveys in Uganda, 56 percent of individuals believed that excessive alcohol consumption was a cause of poverty (USAID (2003)). Strikingly, this percentage was higher than the percentages of individuals that believed “lack of education and skills," “lack of access to financial assistance and credit,” or “idleness and laziness," caused poverty. At the same time, a quarter of individuals viewed excessive alcohol consumption as an outcome of poverty.

---

[14] Some assumptions in this calculation can be questioned. In particular, the WHO (2014) calculates the number of drinks per drinker and day by dividing an estimate of the overall quantity consumed by the estimated fraction of drinkers in the population. Hence, underestimating the prevalence of alcohol consumption among males in India could lead to overestimates of the number of standard drinks per drinker. However, even adjusting for the somewhat higher prevalence according to IIPS and International (2007), 31.9 percent rather than 24.8 percent in WHO (2014), yields just over four standard drinks per drinker and day. In addition, other studies find significantly lower prevalence of drinking in India (e.g. Subramanian *et al.* (2005)).

[15] For instance, an average drinker in Rwanda is estimated to consume 4.2 standard drinks per day. These numbers are similar for Burundi (4.1 standard drinks), Kenya (3.5 standard drinks), and Tanzania (3.4 standard drinks).

1.3 Experimental Design and Balance Checks

The first part of this section consists of a broad overview of the experimental design of my study. Next, I describe the recruitment and screening procedures and, hence, the selection mechanism of potential study participants into the study. I then provide detailed information about the timeline and the treatment conditions, followed by a description of the mechanism used to elicit willingness to pay for sobriety incentives and the outcomes of interest of the experiment. Finally, I discuss summary statistics for the study sample and balance checks.

1.3.1 Overview of Experimental Design

Between April and September 2014, I asked 229 cycle-rickshaw peddlers working in central Chennai to visit a nearby study office every day for three weeks each. During these daily visits, study participants completed a breathalyzer test and a short survey on labor supply, earnings, and expenditure patterns on the previous day, and alcohol consumption both on the previous day and on the same day before coming to the study office. To study the impact of increased sobriety due to financial incentives on savings behavior, all subjects were given the opportunity to save money at the study office. Additionally, participants were randomly assigned to varying conditions with the following considerations. First, to create exogenous variation in sobriety, a randomly selected subsample of study participants was offered financial incentives to visit the study office sober while the remaining individuals were paid for coming to the study office regardless of their alcohol consumption. Second, to examine the interaction between sobriety incentives and commitment savings, a cross-randomized subset of individuals was provided with a commitment savings account, i.e. their savings account did not allow them to withdraw their savings until the end of their participation in the study. Finally, to identify self-control problems regarding alcohol, a randomly selected subset of individuals was given the choice between incentives for sobriety and unconditional payments.

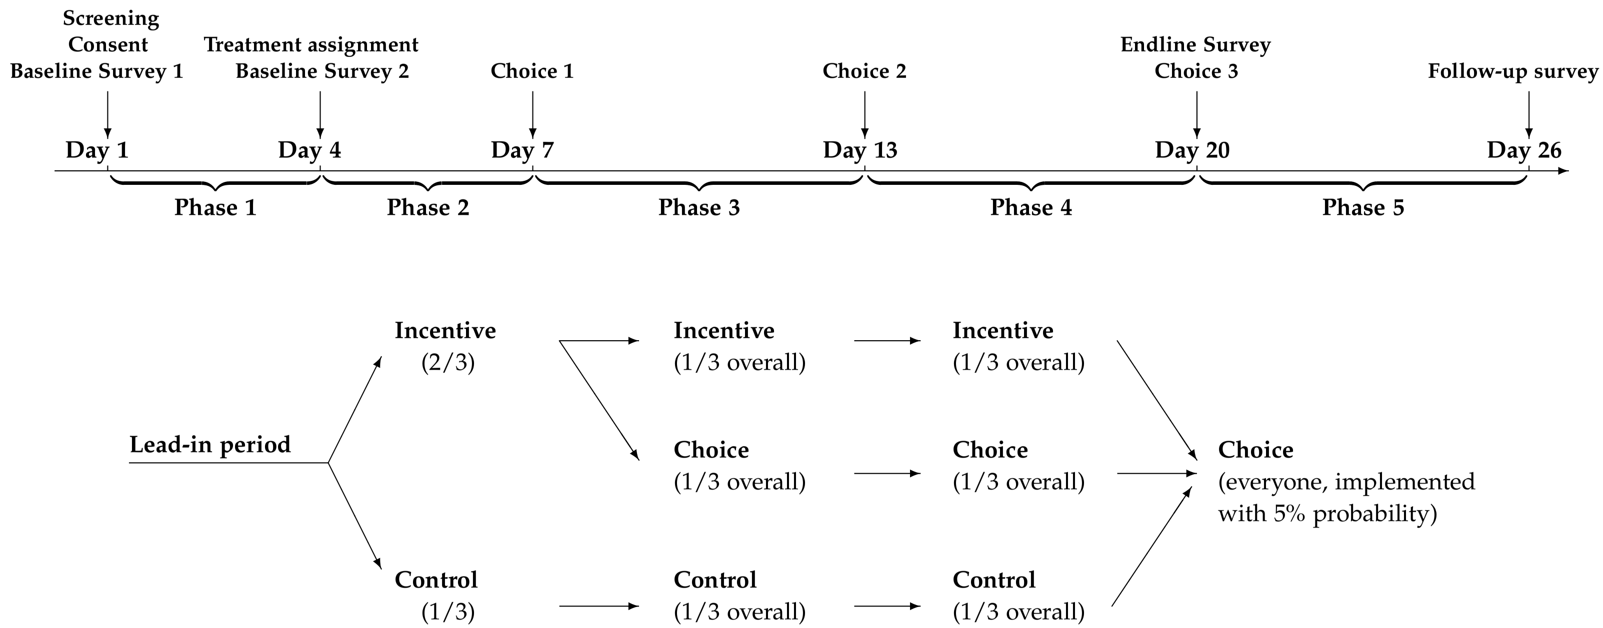

Notes: This figure gives an overview of the experimental design and the timeline of the study.

1. On day 1, individuals responded to a screening survey. Interested individuals then gave informed consent upon learning more about the study. Regardless of the consent decision regarding participation decision in the full study, all individuals were asked to complete a baseline survey, for which a separate consent was elicited.

2. On day 4, individuals who passed the lead-in period (Phase 1) completed a second baseline survey, and were then informed of their treatment status. On this day, individuals were fully informed about their payment structure and the decisions to be made over the course of the study.

3. The payments for the three treatment groups were as follows.

   (i) The Control Group was given the same unconditional payments as in Phase 1 (Rs. 90 regardless of breathalyzer score).

   (ii) Study payments for the Incentive Group depended on the breathalyzer score starting with day 5 of the study (Rs. 60 if BAC $>$ 0, Rs. 120 if BAC $=$ 0).

   (iii) After facing the same payment schedule as the Incentive Group in Phase 2, the Choice Group was asked to choose whether they wanted to continue receiving these incentives, or whether they preferred payments that did not depend on their breathalyzer scores. These choices were made on days 7 and 13, each for the subsequent week.

4. On day 20, all individuals were asked to participate in an endline survey. No incentives for sobriety were given on this day. *All* individuals were then given the same choices between conditional and unconditional payments as individuals in the Choice Group on days 7 and 13. To ensure incentive compatibility, these choices were then implemented for a small subset (5 percent) of study participants.

5. One week after their last day in the study, individuals were visited for a follow-up survey including a breathalyzer test.

Figure 1.1: *Experimental Design*

1.3.2 Recruitment and Screening

The study population consisted of male cycle-rickshaw peddlers aged 25 to 60 in Chennai, India.[16] Individuals enrolled in the study went through a three-stage recruitment and screening process. Due to capacity constraints, enrollment was conducted on a rolling basis such that there were typically between 30 and 60 participants enrolled in the study at any given point in time.

Field recruitment and screening. Field surveyors approached potential participants during work hours near the study office, and asked interested individuals to answer a few questions to determine their eligibility to participate in "a paid study in Chennai." Individuals were eligible to proceed to the next stage if they met the following screening criteria: (i) between 25 and 60 years old, inclusive, (ii) fluent in Tamil, the local language, (iii) worked at least five days per week on average as a rickshaw puller during the previous month, (iv) having lived in Chennai for at least six months, (v) without plans to leave Chennai during the ensuing six weeks, and (vi) reporting an average daily consumption of 0.7 to 2.0 "quarters" of hard liquor (equivalent to 3.0 to 8.7 standard drinks) per day.[17] If an individual satisfied all field screening criteria, he was invited to visit the study office to learn more about the study and to complete a more thorough screening survey to determine his eligibility.

Office screening. The primary goal of the more detailed office screening procedure was to reduce the risks associated with the study, in particular risks related to alcohol withdrawal symptoms. The criteria used in this procedure included screening for previous and current

[16] The study population included both passenger cycle-rickshaw peddlers as in Schofield (2014) and cargo cycle-rickshaw peddlers. Schofield (2014) exclusively enrolled passenger-rickshaw peddlers with a body-mass index (BMI) below 20. To avoid overlap between the two samples, my study only enrolled passenger cycle-rickshaw peddler with a BMI above 20. There was no BMI-related restriction for cargo cycle-rickshaw peddlers.

[17] "Quarters" refer to small bottles of 180 ml each. Nearly 100% of drinkers among cycle-rickshaw peddlers (and most other low-income populations in Chennai) consume exclusively hard liquor, specifically rum or brandy. The drinks individuals consume contain over 40 percent alcohol by volume (80 proof) and they maximize the quantity of alcohol per rupee. One quarter of hard liquor is equivalent to approximately 4.35 standard drinks.

medical conditions such as seizures, liver diseases, previous withdrawal experiences, and intake of several sedative medications and medications for diabetes and hypertension. This thorough medical screening procedure was strictly necessary since reducing one's alcohol consumption (particularly subsequent to extended periods of heavy drinking) can lead to serious withdrawal symptoms. If not adequately treated, individuals can develop delirium tremens, a severe and potentially even lethal medical condition (Wetterling *et al.* (1994), Schuckit *et al.* (1995)).

Lead-in period. Overall attrition and, in particular, differential attrition are first-order threats to the validity of any randomized-controlled trial. In my study, attrition was of particular concern since the study requested participants to visit the study office for three weeks every day with varying payment structures across treatment groups. In early-stage piloting, a non-negligible fraction of individuals visited the study office on the first day, which provided high renumeration to compensate for the time-consuming enrollment procedures, but then dropped out of the study relatively quickly. To avoid this outcome in the actual study, participants were required to attend on three consecutive study days (the "lead-in period") before being fully enrolled in the study and informed about their treatment status. Individuals were informed about this feature of the study during their first visit to the study office. They were allowed to repeat the lead-in period if they missed one or more of the three consecutive days. However, individuals were only allowed to repeat the lead-in period once.

Selection. At each stage, between 64 and 83 percent of individuals were able and willing to proceed to the subsequent stage (Table 1.1). Among individuals who were approached on the street to conduct the field screening survey, 64 percent were eligible and decided to visit the study office to complete the office screening survey. 21 percent were either not willing to participate in the survey when first approached (14 percent), or were not interested in learning more about the study after participating in the survey and being found to be eligible

(7 percent). The majority among the remaining individuals (12 percent) participated in the survey, but did not meet the drinking criteria outlined above, primarily because they were abstinent from alcohol or reported drinking less than 3 standard drinks per day on average (11 percent). During the next stage, the office screening survey, 83 percent of individuals were found eligible. The majority of the remaining, ineligible individuals (13 percent) were not able to participate due to medical reasons. Finally, 66 percent of individuals passed the lead-in period. Importantly, leaving the study at this stage does *not* appear to be related to alcohol consumption as measured by individuals' sobriety during their first visit to the study office.

1.3.3 Timeline and Treatment Groups

Figure 1.1 provides an overview of the study timeline, the different activities, and the treatment conditions. All participants completed five phases of the study as described in more detail below. During the first four phases, consisting of 20 study days in total, individuals were asked to visit the study office every day, excluding Sundays, at a time of their choosing between 6 pm and 10 pm. The office was located in the vicinity of their usual area of work to limit the time required for the visit. During Phase 1, the first four days of the study, all individuals were paid Rs. 90 ($1.50) for visiting the study office, regardless of their blood alcohol content (BAC). This period served to gather baseline data in the absence of incentives and to screen individuals for willingness to visit the study office regularly. On day 4, individuals were randomly allocated to one of the following three experimental conditions for the subsequent 15 days.

(I) **Control Group.** The Control Group was paid Rs. 90 ($1.50) per visit regardless of BAC on days 5 through 19. These participants simply continued with the payment schedule from Phase 1.

(II) **Incentive Group.** The Incentive Group was given incentives for sobriety on days 5 through 19. These payments consisted of Rs. 60 ($1) for visiting the study office, and

Table 1.1: *Eligibility Status at Different Recruitment Stages*

STAGE	**FRACTION**
(1) Field Screening Survey	
Eligible and willing to participate	**64%**
Not willing to conduct survey	14%
Drinks too little to be eligible	11%
Drinks too much to be eligible	1%
Ineligible for other reasons	3%
Eligible, but not interested	7%
(2) Office Screening Survey	
Eligible in Office Screening	**83%**
Ineligible for medical reasons	13%
Ineligible for other reasons	4%
(3) Lead-in Period	
Proceeded to enrollment	**66%**
Didn't proceed and BAC $= 0$ on day 1	19%
Didn't proceed and BAC > 0 on day 1	15%

Notes: This table gives an overview of the three-stage screening process of the study.

1. For each stage, it shows the fraction of individuals who were eligible and willing to proceed to the next stage of the study, the reasons for individuals not to proceed, and the relative frequencies of these reasons (each conditional on reaching the respective stage).
2. The tiers of the selection process are (1) the field screening survey (top panel), (2) the office screening survey (center panel), and (3) the lead-in period (bottom panel).

an additional Rs. 60 if the individual was sober as measured by a score of zero on the breathalyzer test. Hence, the payment was Rs. 60 if they arrived at the office with a positive BAC and Rs. 120 if they arrived sober. Given the reported daily labor income of about Rs. 300 ($5) in the sample, Rs. 60 ($1) was a relatively strong incentive for sobriety.

(III) **Choice Group.** To familiarize individuals with the incentives, the Choice Group was given the same incentives as the Incentive Group in Phase 2 (days 5 to 7). Then, right before the start of Phase 3 (day 7) and Phase 4 (day 13), they were asked to choose for the subsequent week (six study days) whether they preferred to continue receiving the same incentives, or to receive unconditional payments ranging from Rs. 90 ($1.50) to Rs. 150 ($2.50), as described below.

Eliciting willingness to pay for incentives. On days 7 and 13 of the study, surveyors elicited individuals' preferences in each of the three choices shown in Table 1. Each of these choices consisted of a tradeoff between two options. The first option, Option A, was the same for all choices. The payment structure in this option was the same as in the Incentive Group, i.e. a payment of Rs. 60 ($1) for arriving with a positive BAC, and Rs. 120 ($2) for arriving sober. In contrast, Option B varied across the three choices, with unconditional amounts of Rs. 90, Rs. 120, and Rs. 150. To gather as much information as possible while ensuring incentive compatibility, preferences for all three choices were elicited, before one of these choices was randomly selected to be implemented.[18] However, to maintain similar average study payments across treatment groups, Choice 1 was implemented in 90 percent of choice instances (independent over time) so that particularly high payments were only

[18] This is an application of the "random-lottery incentive system" (RLIS), in which a subject is asked to choose in several choice situations, one of which is randomly selected to be implemented once all choices are made. This method is extensively used in the experimental economics literature, for instance, recently by Augenblick *et al.* (2014) or Andreoni and Sprenger (2012). Holt (1986) put forward a theoretical criticism suggesting that subjects may not perceive every choice situation as isolated, but instead treat all choices as a grand meta-lottery. However, in subsequent experimental work, Starmer and Sugden (1991) and Hey and Lee (2005) did not find evidence in support of this concern. For a brief summary of the debate, see Wakker (2007).

actually paid out to a small number of individuals in the Choice Group.[19]

Table 1.2: *Choices between Incentives for Sobriety and Unconditional Payments*

	Option A		Option B
Choice	BAC $>$ 0	BAC $=$ 0	regardless of BAC
(1)	Rs. 60	Rs. 120	Rs. 90
(2)	Rs. 60	Rs. 120	Rs. 120
(3)	Rs. 60	Rs. 120	Rs. 150

I designed these choices with two main objectives in mind: first, to elicit demand for commitment to sobriety and, hence, potential self-control problems regarding alcohol consumption; second, to allow the Choice Group to be part of the evaluation of the impact of incentives for sobriety. In addition, given low literacy and numeracy levels in the study sample, the design seeks to minimize the complexity of decisions while achieving the other two objectives. In particular, Option A was the same across choices and individuals were given three days to familiarize themselves with these incentives during Phase 2. Accordingly, in all three choices, subjects knew Option A from previous office visits, and Option B was simply a fixed payment regardless of BAC as already experienced in Phase 1. To address potential concerns regarding anchoring effects, the order of choices was randomized. Half of participants made their choices in the order as outlined above, and the remaining individuals completed the choices in the opposite order.

[19]Before making their choices, study participants were told to take all choices seriously since each choice had a positive probability of being implemented. Individuals were *not* informed regarding the specific probabilities of implementing each of the choices. One potential concern regarding the procedure to elicit demand for commitment in this study is that subjects' choices may have been affected by the fact that none of the choices were implemented with certainty. Such effects would be a particular concern for this study if they increased the demand for commitment. However, the existing evidence suggests that introducing uncertainty into intertemporal choices *reduces* present bias (as measured by the immediacy effect) rather than increasing it (Keren and Roelofsma (1995); Weber and Chapman (2005)).

Demand for commitment. The choice of the conditional payment (Option A) in Choice 1 is *not* evidence of demand for commitment. An individual who did not prefer to change his drinking patterns may have chosen Option A if he expected to visit the study office sober at least 50 percent of the time and, therefore, to receive higher average study payments than from choosing Option B. In contrast, study payments for Option A were weakly dominated by the ones in Option B for Choice 2. Therefore, choosing Option A in Choice 2 is evidence of demand for commitment to increase sobriety, which reveals underlying self-control problems. Furthermore, study payments in Option A were *strictly* dominated by the ones in Option B for Choice 3. Choosing Option A in Choice 3 implied sacrificing Rs. 30 ($0.50) in study payments per day even during sober visits to the study office, a non-trivial amount given reported labor income of about Rs. 300 ($5) per day.

Endline. On day 20 of the study, all participants were asked to come to the study office once again for an endline visit at any time of the day of their choosing. No incentives for sobriety were provided on this day. During this visit, surveyors conducted the endline survey with individuals, and participants were we given the money they had saved. Moreover, *all* study participants were given the same set of three choices, described above. This allows me to understand whether exposure to incentives for sobriety affected subsequent demand for incentives. Again, preferences for all three choices were elicited, and then one of them was randomly selected to be implemented. However, the choices from day 20 were only implemented for a randomly selected five percent of individuals for budgetary and logistical reasons. These individuals were invited to visit the study office for six additional days. The endline visit was the last scheduled visit to the study office for the remaining study participants.

Follow-up visits. To measure the effects of the intervention beyond the incentivized period, surveyors attempted to visit each study participant about one week after their last scheduled office visit. This visit was announced during the informed consent procedures, and participants were reminded of this visit on day 20 of the study, but they were not informed regarding the exact day of this visit. During the follow-up visit, individuals were breathalyzed and surveyed once again on the main outcomes of interest. The compensation for this visit did *not* depend on the individuals' breathalyzer scores.

1.3.4 Lottery

In addition to the payments described above, study participants were given the opportunity to earn additional study payments in a lottery on days 10 through 18 of the study. The lottery was conducted as follows: If the participant arrived at the study office on a day on which he was assigned to play the lottery, he was given the opportunity to spin a 'wheel of fortune'. This gave him the chance to win a voucher for Rs. 30 or Rs. 60, at a probability of approximately 5 percent each. This voucher was valid only on the participant's subsequent study day, i.e. if the participant came back on the following study day and showed the voucher, he received the equivalent cash amount at the beginning of his visit. The lottery allows me 1) to estimate the impact of increased study payments on labor supply and earnings, 2) to estimate the impact of study payments on attendance and savings at the study office, and 3) to test whether sobriety incentives raised the marginal propensity to save.

1.3.5 Outcomes of Interest and Savings Treatments

The main outcomes of interest in this study are: (i) alcohol consumption and expenditures, (ii) savings behavior, and (iii) labor market participation and earnings. Each of these outcomes is described below.

Alcohol consumption data was collected daily during each study office visit by measuring individuals' blood alcohol content (BAC), and via self-reports regarding drinking times, quantities consumed and amounts spent on alcohol. BAC was measured via breathalyzer tests using devices with US Department of Transportation level of precision.[20] During each visit, after the breathalyzer test, individuals were asked about their alcohol consumption on the same day prior to visiting the study office, and about their overall alcohol consumption on the previous day. To cross-check self-reported drinking patterns, a randomly selected subset of subjects was visited unannounced between 7:30 pm and 10 pm for random breathalyzer tests.[21]

Saving. To study individuals' savings behavior, all individuals were given the opportunity to save money in an individual savings box at the study office. During each office visit, study participants could save up to Rs. 200, using either payments received from the study or money from other sources. Two features of the savings opportunity were cross-randomized to the sobriety incentive treatment groups.

(i) **Matching contribution rate.** Individuals were given a matching contribution ("savings bonus") as an incentive to save. During their endline visit, subjects were paid out their savings plus a matching contribution, randomized with equal probability to be either 10% or 20% of the amount saved. Hence, even in a setting with high daily interest rates, saving money at the study office was a high-return activity for many study participants.[22]

(ii) **Commitment savings.** Half of study participants were randomly selected to have

[20]As in Burghart *et al.* (2013), this study uses the breathalyzer model AlcoHawk PT500 (Q3 Innovations LLC). For more information on the measurement of BAC via breathalyzers, see O'Daire (2009).

[21]Ideally these tests would have been conducted at later times in the night to fully capture individuals' drinking patterns at night. However, staff constraints, safety considerations, and the intrusive nature of visiting individuals late at night at their homes made it infeasible to conduct these tests after 10 pm. The random breathalyzer tests were only conducted for the subset of individuals who consented to be visited unannounced. However, since the renumeration for these visits was deliberately chosen to be high (Rs. 100 for a successful visit regardless of the outcome of the breathalyzer test), the fraction of individuals that agreed to be randomly breathalyzed was nearly 100 percent.

[22]Individuals found the matching contribution easier to understand rather than a daily interest rate on savings during early-stage piloting work. The implied daily interest rate from saving an additional rupee increased for each participant over the course of his participation in the study. However, anecdotal evidence suggests that few individuals were aware of this feature.

their savings account include a commitment feature. Instead of being able to withdraw money during any of their daily visits between 6 pm and 10 pm, they were only allowed to withdraw money at the end of their participation in the study.[23] Notably, the savings option for the remaining individuals also entailed a weak commitment feature. While individuals could withdraw as much as they desired on any given office visit, they were only able to withdraw money in the evenings, i.e. between 6 pm and 10 pm.

The savings option served three purposes. First, it allows me to study the impact of increased sobriety on savings behavior and, more generally, the impact of alcohol on inter-temporal choices and investments in high return opportunities. Second, the cross-randomized commitment savings feature allows to consider the relationship between sobriety and self-control in savings decisions. Third, the savings feature was meant to help study participants avoid using the money received from the study to drink alcohol on the same evening or on subsequent days.

Labor market outcomes included reported earnings, labor supply, and productivity. These outcomes are measured by individuals' self-reports during the baseline survey, daily surveys, and the endline survey. Reported earnings are a combination of income from rickshaw work and other sources such as load work. Labor supply is a combination of the number of days worked per week and the number of hours worked per day. Finally, productivity is measured as income per hour worked.

1.3.6 Sample Characteristics and Randomization Checks

Appendix Tables 5.1 through 5.3 summarize study participants' key background characteristics, and demonstrate balance on these characteristics across treatment groups. Tables 5.1 and 5.2 give an overview of basic demographics, and work- and savings-related variables.

[23]For ethical reasons, all individual had the option to leave the study and withdraw all of their money at any day in the study.

As to be expected with a large number of comparisons, there are imbalances across treatment groups for some characteristics. However, overall only 5 out of 72 coefficients are statistically significantly different at the 10 percent level, and 3 coefficients are significantly different at the 5 percent level.[24] Most notably among these, individuals in the Control Group reported lower savings at baseline than in the Incentive and Choice Groups. Baseline savings are calculated as the sum of amounts saved in a number of different options including savings at home in cash or in gold or silver, with relatives and friends, with self-help groups, or with shopkeepers, as reported in the baseline survey. There is no statistically significant difference in the comparisons between the Incentive and Choice Group with the Control Group individually. However, the difference in reported baseline savings is statistically significant when comparing the Control Group to the Incentive and Choice Groups combined. As illustrated in the Appendix Figure 5.5, this difference is driven entirely by six individuals who reported very high savings, among them one individual in the Choice Group who reported in the baseline survey having Rs. 1 million in cash savings at his home.[25]

Differences in reported baseline savings are *not* driving the savings result shown below. First, there were only small and statistically insignificant differences in savings at the study office across treatment groups in the unincentivized Phase 1 (last row of Table 5.2). Second, controlling for Phase 1 savings and baseline survey variables, including total savings, does not substantially alter the regression results. If anything, the estimated effect of sobriety incentives on savings becomes larger. Third, there is no apparent relationship between reported savings in the baseline survey and savings at the study office. Among the six individuals with total savings above Rs. 200,000 in the baseline survey, four are in the Choice Group, and two are in the Incentive Group.[26] Only two of them, both in the Choice Group,

[24] Among the demographics in Table 5.1, the Control Group reports having lived for a few years longer in Chennai, and they are more likely to have electricity and a TV. In addition, they are somewhat less likely to own a rickshaw. In contrast, the overall fraction of individuals who reports 'lack of money' as a reason for not owning a rickshaw is balanced across treatment groups. Other reasons for not owning a rickshaw include not having a safe place to store it, or getting it provided by an employer.

[25] This amount was confirmed not only in the endline survey, but also during a subsequent follow-up visit.

[26] This outcome is more likely than it may seem. The probability of that none of the six high savers were

saved more than the average study participant in the course of the study.[27] However, their influence on the below results is negligible, in particular because these individuals already saved high amounts in the unincentivized Phase 1, and the below regressions control for savings in Phase. Hence, excluding these two individuals from the analysis does not change the conclusions of this paper.

Table 5.3 shows balance of alcohol consumption at baseline. Only one of the 36 comparisons shows a statistically significant difference at the 10 percent level. Compared to the Control Group, individuals in the Choice Group report somewhat lower alcohol expenditures per day.

allocated into the Control Group is $(2/3)^6 \approx 9\%$.

[27] Three of the remaining four individuals saved a total of Rs. 50 or less, and the fourth individual saved Rs. 500 in the course of the study, i.e. about the average amount in the Control Group.

1.4 Does Alcohol Affect Saving?

Time preferences are a fundamental aspect of decision-making and are critical for consumption-saving decisions. Savings can increase future consumption and serve as a buffer against adverse shocks, such as health emergencies. Accordingly, a growing body of recent research has focused on savings behavior among the poor and the impact of offering different savings accounts to low-income individuals in developing countries (Karlan *et al.* (2014b)). This literature largely focuses on the availability of different savings technologies and their potential impact on savings behavior Ashraf *et al.* (2006) and other outcomes such as investment in health (Dupas and Robinson (2013b)). There is less emphasis on determinants of savings behavior for given technologies and on heterogeneity in take-up or impact. In this section, I present evidence that alcohol distorts intertemporal choice by causing present bias, and hence self-control problems in savings decisions. I show that increasing sobriety can impact individuals' savings behavior beyond effects on income net of alcohol expenditures. I complement this evidence with Section 1.5, which shows that sobriety incentives lower the impact of a commitment savings feature on savings.

Figure 1.2 shows a strong correlation between daily amounts saved at the study office and blood alcohol content (BAC) measured during the same office visits, both across Control Group participants and within the same individuals over time. Individuals who, on average, exhibited higher sobriety also saved more. Moreover, individuals in the Control Group saved more during study office visits with lower levels of inebriation than the same individuals during high-inebriation visits. The remaining part of this section considers whether this correlation reflects a causal impact of alcohol consumption on individuals' savings behavior. Understanding the causal impact of alcohol on savings behavior requires exogenous variation in sobriety. Therefore, I first consider the impact of financial incentives on alcohol consumption. While the outcomes in this section are of interest in and of themselves, they can also be viewed as a first stage for the subsequent analysis of the impact of increased sobriety on savings decisions.

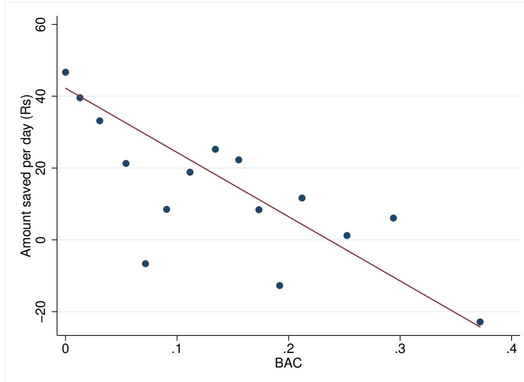

(a) *Daily amount saved and BAC (no individual FE)*

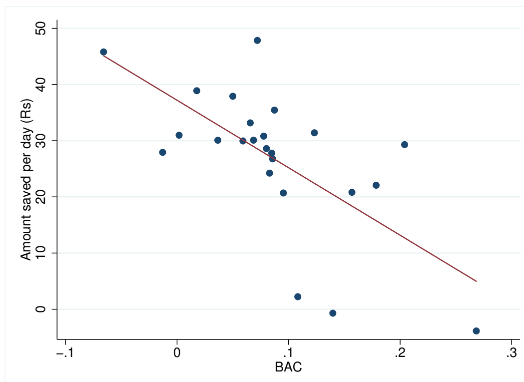

(b) *Daily amount saved and BAC (individual FE)*

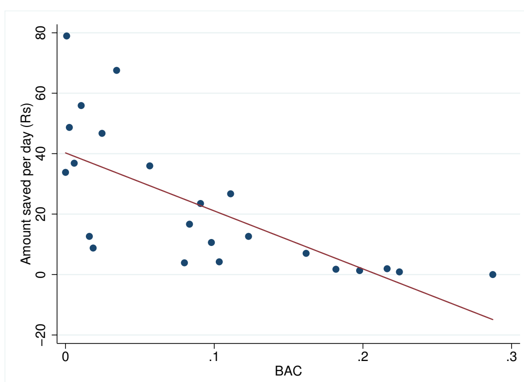

(c) *Mean amount saved and mean BAC*

Notes: This figure shows the correlation between breathalyzer scores during study office visits and amounts saved at the study during the same visits for individuals in the Control Group. The top panel depicts a binned scatter plot (including regression line) for all observations in the Control Group. The center panel shows the same graph, controlling for individual fixed effects. The bottom panel depicts the correlation across study participants by collapsing observations by individual.

Figure 1.2: *Cross-sectional Relationship between Daily Amounts Saved and BAC*

1.4.1 The Impact of Incentives on Alcohol Consumption (First Stage)

Financial incentives significantly reduced daytime drinking, but they had only a moderate effect on overall drinking. Table 1.3 give a summary of the results from this section. Since estimated treatment effects of the Incentive and Choice Conditions on alcohol consumption are remarkably similar, the table shows results from regressions that pool these two groups. Both sobriety incentive treatments lowered daytime drinking (left panel of Table 1.3), as measured by the fraction of individuals showing up sober, measured BAC, and the reported number of standard drinks before coming to the study office. The estimated treatment effects for all three measures correspond to a 33% change relative to the mean in the Control Group. However, this effect translates into only a moderate reduction of overall drinking (right panel of Table 1.3). Reductions in self-reported consumption and expenditures are relatively small (5.0 to 9.5 percent decrease), and, while larger in relative terms, the effect on reported abstinence is only moderate (2 percentage points) and not statistically significant.

The Impact of Sobriety Incentives on Daytime Drinking

The main outcome measure used to assess the impact of incentives on daytime drinking is the fraction of individuals who arrived sober at the study office among *all* participants who were enrolled (as opposed to only among individuals who visited the study office). That is, anyone who did not visit the study office on a particular day is counted as "not sober at the study office," along with individuals for whom a positive BAC was measured when they visited the office. Since attendance in the Incentive Group is lower than in the Control Group, this measure is preferable to other measures of sobriety as it less vulnerable to attrition concerns.

Financial incentives significantly increased sobriety during the day, as measured by the fraction of individuals who visited the study office *and* had a zero breathalyzer test result among all individuals in the respective treatment groups (upper panel of Figure 1.3). In the pre-incentive period, there are only small differences in sobriety across treatment groups. In each group, about half of the individuals visited the study office sober on days 1 through 4.

This fraction gradually decreased in the Control Group over the course of the study to about 35 percent by the end of the study.[28] In contrast, with the start of the incentivized period (day 5), sobriety in the Incentive and Choice Groups increased by about 15 percentage points. Sobriety at the study office declined as well in the course of the study, but individuals in these two groups remained about ten to fifteen percentage points more likely to visit the study office sober than the Control Group through the end of the study.

Table 1.3: *Summary of Estimated Effect of Incentives on Alcohol Consumption*

	Before/during visits			Overall drinking		
	Control	Change	%	Control	Change	%
Breathalyzer scores						
Fraction sober/abstinent	0.39	$+0.13^{***}$	+33.3	0.10	+0.02	+19.0
BAC (%)	0.09	-0.03^{***}	–33.3	–	–	–
Self reports						
# standard drinks	2.96	-0.98^{***}	–33.1	5.65	-0.28	–5.0
Expenditures (Rs/day)	–	–	–	91.2	-8.7^{*}	–9.5

Notes: This table gives an overview of the estimated treatment effects on sobriety before/during the study office visit (left panel) and overall alcohol consumption (right panel).

1. The table includes control means and estimated coefficients, both in absolute terms and as a share of the respective control mean.
2. The coefficients shown are from pooled estimates (i.e. pooling the Incentive and Choice Groups) from Table 1.4 (left panel) and Table 1.5 (right panel), including Phase 1 and baseline survey controls.
3. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

Remarkably, the two treatments had a nearly identical effect on the fraction of individuals who visited the study office sober. This is not a surprise in Phases 1 and 2 since the payment structure was the same in the Incentive and Choice Groups at the beginning of the study. However, overall sobriety levels in these two groups tracked each other even once

[28]The decline in sobriety in the Control Group over the course of the study is in part explained by lower overall attendance in all treatment groups. In addition, individuals may have felt more comfortable visiting the study office inebriated or drunk at later stages of the study.

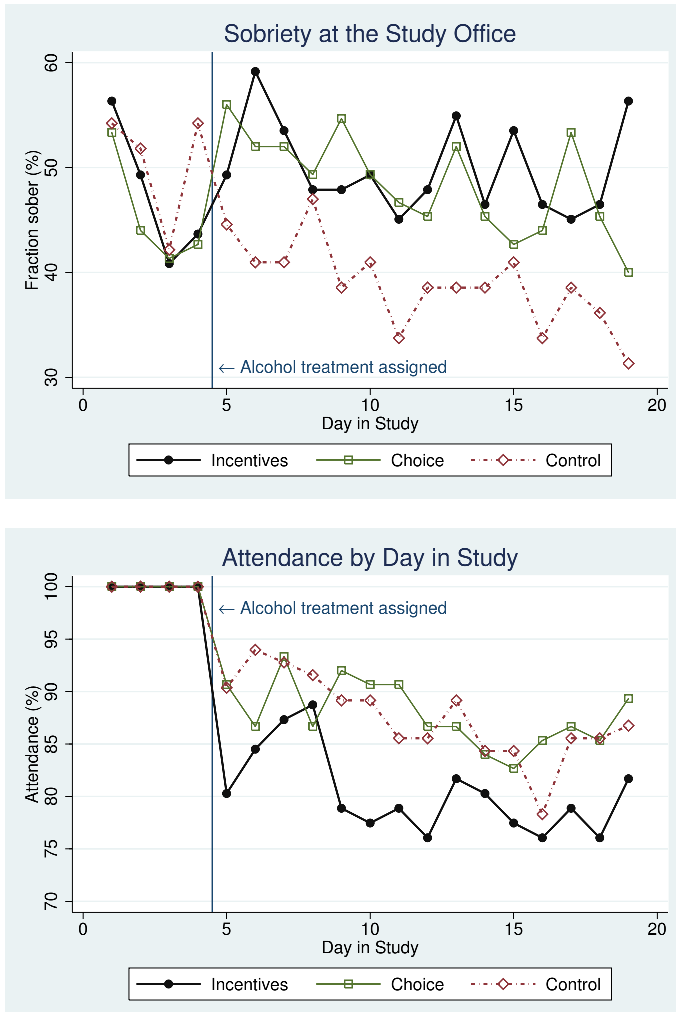

Notes: This figure shows sobriety and attendance over the course of the study for each of the three sobriety incentive treatment groups.

1. The upper panel of this figure shows the fraction of individuals who visited the study office sober. The indicator variable 'sober at the study office' takes on the value '1' for a study participant on any given day of the study if he (i) visited the study office on this day, *and* (ii) his breathalyzer test was (exactly) zero. The variable is, hence, '0' for individuals with a positive breathalyzer or those who did not visit the study office on this day.
2. The lower panel of the figure shows the fraction of individuals who visited the study office. Since only individuals who came to the study office on days 2 through 4 were fully enrolled in the study, by construction, attendance is 100 percent on days 1 through 4.

Figure 1.3: *Sobriety and Attendance by Alcohol Incentive Treatment Group*

individuals were given the choice of whether they wanted to continue receiving incentives at the beginning of Phase 3. The Incentive Group was only slightly more likely to visit the study office sober compared to the Choice Group in Phase 4. The similarity of drinking patterns in the Choice and Incentive Groups suggests sophistication regarding the effect of the incentives on individuals' sobriety. The subset of study participants who would have increased their sobriety during study office visits if they had been provided with incentives also chose to receive the incentives when given the choice.[29]

The corresponding regressions in Table 1.4 confirm the visual results. Individuals in the Incentive and Choice Group were approximately ten percentage points more likely to visit the study office sober, respectively (column 1). The estimates increase to 13 percentage points when regressions include baseline survey and Phase 1 control variables, in particular sobriety in Phase 1 (columns 2 to 4). This estimate corresponds to a 33 percent increase compared to the Control Group. Conditional on visiting the study office, individuals' measured BAC in the Incentive Group was four percentage points lower than in the Control Group (columns 5 through 7). The estimate is smaller for the Choice Group, which translates into a lower pooled estimate (column 8). Nonetheless, the three percentage-point decrease in BAC shown represents a 33 percent reduction compared to the Control Group. Moreover, both treatments reduced the reported number of drinks before visiting the study office by about one standard drink from a base of just under three standard drinks (columns 9 through 12). The point estimate for the pooled treatment effect, 0.98 standard drinks (column 12), corresponds to a reduction of 33 percent as well.

The Impact of Sobriety Incentives on Overall Drinking

The estimated treatment effect on overall alcohol consumption is substantially lower than the estimated effect on daytime drinking (Table 1.5). First, both treatments reduced reported

[29]This assumes that self-imposed and external incentives were equally effective, which may not have been the case. For instance, external incentives may have decreased intrinsic motivation to stay sober (Bénabou and Tirole (2003)).

overall alcohol consumption by about 0.3 standard drinks per day (columns 1 to 4), about a third of the effect on the reported number of drinks before coming to the study office described above. None of these estimates are statistically significant. Second, the reduction at the extensive margin of drinking was small at best (columns 5 to 8). The point estimate for the pooled treatment effect suggests a 2 percentage point increase in reported abstinence from drinking altogether (column 8), but none of the estimates are statistically significant either. Third, the treatment effect on reported overall alcohol expenditures is about Rs. 10 per day (columns 9 to 12), with a point estimate of Rs. 8.7 for the pooled treatment effect, statistically significant at the ten percent level. Taken together, these estimates provide evidence that subjects who responded to the incentives mostly shifted their alcohol consumption to later times of the day rather than reducing their overall consumption, or not drinking at all.

The Role of Differential Attendance

The estimated effect of incentives on sobriety was not caused by differences in attendance across treatment groups. Across all treatment groups and days of the study, attendance was high (lower panel of Figure 1.3).[30] However, compared to the Choice and Control Groups, individuals in the Incentive Group were 7 percentage points less likely to visit the study office post Phase 1. This attendance gap emerged with the start of sobriety incentives, and remained relatively constant thereafter. Anecdotal evidence suggests that this difference in attendance was caused by individuals in the Incentive Group who were not able or willing to remain sober until their study office visit on some days, and, hence, faced reduced incentives to visit the study office on these days. This explanation is consistent with the fact that there was no attendance gap between the Choice and Control Groups because individuals for whom sobriety incentives were not effective or preferable could select out of

[30] Attendance was 88.4 percent overall and 85.4 percent post treatment assignment. By construction, attendance in the lead-in period (Phase 1) was 100 percent.

them.[31]

On average, the Incentive Group was seven to eight percentage points less likely to visit the study office compared to the Control Group (column 1 of Table 1.6). Moreover, though not statistically significant, surprisingly, higher sobriety during the unincentivized Phase 1 *negatively* predicts subsequent attendance (column 2). This appears to be the case in the Incentive and Control Groups, but not in the Choice Group (column 3). Finally, on average, participants with higher savings in Phase 1 exhibited significantly higher subsequent attendance (column 4). However, there is no evidence that the two treatments caused high savers to visit the study office more frequently. If anything, the opposite was the case (column 5). This suggests that differential attendance of high savers does *not* explain the savings results shown below.

1.4.2 Did Increased Sobriety Change Savings Behavior?

Both sobriety incentive treatments increased savings at the study office (upper panel of Figure 1.4). Until day 4, when individuals learnt about their incentive treatment status, average amounts saved were nearly identical across treatment groups. After the start of the incentivized period, individuals in the Incentive and Choice Groups saved 46 percent and 65 percent more until the end of the study (Rs. 446 and Rs. 505 in the Incentive and Choice Groups, respectively, compared to Rs. 306 in the Control Group). The difference in savings across treatment groups did not emerge immediately after the beginning of the incentivized period, but accumulated mainly between days 8 and 15.

The corresponding regression results in Table 1.7 confirm the visual evidence. Individuals in both the Incentive and Choice Groups saved more at the study office, though only the coefficient for the Choice Group is statistically significant at the 10 percent level in the specification without controls (column 1). The pooled estimate shows a treatment effect of Rs. 12.45, corresponding to an increase of 61 percent compared to Control Group savings

[31] However, it remains unclear why there is an attendance gap for the Choice Group on days 5 through 7 of the study.

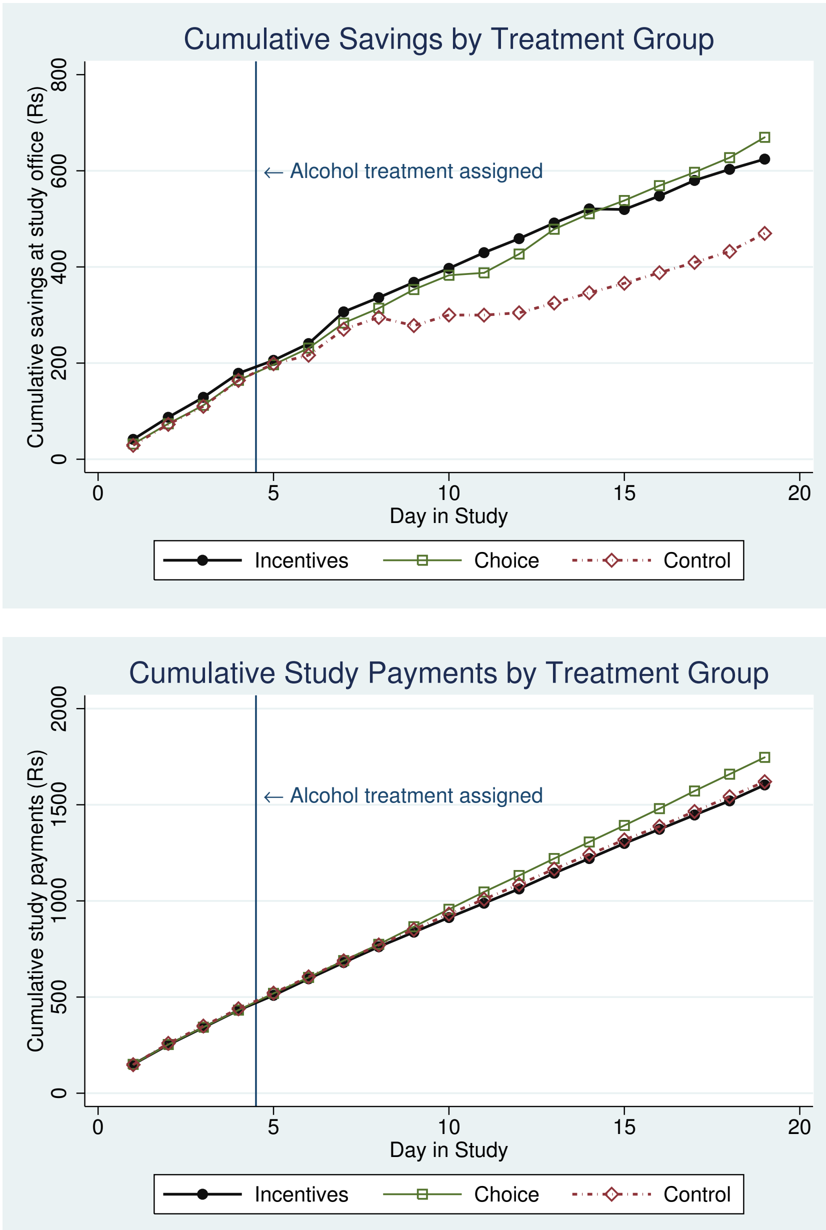

Notes: This figure depicts subjects' cumulative savings at the study office (upper panel) and cumulative study payments (lower panel) by alcohol incentive treatment group.

Figure 1.4: *Cumulative Savings by Day in Study*

of Rs. 20.42 (column 6). This estimate–as well as both individual estimates in column 1–is larger than the coefficients for both the high matching contribution and the commitment savings option. Incentives for sobriety had a larger effect than increasing the matching contribution on savings from 10 to 20 percent, or introducing a commitment feature on the savings option.[32] Importantly, these estimates are ITT estimates, i.e. they measure the impact of *offering* incentives for sobriety. While only effective for a relatively small fraction of individuals as shown above, sobriety incentives increased savings by 61% overall.[33]

[32]As discussed above, even individuals in the "no commitment savings" group were given a weak commitment feature since they were only able to withdraw money during their study visits between 6 pm and 10 pm. Hence, the estimate for "commitment savings" is likely an underestimate of the impact of commitment on savings.

[33]Since BAC levels differed across treatment groups conditional on visiting the study office with a positive blood alcohol content, using the difference in the fraction sober to calculate a ToT is not accurate.

Table 1.4: *The Effect of Incentives on Sobriety Before and During Study Office Visits*

VARIABLES	(1) Sober	(2) Sober	(3) Sober	(4) Sober	(5) BAC	(6) BAC	(7) BAC	(8) BAC	(9) # Drinks	(10) # Drinks	(11) # Drinks	(12) # Drinks
Incentives	0.11*	0.13***	0.13***		-0.04***	-0.04***	-0.04***		-1.09***	-1.22***	-1.14***	
	(0.058)	(0.047)	(0.044)		(0.013)	(0.010)	(0.010)		(0.372)	(0.279)	(0.262)	
Choice	0.10*	0.13***	0.13***		-0.01	-0.02*	-0.02*		-0.76**	-0.86***	-0.84***	
	(0.058)	(0.041)	(0.043)		(0.015)	(0.010)	(0.010)		(0.375)	(0.246)	(0.255)	
Pooled alcohol treatment				0.13***				-0.03***				-0.98***
				(0.038)				(0.009)				(0.221)
Observations	3,435	3,435	3,435	3,435	2,932	2,932	2,932	2,932	2,932	2,932	2,932	2,932
R-squared	0.010	0.248	0.294	0.294	0.019	0.299	0.355	0.352	0.022	0.280	0.306	0.305
Baseline survey controls	NO	NO	YES	YES	NO	NO	YES	YES	NO	NO	YES	YES
Phase 1 controls	NO	YES	YES	YES	NO	YES	YES	YES	NO	YES	YES	YES
Control group mean	0.389	0.389	0.389	0.389	0.0910	0.0910	0.0910	0.0910	2.957	2.957	2.957	2.957

Notes: This table considers the effect of the two sobriety incentives treatments on sobriety before and during study office visits.

1. All regressions use data from day 5 (the first day of sobriety incentives) through day 19 (the last day of sobriety incentives) of the study.
2. The outcome variable in columns 1 through 4, sobriety at the study office, is an indicator variable that is "1" for an individual on a given day if he visited the study office on this day *and* had a zero breathalyzer score on this day, and "0" otherwise. That is, individuals who did not visit the study office on any given day are included in these estimates as "not sober at the study office".
3. Columns 5 through 12 are conditional on visiting the study office. The outcome variable in columns 5 through 8 is individuals' measured blood alcohol content from a breathalyzer test. The outcome variable in columns 9 through 12 is the reported number of drinks *before* visiting the study office on any given day.
4. Standard errors are in parentheses, clustered by individual. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.
5. Phase 1 controls are the fraction of sober days, mean BAC during study office visits, the mean reported number of standard drinks consumed before coming to the study office and overall, and reported overall alcohol expenditures (all in Phase 1). Baseline survey control variables are all baseline survey variables shown in Tables A.1 through A.3.

Table 1.5: *The Effect of Incentives on Overall Alcohol Consumption*

VARIABLES	(1) # Drinks	(2) # Drinks	(3) # Drinks	(4) # Drinks	(5) No drink	(6) No drink	(7) No drink	(8) No drink	(9) Rs Exp	(10) Rs Exp	(11) Rs Exp	(12) Rs Exp
Incentives	-0.34	-0.20	-0.32		0.01	0.01	0.02		-10.27**	-8.12*	-8.01	
	(0.288)	(0.252)	(0.246)		(0.028)	(0.028)	(0.031)		(4.883)	(4.752)	(5.237)	
Choice	-0.35	-0.16	-0.25		0.02	0.01	0.02		-10.10**	-6.70	-9.31*	
	(0.344)	(0.261)	(0.269)		(0.029)	(0.028)	(0.030)		(4.986)	(4.274)	(4.747)	
Pooled alcohol treatment				-0.28				0.02				-8.71*
				(0.217)				(0.025)				(4.485)
Observations	2,932	2,932	2,932	2,932	2,932	2,932	2,932	2,932	2,932	2,932	2,932	2,932
R-squared	0.003	0.147	0.181	0.181	0.001	0.025	0.064	0.064	0.012	0.132	0.172	0.172
Baseline survey controls	NO	NO	YES	YES	NO	NO	YES	YES	NO	NO	YES	YES
Phase 1 controls	NO	YES	YES	YES	NO	YES	YES	YES	NO	YES	YES	YES
Control group mean	5.650	5.650	5.650	5.650	0.105	0.105	0.105	0.105	91.22	91.22	91.22	91.22

Notes: This table shows regressions of measures of *overall* alcohol consumption on indicator variables for the two sobriety incentive treatments.

1. All regressions use data from day 5 (the first day of sobriety incentives) through day 19 (the last day of sobriety incentives) of the study, conditional on visiting the study office.
2. The outcome variables are the reported overall number of standard drinks consumed per day (columns 1 through 4), abstinence from drinking altogether on a given day (columns 5 through 8), and reported alcohol expenditures (Rs. per day, columns 9 through 12).
3. Standard errors are in parentheses, clustered by individual. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.
4. Phase 1 controls are the fraction of sober days, mean BAC during study office visits, the mean reported number of standard drinks consumed before coming to the study office and overall, and reported overall alcohol expenditures (all in Phase 1). Baseline survey control variables are all baseline survey variables shown in Tables A.1 through A.3.

Table 1.6: *The Effect of Incentives on Attendance*

VARIABLES	(1) Present	(2) Present	(3) Present	(4) Present	(5) Present
Incentives	-0.07*	-0.07*	-0.08	-0.08*	-0.06
	(0.043)	(0.043)	(0.053)	(0.042)	(0.069)
Choice	0.00	0.00	-0.05	0.00	0.04
	(0.036)	(0.035)	(0.049)	(0.035)	(0.055)
Fraction of sober days in phase 1		-0.04	-0.08		
		(0.040)	(0.064)		
Incentives X Fraction sober in Phase 1			0.02		
			(0.105)		
Choice X Fraction sober in Phase 1			0.12		
			(0.084)		
Amount saved in Phase 1 (divided by 100)				0.02***	0.04***
				(0.009)	(0.012)
Incentives X Amount saved in Phase 1					-0.01
					(0.025)
Choice X Amount saved in Phase 1					-0.02
					(0.014)
Observations	3,435	3,435	3,435	3,435	3,435
R-squared	0.009	0.011	0.015	0.025	0.027
Baseline survey controls	NO	NO	NO	NO	NO
Phase 1 controls	NO	NO	NO	NO	NO
Control group mean	0.875	0.875	0.875	0.875	0.875

Notes: This table shows regressions of daily attendance at the study office on indicators for the two sobriety incentive treatments.

1. All regressions use data from day 5 (the first day of sobriety incentives) through day 19 (the last day of sobriety incentives) of the study.
2. The outcome variable is an indicator variable for whether an individual visited the study office on any given study day when he was supposed to.
3. Standard errors are in parentheses, clustered by individual. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

1.4.3 Robustness and Potential Confounds

Before examining the potential channels of the described effect of sobriety incentives on savings, this subsection investigates three potential confounds.

Pre-existing differences across treatment groups do not explain the observed differences in savings after day 4. The amounts saved by day 4 are nearly identical across treatment groups (upper panel of Figure 1.4). Moreover, controlling for baseline savings and baseline survey characteristics both decreases standard errors and increases point estimates (columns 2 of Table 1.7). The resulting point estimate for the pooled regression in column 4 is Rs. 13.44 and statistically significant at the 1 percent level (column 7 of Table 1.7).

Differential study payments across treatment groups could have been responsible for the increase in savings in the two treatment groups. Indeed, the Choice Group received slightly higher study payments (Rs. 7 per day) compared to the Control Group. However, the Incentive Group received in fact slightly lower study payments (lower panel of Figure 1.4), which implies that differences in average study payments cannot explain higher savings in both treatment groups. Consistent with this, controlling for study payments does not substantially alter the estimated treatment effects (columns 3 and 8 in Table 1.7). The estimate for the pooled treatment effect decreases slightly to Rs. 11.57 per day.

Differential attendance could have caused the increase in savings. However, as discussed in Section 1.4.1, while attendance was nearly identical in the Choice and Control Groups, it was in fact significantly *lower* in the Incentive Group (lower panel of Figure 1.3). In addition, if anything, the two treatments caused high savers to visit the study office *less* (column 5 of Table 1.6). Accordingly, restricting the sample to days when individuals showed up at the study office increases the estimated treatment effects (columns 4, 5, 9, and 10 of Table 1.7).

Table 1.7: *The Effect of Sobriety Incentives on Savings at the Study Office*

VARIABLES	(1) Rs/day	(2) Rs/day	(3) Rs/day	(4) Rs/day	(5) Rs/day	(6) Rs/day	(7) Rs/day	(8) Rs/day	(9) Rs/day	(10) Rs/day
Incentives	10.10	9.98	10.28*	14.81**	10.34					
	(7.555)	(6.455)	(6.194)	(7.031)	(6.700)					
Choice	14.71*	16.56***	12.77**	19.21***	13.07**					
	(7.772)	(5.679)	(5.382)	(6.288)	(6.208)					
Pooled alcohol treatment						12.45**	13.44***	11.57**	17.18***	11.77**
						(6.262)	(5.030)	(4.801)	(5.529)	(5.293)
High matching contribution	9.40	9.82**	11.41**	12.67**	11.77**	9.29	9.87**	11.45**	12.68**	11.77**
	(6.534)	(4.849)	(4.613)	(5.051)	(4.958)	(6.532)	(4.855)	(4.608)	(5.045)	(4.955)
Commitment savings	7.74	3.15	3.01	4.84	4.64	7.59	2.92	2.92	4.69	4.55
	(6.516)	(5.004)	(4.788)	(5.353)	(5.283)	(6.539)	(5.063)	(4.816)	(5.369)	(5.300)
Daily study payment (Rs)			0.34***		0.49***			0.34***		0.50***
			(0.050)		(0.125)			(0.050)		(0.123)
Observations	3,435	3,435	3,435	2,932	2,932	3,435	3,435	3,435	2,932	2,932
R-squared	0.007	0.114	0.129	0.123	0.131	0.006	0.113	0.129	0.123	0.131
Baseline survey controls	NO	YES	YES	YES	YES	NO	YES	YES	YES	YES
Phase 1 controls	NO	YES	YES	YES	YES	NO	YES	YES	YES	YES
Control mean	20.42	20.42	20.42	20.42	20.42	20.42	20.42	20.42	20.42	20.42

Notes: This table shows the impact of the two sobriety incentive treatments on participants' daily amount saved at the study office (Rs/day).

1. All regressions use data from day 5 (the first day of sobriety incentives) through day 19 (the last day of sobriety incentives) of the study. The outcome variable is the amount saved at the study office. If an individual did not visit the study office on any given day of the study, the amount saved is set to zero on this day. Similarly, the daily study payment is zero for those observations.
2. Regressions include the dummies "high matching contribution" for individuals who were offered a 20 percent matching contribution on their savings as opposed to 10 percent, and "commitment savings" for individuals who were not allowed to withdraw their saving until the last day of the study.
3. Columns (1) through (5) show regressions for the two sobriety incentive treatments separately. Columns (6) through (10) show pooled regressions for the Incentive and Choice Groups. Columns (1) and (6) are without controls, columns (2) and (7) include baseline survey and Phase 1 controls as in the previous tables. Columns (3) and (8) show the same regressions, but additionally control for study payments. The columns (4), (5), (9), and (10) show regressions conditional on attendance.
4. Standard errors are in parentheses, clustered by individual. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively. Phase 1 and baseline survey controls are the same as in the above tables.

1.4.4 The Effect of Changes in Income Net of Alcohol Expenditures

This paper argues increased sobriety caused changes in time preferences, which in turn increased savings. An alternative or complementary channel could be increased income net of alcohol expenditures, either due to reduced overall alcohol expenditures or increased earnings. This section considers the contribution of these channels to the increase in savings. I estimate this contribution to be about one half of the treatment effect on savings, and attribute the remaining share to a change in preferences.

Estimating the Marginal Propensity to Save

Assessing the contribution of increased resources requires knowledge of the marginal propensity to save out of additional resources, which the lottery allows me to estimate. Table 1.8 shows regressions of the daily amounts saved on a dummy for the pooled alcohol treatment as well as the amount won in the lottery on the previous day, and interactions of the treatment dummies with the lottery amount.[34] These regressions show a marginal propensity to save of 0.15 to 0.21 in the Control Group, and 0.36 to 0.37 in the pooled alcohol treatment groups. The below calculations use the marginal propensity to save from the Control Group in the preferred specification in column 4 of Table 1.8.

The estimates in Table 1.8 provide additional suggestive evidence that increasing sobriety affected time preferences. While the difference is not statistically significant, the estimated marginal propensity to save is higher (0.37, statistically significant at the 5 percent level) for the two groups that received sobriety incentives compared to the Control Group (0.21, not significant). Importantly, this difference is unlikely to be explained by the aforementioned confounds or increases in overall resources, since they are conditional on participating in the lottery.

[34]The regressions also control for whether the lottery was conducted on the previous day.

Table 1.8: *The Marginal Propensity to Save out of Lottery Earnings*

VARIABLES	(1) Rs saved	(2) Rs saved	(3) Rs saved	(4) Rs saved
Pooled alcohol treatment	12.32*	11.71*	15.03***	14.44***
	(6.256)	(6.110)	(5.174)	(5.202)
Amount won in lottery on previous study day	0.29*		0.29**	
	(0.166)		(0.143)	
Pooled alcohol treatment X Lottery amount		0.36*		0.36**
		(0.192)		(0.162)
Control Group X Lottery amount		0.15		0.16
		(0.295)		(0.261)
Observations	3,435	3,435	3,435	3,435
R-squared	0.008	0.008	0.117	0.118
Baseline survey controls	NO	NO	YES	YES
Phase 1 controls	NO	NO	YES	YES
Control mean	20.42	20.42	20.42	20.42

Notes: This table shows estimates of the impact of lottery winnings on the amounts saved at the study office.

1. All regressions use data from day 5 (the first day of sobriety incentives) through day 19 (the last day of sobriety incentives) of the study. The outcome variable is the amount saved at the study office. If an individual did not visit the study office on any given day of the study, the amount saved is set to zero on this day. Similarly, the daily study payment is zero for those observations.
2. The lottery was conducted on days 10 through 18 of the study. All regressions control for whether individuals participated in the lottery on any given day. Lottery winnings were Rs. 0 (no win), Rs. 30, or Rs. 60. If an individual won in the lottery, he was given a personalized voucher for the respective amount (Rs. 30 or Rs. 60) that was redeemable *only* by this individual *only* on the subsequent study day.
3. Standard errors are in parentheses, clustered by individual. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively. Phase 1 and baseline survey controls are the same as in the above tables.

The Effect of Reduced Alcohol Expenditures on Savings

Cycle-rickshaw peddlers spend a large fraction of their income on alcohol, on average, about Rs. 100 per day. Hence, even relatively small reductions in alcohol consumption can significantly increase the overall resources available. The above estimates find that the two treatments decreased alcohol expenditures by between Rs. 4.7 (using the implied expenditure reduction based on the reported physical quantities consumed) to Rs. 8.7 per day (using the estimate from reported expenditures). Combining these estimates with the estimated marginal propensity to save from available resources of 0.21 in the Control Group (column 4 of Table 1.8) implies that reduced alcohol expenditures account for Rs. 1.0 to Rs. 1.8 of the increase in savings.[35]

The Effect of Increased Earnings on Savings

Alcohol consumption may interfere with individuals' ability to earn income.[36] In addition to reduced alcohol expenditures, the treatments may have affected available resources via increased earnings. However, while positive, I estimate the effect of sobriety incentives on earnings to be relatively small and statistically insignificant, with a point estimate for the pooled treatment effect of Rs. 17.8 per day (columns 1 through 3 of Table 1.9.) Combined with the marginal propensity to save from above, this estimate implies that increased earnings account for Rs. 3.7 in increased savings. Similarly, the estimates on labor supply

[35] I use the estimated marginal propensity from the Control Group since the purpose of this exercise is to understand the effect of increased resources for *given* preferences, i.e. under the null hypothesis of unchanged preferences.

[36] Irving Fisher (1926) was among the first to investigate the relationship between alcohol and productivity. Based on small-sample experiments by Miles (1924) that showed negative effects of alcohol on typewriting efficiency, he argued that drinking alcohol slowed down the "human machine". He also argued that industrial efficiency was one of the main reasons behind the introduction of alcohol prohibition in the US. While many studies since Fisher (1926) have considered the relationship between alcohol consumption, income, and productivity (for an overview, see of the European Alcohol and Forum (2011)), there is a dearth of well-identified studies of the causal effect of alcohol on earnings and productivity, especially in developing countries. Cook and Moore (2000) summarized the literature as follows: "Modern scholars studying productivity effects have enjoyed larger sample sizes but unlike Fisher have utilized non-experimental data. The typical econometric study estimates the productivity effects of drinking, utilizing survey data in which respondents are asked about their drinking, work, income, and other items. The dependent variable is a measure of earnings or hours worked, while the key independent variable is a measure of the quantity or pattern of contemporaneous drinking, or alcohol-related psychiatric disorder (alcohol dependence or abuse)."

are relatively small and not statically significant (columns 4 through 9 of Table 1.9). In fact, the estimates of the treatment effect on labor supply at the extensive margin (i.e. whether an individual worked at all on any given day) is negative (columns 4 through 6). In contrast, the estimates on hours worked overall are positive in most specifications (columns 7 and 9).

Importantly, the estimates from this paper do *not* imply that alcohol does not have important effects on labor market outcomes for at least three reasons. First, the estimates in Table 1.9 are relatively imprecise. Since, while large in relative terms, the effect of incentives on daytime drinking is only moderate in absolute terms (13 percentage points), I cannot rule out large effects of daytime drinking on labor market behavior. Thus a more powerful intervention to reduce daytime drinking would have caused larger effects. Second, the impact of reduced drinking in the medium or long run might be much larger than the short-run effects considered in this paper. Third, the potentially negative impact of alcohol on productivity and labor supply via reduced physical or cognitive function may have been mitigated by analgesic effects of alcohol, which may not be the case in other settings.

1.4.5 Accounting for Mechanical Effects

Table 1.10 shows a decomposition of the effect of incentives on savings. This composition considers what share of the increase in savings is explained by mechanical effects, i.e. by individuals having increased resources for given preferences. The starting point in this decomposition is the estimate of Rs. 11.57 for the overall pooled treatment effect in column 8 of Table 1.7 (which controls for study payments). From this effect, I subtract the contribution of the two effects described above: (i) the contribution of reduced alcohol expenditures, and (ii) the contribution of increased earnings. This leaves an unexplained treatment effect of Rs. 6.00, i.e. about half of the overall treatment effect, and about 29% of control group savings. I attribute this share of the increase in savings to the effect of increased sobriety on time preferences. This argument is further supported by the next section, which shows evidence that sobriety incentives and commitment savings are substitutes.

Table 1.9: *The Effect of Sobriety Incentives on Labor Market Outcomes*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
VARIABLES	Rs earned	Rs earned	Rs earned	Worked	Worked	Worked	Hours	Hours	Hours
Incentives	20.03	17.46		-0.04	-0.04		0.23	0.27	
	(23.365)	(16.130)		(0.029)	(0.028)		(0.395)	(0.347)	
Choice	-2.83	17.67		-0.05	-0.02		-0.32	0.21	
	(25.363)	(19.991)		(0.032)	(0.030)		(0.401)	(0.330)	
Pooled alcohol treat			17.57			-0.03			0.24
			(15.552)			(0.025)			(0.293)
Observations	3,084	3,084	3,084	3,084	3,084	3,084	3,082	3,082	3,082
R-squared	0.002	0.315	0.315	0.004	0.070	0.069	0.003	0.163	0.163
Baseline survey controls	NO	YES	YES	NO	YES	YES	NO	YES	YES
Phase 1 controls	NO	YES	YES	NO	YES	YES	NO	YES	YES
Control group mean	287.4	287.4	287.4	0.894	0.894	0.894	6.829	6.829	6.829

Notes: This table shows the impact of the two sobriety incentive treatments on labor market outcomes.

1. All regressions use data from day 5 (the first day of sobriety incentives) through day 19 (the last day of sobriety incentives) of the study.
2. The outcome variables are (i) reported earnings (Rs. per day; columns 1 through 3) (ii) whether an individual worked on a particular day (columns 4 through 6), and (iii) the number of hours worked on this day (columns 7 through 9). If an individual did not work on any given day, this is counted as zero hours worked.
3. The data used in the regressions is from retrospective surveys on the consecutive study days, during which individuals are asked about earnings and hours worked on the previous day. In addition, if individuals missed a day or two (and on Mondays), they were asked about the same outcomes two or three days ago, respectively.
4. Standard errors are in parentheses, clustered by individual. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively. Phase 1 and baseline survey controls are the same as in the above tables.

Table 1.10: *Decomposing the Impact of Incentives on Savings*

Estimated overall treatment effect	**Rs. 11.57**
Resource effect 1: reduced expenditures	Rs. 1.83
Resource effect 2: increased earning	Rs. 3.74
Remaining treatment effect	**Rs. 6.00**

1.4.6 Household Resources and Complementary Consumption

This subsection addresses two additional concerns regarding the above findings. First, the increase in savings at the study office due to increased sobriety may have come at the cost of reduced household resources. Second, reduced alcohol consumption during the day or overall may have lowered complementary consumption such as smoking.

Household Resources

The increase in savings due to the incentives treatments does not appear to have crowded out money spent on family resources (Table 5.4). While not statistically significant, I find that sobriety incentives *increased* money given to wives by about Rs. 17.4 (columns 1 through 3). In contrast, resources spent on other family expenses decreased by about Rs. 8.9 (columns 4 through 6) such that reported resources spent on family expenses overall increased by about Rs. 8.6 (columns 7 through 9).

Food Expenditures and Complementary Consumption

I find no evidence of the treatment affecting expenditures on other goods (Table 5.5). Expenses on food outside of the household increased slightly by about Rs. 4 (columns 1 through 3), and reported expenditures on coffee and tea remained constant (columns 4 through 6; these may be underreported altogether). Of particular interest are expenses on tobacco products as they are often thought of as complements to alcohol (Room (2004)). However, there is no evidence of such effects (columns 7 through 9). This is not particular surprising in the light of the facts that reported expenditures on tobacco and paan[37] products are low to start with, and the incentives reduced overall alcohol expenditures only moderately, hence limiting the scope of effects through complementarities in consumption.

1.5 Are Sobriety and Commitment Savings Substitutes?

The structure of the experiment allows for an additional test of the hypothesis that increasing sobriety lowers self-control problems. The intuition for this test is straightforward. If self-control problems prevent individuals from saving as much as they would like to, and if commitment savings products help sophisticated individuals overcome these problems, then commitment savings should have a larger effect for individuals with more severe self-control problems. Hence, if alcohol reduces self-control, then increasing sobriety should lower the effect of commitment savings. However, this intuition overlooks an additional, opposing effect. While commitment savings products may help individuals overcome self-control problems in future savings decisions by preventing them from withdrawing their savings prematurely, the immediate decision to save always requires incurring instantaneous costs. A sophisticated individual with severe self-control problems may not save (much) even if a commitment savings product is offered, simply because he does not put much weight on future consumption. In the extreme case, for β close to zero, the individual will not save regardless of the availability of a commitment option.

This section shows a simple model that formalizes this intuition. I then consider a specific case (isoelastic utility) to demonstrate two features of this model. First, the impact of commitment savings is an inverse-U shaped function in present bias for sophisticated individuals. The impact of commitment savings devices on savings is lowest for individuals without present bias ($\beta \approx 1$) and for the most present-biased individuals ($\beta \approx 0$). At

[37]Paan is a mixture of ingredients including betel leaf, areca nut, and often tobacco. Chewing paan is popular in many parts of India.

least in theory, for individuals with the greatest need to overcome self-control problems, commitment savings devices in the form in which they are often offered may only be moderately helpful (if at all).[38] Second, for the empirically relevant parameter range of $\beta > 0.5$, an increase in β lowers the impact of commitment savings on savings. Accordingly, a decrease in the impact of commitment savings due to increased sobriety, as demonstrated in Section 1.5.2, can be viewed as evidence for increased self-control due to increased sobriety.

1.5.1 A Simple Model

Consider a simple consumption-saving problem. A consumer lives for three periods. In Period 1 he receives an endowment Y_1. There are no other income sources in Periods 2 and 3, but the consumer is paid a matching contribution of M times the amount saved by the start of Period 3. In Periods $t = 1, 2$, he has to decide how to allocate his available resources into instantaneous consumption c_t or savings. The instantaneous utility function $u(c_t)$ is increasing and concave: $u'(\cdot) > 0$ and $u''(\cdot) < 0$. The consumer has β-δ time preferences as in Laibson (1997), with $\delta = 1$ for simplicity and $\beta \in (0, 1]$. The individual is sophisticated in the O'Donoghue and Rabin (1999) sense. He understands the extent of future self-control problems, i.e. he knows his future β. There is no uncertainty. In Period 1, he maximizes $U_1(c_1, c_2, c_3) \equiv u(c_1) + \beta[u(c_2) + u(c_3)]$ and in Period 2 he maximizes $U_2(c_2, c_3) \equiv u(c_2) + \beta u(c_3)$.

No commitment savings. Consider first a situation without commitment savings. We solve the problem recursively. In Period 3, the individual will consume the entire amount saved plus the matching contribution: $c_3 = (Y_1 - c_1 - c_2)(1 + M)$. In Period 2, the individual

[38]Note that interventions designed along the lines of the Save More Tomorrow program (Thaler and Benartzi (2004)) overcome this problem, since it allows individuals to commit to saving more without reducing today's consumption.

takes c_1 as given and maximizes

$$\max_{c_2} \; u(c_2) + \beta u((Y_1 - c_1 - c_2)(1+M)) \tag{1.1}$$

The associated FOC is $u'(c_2) = \beta(1+M)u'((Y_1 - c_1 - c_2)(1+M))$. This choice is anticipated in Period 1 such that the individual chooses c_1 to solve the following problem:

$$\max_{c_1} \; u(c_1) + \beta[u(c_2) + u(c_3)] \tag{1.2}$$

$$\text{s.t.} \quad c_3 = (Y_1 - c_1 - c_2)(1+M) \tag{1.3}$$

$$u'(c_2) = \beta(1+M)u'(c_3) \tag{1.4}$$

$$c_1, c_2, c_3 \geq 0 \tag{1.5}$$

Defining $Y_2 \equiv Y_1 - c_1$, the solution is described by the following three equations.

$$u'(c_1) = \beta \left[u'(c_2)\frac{dc_2}{dY_2} + u'(c_3)\frac{dc_3}{dY_2} \right] \tag{1.6}$$

$$u'(c_2) = \beta(1+M)u'(c_3) \tag{1.7}$$

$$c_3 = (Y_2 - c_2)(1+M) \tag{1.8}$$

Combining these equations yields a version of the familiar modified Euler equation (Harris and Laibson (2001)):[39]

$$u'(c_1) = \left[\beta\frac{dc_2}{dY_2} + \left(1 - \frac{dc_2}{dY_2}\right) \right] u'(c_2) \tag{1.9}$$

Commitment savings. Consider now the situation in which a commitment savings account is available. That is, any money that is saved in Period 1 cannot be withdrawn until Period 3. Period 1 self would like to set $u'(c_2) = (1+M)u'(c_3)$. However, in the absence of commitment savings, Period 2 self deviates from this, i.e. chooses c_2 such that $u'(c_2) = \beta(1+M)u'(c_3)$ and, hence, consumes more than the Period 1 self would like him to. This creates a demand for commitment for Period 1 self. Since the Period 1 self is always

[39] In contrast to Harris and Laibson (2001), there is no interest rate in this equation since M is a matching contribution rather than an interest rate.

(weakly) more patient than the Period 2 self, this implies that the solution to this problem is simply the case in which the Period 1 self determines consumption in all three periods. The individual will consume c_1 and deposit c_3 into the commitment savings account such that $u'(c_1) = \beta u'(c_2) = \beta(1+M)u'(c_3)$, subject to the above budget constraint. Hence, the solution is described by the following equations:

$$u'(c_1) = \beta u'(c_2) \tag{1.10}$$

$$u'(c_2) = (1+M)u'(c_3) \tag{1.11}$$

$$c_3 = (Y_2 - c_2)(1+M) \tag{1.12}$$

Comparing the two above solutions clarifies the relationship between present bias and commitment savings. Introducing a commitment savings option increases savings iff $0 < \beta < 1$, since the commitment savings device makes both the Period 1 and 2 selves consume a smaller share of their available resources Y_1 and Y_2, respectively. If $\beta = 1$, commitment savings has no effect as there is no discrepancy between the Period 1 and Period 2 preferences. At the other extreme, if $\beta \to 0$, there are no savings even if commitment is available such that there is no impact of the commitment device on savings choices either.[40] Taken together, this implies that the impact of commitment savings is non-monotonic in present bias.

For $\beta \in (0,1)$, changing β has two opposing effects on the impact of commitment on savings. The first effect is that, in the absence of commitment, the Period 2 self will deviate more from the allocation that maximizes Period 1 self's utility (by increasing c_2 relative to c_3). This not only reduces Period 2 self's savings for given resources, but it also reduces Period 1 self's saving as he anticipates this effect. In contrast, in the presence of the commitment device, the Period 1 self can prevent this from happening by saving the desired amount using the commitment device. Hence, the impact of the commitment device on savings is larger for increased present bias due to this effect. However, there is a second, opposing effect. Since Period 1 self's β also decreases, the desire to allocate resources to Periods 2 and

[40]Subsistence levels in consumption could change this in the absence of income sources in Periods 2 and 3.

3 falls even if a commitment savings option is available. This lowers the impact of offering the commitment savings option. In the extreme case for $\beta \to 0$, there is no effect.

Solving for the isoelastic case. Consider the case of the commonly used isoelastic utility function.

$$u(c_t) = \begin{cases} \frac{c_t^{1-\gamma}}{1-\gamma} & \text{if } \gamma \neq 1, \\ \log(c_t) & \text{if } \gamma = 1. \end{cases} \tag{1.13}$$

The impact of commitment savings on savings is given by the difference in consumption levels in period 3 with and without commitment (see Appendix Section 5.2.1 for details).

$$\Delta \equiv c_3^{\mathrm{C}} - c_3^{\mathrm{NC}} = \frac{Y(1+M)}{1+\theta+\theta\left[\frac{1+\beta\theta}{1+\theta}\right]^{\frac{-1}{\gamma}}} - \frac{Y(1+M)}{1+\theta+(1+M)^{1-\frac{1}{\gamma}}}. \tag{1.14}$$

Figure 1.5 depicts Δ as a function of β for different values of γ. For the empirically relevant ranges of $\beta \in [0.5, 1]$ and $\gamma > 0.5$, a decrease in present bias, i.e. an increase in β, lowers the impact of commitment savings devices on savings.[41] This implies that an increase in sobriety (which lowers the use of commitment savings in my experiment) is effectively equivalent to an increase in β.

1.5.2 Empirical Evidence

In my study, increasing sobriety and commitment savings are substitutes in terms of their impact on savings. Figure 1.6 shows cumulative savings by the (pooled) sobriety treatment and the cross-randomized savings conditions.[42] In the upper panel of the figure, individuals are divided into four groups according to whether they were offered sobriety incentives—

[41] See, for instance, Frederick *et al.* (2002) for a review of estimates of present bias, and Chetty (2006) for estimates of γ.

[42] The two sobriety treatments are pooled solely for expositional purposes. The equivalent graphs without pooling the sobriety treatment groups show only very minor differences in savings behavior between the Incentive and Choice Groups (Figure 5.6).

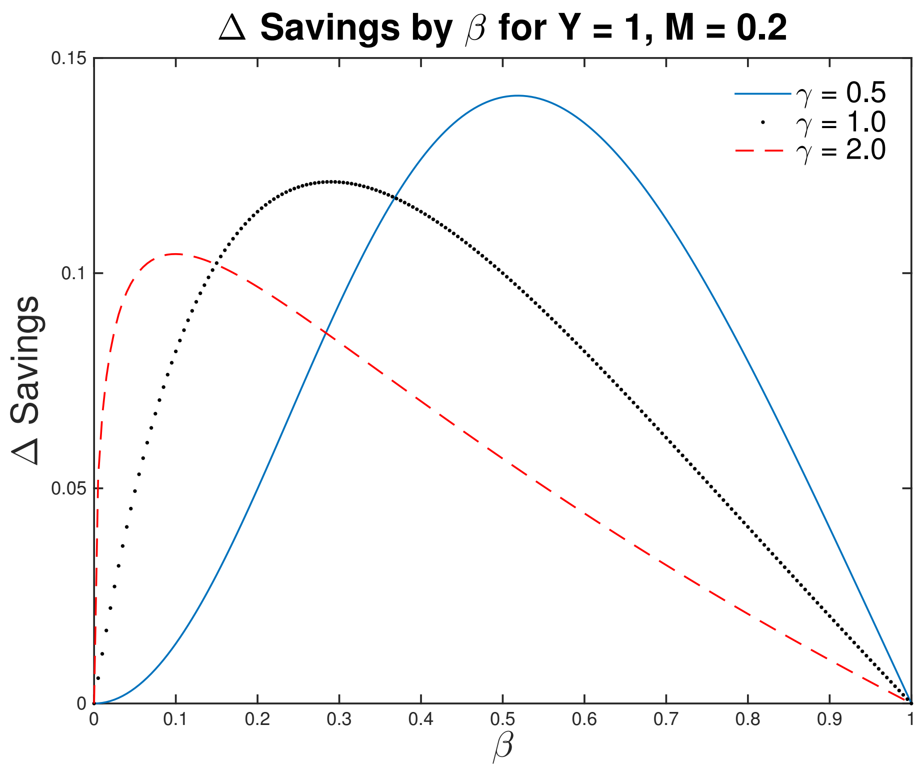

Notes: This figure shows the relationship between present bias and the effect of commitment savings in the model described in Sections 1.5.1 and 5.2.1.

1. The figure shows the present bias (as measured by $\beta \in [0, 1]$) on the horizontal axis and the increase in savings due to offering a commitment savings option on the vertical axis for the isoelastic utility case.
2. This increase in savings is given by the difference in consumption in period 3 between the two cases described in my model, i.e. $\Delta = c_3^{C} - c_3^{NC}$ as shown in equation (1.14).
3. The figure depicts the relationship between Δ and β for $\gamma = 0.5$ (the solid line), $\gamma = 1$ (the dotted line), and $\gamma = 2$ (dashed line).
4. In the specific figure shown here, $Y = 1$ and $M = 0.2$. The relationship is very similar, if not identical, for different parameter values. An explicit solution for Δ in the log case ($\gamma = 1$) is given in the Supplementary Appendix below.

Figure 1.5: *Effect of Commitment Savings as Function of β*

pooling the Incentive and Choice Groups—and whether their savings option included the cross-randomized Commitment Savings feature.[43] Cumulative savings for the four groups are nearly identical through the pre-incentive period until day 4, and throughout the study, three of the four lines in the graph remain nearly indistinguishable. However, the group that received neither commitment savings nor the alcohol treatment (as represented by the green line with solid circles) saved much less than each of the remaining groups subsequently. While both incentives for sobriety and the commitment savings option have a large impact on savings, being assigned to both does not further increase savings.

These differences across treatment groups are due to differences in both deposits and withdrawals (Figure 1.7). Compared to the group without either incentives for sobriety or commitment savings, sobriety incentives and commitment savings each on their own increased deposits (upper panel), and reduced withdrawals (lower panel). The magnitudes of these effects vary slightly. The effect of sobriety incentives on deposits is somewhat larger than the effect of commitment savings, but this difference is offset by an equivalent difference in withdrawals resulting in nearly identical overall savings.

These results suggest that increasing sobriety reduced self-control problems. An alternative interpretation could be that alcohol is a key temptation good for this population such that reducing alcohol consumption mitigates the need for commitment savings. However, given that the intervention only moderately reduced overall alcohol consumption and expenditures, this channel is unlikely.

A second competing explanation could be that there was an upper bound of how much individuals were able to or wanted to save. However, average daily savings are well below the savings limit of Rs. 200 per day. Moreover, in the course of the study, all individuals received relatively large study payments in addition to their earnings outside of the study, which appear to have been largely unaffected by the study. This suggests that the majority of individuals would have been able to increase their savings if they had preferred to

[43]For instance, the blue line with squares shows cumulative savings for individuals who were not offered incentives for sobriety, but who were given the commitment savings options.

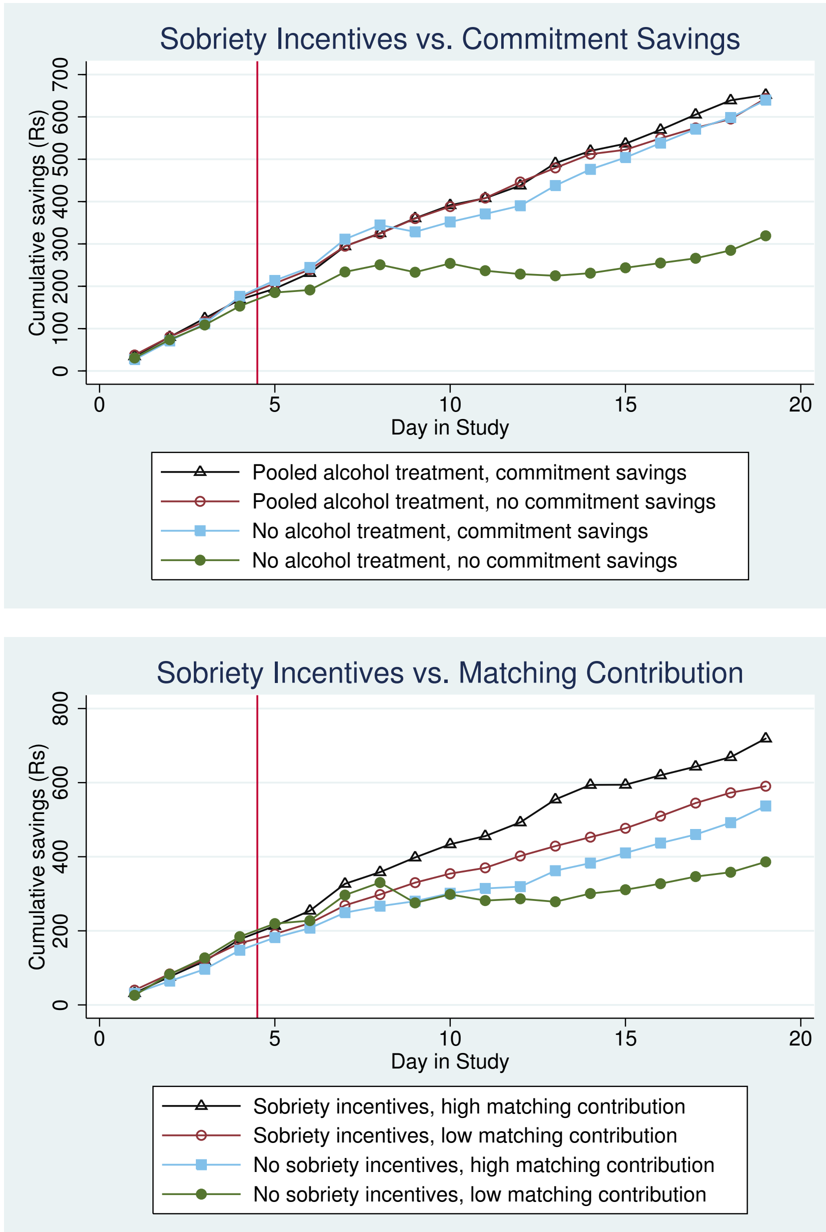

Notes: This figure shows the interaction between the cross-randomized sobriety incentives and savings treatments. The upper panel shows cumulative savings for four different groups: individuals who were offered

(i) neither sobriety incentives nor commitment savings (green line with solid circles),

(ii) no sobriety incentives, but commitment savings (blue line with squares),

(iii) sobriety incentives, but not commitment savings (red line with hollow circles), and

(iv) both sobriety incentives and commitment savings (black line with triangles).

The lower panel of the figure shows the equivalent graph for the interaction between receiving sobriety incentives and a matching contribution (20 percent instead of 10 percent on the amount saved by day 20).

Figure 1.6: *Interaction between Sobriety Incentives and Savings Treatments*

do so. Consistent with this, increasing the matching contribution rate did *not* serve as a complement to increased sobriety, i.e. the effects of incentives for sobriety and a high matching contribution appear to have been additive (lower panel of Figure 1.6).

Table 1.11: *Interaction between Sobriety Incentives and Savings Treatments*

VARIABLES	(1) Rs/day	(2) Rs/day	(3) Rs/day	(4) Rs/day
Either Incentives or Commitment Savings	19.77**	15.48*		
	(9.037)	(8.679)		
Sobriety Incentives only	0.49	0.06		
	(9.745)	(9.048)		
Both Incentives and Commitment Savings	1.43	2.36		
	(9.562)	(9.997)		
Either Incentives or High Matching Contribution			12.43	12.23
			(8.841)	(9.489)
Sobriety Incentives only			2.42	0.15
			(8.957)	(9.851)
Both Incentives and High Matching Contribution			10.16	8.30
			(9.468)	(9.731)
Observations	3,435	3,435	3,435	3,435
R-squared	0.006	0.037	0.005	0.037
Baseline survey controls	NO	YES	NO	YES
Phase 1 controls	NO	NO	NO	NO
Control mean	20.42	20.42	20.42	20.42

Notes: This table shows estimates of the impact of lottery winnings on the amounts saved at the study office.

1. All regressions use data from day 5 (the first day of sobriety incentives) through day 19 (the last day of sobriety incentives) of the study. The outcome variable is the amount saved at the study office. If an individual did not visit the study office on any given day of the study, the amount saved is set to zero on this day. Similarly, the daily study payment is zero for those observations.
2. Columns (1) and (2) show the relationship between the effects of offering sobriety incentives and commitment savings. Columns (3) and (4) show the relationship between the effects of offering sobriety incentives and a high matching contribution.
3. Standard errors are in parentheses, clustered by individual. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively. Baseline survey controls are the same as in the above tables.

1.6 Do Individuals Want to Reduce Their Drinking?

Given the above short-term costs and other longer-run costs of alcohol consumption, a natural question to ask is whether individuals are aware of the costs of alcohol consumption.

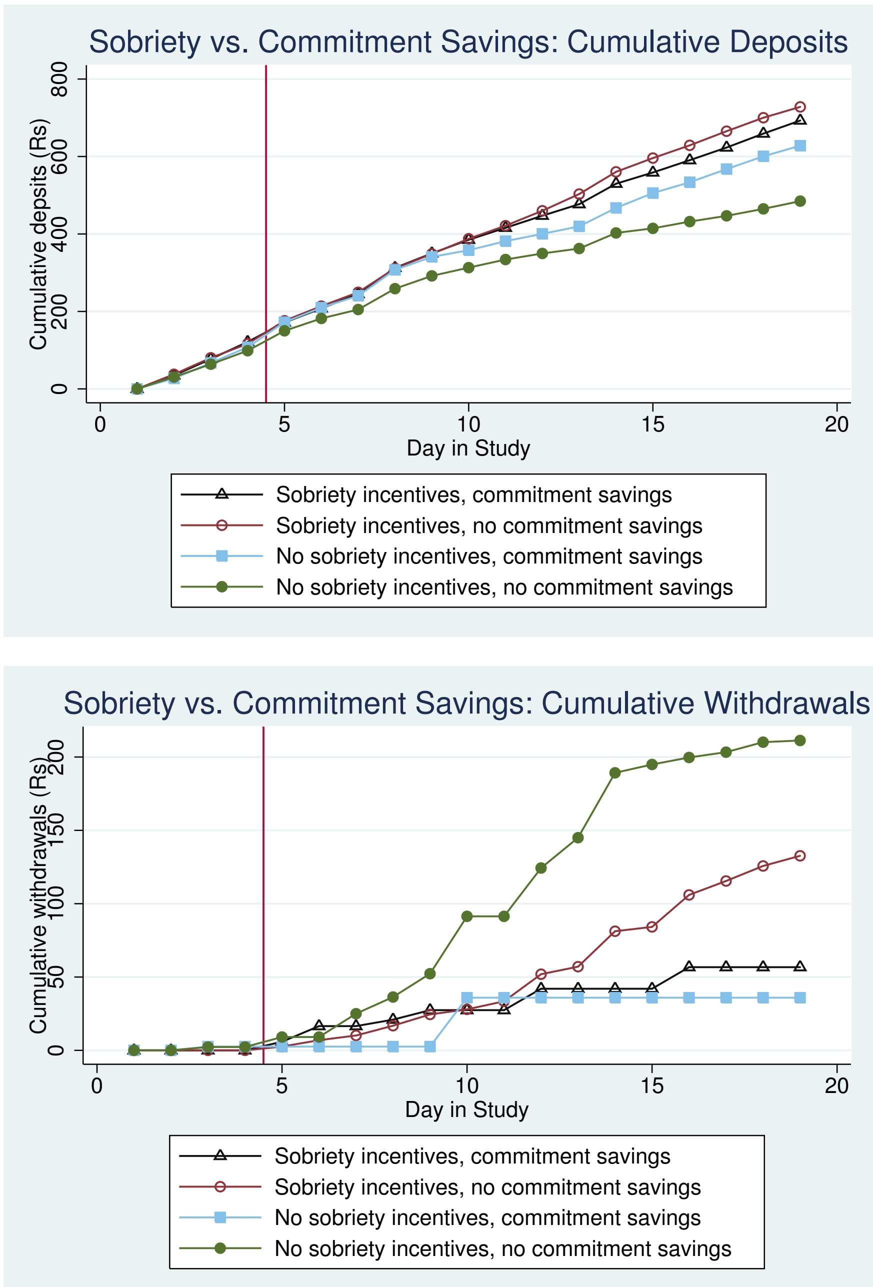

Notes: This figure splits up the results shown in the upper panel of Figure **??** into cumulative deposits (upper panel) and cumulative withdrawals (lower panel).

Figure 1.7: *Sobriety Incentives vs. Commitment Savings: Deposits and Withdrawals*

In particular, if these costs exceed the benefits of drinking, why are individuals not reducing their consumption? This section considers the extent to which self-control problems contribute to individuals' demand for receiving incentives for sobriety. After receiving incentives for three days, individuals in the Choice Group were asked to choose between incentives to arrive sober and different amounts of unconditional payments. Individuals in the Choice Group first made these choices at the beginning of Phase 3 (day 7), and then again at the beginning of Phase 4 (day 13). Finally, regardless of experimental condition, all study participants were given the same choices at the end of Phase 4 (day 20). This structure allows me to investigate whether individuals in the Choice Group changed their choices over time, and whether receiving incentives in earlier phases of the study affected individuals' demand for commitment. During each choice session, individuals chose their incentive structure for the subsequent six study days.[44]

The demand for incentives was high, even when choosing incentives entailed a potential (Choice 2) or certain (Choice 3) reduction in overall study payments (upper panel of Figure 1.8 and Table 5.7). More than one third of individuals in the Choice Group preferred sobriety incentives over receiving Rs. 150 regardless of their breathalyzer scores, and in each week, over 50 percent of individuals chose incentives over receiving Rs. 120 unconditionally. Holding attendance constant, this choice implied losses of Rs. 30 ($0.50) in study payments at the minimum (on days when the individual visits the study office sober) and Rs. 90 ($1.50) at the maximum (on days when the individual visits the study with a positive breathalyzer score). These amounts are economically meaningful, representing between 10 and 30 percent of reported daily labor earnings. Moreover, the fraction of individuals choosing sobriety

[44] Attrition and inconsistencies of preferences during the choice session cause relatively minor concerns for the below analysis (Table 5.6). In the Choice Group, less than 7 percent of individuals missed their choices in any given week, and, in each week, less than 7 percent of individuals stated inconsistent preferences. Furthermore, over 88 percent of all study participants completed the endline choices with consistent choices. This fraction varies only slightly across treatment groups (90.1 in the Incentive Group and 88.0 in the Choice Group vs. 86.7 in the Control Group). In an attempt to be conservative regarding the demand for commitment in Figure 1.8 and Table 5.7, an individual is counted as not choosing incentives in any given choice when he did not attend the respective choice session or when he attended, but made inconsistent choices. The below regressions in Tables 1.12 and 1.13 are conditional on attendance. The analysis is robust to alternative specifications.

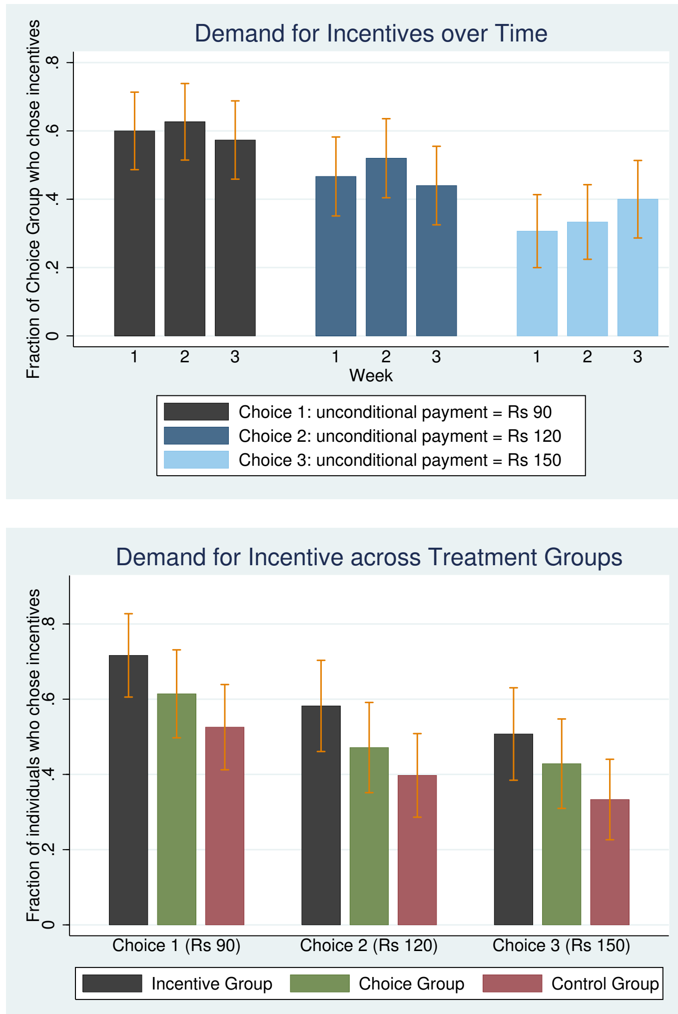

Notes: This figure depicts the fraction of individuals who preferred incentives for sobriety over unconditional payments.

1. All choices were made for the subsequent week, i.e. for the next six days in the study. Under incentives for sobriety, if an individual visited the study office, he received Rs. 60 ($1) if his breathalyzer score was positive, and Rs. 120 ($2) if his breathalyzer score was zero.

2. Unconditional payments are Rs. 90 (Choice 1), Rs. 120 (Choice 2), and Rs. 150 (Choice 3). Hence, an individual exhibited demand for commitment to sobriety if he chose incentives in Choices 2 and/or 3. At any point in time, individuals made all rickshaw peddlers three choices before one of these choices was randomly selected to be implemented.

3. If an individual did not complete the set of choices, or if he chose inconsistently, the observation is counted as *not* preferring incentives. During a given choice session, an individual chose inconsistently if he chose Option B for the unconditional amount Y_1, but Option A for the unconditional amount Y_2 with $Y_2 > Y_1$.

4. The upper panel of the figure shows how the fraction of individuals in the Choice Group who chose incentives evolved over time (i.e. on days 7, 13, and 20 of the study). The lower panel of the figure depicts the fraction of individuals who chose incentives on day 20 in the three treatment groups, i.e. it shows how previous exposure to incentives affected the demand for incentives. Error bars show 95 percent confidence intervals.

Figure 1.8: *Choices Across Treatment Groups and Over Time*

incentives over Rs. 150 unconditionally did not decline over time. Instead, though not statistically significant, it in fact increased slightly over the course of the study.

Subjects' choices provide clear evidence of self-control problems. In particular, the fraction of individuals who exhibited costly demand for commitment was larger than found previously for smoking (Gine *et al.* (2010)) or exercising (Royer *et al.* (2014)). A growing literature has demonstrated demand for commitment in a number of domains.[45] However, with the exceptions of Beshears *et al.* (2011) and Milkman *et al.* (2014), there is little existing evidence that individuals are willing to pay for commitment beyond the potential costs of failing to achieve the behavior they are committing to.[46] In my study, about a third of subjects made choices that implied significant losses in study payments even in the best case of visiting the study office sober every day.

Moreover, Table 1.12 shows the relationship between the number of sober days in each phase of the study and the demand for sobriety incentives. Individuals who visited the study office sober more often in the incentivized Phase 2 were subsequently more likely to choose incentives for all three unconditional amounts. This is not surprising since expected study payments from choosing incentives were higher if a study participant was more likely to visit the study office sober. In contrast, the difference in sobriety between Phase 2 (when some individuals were receiving incentives) and Phase 1 (the pre-incentive period) positively predicts demand *only* for costly incentives (i.e. when the unconditional payment is Rs. 150). This is reassuring since individuals should have chosen costly incentives only when they

[45]For instance, Ashraf *et al.* (2006) and Beshears *et al.* (2011) on commitment savings; Gine *et al.* (2010) on smoking cessation; Kaur *et al.* (2014) on self-control at the workplace; Ariely and Wertenbroch (2002), Augenblick *et al.* (2014), and Houser *et al.* (2010) on effort tasks; and Royer *et al.* (2014) and Milkman *et al.* (2014) for gym attendance. See Bryan *et al.* (2010) and Augenblick *et al.* (2014) for overviews.

[46]A large number of studies in the psychology literature have associated excessive alcohol consumption with survey measures of (lack of) self-control, behavioral undercontrol, and susceptibility to temptation (Hull and Slone (2004)). In addition, the existence of and demand for disulfiram (Antabuse) can be viewed as evidence of self-control problems causing alcohol consumption (Glazer and Weiss (2007), Bryan *et al.* (2010)). However, evaluations of disulfiram treatment for alcohol dependence have shown inconsistent findings, in a large part because of low treatment adherence as in Fuller *et al.* (1986). Studies evaluating incentives to increase compliance (O'Farrell *et al.* (1995)) and a combination of disulfiram with other medication to reduce cravings or withdrawal symptoms such as naltrexone or acamprosate have found more promising results (Suh *et al.* (2006)), but do not necessarily show evidence of demand for commitment and, hence, self-control problems.

expected them to help increase their sobriety, which in turn should have been informed by their own experience in the study.

Exposure to incentives for sobriety increased the demand for the incentives (lower panel of Figure 1.8). For all three choices, the Incentive Groups were more likely to choose incentives than the Control Group. The fraction of individuals choosing incentives in the Choice Groups (on day 20) was in between the corresponding fractions in the Incentive and Control Groups. The corresponding regressions show significant differences between the fraction choosing incentives in the Incentive and Control Groups for all three choices (Table 1.13). These differences are not explained by differences in sobriety while making these choices, or by differences in expectations of future sobriety under incentives. Before preferences were elicited, individuals were asked how often they expected to visit the study office sober if they were to be given incentives for sobriety. Reassuringly, subjects' beliefs about their expected sobriety under incentives strongly predicts demand for incentives. Finally, higher sobriety during the time of choosing predicts a higher probability of choosing incentives.

The above findings raise the question why so many study participants exhibited the demand for commitment despite the fact that overall drinking only fell moderately. Several, not mutually exclusive explanations are possible. First, the above estimates suggest that incentives for sobriety caused several small benefits, which taken together may well exceed Rs. 30. On average, though not statistically significant, sobriety incentives increased reported earnings by about Rs. 17.6), and reduced reported alcohol expenditures by about Rs. 8.7. Moreover, as shown above, savings increased significantly. Increasing sobriety may have also improved other decisions, and individuals may have valued daytime sobriety on its own despite potentially increased disutility of work due to increased physical pain.

Second, partial naïveté may have contributed to the demand for commitment. On the one hand, underestimating the extent of their self-control problems due to partial or full naïveté as in O'Donoghue and Rabin (1999) may lower the demand for (costly) commitment by

Table 1.12: *Demand for Incentives over Time*

VARIABLES	(1) Rs 90	(2) Rs 90	(3) Rs 90	(4) Rs 120	(5) Rs 120	(6) Rs 120	(7) Rs 150	(8) Rs 150	(9) Rs 150
Week 2	0.01	0.04	0.01	0.05	0.07	0.04	0.02	0.03	0.01
	(0.060)	(0.060)	(0.063)	(0.070)	(0.070)	(0.070)	(0.067)	(0.068)	(0.068)
Week 3	0.01	-0.01	-0.04	0.00	-0.02	-0.03	0.12	0.11	0.10
	(0.082)	(0.081)	(0.075)	(0.081)	(0.079)	(0.076)	(0.081)	(0.081)	(0.079)
BAC during choice	-1.63***			-1.12***			-0.67**		
	(0.318)			(0.322)			(0.279)		
Days sober in Phase 1		0.06			0.02			-0.05	
		(0.043)			(0.045)			(0.049)	
Days sober in Phase 2		0.09**			0.07			0.07	
		(0.042)			(0.043)			(0.045)	
Incentives increased sobriety			0.04			0.08			0.15**
			(0.065)			(0.077)			(0.071)
Exp frac sober under incentives			0.56***			0.40***			0.21**
			(0.083)			(0.085)			(0.086)
Constant	0.76***	0.40***	0.22**	0.58***	0.36***	0.18**	0.37***	0.27***	0.12
	(0.057)	(0.080)	(0.085)	(0.065)	(0.082)	(0.080)	(0.062)	(0.079)	(0.077)
Observations	211	211	211	211	211	211	211	211	211
R-squared	0.122	0.147	0.205	0.057	0.046	0.109	0.028	0.024	0.062

Notes: This table considers the relationship between the demand for incentives and sobriety for the Choice Group at different points in the study.

1. In all columns, the outcome variable is whether the individual chose incentives over unconditional payments. The unconditional amounts are Rs. 90 in columns (1) through (3), Rs. 120 in columns (4) through (6), and Rs. 150 in columns (7) through (9).
2. "BAC during choice" refers to the subjects' blood alcohol content measured before making choices between incentives and unconditional amounts. "Exp sober days under incentives" are subjects' answers to asking how many days they expected to show up sober if they were to receive incentives for sobriety during the subsequent six days (always asked before choices were made). "Days sober in Phase 1" and "Days sober in Phase 2" refer to the number of days the individual visited the study office sober during Phase 1 and 2, respectively. "Incentives increased sobriety" indicates whether the difference in the fraction of sober days in the phase before choosing and the fraction of sober days in Phase 1 is positive.
3. Standard errors are in parentheses, clustered by individual. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

Table 1.13: *Demand for Incentives Across Treatment Groups*

VARIABLES	(1) Rs 90	(2) Rs 90	(3) Rs 90	(4) Rs 120	(5) Rs 120	(6) Rs 120	(7) Rs 150	(8) Rs 150	(9) Rs 150
Incentives	0.13*	0.15**	0.13*	0.15*	0.16**	0.15**	0.14*	0.16**	0.14*
	(0.075)	(0.070)	(0.070)	(0.082)	(0.075)	(0.076)	(0.081)	(0.077)	(0.078)
Choice	0.10	0.07	0.08	0.09	0.06	0.07	0.11	0.09	0.10
	(0.079)	(0.074)	(0.074)	(0.081)	(0.078)	(0.078)	(0.079)	(0.078)	(0.078)
BAC during choice	-1.70***		-0.85**	-1.10***		-0.32	-1.10***		-0.52
	(0.315)		(0.358)	(0.323)		(0.355)	(0.304)		(0.349)
Exp sober days under incentives		0.10***	0.08***		0.08***	0.07***		0.06***	0.06***
		(0.011)	(0.014)		(0.011)	(0.013)		(0.011)	(0.012)
Observations	215	215	215	215	215	215	215	215	215
R-squared	0.144	0.251	0.275	0.070	0.170	0.173	0.071	0.122	0.130
Control mean	0.494	0.494	0.494	0.373	0.373	0.373	0.313	0.313	0.313

Notes: This table considers how the two sobriety incentives treatments affected the demand for incentives.

1. In all columns, the outcome variable is whether the individual chose incentives over unconditional payments. The unconditional amounts are Rs. 90 in columns (1) through (4), Rs. 120 in columns (5) through (8), and Rs. 150 in columns (9) through (12).
2. "BAC during choice" refers to the subjects' blood alcohol content measured during the visit to the study office when he was choosing between incentives and unconditional amounts. Before making these choices, individuals were asked on how many days they expected to show up sober if they were to receive incentives for sobriety during the subsequent six days. The variable "Expected sober days under incentives" refers to subjects' answer to this question.
3. Standard errors are in parentheses, clustered by individual. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

decreasing the perceived benefits of commitment (Laibson (2015)). On the other hand, partial naïveté can also increase the demand for commitment by causing individuals to overestimate the effectiveness of commitment devices in overcoming their self-control problems.[47] In the context of my study, while being aware of their own self-control problems, some individuals may have overestimated the usefulness of the incentives for sobriety in reducing their daytime or overall drinking.

1.7 Conclusion

This paper provides evidence that self-control problems may not only cause undesired alcohol consumption, but that alcohol itself exacerbates present bias, and hence creates further self-control problems in other domains. Increasing sobriety during the day causes a stark increase in individuals' savings at the study office. I provide evidence that this increase was not just the result of mechanical effects from increased resources, but due to lowered self-control problems in savings decisions as a consequence of decreased myopia. Taken together, these results imply that effective commitment devices for sobriety not only help individuals reduce undesired alcohol consumption, but also lessen self-control problems caused by alcohol. More generally, the results suggest that alcohol changes decision processes in a way that may reinforce poverty.

A significant fraction of cycle-rickshaw peddlers in a large Indian city were willing to sacrifice money for commitment to increase sobriety during the day, indicating a greater awareness of and willingness to overcome self-control problems than found in most other settings. This high prevalence of self-control problems suggests that "sin taxes" could be an attractive policy option (Gruber and Kőszegi (2001), O'Donoghue and Rabin (2006)). Given the negative correlation of alcohol consumption and income, such taxes may be regressive. However, the regressiveness of taxation may be mitigated if consumers have self-control problems. Gruber and Kőszegi (2004) show that "sin taxes" can even be progressive (in

[47]For a more detailed treatment of this argument and an application in the savings domain, see John (2014).

particular in the utility domain) if poor individuals are more price-elastic and/or are more present-biased compared to rich individuals. The results from this study suggest that the regressiveness of taxing alcohol may be further lessened due to effects of reduced drinking on earnings and savings. However, given that the price elasticity of the demand for alcohol in this setting is below unity, increasing taxes would further reduce individuals' – and therefore many families' – already low income net of alcohol expenditures, unless the effects of reduced drinking on earnings turn out to be particularly large.[48]

A second, more extreme policy option could be prohibition, as already implemented in several Indian states such as Gujarat. Prohibition may be a particularly attractive policy option for India and other developing countries compared to developed countries since the distribution of alcohol consumption is heavily skewed, with the majority of the population abstaining from alcohol and a relatively large share among the drinkers consuming alcohol excessively. However, enforcement of prohibition is known to be difficult and may result in other unintended consequences such as crime and corruption (Thornton (1991)). Moreover, many Indian state governments heavily depend on excise taxes, which makes the implementation of prohibition difficult. Given these concerns, second-best policies aimed at reducing the costs of inebriation by shifting critical decision away from drinking times could be welfare-improving even if they do not change overall drinking levels.

[48]In most other studies, the price elasticity of alcohol consumption has been found to be below unity, and heavy drinkers' price response tends to be particularly small Manning *et al.* (1995). For an overview, see Wagenaar *et al.* (2009).

Chapter 2

Technology Diffusion and Appropriate Use: Evidence from Western Kenya[1]

2.1 Introduction

Low technology adoption has hampered growth in agricultural productivity in Sub-Saharan Africa (Bank (2008), Evenson and Gollin (2003)). Insufficient knowledge of appropriate use may inhibit technology diffusion since information regarding appropriate use of often risky new technologies is scarce in many contexts. In the absence of reliable sources of information, sharing experiences and learning from each other can be important ways to increase knowledge for small-scale farmers.[2] If input decisions are unobserved, the ability and willingness of farmers to communicate with each other become crucial aspects of knowledge diffusion. However, lack of ability or willingness to communicate among small-scale farmers may curb knowledge diffusion and, hence, technology adoption (BenYishay

[1]Co-authored with Esther Duflo, Michael Kremer, and Jon Robinson

[2]See, for example, Foster and Rosenzweig (1995), Munshi (2004), Bandiera and Rasul (2006), Conley and Udry (2010); see Foster and Rosenzweig (2010) for a review.

and Mobarak (2014)).

Using evidence from a field experiment with over 30,000 small-scale maize farmers, this paper considers the impact of introducing a simple measurement tool for fertilizer on knowledge diffusion and fertilizer use in Western Kenya. In the experiment, we randomly selected 15% of farmers to receive a “Bluespoon", a small measuring spoons painted blue. These spoons measure the physical quantity of one-half teaspoon, which was found to result in the highest profits on average among four different options in earlier work (Duflo *et al.* (2008)). The Bluespoons could be used to apply fertilizer to maize and they were also made available in local shops. Since the spoons were commercially unavailable before our project, we identify spillovers by examining whether the other 85% of farmers not offered spoons learned about or purchased them. Moreover, the experiment included two cross-randomized school-level treatment conditions. First, we attempted to reduce the cost of sharing information by encouraging farmers to form cooperatives. While we helped organizing the groups and coordinated the first few meetings, no information was provided directly. Second, motivated by previous work (Duflo *et al.* (2011)), we provided farmers with small, time-limited discounts, valid within a short window right after harvest, redeemable at a fertilizer shop in local market centers. This treatment condition was supplemented by text message reminders for a randomized subsample of farmers.

The main results of our paper concern the diffusion of Bluespoons and knowledge among farmers. First, as predicted by a target input model as in Jovanovic and Nyarko (1996), being assigned to receive a Bluespoon significantly increases fertilizer use among farmers, as measured by administrative data on discount coupon redemption and self-reported fertilizer use by farmers. Second, we consider the diffusion of Bluespoons through social networks. Bluespoons were extremely popular among farmers: about one year after their introduction, 59% of farmers had heard of the Bluespoon, and 23% owned one, which is a remarkable level of diffusion, given that the Bluespoons had not been not previously available to farmers before the experiment. We relate the diffusion of Bluespoons to two treatment conditions. Contacts of Bluespoon farmers were 14 percentage points more likely to have heard of the Bluespoon and 7 percentage points more likely to own one. Moreover, individuals in cooperative schools were 7 percentage points more likely to own a Bluespoon. Third, the diffusion of Bluespoons was paralleled by diffusion of knowledge. Farmers who were assigned to receive a Bluespoon are significantly more likely to identify one-half teaspoon – the quantity measured by a Bluespoon – as the optimal quantity of fertilizer to be used per planting hole. Finally, while we also find some evidence that the diffusion of knowledge and Bluespoon is increased by the coupon intervention, we do not find any evidence of the diffusion of fertilizer adoption through social networks. In preliminary results, none of the different treatment conditions affected fertilizer use among contacts of farmers who participated in the program.

The paper also provides evidence on the impact of a large-scale discount coupon intervention on fertilizer use. The time-limited discount around harvest time increased reported fertilizer use by 9 to 13 percentage points (on a base of about 50%). These results suggest that, given the low fiscal costs and potentially lower distortions associated with small, time-limited discounts, such discounts may be an attractive policy alternative to larger subsidies at the time of planting for governments intending to increase fertilizer use on a broad scale. While small discounts have a more muted effect on adoption than larger subsidies, they are clearly much less expensive. Our results on text message reminders supplementing such a program are mixed. On the one hand, the text message reminders significantly increased fertilizer coupon redemption, as others have found in regards to ARV adherence (Lester *et al.* (2010)) and savings (Karlan *et al.* (2014a)). However, on the other hand, we do not find evidence of the text message reminders increasing fertilizer use, suggesting that the reminders only had an impact on coupon redemption of infra-marginal farmers, i.e. the reminders increased coupon redemption among individuals who would have used fertilizer anyway.

The remaining parts of this paper are organized as follows. First, Section 2.2 describes the experimental design and data. Section 2.3 then describes the impact of the different treatment conditions on fertilizer use. Section 2.4 considers the impact of the different treatments on Bluespoon and knowledge diffusion. Section 2.5 concludes.

2.2 Experimental Design and Data

2.2.1 Experimental Design

The experiment was conducted with small-scale farmers over two agricultural seasons (short rains 2010 and long rains 2011).[3] To reach a high number of farmers at affordable costs, we leveraged a large social network: the parents of school children in 184 rural primary schools in Western Kenya. Children in these schools were given a letter inviting their parents inviting them to a meeting at the school. Every parent who participated in a meeting at a particular school was eligible for the treatment administered at this school. At the meeting, the experimental treatments (if any) were explained, and enumerators completed a short baseline survey with each participant. We enrolled 26,856 farmers (on average about 146 farmers per school) into the program, whom we refer to as "original respondents" below.

Coupons, Cooperatives, and SMS Reminders. We cross-randomized farmers at the school level into two main treatment conditions. First, to lower costs of communicating and sharing knowledge about fertilizer and other agricultural practices, we encouraged farmers to form farmers' cooperatives to talk about agriculture. While we facilitated organizing the groups and coordinated the first few meetings, we did not provide any information directly to farmers. Second, in previous work, we find that providing farmers with small incentives to invest in fertilizer when they have money (right after harvest) can substantially increase usage (Duflo *et al.* (2011)). Hence, to increase usage exogenously, we implemented a scalable version of a program to provide farmers with small, time-limited discounts which were valid within a short window (3 to 4 weeks) right after harvest, redeemable at a local shop. Farmers received coupons for a discount of 15% of the price of fertilizer. The coupon was

[3] In our study area, there are two agricultural seasons for growing maize each ear: the"long rains" from March/April to July/August, and the "short rains" from July/August until December/January.

valid for discounts to either diammonium phosphate fertilizer (DAP), used at planting, and calcium ammonium nitrate fertilizer (CAN), used at top dressing, when the maize plant is knee high, approximately one to two months after planting. Farmers could choose any combination of DAP and CAN up to 25 kilograms in total. Moreover, to evaluate the effectiveness of text message reminders, we randomly selected half of the schools that received fertilizer discount coupons into a text message reminder program. In these schools, we randomized a subset of individuals who either owned a cellphone or had access to a cellphone to receive a text message reminder two days before the expiration date of their time-limited discount. We conducted a distinct individual- and school-level randomization of text message reminders in the second season, i.e. individuals who received a text message reminder in the first season did not necessarily receive one in the second season as well. This allows us to also consider the effect of receiving a text message reminder on coupon redemption and fertilizer use in the consecutive season.

Social Networks and Bluespoons. Due to staff and time constraints at the school meetings, we restricted the collection of social network information to a randomly selected subset of individuals. At each school meeting, we asked about 25% of farmers in each school to provide names and contact information of up to 3 individuals outside their own household with whom they discussed agriculture whom we refer to as “agricultural contacts" below. About two thirds of these individuals were randomly selected to receive a "Bluespoon," a half-teaspoon, painted blue, which farmers could use to apply fertilizer to their maize. We chose this simple technology because in earlier work we had found 1/2 teaspoon of CAN to yield the highest profits on average among four different quantities (Duflo *et al.* (2008)). We delivered these spoons in an additional short meeting (at the same schools) to which we had invited the randomly selected subset of farmers. In addition, farmers were given the information that in earlier work we had found that this quantity of CAN resulted in the highest profits on average. Again due to logistical constraints, the timing of these meetings was randomized. One third of individuals who were randomized to receive a Bluespoon (group 1) was visited about one week after the large school meeting, one third (group 2) was visited about three weeks after the meeting, and the remaining individuals (group 3) were visited about six weeks after the meeting. Bluespoons were also made available to anyone for a nominal fee ($0.05) in the local market center at the same shops which handled the coupon redemption.

Tracking Bluespoon and Knowledge Diffusion. The diffusion of Bluespoons and knowledge in social networks is a key outcome of this paper, which we measured in several ways. First, farmers who received a Bluespoon were also given 10 vouchers for Bluespoons along with an encouragement to share them with their friends. The vouchers did not entitle farmers to purchase spoons at a discounted price or for free, i.e. there was no financial incentive to use the vouchers. However, we asked farmers to encourage their friends to use these vouchers when purchasing spoons. Individual-specific identifiers printed on all vouchers allow us to trace Bluespoon diffusion. Second, we administered a short survey at endline (similar to the baseline survey) to which we invited all farmers who had been surveyed at baseline. In addition, we invited one randomly selected agricultural contact of all farmers who had provided us with information of at least one individual (as described above). In this short endline survey, we asked all individuals about fertilizer use, knowledge and ownership of Bluespoons, as well as their belief regarding the optimal quantity of fertilizer per planting hole. Third, we conducted in-depth endline surveys with a randomly selected subset of original respondents and agricultural contacts. In each school, we randomly selected 16 individuals among all respondents who had been asked about their agricultural contacts during the baseline survey. The sample was chosen such that half of these 16 individuals had been selected to receive a Bluespoon, and, accordingly, the other half had been not been selected to receive a Bluespoon. Furthermore, for each of the 16 selected individuals from each school, we again selected one of their randomly selected contact (if available) to be visited as well for the in-depth endline survey.

2.2.2 Data

Our analysis uses four main sources of data.

Short Baseline Surveys. First, we conducted short baseline surveys during school meetings. These surveys elicited information on previous use of CAN and DAP, as well as on fertilizer quantities and the application method used. In addition, we asked respondents how much of each of the two types of fertilizer they thought should be used per maize planting hole. To get an objective measure of these quantities, subjects were asked to pick one of four small containers that contained different amounts of fertilizer, one of which contained half a teaspoon of fertilizer. In addition, we asked a randomly selected subsample for contact information of up to three individuals outside their household with whom they discussed agriculture on a regular basis.

Short Endline Surveys. Second, we conducted similarly short follow-up surveys at endline school meetings, including the same questions as at baseline, plus a verbal measure of optimal qualities per planting hole. In addition to usage and quantities, endline surveys also included questions on group membership, as well as Bluespoon knowledge, usage, and ownership.

Long Endline Surveys. Third, to get more detail on fertilizer use and communication among farmers, we conducted a longer one-on-one survey with a randomly selected subset of respondents at their homes. This survey included several demographic questions, as well as more detailed questions on fertilizer and Bluespoon usage and knowledge. The survey also included a number of questions to measure social interactions.

Administrative Redemption Data. Fourth, we collected administrative data on fertilizer coupon redemption as well as Bluespoon voucher redemption and sales at local fertilizer shops.

2.2.3 Summary Statistics, Randomization, and Attrition

In an effort to reach as many farmers as possible during the school meetings, we kept the baseline survey conducted at these meetings as short as possible. Hence, only a few baseline measures other than prior usage of fertilizer are available for balance checks. We collected additional covariates during the endline, but since these surveys were conducted after the program had been implemented, potential endogeneity concerns arise. We therefore look at outcomes which are unlikely to be affected by the program, namely household demographics.

We check for orthogonality of treatments by running the following regression for each background characteristic y_i:

$$y_i = \beta_0 + T_i'\beta + \varepsilon_i \tag{2.1}$$

where T_i is the set of treatment indicators for individual i. At the school level, the treatment conditions are whether the school was sampled for the coupon and cooperative treatments, and whether the school was sampled to receive SMS reminders (to redeem their fertilizer coupon). We also include an interaction between the two school-level treatments to check for potential interaction effects. At the individual level, the treatments are receiving a Bluespoon and being sampled to receive an SMS reminder. Since only farmers with a cell phone or access to a cell phone could receive an SMS (which, from Table 2.1, makes up 54% of the sample) we also include a control for eligibility to the SMS treatment (coefficients omitted).

Randomization Checks. Table 2.1 shows basic summary statistics and randomization checks for the study sample. The regression results shown in this table reveal little evidence of imbalance across treatment groups. Columns 1 through 5 show sample characteristics based on the large baseline school meetings. Importantly, prior fertilizer usage is similar across all treatment groups (columns 1 through 3). There are slight imbalances in reported qualities of fertilizer to be used per planting hole (column 4) and phone availability (column 5). Whenever possible, we control for these variables in our analyses below. The remaining columns (6 through 11) are based on the in-depth endline surveys at respondents' homes. While there are a few variables which differ significantly across treatments, these differences seem likely due to chance.

Attrition patterns. We also check for differential attrition patterns in Table 2.2. We find that farmers in cooperative-only schools are less likely to be in the school endline and Bluespoon farmers are more likely to be in the school endline. However, we find no evidence of differential attrition in the home endline. We also find no evidence of attrition among sampled contacts, other than that the friends of farmers who received SMS reminders are more likely to be in the sample, a result which is likely due to chance given the limited overall effect of that program.

2.3 Fertilizer Use

This section considers the impact of the different treatments on fertilizer use. Section 2.3.1 first presents results based on administrative data (coupon redemption) before Section 2.3.2 considers self-reported fertilizer use based on short and long endline surveys.

2.3.1 Fertilizer Coupon Redemption

Overall coupon redemption, summarized in Table 2.3, was relatively low: 18% and 12% of farmers redeemed their discount coupons in season 1 (short rains 2010) and season 2 (long rains 2011), respectively. The fraction of farmers taking advantage of the program is significantly lower than in Duflo *et al.* (2011), in which take-up was between 31 and 39 percent in the two seasons, respectively. On the one hand, it is still a substantial fraction given that the subsidy is small (15% of the price of fertilizer), and that farmers had to redeem the coupon at a local shop rather than at their homes, e.g. leaving room for the possibility of individuals forgetting about the redemption period. Moreover, while the

Table 2.1: *Summary Statistics and Randomization Checks*

VARIABLES	(1) DAP	(2) CAN	(3) Fert	(4) 1/2 Tsp	(5) Phone	(6) Male	(7) Age	(8) Literate	(9) Education	(10) Married	(11) Farming
Coupon School	0.05	0.05	0.06	0.02*	-0.05*	0.01	-0.21	0.01	0.12	0.02	0.01
	(0.060)	(0.045)	(0.059)	(0.014)	(0.030)	(0.036)	(0.855)	(0.033)	(0.259)	(0.027)	(0.032)
Cooperative School	-0.04	-0.04	-0.03	0.00	-0.03	-0.02	-1.55*	0.04	0.71***	0.03	-0.04
	(0.059)	(0.039)	(0.058)	(0.013)	(0.029)	(0.033)	(0.805)	(0.031)	(0.248)	(0.023)	(0.028)
Coupon Coop Interaction	0.02	0.07	0.02	-0.01	0.02	0.03	1.70	-0.03	-0.52	-0.02	0.02
	(0.078)	(0.057)	(0.077)	(0.019)	(0.039)	(0.045)	(1.058)	(0.040)	(0.337)	(0.033)	(0.039)
Individual selected to receive a Bluespoon	-0.01	0.00	-0.01	0.01	0.02*	0.03	-0.33	0.02	0.04	0.00	0.00
	(0.010)	(0.011)	(0.010)	(0.011)	(0.013)	(0.018)	(0.551)	(0.016)	(0.156)	(0.016)	(0.016)
School sampled to receive SMS	-0.06	-0.06	-0.07	0.00	0.02	0.02	-0.26	-0.02	-0.27	-0.01	-0.00
	(0.052)	(0.042)	(0.051)	(0.012)	(0.024)	(0.034)	(0.768)	(0.028)	(0.248)	(0.026)	(0.029)
Individual sampled to receive SMS	0.02	-0.02	0.01	-0.00		-0.05	0.31	0.02	0.06	0.05**	0.03
	(0.021)	(0.016)	(0.021)	(0.013)		(0.040)	(1.065)	(0.026)	(0.291)	(0.025)	(0.036)
Observations	26,856	26,856	26,856	26,856	26,856	2,733	2,684	2,733	2,724	2,733	2,733
Adjusted R-squared	0.023	0.021	0.023	0.002	0.002	0.019	0.037	0.154	0.170	0.021	0.012
Mean of dependent variable	0.624	0.374	0.652	0.241	0.541	0.353	43.20	0.685	6.419	0.814	0.777
Standard dev of dep variable	0.484	0.484	0.476	0.427	0.498	0.478	13.56	0.465	4.082	0.389	0.416
Number of clusters	184	184	184	184	184	184	184	184	184	184	184

Notes: This table shows basic summary statistics for the study sample, and their relationship to the randomized treatment conditions.

(1) Columns 1 through 5 show summary statistics based on the short baseline surveys during school meetings. 'DAP', 'CAN', and 'Fert' are indicator variables for reported previous use of DAP, CAN, or either type of fertilizer, respectively. '1/2 Tsp' indicates whether the respondent chose the container that included the quantity of 1/2 teaspoon when asked what quantity of CAN should be used per planting hole. 'Phone' indicates whether the respondent owned or had access to a cell phone (which determined eligibility for the SMS reminder treatment).

(2) Columns 6 through 11 show summary statistics for demographics based on in-depth one-to-one endline surveys at respondents' homes. 'Literate' is an indicator for ability to read and write, 'Education' denotes the number of years of education, and 'Farming' indicates whether the respondent reported farming as her main occupation.

(3) Standard errors are in parentheses, clustered by school. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

Table 2.2: *Attrition*

	In short endline	In long endline	
	(1)	(2)	(3)
Coupon School	-0.02	0.01	-0.00
	(0.015)	(0.015)	(0.021)
Cooperative School	-0.04**	0.01	0.00
	(0.020)	(0.013)	(0.019)
Coupon Coop Interaction	0.04*	-0.02	0.01
	(0.025)	(0.019)	(0.027)
Individual selected to receive a Bluespoon	0.03***	0.00	0.00
	(0.009)	(0.008)	(0.013)
School sampled to receive SMS	-0.03*	0.01	-0.02
	(0.017)	(0.015)	(0.018)
Individual sampled to receive SMS	0.02*	0.01	0.06**
	(0.012)	(0.022)	(0.024)
Observations	26,856	2,914	2,681
R-squared	0.004	0.001	0.002
Overall mean	0.828	0.938	0.863
Control mean	0.848	0.933	0.874

Notes: This table shows attrition patterns in the study.

(1) The outcome variable in columns 1 and 2 is whether we were able to conduct a short endline with an individual (among everyone surveyed at baseline). The outcome variable in columns 3 is whether we were able to conduct a long endline survey with an agricultural contact named by the original respondents (among everyone sampled for the long endline survey).

(2) The independent variables in columns 5 and 6 correspond to the original respondent who named the particular friend during the baseline meeting.

(3) Standard errors are in parentheses, clustered by school. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

previous intervention entailed a personal explanation of the program to each farmer at his or her home, the current intervention was conducted in large public meetings with almost 150 farmers on average, such that it is possible that a fraction of farmers may not have fully understood all details of the coupon program. On the other hand, given that over half of the respondents individuals reported fertilizer use in each of the seasons, the relatively low coupon redemption rate is puzzling, especially since liquidity constraints are an unlikely explanation due to the timing of the intervention.[4] Finally, most coupons were used to purchase DAP (planting) fertilizer, and, as in previous work, the quantities purchased conditional on making use of the program were small. In each season, farmers who redeemed their coupon purchased about 9 kgs of fertilizer with their coupon.[5]

Table 2.4 examines the effect of the different treatment conditions on fertilizer use, as measured by coupon redemption. It includes regressions similar to Tables 2.1 and 2.2, though restricted to coupon schools only. The outcome variable in the first five columns is whether a particular individual redeemed their coupon; the outcome variable in the remaining five columns is the (unconditional) quantity redeemed using the coupon (if any). In each half of the table, the first two columns consider season 1, and the remaining three columns consider season 2. The table shows three sets of results.

First, we do not find evidence that the cross-randomized school-level treatments affected coupon redemption. None of the 20 coefficients in the entire table show statistically significant effects of the cooperative or school-level text message treatments. This is not particularly surprising given that one would expect the effect of the cooperative treatment to come with some lag, and – given the estimated effect size of text message reminders at

[4] Similarly to our study, Carter *et al.* (2013) found low coupon redemption rates in a recently completed randomized evaluation with maize farmers in Mozambique. In this intervention, despite a 73% subsidy (worth USD 117) and demanding eligibility criteria (access to agricultural extension, ability to pay USD 32, being a “progressive farmer" and land size between 0.5 and 5 hectares), coupon redemption for a package of improved seeds and fertilizer was only 16% higher in the treatment compared to the control group (22% vs. 6%).

[5] The distribution of redeemed quantities was also highly skewed. In season 1, the median farmer who redeemed her coupon purchased 4 kgs of DAP and only ten percent of farmers bought at least 5 kgs of CAN conditional on redeeming the coupon. Given that using half a teaspoon on a one acre plot requires about 50 kgs of fertilizer, these quantities are relatively small.

Table 2.3: *Fertilizer Coupon Redemption*

	Season 1 (SR 10)	Season 2 (LR 11)
Fraction of individuals who redeemed their coupon		
Percentage of coupons redeemed	0.18	0.12
	(0.38)	(0.32)
Percentage of coupons redeemed for DAP	0.15	0.10
	(0.36)	(0.30)
Percentage of coupons redeemed for CAN	0.06	0.04
	(0.24)	(0.21)
Quantities purchased using coupons		
Quantity purchased (unconditional)	1.62	1.10
	(4.86)	(4.37)
Quantity purchased (conditional on redeeming)	9.00	9.19
	(8.08)	(9.27)
Quantity DAP purchased (conditional on redeeming)	7.40	7.23
	(7.77)	(8.40)
Quantity CAN purchase (conditional on redeeming)	1.60	1.96
	(3.93)	(4.47)
Total number of coupons	9,505	7,902

Notes: This table shows summary statistics for fertilizer coupon redemption in the two seasons of the program. Standard errors are in parentheses.

1. Standard errors are in parentheses, clustered by school. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

Table 2.4: *Fertilizer Coupon Redemption: Effect of Bluespoons and Text Message Reminders*

	Redeemed coupon					Quantities redeemed				
VARIABLES	Short rains 2010		Long rains 2011			Short rains 2010		Long rains 2011		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Current Season Treatments										
Cooperative School	0.01	0.01	-0.01	-0.02	-0.02	0.11	0.12	0.06	0.04	0.03
	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.35)	(0.35)	(0.26)	(0.24)	(0.24)
SMS school	-0.02	-0.01	-0.00	-0.00	-0.00	0.10	0.11	-0.11	-0.07	-0.07
	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.32)	(0.32)	(0.23)	(0.20)	(0.20)
Individual sampled to receive SMS	0.03**	0.03**	0.06***	0.06***	0.06***	0.14	0.14	1.30***	1.12***	1.12***
	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.21)	(0.21)	(0.39)	(0.36)	(0.36)
Individual sampled to received Bluespoon	0.04***		0.03*	0.02	0.02	0.37**		-0.06	-0.00	-0.00
	(0.01)		(0.02)	(0.01)	(0.01)	(0.15)		(0.13)	(0.11)	(0.11)
Bluespoon group 1 (1 week after school meeting)		0.10***					0.80**			
		(0.03)					(0.35)			
Bluespoon group 2 (3 to 4 weeks after school meeting)		0.03					0.30			
		(0.04)					(0.42)			
Bluespoon group 3 (6 weeks after school meeting)		0.01					0.05			
		(0.02)					(0.26)			
Previous Season Treatments (Season 2 only)										
SMS school				0.02	0.02				0.20	0.22
				(0.02)	(0.02)				(0.31)	(0.31)
Individual sampled to receive SMS				0.03*	0.03*				0.45	0.45
				(0.02)	(0.02)				(0.29)	(0.29)
Individual sampled to receive a Bluespoon				-0.01					-0.32*	
				(0.01)					(0.16)	
Bluespoon group 1 (previous season)					-0.03					-0.71***
					(0.03)					(0.25)
Bluespoon group 2 (previous season)					-0.02					-0.27
					(0.02)					(0.28)
Bluespoon group 3 (previous season)					-0.00					-0.04
					(0.02)					(0.37)
Observations	9,505	9,505	7,902	7,902	7,902	9,505	9,505	7,902	7,902	7,902
Adjusted R-squared	0.008	0.010	0.009	0.015	0.014	0.010	0.010	0.010	0.014	0.014
Control mean	0.169	0.169	0.119	0.119	0.119	1.432	1.432	0.894	0.894	0.894

Notes: This table shows the impact of the different treatment conditions on coupon redemption. Standard errors in parentheses, clustered by school. Standard errors are in parentheses, clustered by school. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

the individual level – statistical power to detect effects at the school level was low.

Second, receiving a text message increased the probability of a farmer redeeming her coupon by approximately 3 and 6 percentage points in the two seasons at the individual level, respectively (columns 1 through 5). The stronger individual-level effect in season 2 may be due to differences in the nature of the text message intervention. In season 1, we sent text messages to everyone who provided us with a phone number, while in season 2 we sent messages only to respondents who gave us a phone number *and* who indicated that a household member was the owner of the phone (information that we did not collect in season 1). Moreover, we find a small positive effect of text messages sent in the first season on coupon redemption in the schools visited in both seasons (columns 4 and 5). While the estimated effects of text message reminders are modest in absolute terms, they are large in relative terms (18 percent and 50 percent, respectively), and, conditional on having cellphone numbers available, the costs of sending mass text messages are extremely low.

Third, being sampled to receive a Bluespoon increased fertilizer coupon redemption by 2 to 4 percentage points (columns 1 through 5). As described above, in the first season the timing of the Bluespoon treatment was randomized, such it could have only possibly impacted coupon redemption for group 1 and possibly group 2, and there could not have been a causal effect on coupon redemption for group 3. This is what we find: Being randomized to group 1 increased coupon redemption by as much as 10 percentage points, while there is only a small, but insignificant effect for group 2, and no effect for group 3. Furthermore, we do not find any evidence for a persistent effect of the Bluespoon intervention. Respondents who had received a Bluespoon in the previous season were in fact somewhat *less* likely to redeem their coupon. This appears to be particularly the case for individuals in Bluespoon group 1. One interpretation of this is that the Bluespoon intervention encouraged respondents to experiment with fertilizer and the Bluespoon, but they got disappointed by the results.

2.3.2 Self-reported Fertilizer Use

Table 2.5 examines the impact of the different treatment conditions on self-reported fertilizer use, based on survey answers in the short school endline and longer endline surveys at respondents' homes. Since some schools participated in two seasons of the program, we stack observations in these regressions and control for the season, again clustering standard errors by school. Results are presented in the season of the program (columns 1 and 2) and the season after the program (columns 3 and 4). The table shows three sets of results. First, we find evidence that the coupon increased overall fertilizer usage (columns 1 and 2). However, the estimated effect varies substantially across types of survey, i.e. we estimate a larger effect in the school survey (13 percentage points; statistically significant at the one-percent level) compared to the home survey (6 percentage points; not statistically significant), a discrepancy that may be due to social desirability bias. We also find suggestive evidence of an effect in the season after the treatment (columns 3 and 4). Second, we find no impact of the cooperative or text message treatments on fertilizer use. In contrast to the results on coupon redemption, neither the school-level nor the individual-level text message treatment affected self-reported fertilizer use. Given the effect of text messages on coupon redemption at the individual level, the latter suggests that the text message reminders increased coupon redemption among a subsample of individuals who would have used fertilizer already anyway, rather than encouraging some individuals to use fertilizer who would not have done so anyway. Third, similarly to the above results on coupon redemption, we find that the Bluespoon treatment significantly increased fertilizer use, both in the same and in the subsequent season of the program.

2.4 Bluespoon and Knowledge Diffusion

The section considers the impact of the different treatment conditions on the diffusion of Bluespoons and knowledge across social networks.

Table 2.5: *Self-reported Fertilizer Use*

	Season of Program				Season after Program			
	School Survey		Home Survey		School Survey		Home Survey	
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Coupon School	0.16***	0.13***	0.10	0.06	0.04	0.03	0.09	0.05
	(0.05)	(0.03)	(0.06)	(0.05)	(0.04)	(0.03)	(0.06)	(0.05)
Cooperative School	0.01	0.03	-0.00	0.01	-0.04	-0.03	0.02	0.03
	(0.05)	(0.04)	(0.05)	(0.04)	(0.04)	(0.04)	(0.05)	(0.04)
Coupon Coop Interaction	-0.03	-0.05	-0.05	-0.05	0.01	0.00	-0.02	-0.01
	(0.06)	(0.04)	(0.07)	(0.05)	(0.05)	(0.05)	(0.07)	(0.05)
Individual sampled to receive SMS	0.00	-0.00	0.06	0.04	-0.03	-0.03	0.01	0.01
	(0.02)	(0.01)	(0.04)	(0.04)	(0.02)	(0.02)	(0.04)	(0.04)
School sampled to receive SMS	-0.02	0.01	-0.06	-0.01	0.04	0.04	-0.05	-0.02
	(0.04)	(0.02)	(0.05)	(0.04)	(0.03)	(0.03)	(0.05)	(0.04)
Individual selected to receive a Bluespoon	0.05***	0.06***	0.06***	0.05***	0.05***	0.05***	0.03	0.03
	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
R-squared	0.036	0.177	0.029	0.142	0.033	0.094	0.028	0.167
Control mean	0.606	0.606	0.478	0.478	0.703	0.703	0.547	0.547
# of schools	183	183	184	184	60	60	184	184
# of farmers	21939	21939	2509	2509	8250	8250	2622	2622
# of observations	27351	27351	3211	3211	8250	8250	2622	2622
Baseline survey controls	No	Yes	No	Yes	No	Yes	No	Yes

Notes: This table shows the impact of the different treatment conditions on self-reported fertilizer use.

1. Standard errors are in parentheses, clustered by school. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

2.4.1 Group Membership

Table 2.6 shows the effect of the different treatment conditions (and in particular the cooperative treatment) on group membership, as reported during the short endline surveys (columns 1 and 2) and the long endline meetings (columns 3 through 5). Even though the only assistance provided was coordination on forming a group and on organizing the first meetings (as well as the provision of sodas at the first meeting), the cooperative treatment had a strong impact on reported group membership overall (column 1), and on membership of groups that discuss agriculture (column 2).

In the school endline survey, 49% of farmers in control schools report participating in any group (cooperative, ROSCA, self-help group, or any other social group in which members meet regularly). This percentage increases by 25 percentage points in cooperative schools (column 1). The effect is slightly larger when examining cooperatives which discuss agriculture: farmers in cooperative schools are 30 percentage points more likely to participate in such a group, on a base of 33% in the control schools. These effects are attenuated but still sizable in the home survey (columns 3 and 4). While the baseline means are similar (at 59% and 39%, respectively) the estimated effect of the cooperative treatment is 11 to 13 percentage points (though still significant at the 1% level).[6]

2.4.2 Bluespoon Diffusion

Table 2.8 considers the diffusion of Bluespoons among original respondents, i.e. individuals who attended the baseline school meetings. Diffusion was substantial even among individuals for whom we did *not* collect social network information and who, therefore, also did not receive a Bluespoon: in control schools, 42% of these individuals had heard of the Bluespoon (column 1), 24% owned at least one Bluespoon, and 28% were able to name its correct price (Ksh 5). As expected, diffusion is near complete among individual who were

[6]A possible explanation for the lower estimates in home survey is timing: home surveys were conducted approximately half a year after school surveys.

Table 2.6: *Group Membership (Among Original Respondents)*

	School Survey		Home Survey		
	Respondent reported being a member of ...				
VARIABLES	any group (1)	a group that discussed agric. (2)	any group (3)	a group that discussed agric. (4)	a group that discussed fert. (5)
Coupon School	-0.00	0.00	0.00	-0.03	-0.01
	(0.018)	(0.018)	(0.034)	(0.028)	(0.022)
Cooperative School	0.25***	0.30***	0.11***	0.13***	0.06**
	(0.022)	(0.027)	(0.033)	(0.030)	(0.026)
Coupon Coop Interaction	-0.05	-0.05	0.01	0.00	0.01
	(0.031)	(0.035)	(0.044)	(0.040)	(0.034)
Individual selected to receive a Bluespoon	-0.01	-0.02	0.01	0.02	0.02
	(0.014)	(0.014)	(0.018)	(0.019)	(0.016)
Observations	21,872	21,872	2,733	2,733	2,733
R-squared	0.060	0.081	0.019	0.031	0.028
Control mean	0.486	0.327	0.592	0.390	0.179

Notes: This table shows the impact of the different treatment conditions on group membership among individuals who had participated in the baseline school meetings.

1. An individual is a member of a group if she “is a member of a cooperative, ROSCA, self-help group, or any other social group in which members meet regularly."
2. Standard error are in parentheses, clustered by school. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

selected for the Bluespoon treatment. The cooperative treatment increased the diffusion of Bluespoons, as measured by the school endline surveys. Individuals in cooperative schools were 5 percentage points more likely to have heard of the Bluespoon, and 7 percentage points more likely to own a Bluespoon and to know the price of a Bluespoon. We also find suggestive evidence of a treatment effect of the cooperative treatment on Bluespoon diffusion in the longer home surveys, but the estimated coefficients are not statistically significant.

Table 2.9 considers Bluespoon diffusion among friends of original respondents, based on the home endline surveys. As expected, knowledge and ownership of Bluespoons is higher among individuals whose friends were selected to receive a Bluespoon. In addition, among individuals who were not friends of farmers who had received a Bluespoon, knowledge and ownership were higher in cooperative, but not in coupon schools. Furthermore, being a friend of a farmer who had received a Bluespoon appears to matter more in coupon schools (column 1), but this is only the case for knowledge of Bluespoons (column 1), not for ownership (column 2). Columns 3 through 9 of the table show essentially the same information as the first two columns.

2.4.3 Knowledge Diffusion

Tables 2.7 and 2.10 consider knowledge diffusion among original respondents and friends, respectively. The outcome variable of interest in these tables is respondents' answers to the question, "How much fertilizer they thought should be used per planting hole?" This question was asked in two ways. First, we asked respondents to choose between five different verbal options, one of which was 'one half teaspoon'.[7] Second, we showed respondents four different containers with fertilizer, and asked them to choose the container which most closely represented the quantity of fertilizer per planting hole which they thought was

[7]The remaining answers were 'one quarter teaspoon', 'one teaspoon', 'one tablespoon', and 'don't know'.

Table 2.7: *Knowledge Diffusion Among Original Respondents (School Survey)*

	Quantity per planting hole (verbal)						Quantity per planting hole (visual)					
	CAN			DAP			CAN			DAP		
VARIABLES	1/4 Tsp	1/2 Tsp	$\geq$ 1 Tsp	1/4 Tsp	1/2 Tsp	$\geq$ 1 Tsp	1/4 Tsp	1/2 Tsp	$\geq$ 1 Tsp	1/4 Tsp	1/2 Tsp	$\geq$ 1 Tsp
Coupon School	0.02***	0.01	-0.04*	0.01	0.03*	-0.04*	-0.00	0.02*	-0.02	-0.00	0.03**	-0.02
	(0.007)	(0.018)	(0.021)	(0.006)	(0.017)	(0.019)	(0.017)	(0.012)	(0.020)	(0.017)	(0.013)	(0.019)
Cooperative School	0.01*	0.01	-0.02	0.00	0.01	-0.02	0.01	0.01	-0.02	0.01	0.01	-0.02
	(0.006)	(0.021)	(0.023)	(0.005)	(0.019)	(0.020)	(0.016)	(0.013)	(0.020)	(0.017)	(0.014)	(0.018)
Coupon Coop Interaction	-0.02**	-0.02	0.04	-0.00	-0.03	0.04	-0.01	-0.02	0.03	-0.01	-0.02	0.03
	(0.010)	(0.026)	(0.031)	(0.008)	(0.024)	(0.027)	(0.023)	(0.018)	(0.026)	(0.024)	(0.019)	(0.025)
Individual selected to receive a Bluespoon	0.02***	0.17***	-0.19***	0.02***	0.17***	-0.20***	0.06***	0.05***	-0.11***	0.07***	0.05***	-0.12***
	(0.007)	(0.014)	(0.014)	(0.006)	(0.015)	(0.015)	(0.013)	(0.013)	(0.014)	(0.012)	(0.012)	(0.013)
Observations	22,230	22,230	22,230	22,230	22,230	22,230	22,230	22,230	22,230	22,230	22,230	22,230
R-squared	0.002	0.042	0.040	0.001	0.046	0.048	0.006	0.002	0.012	0.009	0.002	0.013
Control mean	0.0508	0.468	0.481	0.0508	0.426	0.523	0.225	0.316	0.458	0.225	0.317	0.459

Notes: This table shows the impact of the different treatment conditions on knowledge diffusion, as measured by respondents' answers to the question "How much fertilizer [CAN/CAN] should be used per planting hole?" during short endline surveys.

1. Columns 1 through 6 use respondents' verbal answers. Columns 7 through 12 use respondents' answers to the visual equivalent of the question, in which we showed respondents four containers with different quantities of fertilizer.
2. Columns 1 through 3 and 7 through 9 show answers to questions about CAN fertilizer, the remaining columns show questions about DAP fertilizer.
3. '1/4 Tsp' corresponds to the answer 'one-quarter teaspoon' (only), '1/2 Tsp' to 'one-half teaspoon' (only), and '$\geq$ 1 Tsp' to the remaining answers (i.e. 'one teaspoon', 'one tablespoon', or 'don't know').
4. Standard errors are in parentheses, clustered by school. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

Table 2.8: *Bluespoon Diffusion (Among Original Respondents)*

	School Survey			Home Survey	
VARIABLES	Heard of BS (1)	Owns BS (2)	Knows price (3)	Heard of BS (4)	Owns BS (5)
Coupon School	0.02	0.04	0.00	-0.06**	-0.03
	(0.027)	(0.024)	(0.028)	(0.024)	(0.023)
Cooperative School	0.05**	0.07***	0.07***	0.03	0.02
	(0.022)	(0.020)	(0.022)	(0.020)	(0.022)
Coupon Coop Interaction	0.01	-0.03	-0.03	0.04	0.01
	(0.030)	(0.028)	(0.032)	(0.030)	(0.029)
School sampled to receive SMS	-0.03	-0.03	-0.02	0.00	-0.00
	(0.022)	(0.021)	(0.024)	(0.023)	(0.021)
Individual sampled to receive SMS	0.01	0.02*	0.03***	-0.03	-0.00
	(0.014)	(0.011)	(0.012)	(0.033)	(0.031)
Individual selected to receive a Bluespoon	0.41***	0.53***	0.45***	0.31***	0.66***
	(0.014)	(0.014)	(0.014)	(0.016)	(0.016)
Observations	22,230	22,230	22,230	2,733	2,733
R-squared	0.118	0.199	0.143	0.168	0.455
Control mean	0.418	0.244	0.279	0.670	0.235

Notes: This table shows the impact of the different treatment conditions on Bluespoon diffusion among original respondents, based on home surveys.

1. 'Heard of BS' indicates whether the respondent had heard of the Bluespoon at the time of the survey. 'Owns BS' indicates whether the respondent owned a Bluespoon at the time survey. 'Knows price' indicates whether the respondent named the correct price of a Bluespoon (Ksh 5).
2. Standard errors are in parentheses, clustered by school. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

Table 2.9: *Bluespoon Diffusion (Agricultural Contacts)*

	Type of Schools									
	All		No Coupon, No Coop		Coupon, No Coop		No Coupon, Coop		Coupon, Coop	
VARIABLES	Heard of BS (1)	Owns BS (2)	Heard of BS (3)	Owns BS (4)	Heard of BS (5)	Owns BS (6)	Heard of BS (7)	Owns BS (8)	Heard of BS (9)	Owns BS (10)
Friend of Bluespoon Farmer	0.09***	0.07**	0.10**	0.09***	0.18***	0.07*	0.06	0.03	0.16***	0.09**
	(0.035)	(0.029)	(0.041)	(0.031)	(0.040)	(0.037)	(0.045)	(0.041)	(0.038)	(0.041)
Friend of Coupon Farmer	-0.06*	-0.02								
	(0.031)	(0.029)								
Friend of Cooperative Farmer	0.07**	0.07**								
	(0.031)	(0.029)								
Coupon X Individual selected for BS	0.09**	0.02								
	(0.041)	(0.037)								
Cooperative X Individual selected for BS	-0.03	-0.02								
	(0.041)	(0.037)								
Observations	2,314	2,314	579	579	600	600	573	573	562	562
R-squared	0.022	0.011	0.010	0.011	0.033	0.006	0.004	0.001	0.027	0.011
Control mean	0.624	0.281	0.606	0.247	0.588	0.255	0.654	0.321	0.648	0.302

Notes: This table shows the impact of the different treatment conditions on Bluespoon diffusion among friends of original respondents, based on in-depth home surveys.

1. 'Heard of BS' indicates whether the respondent had heard of the Bluespoon at the time of the endline survey. 'Owns BS' indicates whether the respondent owned a Bluespoon at the time of the endline survey.
2. Standard errors are in parentheses, clustered by school. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

appropriate. Again, one of the options included 'one half teaspoon'.

The Bluespoon treatment significantly shifted respondents' beliefs regarding how much fertilizer to use (Table 2.10). In the verbal question, the fraction of individuals who chose one-half teaspoon as the quantity of fertilizer to be used per planting hole increased by about 17 percentage points (on a base of 47 percent). While the intervention was targeted to CAN knowledge and usage, we find similar results for DAP fertilizer (columns 7 through 12). We find similar results using answers to the longer endline surveys with original respondents (columns 1 through 6 of Table 2.10. In the verbal question and in a choice between four spoons of different sizes, respondents are 23 to 25 percentage points more likely to choose the option indicating one-half teaspoon. Columns 7 through 12 show the equivalent results for agricultural contacts named by the original respondents during the baseline survey. We only find limited evidence of an impact of the different conditions on beliefs regarding optimal fertilizer use.

2.5 Conclusion

Insufficient knowledge of appropriate use may have been a key cause of low agricultural technology adoption and productivity in Sub-Saharan Africa. In particular, if farmers do not observe each others' inputs, diffusion of both information on the optimal input mix and of the technology itself may be slow. This paper reports results from a large-scale field experiment, which introduced a simple and salient tool, a blue measuring spoon, to help farmers remember how much fertilizer to use. The main result of this paper is that farmers who were randomly assigned to receive a spoon subsequently used more fertilizer, a result that suggests that insufficient knowledge is a relevant barrier to fertilizer adoption. Moreover, spoon purchases among the remaining farmers were higher when these were more likely to use fertilizer due to a randomly assigned fertilizer discount program, and when communication about agriculture was encouraged. Unlike fertilizer adoption itself, purchase and use of measuring spoons diffused rapidly through social networks.

Table 2.10: *Knowledge Diffusion among Original Respondents and Agricultural Contacts (Long Endline Survey)*

VARIABLES	Original respondents						Agricultural contacts					
	1/4 Tsp		1/2 Tsp		$\geq$ 1 Tsp		1/4 Tsp		1/2 Tsp		$\geq$ 1 Tsp	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A: Verbal question (CAN)												
Coupon School	-0.00	0.01	0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	0.02	0.03
	(0.010)	(0.012)	(0.022)	(0.025)	(0.023)	(0.027)	(0.009)	(0.013)	(0.023)	(0.029)	(0.025)	(0.031)
Cooperative School	-0.01	-0.01	-0.04*	-0.02	0.05**	0.03	-0.00	-0.00	0.03	0.02	-0.03	-0.02
	(0.010)	(0.012)	(0.022)	(0.025)	(0.023)	(0.027)	(0.009)	(0.013)	(0.023)	(0.029)	(0.025)	(0.031)
Individual selected to receive a Bluespoon	0.02*	0.03**	0.23***	0.24***	-0.24***	-0.27***	0.00	-0.00	0.03	0.00	-0.03	0.00
	(0.009)	(0.016)	(0.016)	(0.024)	(0.016)	(0.024)	(0.008)	(0.014)	(0.017)	(0.029)	(0.017)	(0.030)
Coupon X Individual selected for BS		-0.03*		0.04		0.00		0.00		0.02		-0.03
		(0.017)		(0.032)		(0.033)		(0.017)		(0.034)		(0.034)
Cooperative X Individual selected for BS		-0.01		-0.04		0.05		0.00		0.03		-0.03
		(0.017)		(0.032)		(0.033)		(0.017)		(0.034)		(0.034)
Observations	2,733	2,733	2,733	2,733	2,733	2,733	2,314	2,314	2,314	2,314	2,314	2,314
R-squared	0.002	0.003	0.061	0.062	0.065	0.065	0.000	0.001	0.002	0.003	0.002	0.003
Control mean	0.0506	0.0506	0.226	0.226	0.717	0.717	0.0311	0.0311	0.242	0.242	0.713	0.713
Panel B: Spoon choice (CAN)												
Coupon School	0.01	0.01	-0.02	-0.04	0.01	0.03	-0.01	-0.03	0.01	-0.04	0.01	0.08**
	(0.017)	(0.019)	(0.020)	(0.027)	(0.018)	(0.029)	(0.016)	(0.022)	(0.023)	(0.033)	(0.025)	(0.034)
Cooperative School	-0.00	-0.01	-0.03	0.01	0.03*	0.00	-0.01	0.00	0.04	0.03	-0.03	-0.04
	(0.017)	(0.019)	(0.020)	(0.027)	(0.018)	(0.029)	(0.015)	(0.022)	(0.023)	(0.033)	(0.025)	(0.033)
Individual selected to receive a Bluespoon	0.03**	0.02	0.24***	0.25***	-0.27***	-0.27***	0.00	-0.00	0.04*	-0.01	-0.04*	0.02
	(0.012)	(0.020)	(0.019)	(0.032)	(0.019)	(0.031)	(0.014)	(0.026)	(0.023)	(0.039)	(0.022)	(0.038)
Coupon X Individual selected for BS		-0.00		0.05		-0.05		0.04		0.09**		-0.13***
		(0.024)		(0.038)		(0.038)		(0.028)		(0.046)		(0.043)
Cooperative X Individual selected for BS		0.02		-0.07*		0.06		-0.02		0.01		0.01
		(0.024)		(0.038)		(0.038)		(0.028)		(0.046)		(0.043)
Observations	2,733	2,733	2,733	2,733	2,733	2,733	2,314	2,314	2,314	2,314	2,314	2,314
R-squared	0.002	0.002	0.059	0.061	0.078	0.079	0.001	0.002	0.003	0.005	0.003	0.007
Control mean	0.122	0.122	0.357	0.357	0.506	0.506	0.121	0.121	0.384	0.384	0.471	0.471

Notes: This table shows the impact of the different treatment conditions on knowledge diffusion. as measured by respondents' answers to the question "How much CAN fertilizer should be used per planting hole?" during short endline surveys.

1. '1/4 Tsp' corresponds to the answer 'one-quarter teaspoon' (only), '1/2 Tsp' to 'one-half teaspoon' (only), and '$\geq$ 1 Tsp' to the remaining answers (i.e. 'one teaspoon', 'one tablespoon', or 'don't know').
2. Panel A shows the answer to a verbal question. Panel B shows the answer to a question that asks respondents to pick between four different spoons.
3. Columns 1 through 6 show regressions with answers from original respondents. Columns 7 through 12 show regressions with answers from original respondents' agricultural contacts. Regressors in columns 7 through 12 all correspond to the respective original respondents.
4. Standard errors are in parentheses, clustered by school. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

Chapter 3

Trading Decisions as a Function of Stock Return Quantiles: Implications for the Disposition Effect[1]

3.1 Introduction

A long-standing puzzle in finance research is investors' tendency to hold losing investments too long and to sell winning investments too soon, a phenomenon that has been called the "disposition effect" by Shefrin and Statman (1985). Using individual-level trading data, Odean (1998) rules out various competing explanations. The disposition effect is persistent, even when controlling for higher transaction costs of selling losers or tax incentives.[2] Other explanations, such as re-diversification or private information, similarly fail to capture important features of the data.

Several theories have been proposed to explain the disposition effect. The most prominent explanation invokes prospect theory, developed by Kahneman and Tversky (1979). Odean (1998) argues that with reference-dependent preferences, investors are risk-averse over gains

[1]Co-Authored with Tom Zimmermann

[2]Ivkovic *et al.* (2005) show that the disposition effect interferes with a "lock-in effect" for capital gains.

and risk-seeking over losses, such that they sell winning stocks prematurely and gamble on stocks that lost value in the past.[3] Barberis and Xiong (2009) challenge this informal argument by presenting a stylized asset pricing model involving investors with prospect theory preferences.[4] This model does *not* predict a disposition effect for many reasonable parameter values, since loss aversion pushes investors to only purchase stocks with high enough expected returns. In particular, if an investor is in the gain region, she will in expectation be further away from the reference point than when she is in the loss region, such that a disposition effect may not arise for mild curvature of the value function. In a related paper, Barberis and Xiong (2012) introduce the concept of realization utility and apply it to the disposition effect. In their model, investors derive utility from realizing gains and losses, implying that an investor only sells a stock once the return exceeds a certain, potentially investor-specific, threshold.

The main contribution of this paper is to contrast these competing explanations based on their predictions for realizing different sizes of gains and losses. We use the Barber and Odean (2000) data to establish two empirical facts. First, for all holding periods longer than one month and for both gains and losses, the probability to sell a stock declines monotonically with the size of the absolute return. That is, individual investors are not only more likely to sell gains than to sell losses, but they are also more likely to sell stocks with small absolute returns (i.e. stocks with prices close to the purchase price) than to sell stocks with large absolute returns (for given holding periods). Second, the disposition effect is more pronounced for relatively high returns, i.e. the difference in selling probability between gains and losses of similar magnitude is more pronounced for large returns (i.e. comparing large gains to large losses) than for small returns (i.e. comparing small gains to small losses). Theories that attempt to explain the disposition effect also need to be consistent with these

[3] Odean (1998) notes that an irrational belief in mean reversion of stock returns could also explain the disposition effect, but he is not able to separate the two hypotheses. He speculates that investors themselves might not make a clear distinction: "For example, an investor who will not sell a stock for a loss might convince himself that the stock is likely to bounce back rather than admit his unwillingness to accept a loss." Weber and Camerer (1998) control for individuals' beliefs in an experimental setting and find the disposition effect as well.

[4] Barberis and Xiong (2009) do not include probability weighting.

additional facts. In particular, we argue that our first fact is not consistent with the model of realization utility as proposed by Barberis and Xiong (2012), but it is consistent with a version of prospect theory that we outline below.

Our empirical analysis employs two methodologies. First, in our preferred approach, we follow Ivkovic *et al.* (2005) and construct portfolios of investors' stock holdings for each month. This method allows us to measure the survival time of stocks in investors' portfolios, conditional on the stock being in the gain or in the loss region. We extend this approach by splitting up the sample into gain and loss quantiles, and then consider the selling patterns across quantiles, controlling for stock holding periods, which generates the results described above. Second, following the original estimation approach of Odean (1998), we construct an investor's portfolio for every day at which an investor made at least one trade. Replicating Odean's results, we compute the fraction of stocks valued at a gain (loss) that were actually sold relative to all stocks that could have been sold at a gain (loss). Extending this approach to stock returns of different magnitudes, we find that the disposition effect persists for all return sizes, again controlling for the holding period. However, in contrast to the results from our first approach, the selling probability *increases* with the size of the return for gains and is constant in returns for losses.

We trace the apparent tension between results from the two approaches to different conditioning sets of the estimates. The duration model of Ivkovic *et al.* (2005) computes an *unconditional* probability of selling (for given holding periods), while the Odean (1998) methology estimates a probability of selling a stock *conditional on investor activity*. To reconcile the results from the two estimation approaches, we establish that an investor's propensity to make a trade is largest for small absolute portfolio returns, using several parametric and non-parametric estimators.

Our paper is closely related to a recent contribution by Ben-David and Hirshleifer (2012), who investigate the relationship between past security returns and subsequent sales using an approach similar to ours. Ben-David and Hirshleifer find that the probability to sell a security is "asymmetrically V-shaped". That is, larger returns are more likely to be sold, this effect is more pronounced for positive returns compared to negative returns, and the selling probability does *not* have a discontinuity at a stock return of 0. Since these results appear to be in contrast to our findings, we seek to reconcile these findings with our results. We document that the probability to sell is asymmetrically V-shaped only for short holding periods. In particular, if results are pooled over different horizons, we do find a pronounced discontinuity at 0 returns. We also find that the selling probability decreases in the absolute value of the return, in line with our previous findings.

The remainder of the paper is organized as follows. Section 3.2 derives simple theoretical predictions for prospect theory and realization utility. Section 3.3 describes the data and our methodological approach. In Section 3.4 we present our main empirical results and robustness checks, and in Section 3.5 we relate our results to Ben-David and Hirshleifer (2012). Section 3.6 concludes.

3.2 Theoretical Background

In this section, we provide a sketch of models of realization utility and of prospect theory that have been suggested as explanations for the disposition effect. We are particularly interested in the predictions of these models for the probability to sell a stock for different returns, and we show that the models we consider generate quite different predictions in this regard. This allows us to disentangle the explanations by comparing their predictions to the actual probability to sell a stock for different returns.

3.2.1 Realization Utility

Barberis and Xiong (2012) argue that investors derive utility from realizing gains and losses of assets that they own (as opposed to from consumption of the proceeds). The authors set up a dynamic optimization problem and show that investors with realization utility will sell a stock when the return exceeds some (positive) liquidation point. In other words, "if the investor buys a stock, he voluntarily sells it only if its price rises a sufficient amount above the purchase price". In particular, an investor only sells at a loss if he is forced to do so by a liquidity shock.

For illustrative purposes denote by P_0 the investor's purchase price and by P_t the price in period t. Define $g_t := P_t/P_0$. An investor will then sell the stock if $g_t \geq g^i_*$, where $g^i_* \geq 1$ is the investor's liquidation point which depends on individual and stock characteristics.[5] Figure 3.1 shows the trading behavior of the individual investor depending on g_t. The probability to sell a stock below g_t is drawn to be greater than zero to take liquidity shocks into account. If g_t exceeds the investor's liquidation point, the stock is sold in t. While realization utility predicts a step function for the individual's probability to sell a stock, the function is smoothed out when we consider the aggregated prediction. The liquidation point depends on individual characteristics, for instance, on the time discount rate, transaction cost, and the likelihood of a liquidity shock. If these factors vary across individuals, there there will be heterogeneity of liquidation points across investors.

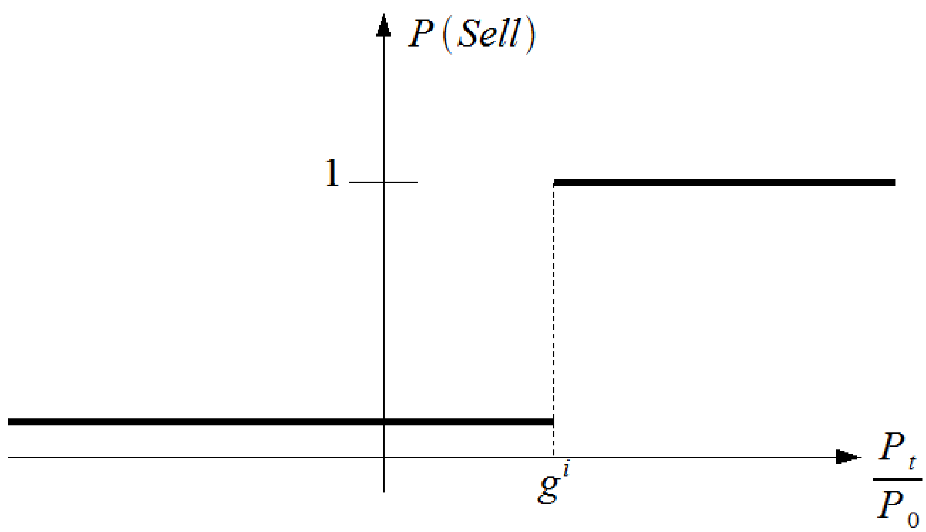

Figure 3.1: *Realization utility prediction for individual trading*

Figure 3.2 illustrates the implications for the aggregate probability that a stock is sold. As investors do not realize losses, the probability to sell is constant in the loss region

[5] g^i_* is strictly greater than 1 if the investor faces some transaction cost.

(liquidity shocks). When the return is slightly positive, some investors start selling the asset because they have a relatively low liquidation point. A higher return implies that additional liquidation points are exceeded. The exact shape of the function obviously depends on the distribution of liquidation points among investors. Of course, there is no reason to expect this relationship to be linear, but the probability to sell should be an increasing function of the realized return.[6]

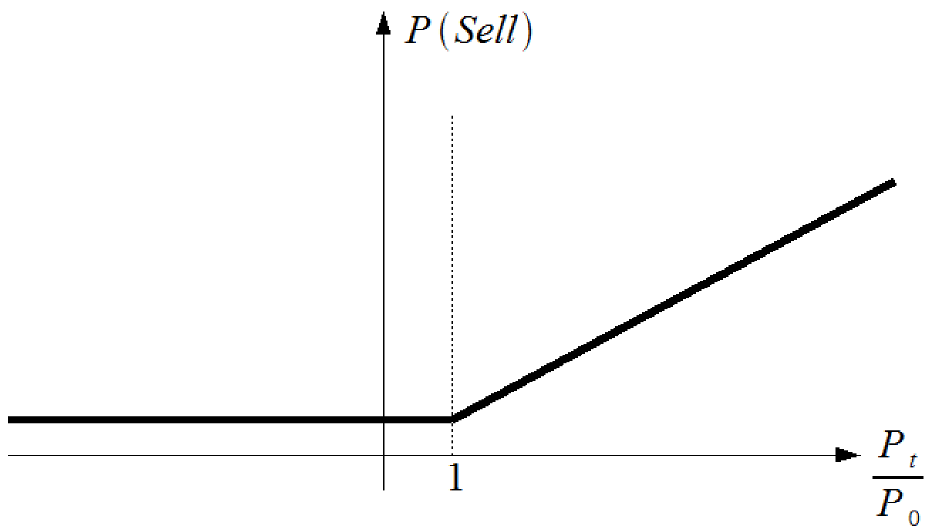

Figure 3.2: *Implied trading pattern in the realization utility model*

3.2.2 Prospect Theory

In their pioneering work, Shefrin and Statman (1985) and Odean (1998) linked prospect theory to the disposition effect using informal arguments. In a formal analysis, however, Barberis and Xiong (2009) find that a model including investors with prospect theory preferences (without probability weighting) does not generate the disposition effect for many combinations of parameter values.[7] Prior arguments had neglected the effect of the

[6] For instance, a positively skewed distribution (many observations close to 1) would yield a concave function.

[7] This has also been studied by Hens and Vlcek (2011).

kink on the initial buying decision. Because of the kink, an investor will only buy a stock in the first place if the initial expected return is high enough. This in turn implies that, in the next period, the investor is relatively far away from the reference point when the stock trades in the gain region, whereas she is closer to the reference point when the stock is in the loss domain. When the value function has only mild curvature, the investor is almost risk-neutral in the gain region when she is relatively far away from the reference point, and she will therefore hold the stock after a gain. On the other hand, after a loss, she is still close to the reference point and sells the stock for many parameter values of the value function. Meng (2013) argues that a modified reference point can help to get around this result. Using expectations as a reference point, she proposes a simple model in which prospect theory generates a disposition effect for investors.

We present two views of prospect theory, both of which are able to generate the disposition effect. First, we consider Barberis' (2012) implementation of prospect theory in a model of casino gambling. This model, very different from most other models in this area of research, takes nonlinear probability weighting into account. It features some interesting predictions that we think are transferable to a model of individual investment decisions.

Consider an asset that in every period with probability 0.5 either increases or decreases by h. Figure 3.3 shows the possible prices after six periods. Barberis' (2012) model implies that even though this stock has an expected return of 0, a prospect theory agent might buy it because she can give its return a favorably skewed distribution by overweighting small probabilities and by choosing a suitable ex-ante exit strategy (e.g. sell the stock as soon as you acquire losses). In stark contrast, Barberis and Xiong's (2012) model implies that prospect theory agents only buy stocks with high expected return. Barberis (2012) then investigates implications for subsequent gambling behavior. He finds that naïve prospect theory agents (that is, those who are not aware of their nonlinear probability weighting) almost never exit after making losses and stop gambling too early after making gains; in other words, they do not stick to their original plans. The actual exit pattern is illustrated by the curved line in Figure 3.3. Intuitively, the pattern comes from a trade-off between the nonlinear probability weighting (which pulls towards keeping on gambling, but less so if the gain is already large) and the concavity of the utility function in the gain region.

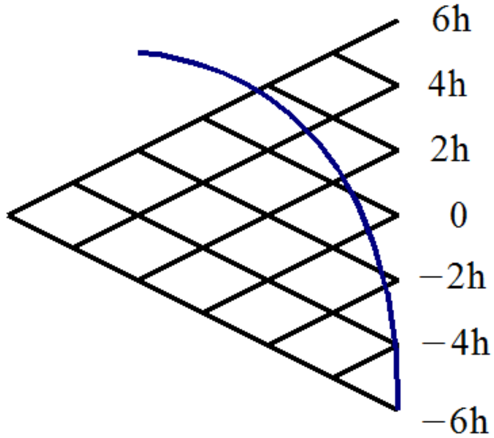

Figure 3.3: *Stock price for six periods and resulting exit behavior*

Carrying this result over to the stock market, it implies that higher positive returns are more likely to be realized. Figure 3.4 shows the trading pattern that can be derived from the model. In particular, investors never realize losses under this specification, but are increasingly willing to realize gains.

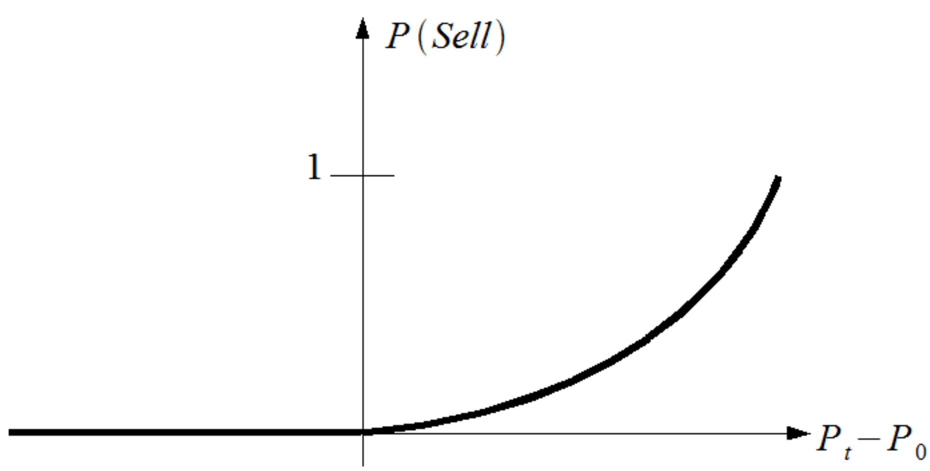

Figure 3.4: *Implied trading pattern in the casino gambling model*

A different, quite natural prediction arising from prospect theory focuses on the kink in the value function. As is well-known (see e.g. Barberis *et al.* (2006)), the kink induces first-order risk aversion around the reference point (Figure 3.5 illustrates this). Even though investors are risk-seeking in the loss domain, the kink induces them to reject a fair bet involving small amounts if the price is close to the reference price.[8] Being far away from the reference return, on the other hand, implies only mild curvature of the utility function (i.e. mild concavity/convexity) and, therefore, very similar behavior in these regions.

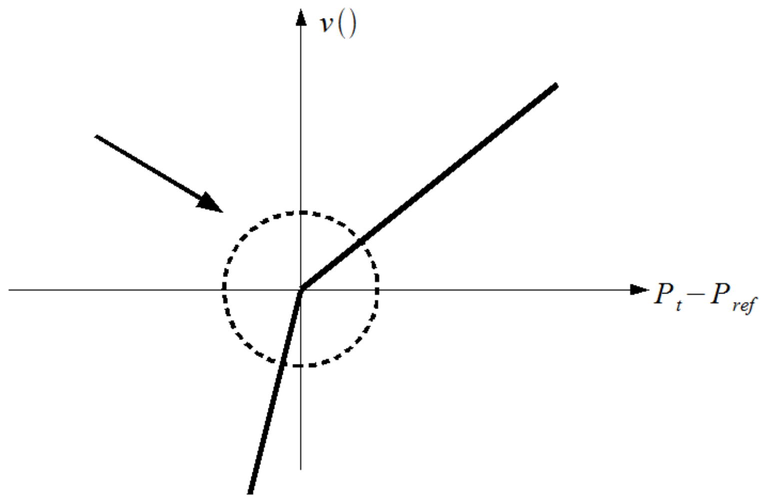

Figure 3.5: *Local risk aversion in the value function*

This gives rise to the prediction that the probability to sell will be higher around the reference point, a phenomenon that we label bunching, and lower for larger gains and larger losses. If, in addition, we assume that some degree of concavity/convexity away from the reference point is preserved, we can also conclude that the probability to sell gains is generally higher than the probability to sell losses (risk aversion vs. risk loving). This is summarized in Figure 3.6. Note that the plot is qualitatively similar to Kaustia (2010) or Meng (2013) who develop models along these lines.

[8]Note that a fundamental problem of this explanation is the question of why they would buy stocks in the first place.

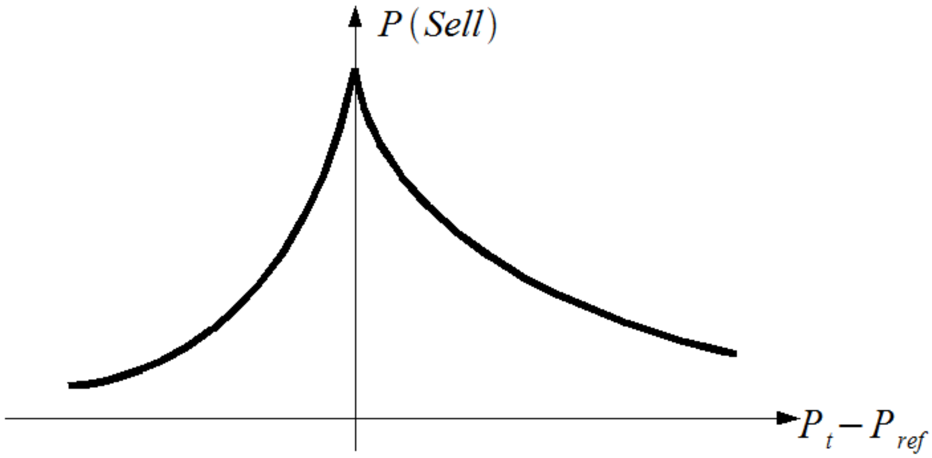

Figure 3.6: *Implied trading pattern in the bunching model*

3.3 Data and Methodology

In this section, we describe the data that we use for our analysis. We then describe two different approaches to estimation of the disposition effect that have been used in the related literature. Section 3.3.2 discusses the duration model approach of Ivkovic *et al.* (2005) and how it can be applied in our context. Section 3.3.3 describes the original approach of Odean (1998). We apply both approaches and reconcile their results later in section 3.4.4.

3.3.1 Data

Terrance Odean kindly provided us with the dataset used in Barber and Odean (2000), which is very similar to the one used in Odean (1998). This dataset from a large discount brokerage house contains all trades as well as end-of-the-month positions for 158,034 US accounts (belonging to 78,000 households) for the time period 1991-1996. Among other variables, the data comprises household and account identifiers, dates, selling and purchase prices, quantities and security identifiers.[9] The data also feature some demographic information that we use to control for investor characteristics in robustness checks of our analysis below.

[9] A detailed description of the data can be found in Barber and Odean (2000).

Table 3.1 shows the the main demographic variables used in our analysis. Our data contain relatively few female individuals, more than 50% of the individuals say they have some knowledge about the stock market (self-reported), and 15% of the households are labeled as frequent traders by the brokerage house (i.e. they trade more than 48 times per year).

Table 3.1: *Demographic characteristics of investors*

Variable	Obs	Mean	Std	Min	Max
Married	37642	0.78	0.41	0	1
Female	47586	0.13	0.34	0	1
Age	41654	50.63	12.76	18	94
Home owner	54914	0.77	0.42	0	1
Knowledge	27179	0.56	0.5	0	1
Equity ('000s)	77981	53	277	-.97	51900
Frequent trader	77984	0.15	0.36	0	1
Taxable account	77984	0.63	0.48	0	1

We match the trades file with specific information for each account and trade (e.g. account type, trading activity, product type). We get monthly and daily price data of securities from the Center for Research in Security Prices (CRSP). Prior to our analysis, we eliminate all trades other than trades of common stocks (e.g. foreign stocks), all trades that involve short-selling and all trades including securities purchased before 1991. Also, we drop all observations for which price data are not available and accounts that own only one stock. This procedure closely follows Odean (1998). To get a sense of the data, and because it is going to be important for our subsequent analysis, Figure 3.7 shows the histogram of security holding durations in the sample. As one would expect, a lot of stocks are being held for a short period of time, and the number of observations declines monotonically in the holding duration.

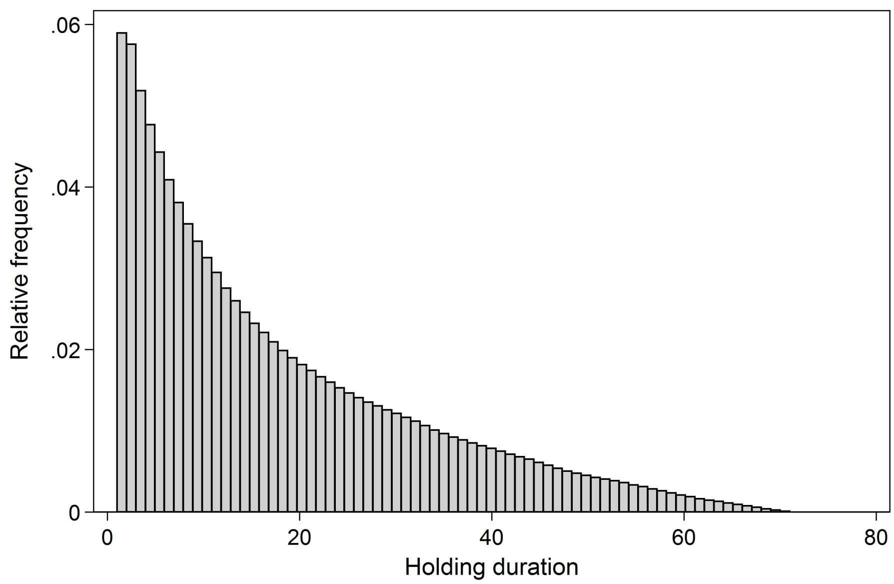

Figure 3.7: *Histogram of stock holding durations*

3.3.2 Duration Model

Ivkovic *et al.* (2005) expand the same dataset that we use by including every month between an initial purchase of a stock and the first sale of the stock (or the end of the observation period, if the security had not been sold by then). For every month they match price data to their dataset and determine whether the stock was sold/could have been sold for a gain or a loss by comparing the buying price to the daily price, just like in Odean (1998). For a given duration they calculate the hazard rate for both gains and losses (i.e. the share of gains/losses that are sold). The resulting figure plots the frequency of selling (and leaving the sample) given that an investor has not sold the security yet, conditional on whether it is a gain or a loss (relative to the buying price), an estimator that is commonly referred to as the nonparametric Kaplan-Meier estimator of the hazard rate.[10]

We take this analysis further and, for every holding period, compute the hazard rate as a function of return magnitude. For illustrative purposes, we show a graph with 3 quantiles of returns for all holding periods (small, medium and large) below, but due to the large sample we can focus on a much finer grid in our analysis when we keep the holding period fixed.

As in Ivkovic *et al.* (2005), we then turn to the estimation of a Cox-proportional hazard model of the form

$$\lambda(t, x_i) = \lambda_0(t) exp(x_i'\beta), \tag{3.1}$$

where $\lambda_0(t)$ is the baseline hazard rate and x_i contains individual-specific information which will allow us to control for various stock-holder (and stock) characteristics, such as those presented in Table 3.1.

3.3.3 Proportion of Realized Gains and Losses: The Odean Approach

To avoid spurious results in an upward-trending market, Odean (1998) does not simply compare realized gains and losses. Instead, he constructs a more sophisticated measure. He calculates portfolios for each account at each trading date by adding up trading records in chronological order. Every time a sale takes place for a particular investor, he compares

[10] The estimate of the hazard rate is given by

$$\hat{\lambda}(t_k) = \frac{s_k}{n_k}$$

where n_k is the number of observations in the sample in period t_k and s_k is the number of observations that leaves the sample in t_k. Both numbers can be conditioned on gains and losses.

the average buying price (i.e. the average price of a security for all purchases up to that date) to the selling price for each stock in the investor's portfolio at this point in time. An observation counts as a realized gain if the selling price is higher than the average buying price. It counts as a realized loss if the average buying price exceeds the selling price. The observation is omitted if the prices are equal. Paper gains and paper losses are defined similarly. Consider, for example, a security that is in an investor's portfolio at the beginning of the day, and is not sold. The observation counts as a paper gain if the average buying price is lower than both the high and low price on that day. It is omitted if the average buying price lies between the high and low price for the day. Paper losses are defined equivalently. Odean (1998) then computes the proportion of realized gains of all gains (PGR):

$$PGR = \frac{\text{\# of realized gains}}{\text{\# of realized gains + \# of paper gains}} \tag{3.2}$$

The proportion of realized losses (PLR) can be calculated in the same way. Odean then tests for the presence of a disposition by testing whether PGR>PLR.

The original implementation of the approach suffers from a bias which comes from the differential treatment of paper gains/losses relative to realized gains/losses. Note that paper gains/losses are only counted when the average buying price is less than/exceeds both the high and low price of that day. Otherwise they are not counted. Realized gains and losses, on the other hand, are determined relative to the actual selling price regardless of whether the average buying price lies within the daily high and low price of that stock. Since it is more likely that small returns are between the daily high and low prices, this procedure systematically overstates small realized gains/losses relative to small paper gains/losses. Figure 3.8 illustrates this. The lower panel shows an observation that counts as a realized gain although $P_b \in [low, high]$, whereas the observation would not have been considered a paper gain had it not been sold (upper panel).

This bias can, of course, be avoided by applying the same rule to both paper gains/losses and realized gains/losses (that is, by dropping realized gains when their average buying

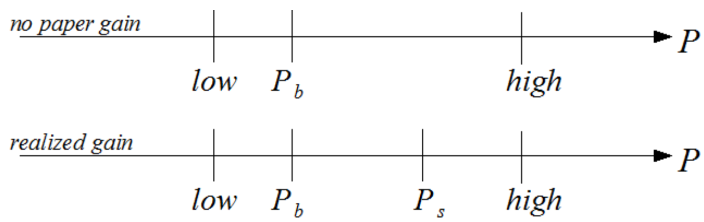

Figure 3.8: *Bias towards realized returns in the Odean approach*

price falls in the interval between the low and high price of that day), a convention that we follow throughout our analysis. While we think that our approach is more rigorous than the original treatment of gains and losses in Odean (1998), we have confirmed that it does not have a substantial effect on any of the results.

3.4 Results

Our main analysis consists of evaluating the Kaplan-Meier estimates for different returns sizes, but given holding periods. Before we present these estimates in Section 3.4.2, we first illustrate the main results using simple histograms. In Section 3.4.3, we present results using the Odean (1998) methodology.

3.4.1 A First Look at the Data using Histograms

Figure 3.9 shows histograms of the returns for a holding period of 12 months. Panel 3.9a shows the histogram for all returns, whereas Panel 3.9b conditions on returns that were realized.[11] Comparison of the two histograms reveals that the conditional histogram contains is much more concentrated around 0, i.e. there are relatively fewer observations in the tails.[12] This observation illustrates our main empirical result: on average, investors are more likely to sell stocks with smaller absolute returns compared to stocks with larger absolute returns (for given holding periods).

[11] We truncate the histograms at a return of 200%.

[12] Note the difference in scales between the two histograms.

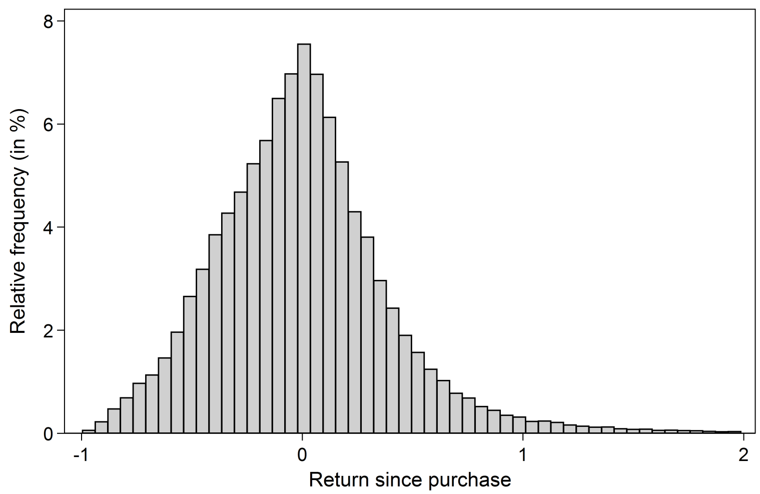

(a) *Unconditional*

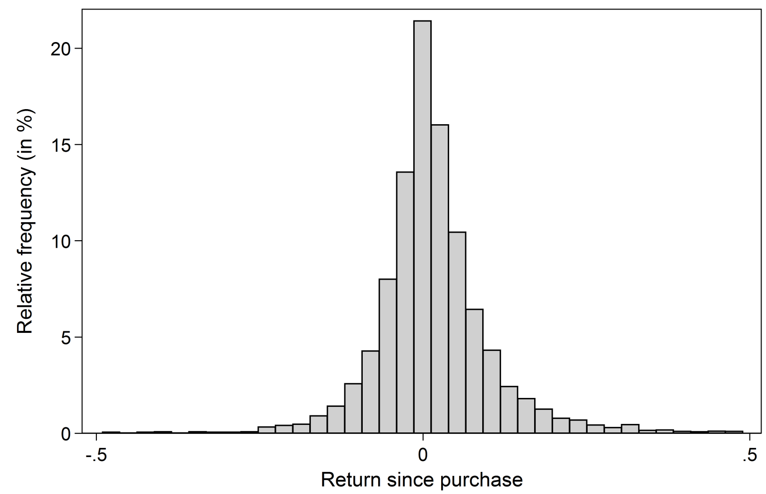

(b) *Conditional on selling*

Figure 3.9: *Histograms of returns for a holding duration of 12 months*

The corresponding kernel density estimates illustrate the same fact in a more compact way (Figure 3.10). The dashed line in this figure represents the kernel density conditional on selling, and the solid line corresponds to the unconditional kernel density function. As in Figure 3.9, the probability mass is more concentrated around 0 conditional on selling, which again implies that investors are more likely to sell stocks with smaller absolute returns. Figures 3.9 and 3.10 use returns for a holding period of 12 months for illustration purposes, but the describe patterns are robust and holds for shorter and longer holding periods as well.

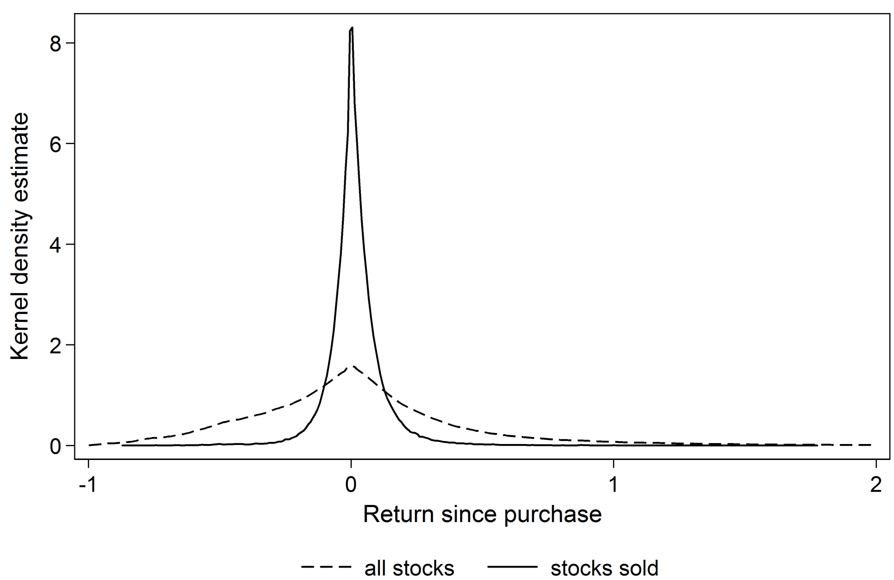

Figure 3.10: *Kernel density estimates*

3.4.2 Duration Model

We now turn to a more rigorous analysis of the patterns shown in the previous section using Kaplan-Meier estimates. Figure 3.11 plots the estimated hazard rate as a function of the holding period for both gains and losses for the entire sample for holding periods between 1 and 30 months. The results are qualitatively and quantitatively in line with Ivkovic *et al.* (2005). In particular, the probability to sell a stock declines with the holding period, and the probability to sell gains is greater than the probability to sell losses for all holding periods. That is, we observe a disposition effect for all holding periods.

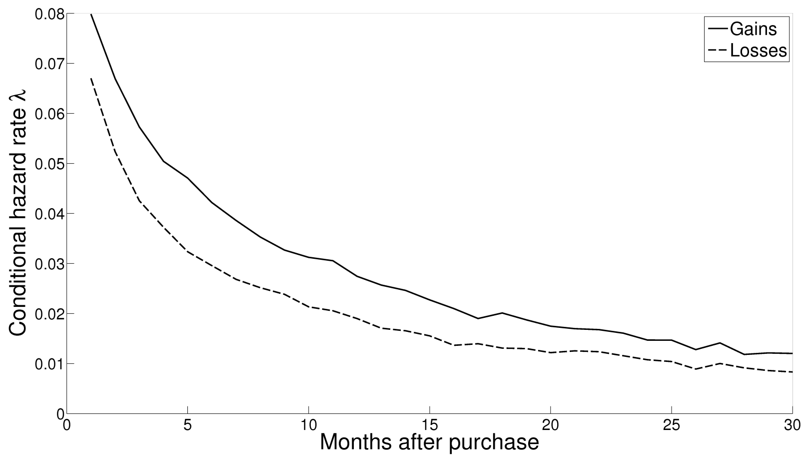

Figure 3.11: *Conditional hazard rate for gains and losses*

Figure 3.12 extends the analysis to the case of different return quantiles. We use quantiles instead of return intervals for two reasons. First, since we are interested in comparing differently-sized gains and losses, we want to make sure that the respective classes of returns that we compare contain equal numbers of observations. Second, we cannot match very high returns with same-sized low returns, because losses cannot exceed 100 percent.[13] The quantile procedure has proven valuable when trying to estimate non-linear functions and has been used by others before (DellaVigna and Pollet (2009)).

The rich dataset enables us to split the data into many quantiles in the subsequent analysis. For illustrative purposes, Figure 3.12 considers only 3 quantiles: small, medium, and large gains and losses, respectively. Solid lines correspond to gains relative to purchase price and dashed lines correspond to losses. Circles denote the smallest gains and losses, triangles denote medium-sized ones and squares denote the largest tercile of gains and losses, respectively. The disposition effect is apparent for all three quantiles of returns, i.e.

[13]If we restrict ourselves to gains smaller than 100% and use intervals instead of quantiles, we do get qualitatively similar results.

the probability to sell gains exceeds the probability to sell losses for each respective pair of quantiles. Moreover, the probability to sell is lower for the larger gains and losses compared to small gains and losses, respectively. While Figure 3.12 is suggestive, it is not conclusive. The figure plots the probability to sell as a function of the holding period. However, we are interested in the probability to sell a stock as a function of the return in order to compare the empirical results to the theories that we considered in Section 3.2. Therefore, we repeat the calculations behind Figure 3.12, split the sample into 25 gain and 25 loss quantiles and plot the probability to sell as a function of those quantiles keeping the holding period fixed.

Figure 3.13 presents our main finding. It shows the hazard rate as a function of return quantiles for four different holding periods: 1, 6, 12, and 24 months.[14] The disposition effect is apparent in all four panels: The probability to sell gains is larger than the probability to sell losses throughout most holding periods and return magnitudes. Small returns with short holding periods and large returns with long holding periods do not display a disposition effect.[15]

More novel and remarkable, however, is the pattern of the selling probability as a function of the return size. For both gains and losses, the probability of selling is largest for small returns and steadily declines for higher (absolute) returns, a consistent pattern throughout all panels. This finding stands in stark contrast to most model predictions that we have derived in Section 3.2. The realization utility view and the casino gambling prospect theory view imply an increasing probability to sell in the gain region. However, we find the exact opposite pattern: smaller returns are more likely to be sold than larger returns. In addition, the probability to sell in the loss region is not constant, which does not support the realization utility view. On the other hand, the bunching view of prospect theory at the end of Section 3.2 appears to be consistent with this evidence.

To further explore the relationship between the propensity to sell a stock and its return

[14]The findings are robust for other holding periods.

[15]Note that we have relatively few observations for holdings of exactly 24 months which makes the estimates less precise.

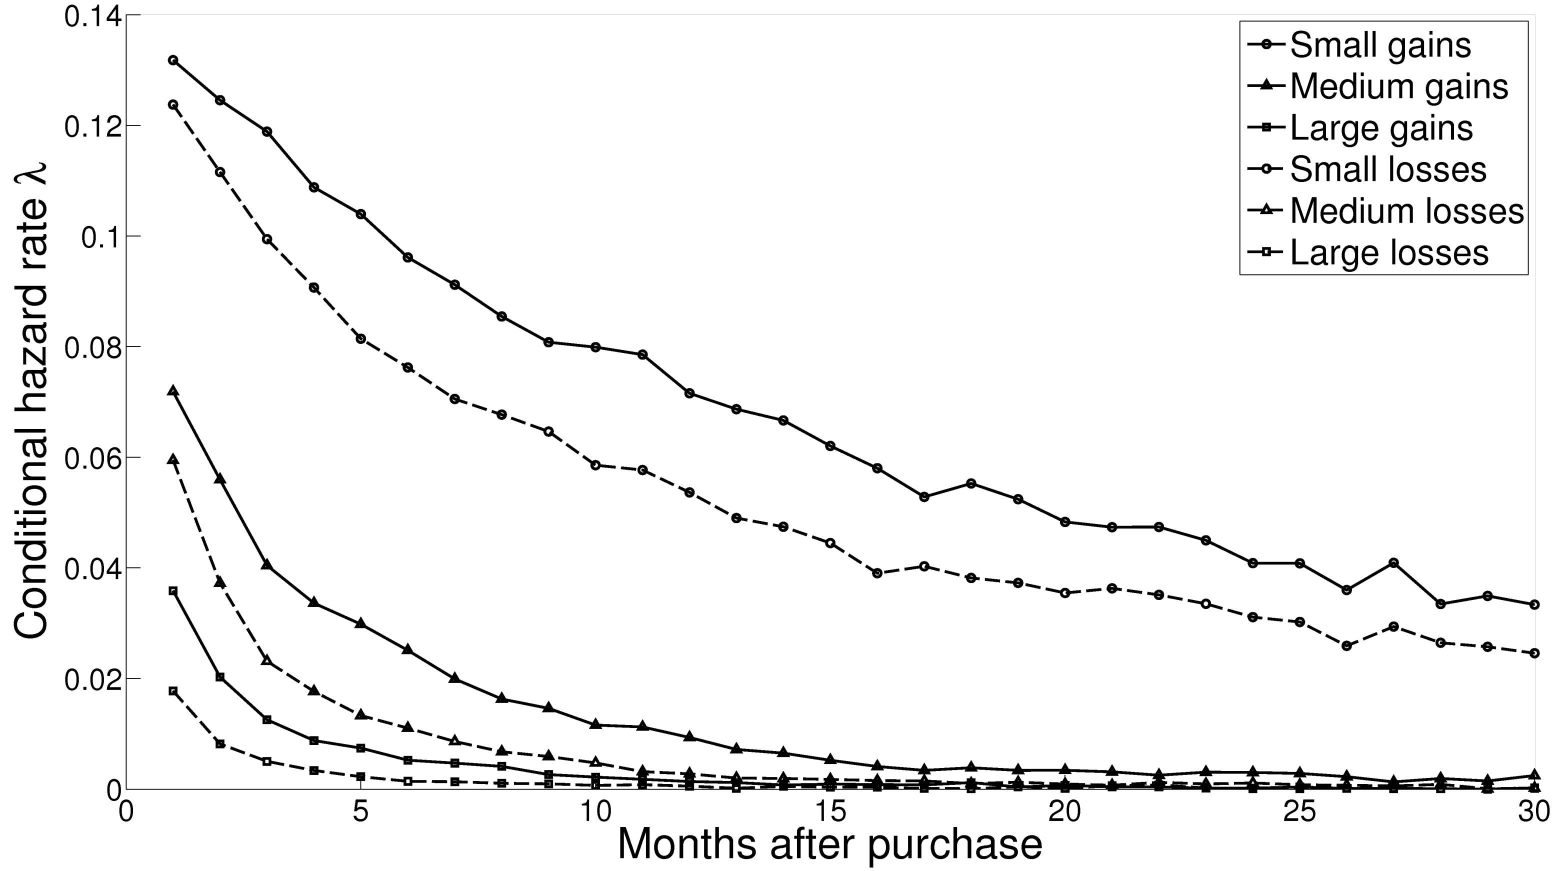

Figure 3.12: *Conditional hazard rate for gains and losses by return tercile*

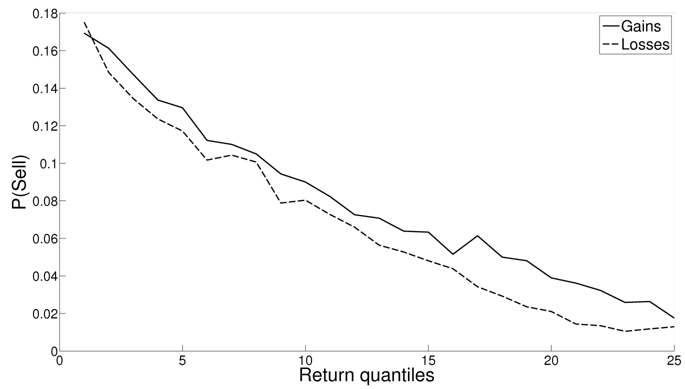

(a) *Holding period: 1 month*

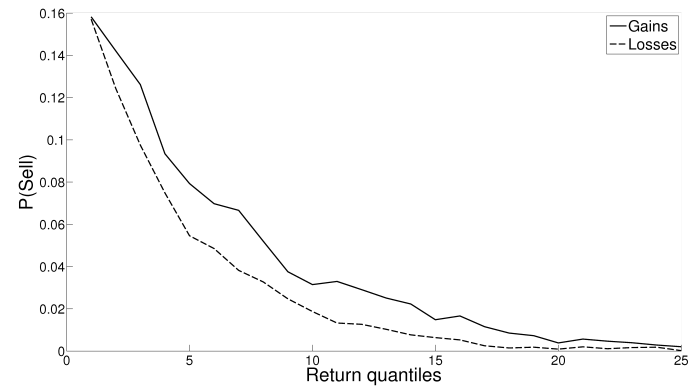

(b) *Holding period: 6 months*

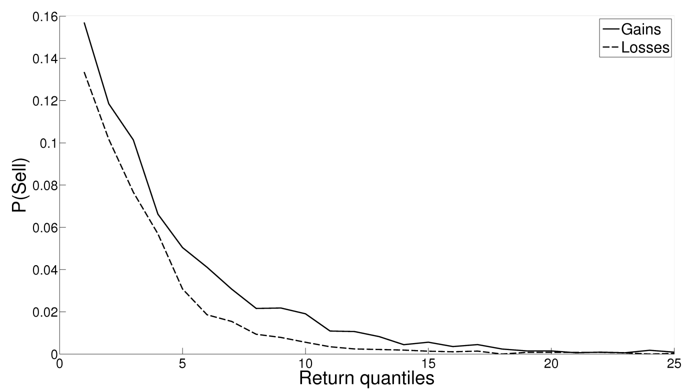

(c) *Holding period: 12 months*

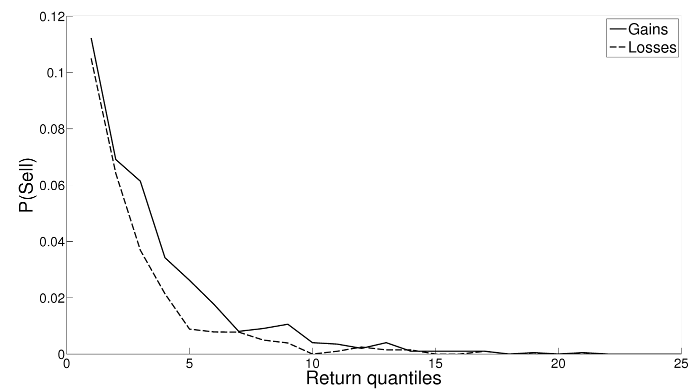

(d) *Holding period: 24 months*

Figure 3.13: *Conditional hazard rate for different holding periods as function of stock returns. Return quantile 1 denotes small returns and return quantile 25 denotes large retunrs.*

relative to the purchase price, we estimate a Cox-proportional hazard model. The parametric approach has the advantage that we can control for other investor-specific characteristics. Every investor-security combination is treated as an observation i, such that the hazard rate is given by:

$$\lambda_i(t,x) = \lambda_0(t) exp\{x'_{it}\delta\}, \text{ where} \tag{3.3}$$

$$x'_{it}\delta = g_{l,it}\beta_l + g_{m,it}\beta_m + g_{h,it}\beta_h + l_{l,it}\gamma_l + l_{m,it}\gamma_m + l_{h,it}\gamma_h + demo_i \tag{3.4}$$

Here, $g_{l,it}$ is a dummy that is equal to 1 if individual-security combination i trades at a low gain in holding period t. The other dummies are interpreted likewise with m being medium gains and losses and h being large gains and losses. Finally, $demo_i$ denotes the additional individual-level demographic control variables from Table 3.1 such as age, gender and wealth.

For both gains and losses, we divide the stock returns into three groups each.[16] Omitting the intermediate middle group enables us to compare effects for small and large gains respectively. For instance, β_1 is the effect of a small gain on the probability to sell a security that has not been sold yet (relative to a medium-size gain). Similarly, γ_1 is the equivalent effect for small losses (relative to a medium-size loss). Using this notation, the previous findings can be rewritten in terms of the model's coefficients: For instance, if the probability of selling declines away from a 0 return, we should observe $\gamma_1 \geq \gamma_2 \geq \gamma_3$ and $\beta_1 \geq \beta_2 \geq \beta_3$.

Columns (1) and (2) in Table 3.2 report regression results with and without control variables, respectively. As in the above figures, we find that the probability of selling declines with the size of returns. Small gains are more likely to be sold than medium-sized gains (the omitted category), and large gains are less likely to be sold than medium-sized gains. The same holds true for small and large losses relative to medium-sized losses. The coefficients between quantiles differ significantly, rejecting the hypothesis that the probability of selling a security is constant across returns. The size of coefficients is in line with the hypothesized order from the preceding paragraph. Adding demographic control

[16]We have tried other splits into 2 to 10 groups with no qualitative change in results or interpretation.

Table 3.2: *Cox-proportional hazard model for the probability to sell a stock given its return relative to purchase price.*

	Dependent variable: Dummy for selling			
Variable	(1)	(2)	(3)	(4)
g_l	1.613***	1.630***	1.640***	1.628***
g_h	-1.095***	-1.107***	-1.268***	-1.297***
l_l	1.332***	1.345***	1.420***	1.415***
l_h	-1.796***	-1.816***	-1.730***	-1.722***
Worst			-0.110***	-0.151***
Best			0.045***	0.027
Controls		X		X
N	4,439,680	1,978,133	1,418,624	578739

***, ** and * denote significance at the 1%, 5% and 10% significance level, respectively.

variables does not change results significantly.

Hartzmark (2015) shows that investors are more likely to sell stocks with more extreme returns among the stocks in a portfolio. In columns (3) and (4) of Table 3.2, we therefore add two dummies as additional control variables: "Best" is equal to 1 if a stock is the best performing stock in an investor's portfolio in a given month, and "Worst" is equal to 1 if a stock is the worst performing stock in an investor's portfolio in a given month. Adding these controls does not alter our results. Given that the number of observations varies widely across specifications, the stability of coefficients is remarkable.[17]

3.4.3 Odean Approach

Table 3.3 replicates Odean's (1998) main result. We use the proposed correction of the bias towards realized gains and losses, as discussed in Section 3.3.3. As a result, all estimates are slightly lower than without bias correction, but qualitatively they are the same. We view this as evidence that the disposition effect cannot be explained by this simple estimation bias. For all months except for December, we find a disposition effect. PGR exceeds PLR in magnitude and the difference is significantly different from zero.[18] From January to November, investors realize 12% of their gains but only 6.7% of their losses. In December, the disposition effect is reversed and investors realize 9.9% of their gains but 10.5% of their losses. As discussed in Odean (1998) and Ivkovic *et al.* (2005), investors face a trade-off between realizing their losses and foregoing tax benefits. Since December is the last month for realizing tax-loss savings, investors choose more often to sell their losers in that month. Odean (1998) shows that the ratio of PGR and PLR declines over the year, implying that

[17]The sample size varies for two reasons: First, demographic control variables are only available for a subset of regressors. Second, Hartzmark (2015) only includes observations that have at least five stocks in their portfolio, and we follow this approach in columns (3) and (4) of Table 3.2.

[18]The standard error is computed in the same way as in Odean (1998). That is, the standard error of the respective difference is given by:

$$\sqrt{\frac{PGR(1-PGR)}{n_{rg}+n_{pg}}+\frac{PLR(1-PLR)}{n_{rl}+n_{pl}}}$$

where $n_{rg}, n_{pg}, n_{rl}, n_{pl}$ are the numbers of realized gains, paper gains, realized losses and paper losses.

tax motives become more important in the course of the year. Our estimates have higher t-statistics than Odean (1998)'s results, mainly because our dataset is much larger.

Table 3.3: *PGR and PLR for the Entire Data Set (bias corrected)*

Variable	Entire Year	Dec	Jan-Nov	Entire Year (Odean)
PLR	0.070	0.105	0.067	0.098
PGR	0.118	0.099	0.120	0.148
Difference	-0.048	0.007	-0.053	-0.05
t-stat	-151.773	5.551	-160.598	-35

Figure 3.14 plots the proportions of realized gains and losses as functions of stock returns for different return quantiles and four different holding periods.[19] For all panels, the proportion of realized gains exceeds the proportion of realized losses for all return quantiles. That is, the disposition effect is apparent for all sizes of returns and for all holding periods. Furthermore, for short holding periods, small gains are less likely to be sold than large gains, while for larger holding periods, stocks with small and large returns are equally likely to be sold.

At first sight, these results appear to be quite different from our previous findings. However, compared to the Cox-proportional hazard model, the Odean approach is biased towards trading activity. That is, the computed probability is not the unconditional probability to sell, but the conditional probability to sell, given some trading activity takes place

[19]For the same reasons as in the previous section, we use quantiles rather than return intervals.

in the portfolio. In the next section, we show that the Odean graphs in Figure 3.14 display the behavior that one would expect if small returns are (unconditionally) more likely to be sold.

3.4.4 The Propensity to Trade

This section reconciles the seemingly contradictory findings from the two above empirical approaches. The duration model computes a conditional probability of selling a stock, given that it was not sold before. In contrast, the Odean approach estimates the probability to sell given that *any* activity in the portfolio takes place. The difference in conditioning sets leads to differences in results if the probability of an investor becoming active depends on the return of her portfolio. A simple way to see this is using Bayes' rule. If $P(Sell)$ is the hazard rate, then

$$P(Sell|active) = \frac{\overbrace{P(active|Sell)}^{=1} P(Sell)}{P(active)}. \tag{3.5}$$

The left-hand side represents the estimated object of the Odean approach, while the numerator of the right-hand size shows represent the estimand of the Cox-proportional hazard model. Therefore, if $P(active)$ varies with the size of the return, then we should expect differences in results from the two approaches.

We estimate $P(active)$ using parametric probit regressions as well as a nonparametric approach. To start, we construct stock portfolios for each investor at the investor-month level as in Section 3.3.2. We then collapse the data, compute the portfolio return for each investor-month combination and construct a dummy that is equal to one if the investor traded at all (sold or bought any stock) in a particular month and that is zero otherwise. Our resulting dataset therefore has one observation for each investor and each month.[20] We then regress this dummy on the portfolio return to get an estimate of the propensity of trading as a function of the portfolio return. Since the shape of this relationship is unknown, we

[20]We require that an investor holds at least two stocks in a given month to be included in the sample.

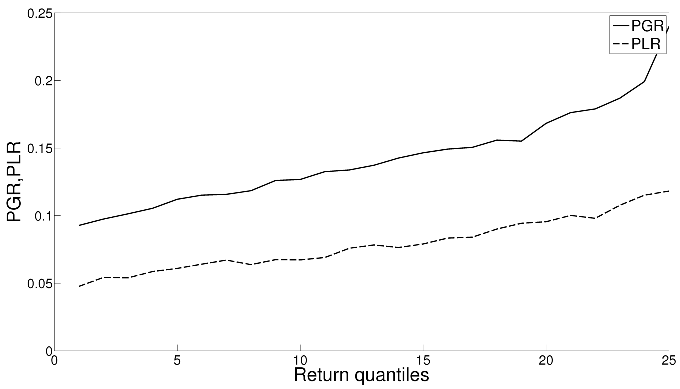

(a) *Holding period: 1 to 3 months*

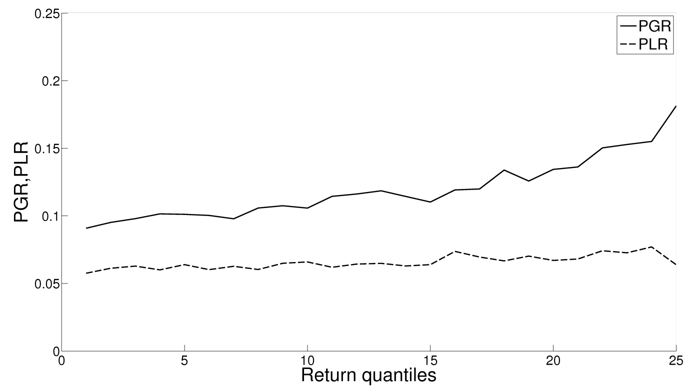

(b) *Holding period: 5 to 7 months*

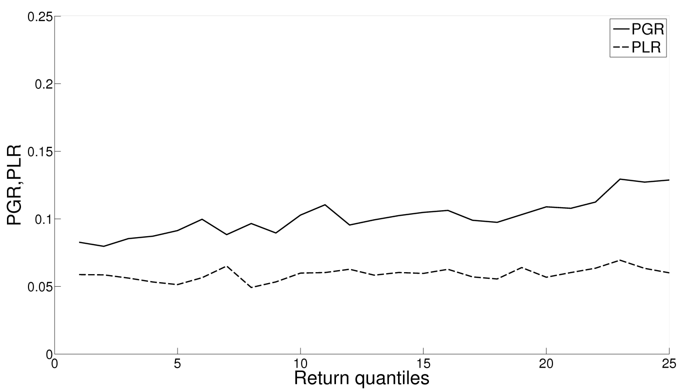

(c) *Holding period: 11 to 13 months*

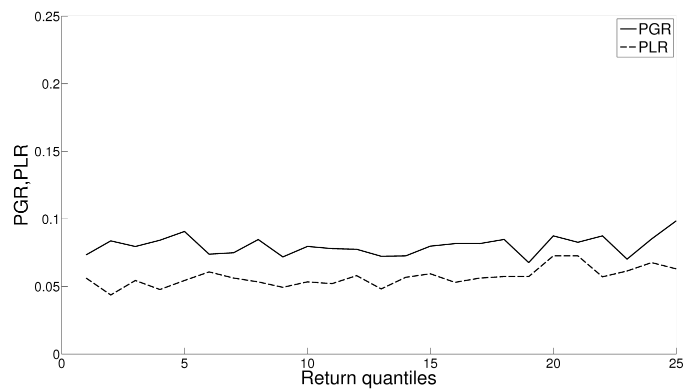

(d) *Holding period: 22 to 24 months*

Figure 3.14: *Proportions of realized gains and losses as in Odean (1998) as a function of stock returns for different holding periods. Return quantile 1 denotes small returns and return quantile 25 denotes large returns.*

start with a nonparametric regression using the Nadaraya-Watson estimator, as described in Haerdle (1990).

The nonparametric model that we want to estimate is given by:

$$g(x) = E[y|X = x], \tag{3.6}$$

where x denotes the portfolio return, and y is our activity dummy. Therefore, an estimate of $g(x)$ can be obtained from the regression

$$y = g(x) + \epsilon, \quad \epsilon \sim (0, \sigma^2(x)) \tag{3.7}$$

Figure 3.15 shows the estimated regression function and 95% confidence bands. The propensity to trade is highest around 0 returns (with the peak slightly above 0), and declines in absolute size of returns. This is exactly what one would expect given the estimates of the two above approaches.

Table 3.4 confirms this result parametrically. We regress the portfolio activity dummy on a (first and third order) polynomial of the portfolio return, allowing for a structural break in the relationship at a portfolio return of 0. The first two columns show results for a linear regression while the last two show results for a probit model. All models confirm the significant impact of the portfolio return on the propensity to trade at all, and the regression function is in line with the non-parametric model above, that is, it is upward-sloping for negative returns and downward-sloping for positive returns.

To summarize, the results of the duration model and the Odean approach can be reconciled by taking into account that the conditioning sets of the two estimates differ. The link is given by the probability that an investor makes any trade as a function of the return of her portfolio. This probability declines with the size of absolute returns, which reconciles our seemingly different results from Sections 3.4.2 and 3.4.3.

3.5 Relation to Ben-David and Hirshleifer (2012)

Ben-David and Hirshleifer (2012) also investigate the relation between past security returns

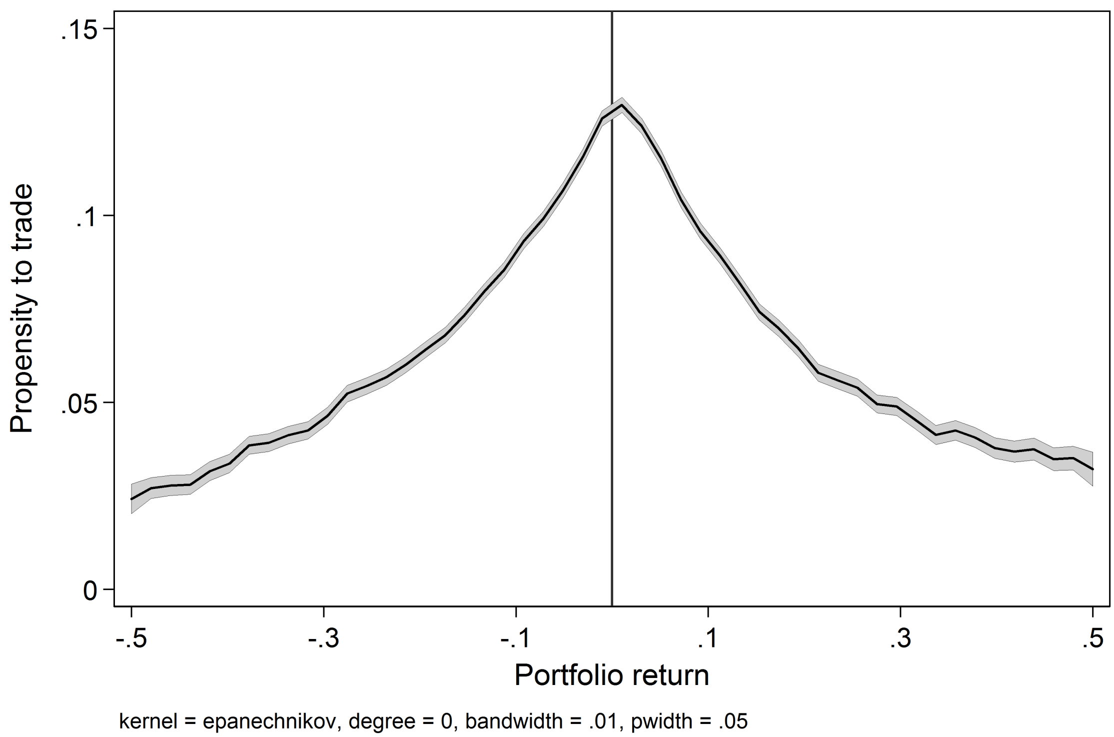

Figure 3.15: *Nonparametric estimate of propensity to trade. Grey area denotes the 95% confidence bands.*

Table 3.4: *Propensity to trade as a function of portfolio return*

	Dependent variable: Dummy for any trade			
Regressors	Linear model		Probit model	
$I(Ret_p > 0)$	0.005***	0.007***	0.022***	0.031***
Ret_p	0.222***	0.509***	1.753***	2.600***
$I(Ret_p > 0)Ret_p$	-0.452***	-1.051***	-3.497***	-5.106***
Ret_p^2		1.026***		3.463***
Ret_p^3		0.887***		3.447***
$I(Ret_p > 0)Ret_p^2$		-0.062		-2.261**
$I(Ret_p > 0)Ret_p^3$		-1.468***		-1.553
Constant	0.115***	0.131***	-1.158***	-1.119***
N	1254918	1254918	1254918	1254918

***, ** and * denote significance at the 1%, 5% and 10% significance level, respectively.

and subsequent sales using the same data that we used. The main result of this paper is that the probability to sell a security is "asymmetrically V-shaped". That is, larger returns are more likely to get sold and more so for positive returns, and the selling probability does not have a discontinuity at a stock return of 0. While the latter is consistent with our results (recall that we find a stronger disposition effect for larger returns), the former fact seemingly stands in contrast to our results. In this section, we investigate the causes of this discrepancy.

In contrast to the approach in Ivkovic *et al.* (2005), Ben-David and Hirshleifer (2012) follow the holding of a stock on each day after it was purchased without conditioning on portfolio activity as in Odean (1998). This approach hugely expands the data set and most of our results are based on a random sample of 25% of accounts.[21] For the short holding horizons in Figure 3.16 below, we are able to use the entire data set.

We start by replicating the main results in Ben-David and Hirshleifer (2012). First, we follow their procedure to construct each investor's portfolio on each possible trading day. This enables us to follow each stock from the purchase day to the selling day. Table 3.5 reports the unconditional probability that a stock gets sold for different holding periods. In general, stocks are sold infrequently: The probability that a stock is sold within the first 30 (trading) days after purchase is 0.79 percent. The monthly selling probability monotonically decreases for longer holding horizons.

Table 3.5: *Unconditional probability (in %) of selling for different numbers of holding days. Note: Days are trading days.*

Holding days	1-30	31-60	61-90	91-120	121-150	151-180	181-210	211-240
Probability to sell	.790	.508	.391	.309	.258	.230	.201	.174

[21] Ben-David and Hirshleifer (2012) follow the same strategy and base their estimates on a random sample of 10,000/77,000 ≈ 12% of accounts.

Figure 3.16 shows our replication of Ben-David and Hirshleifer's (2012) motivating results (Figure 1 of their paper). We document a sharp decrease of the selling probability around small stock returns, and we find that the difference between small positive and small negative returns is negligible.[22] These results support an asymmetrically V-shaped selling schedule. That is, stocks with small returns are unlikely to be sold and stocks with higher returns are more likely to be sold, and more so for positive returns. However, the results in Figure 3.16 are conditional on the specific holding periods of 1 or 5 days, both of which are extremely short.

Figure 3.17 illustrates the results of the same exercise when data are pooled over different holding horizons. The left panel pools all holding periods less than 30 days and is generally in line with the results in Figure 3.16, although there is an apparent small discontinuity around 0. The right panel pools all holding periods of less than 250 trading days and looks very different: The discontinuity of the selling probability around 0 is more apparent and the probability to sell appears to decrease with the absolute value of the return.

Table 3.6 provides a more detailed account of the relation between the discontinuity result around a return of 0 and the holding period. We regress an indicator for stock sales (multiplied by 100) on an indicator of whether a stock's return was greater than zero, on a third-order polynomial of a stock's return and on interactions of the indicator and the polynomial terms. Each column reports the regression results for a different stock holding period. For ease of interpretability, we report results for the linear probability model here, but none of the results change when a logistic regression model is used instead. The indicator for a positive return is statistically significant in each regression. For instance, for a holding period of up to 30 days, a positive return increases the likelihood to sell the stock by .367 percentage points. To assess whether this is large or small, we scale the coefficient by the unconditional probability of selling for each holding interval (from Table 3.5 above)

[22]Our replication is based on a quadratic polynomial in the stock return on each side of the threshold of 0, while Ben-David and Hirshleifer (2012) use higher-order polynomials. We produced additional results using nonparametric kernel density regressions (available on request) that looked even more similar to the original results.

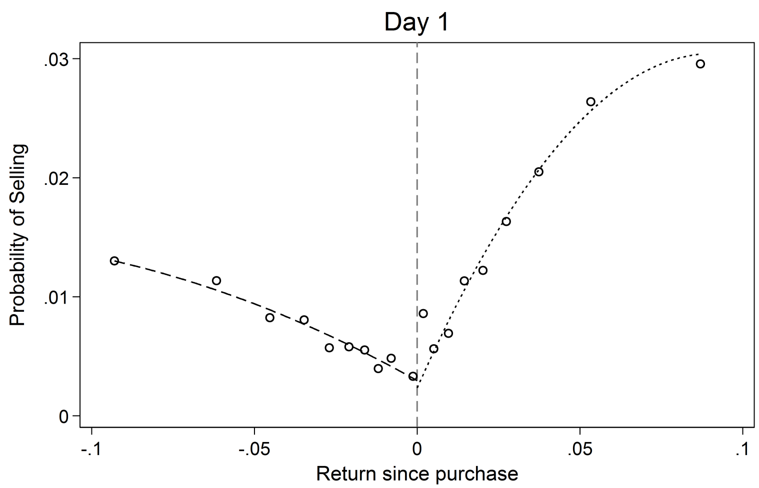

(a) *For holding period of 1 day*

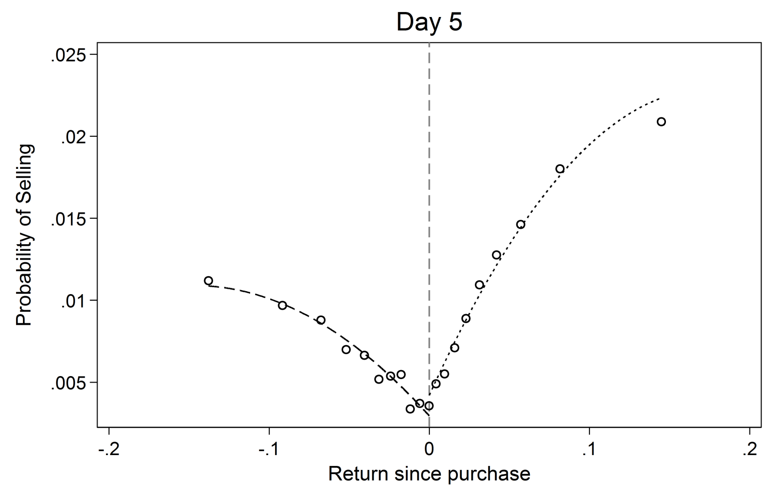

(b) *For holding period of 5 days*

Figure 3.16: *Replication of Ben-David and Hirshleifer (2012)*

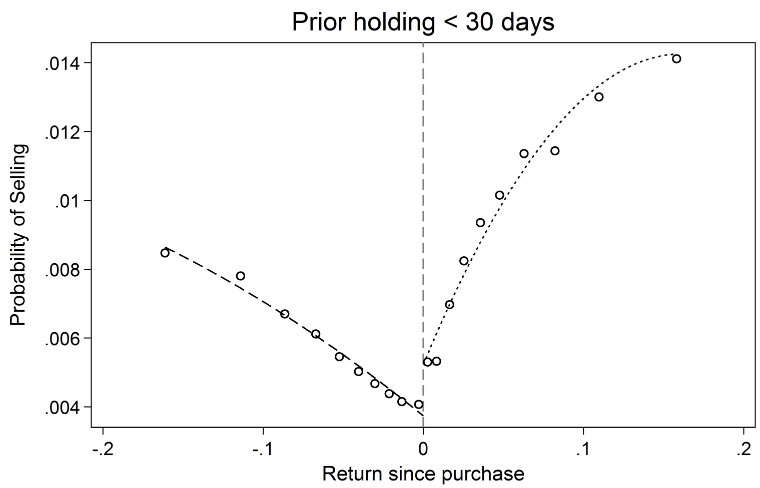

(a) *Holdings less than 30 days*

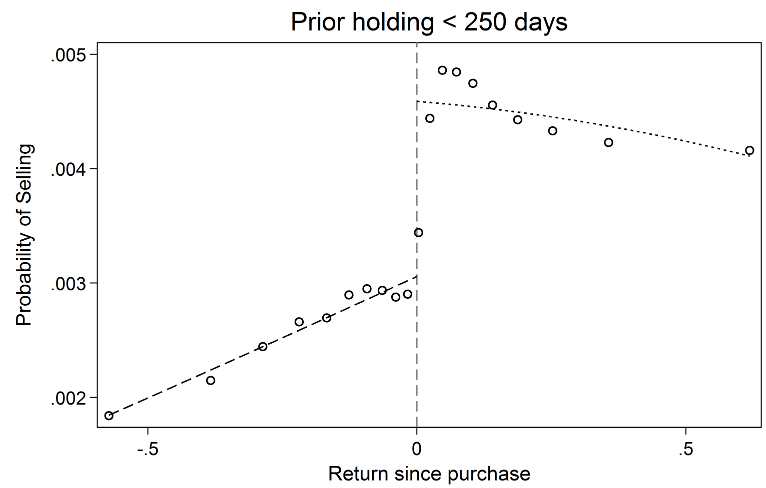

(b) *Holdings less than 250 days*

Figure 3.17: *Pooled selling schedules*

in the last row of the table. For shorter holding horizons, the discontinuity is about 40% of the unconditional probability to sell, while for longer holding horizons it is about 20%.

The effect of positive returns relative to purchase price on the probability to sell a stock is sizeable for all horizons within one year of stock purchase. The long-lasting effect is in line with other research that investigates how investment decisions depend on past prices. For instance, Heath *et al.* (1999) show that employees are more likely to exercise stock options when the stock price is greater than the maximum achieved over the previous year, and Baker *et al.* (2012) show that the 52-week high is an important anchor for merger offers. Here, we show that the purchase price of a stock affects trading decisions for (at least) one year, an unsurprising finding in light of the aforementioned studies.

To summarize, we document that the probability to sell is asymmetrically V-shaped only for very short holding periods. If results are pooled over different horizons, we find a pronounced discontinuity at 0 returns and we also find that the selling probability decreases in the absolute value of the return, in line with our previous findings.

It seems plausible that investors trading at very short horizons are different from those that trading at longer horizons. Indeed, Ben-David and Hirshleifer (2012) report weaker results for infrequent traders (their Figure 5). As our study generally focuses on holding periods greater than one month (and up to two years), it can also be read as a study of the behavior of those less frequent traders.

3.6 Conclusion

What drives the disposition effect? In this paper, we consider two leading explanations, prospect theory and realization utility. We derive implications of the theories for the probability to sell as a function of the return size, and contrast the predictions to new empirical findings of this relationship. Our main empirical finding is that, for all but very short holding horizons, investors are more likely to sell stocks with small returns (i.e. stocks with prices close to the purchase price) than to sell stocks with large returns. We use two different empirical approaches. Using the duration model of Ivkovic *et al.* (2005), we find

Table 3.6: *Probability of selling as a function of stock returns for different holding periods*

Holding days	1-30	31-60	61-90	91-120	121 -150	151-180	181-210	211-240
$I(Ret > 0)$	0.367***	0.211***	0.143***	0.106***	0.039***	0.055***	0.035**	0.031***
Ret	-4.892***	-1.344***	-0.048	-0.071	0.216*	0.090	0.159	0.187*
Ret^2	-14.023***	-2.689***	0.186	-0.165	0.442	0.140	0.338	0.397
Ret^3	-10.649***	-0.915	0.348	-0.104	0.378	0.073	0.277	0.314
$I(Ret > 0)Ret$	9.376***	2.710***	0.500***	0.367**	0.045	0.078	0.074	-0.086
$I(Ret > 0)Ret^2$	10.909***	1.979***	-0.222	0.141	-0.483	-0.194	-0.528	-0.417
$I(Ret > 0)Ret^3$	11.089***	1.002	-0.347	0.104	-0.377	-0.070	-0.245	-0.313
Constant	0.350***	0.278***	0.279***	0.226***	0.225***	0.192***	0.176***	0.158***
N	2808535	2857042	2869490	2888734	2781890	2719315	2712780	2598521
$\frac{\hat{\beta}_{I(Ret>0)}}{\text{Uncond P(Sell)}}$	.465	.416	.366	.343	.151	.239	.174	.178

that larger absolute returns are less likely to be sold than small returns. Using the Odean (1998) methodology, we find that small gains are less likely to be sold than large gains for small holding periods, while there is no relationship between the return size and selling probability for larger holding periods. Realization utility cannot explain important features of the data. In contrast, a version of prospect theory, that puts emphasis on "bunching" around the kink appears to be consistent with the facts.

We then reconcile these seemingly contradictory empirical findings by pointing out that the two approaches estimate different probabilities. While Ivkovic *et al.* (2005) consider the probability of selling a stock with a given return in a given month (conditional on still holding it), Odean considers the probability of selling a stock, given that the investor sells or buys a stock in her portfolio. Therefore, a comparison of the two approaches needs to take into account that the probability of undertaking any transaction in the portfolio is a function of the individual returns of the stocks in the portfolio. We find that individuals are more likely to engage in transactions for stocks with small returns than for stocks with large returns. Jointly, these findings pose yet another challenge: An investor is more likely to perform any action (i.e. sell or buy stocks, look into her portfolio) if returns of stocks in her portfolio are small. Once the investor decides to act, however, she is more likely to sell large returns. It is hard to think of a theory that would predict this.

Of course, consistency with the facts alone does not make a theory the true explanation. For instance, an explanation that combines elements of prospect theory, realization utility, and overconfidence might very well be at the heart of the disposition effect. Moreover, despite the fact that the choice of the reference point is crucial for any analysis of reference-dependent utility, with the exception of Meng (2013), the existing literature on the disposition effect has not thoroughly investigated this choice. Instead, it solely focuses on a stock's buying price as "a noisy proxy for the investor's true reference point" (Odean, 1998).[23]

[23] While Odean discusses the possibility of other determinants of the reference point (in particular, expectations), his focus remains on variants of the purchase price. For investors who buy the same stock several times, Odean considers the average purchase price, the highest purchase price, the first purchase price, and the most recent purchase price

Koszegi and Rabin (2006) argue that expectations, rather than the status quo, play a key role in the formation of individuals' reference points. In particular, when individuals do not plausibly expect to maintain the status quo, "equating the reference point with expectations generally makes better predictions." This suggests that the purchase price might not be a good choice of a reference point for trading decisions. While it may be a good proxy in times of low returns, with soaring stock prices like in the 1990s – our dataset covers transaction from 1991 through 1996, a time period during which on average the S&P 500 index rose over 15% annually – investors may have higher expectations, and, hence, a higher reference point than the status quo. Future research should take into account the possibility that investors' reference points may be driven by expectations or might generally deviate from a stock's purchase price.

Chapter 4

References

Bibliography

ANDERSON, P., CHISHOLM, D. and FUHR, D. (2009). Effectiveness and cost-effectiveness of policies and programs to reduce the harm caused by alcohol. *Lancet*, **373**, 2234–2246.

ANDREONI, J. and SPRENGER, C. (2012). Estimating time preferences from convex budgets. *American Economic Review*, **120** (7), 3357–3376.

ARIELY, D. and WERTENBROCH, K. (2002). Procrastination, Deadlines, and Performance: Self-Control By Recommitment. *Psychological Science*, **13** (3), 219–224.

ASHRAF, N., KARLAN, D. and YIN, W. (2006). Tying odysseus to the mast: Evidence from a commitment savings product in the philippines. *Quarterly Journal of Economics*, **121** (2), 635–672.

AUGENBLICK, N., NIEDERLE, M. and SPRENGER, C. (2014). Working over time: Dynamic inconsistency in real effort tasks. *mimeo*.

BAKER, M., PAN, X. and WURGLER, J. (2012). The effect of reference point prices on mergers and acquisitions. *Journal of Financial Economics*, **106** (1), 49–71.

BANAJI, M. R. and STEELE, C. M. (1989). Alcohol and self-evaluation – is a social cognition approach beneficial? *Social Cognition*, **7** (2), 137–151.

BANDIERA, O. and RASUL, I. (2006). Social networks and technology adoption in northern mozambique. *Economic Journal*, **116** (514), 869–902.

BANERJEE, A. and MULLAINATHAN, S. (2010). The shape of temptation: Implications for the economic lives of the poor. *NBER Working Paper 15973*.

BANK, W. (2008). *World Development Report 2008: Agriculture for Development*. Washington, DC: The World Bank.

BARBER, B. and ODEAN, T. (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance*, **55**, 773–806.

BARBERIS, N. (2012). A model of casino gambling. *Management Science*, **58** (1), 35–51.

—, HUANG, M. and THALER, R. H. (2006). Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American Economic Review*, **96** (4), 1069–1090.

— and Xiong, W. (2009). What drives the disposition effect? an analysis of a long-standing preference-based explanation. *Journal of Finance*, **64** (2), 751–784.

— and — (2012). Realization utility. *Journal of Financial Economics*, **104** (2), 251–271.

Becker, G. S. and Mulligan, C. B. (1997). The endogenous determination of time preferences. *Quarterly Journal of Economics*, **112** (3), 729–758.

— and Murphy, K. M. (1988). A theory of rational addiction. *Journal of Political Economy*, **96** (4), 675–700.

Ben-David, I. and Hirshleifer, D. (2012). Are investors really reluctant to realize their losses? trading responses to past returns and the disposition effect. *Review of Financial Studies*, **25** (8), 2485–2532.

Bénabou, R. and Tirole, J. (2003). Intrinsic and extrinsic motivation. *Review of Economic Studies*, **70**, 489–520.

Benegal, V. (2005). India: Alcohol and public health. *Addiction*, **100**, 1051–1056.

BenYishay, A. and Mobarak, M. (2014). Social learning and communication. *mimeo*.

Bernheim, D., Meer, J. and Novarro, N. (2012). Do consumers exploit precommitment opportunities? evidence from natural experiments involving liquor consumption. *NBER Working Paper*, **17762**.

— and Rangel, A. (2004). Addiction and cue-triggered decision processes. *American Economic Review*, **94** (5), 1558–1590.

—, Ray, D. and Yeltekin, S. (2014). Poverty and self-control. *mimeo*.

Beshears, J., Choi, J., Liaison, D., Madrian, B. C. and Sakong, J. (2011). Self control and liquidity: How to design a commitment contract. *mimeo*.

Brune, L., Gine, X., Goldberg, J. and Yang, D. (2014). Facilitating savings for agriculture: Field experimental evidence from malawi. *mimeo*.

Bryan, G., Karlan, D. and Nelson, S. (2010). Commitment devices. *Annual Review of Economics*, **2**, 671–698.

Burghart, D., Glimcher, P. and Lazzaro, S. (2013). An expected utility maximizer walks into a bar... *"Journal of Risk and Uncertainty"*, **46** (3), 215–246.

Carter, M., Laajaj, R. and Yang, D. (2013). The impact of voucher coupons on the uptake of fertilizer and improved seeds: Evidence from a randomized trial in mozambique. *American Journal of Agricultural Economics*, **95** (5), 1345–1351.

Carvalho, L. S., Prina, S. and Sydnor, J. (2014). The effect of saving on risk attitudes and intertemporal choices. *mimeo*.

Charness, G. and Gneezy, U. (2009). Incentives to exercise. *Econometrica*, **3**, 909–931.

Chetty, R. (2006). A new method of estimating risk aversion. *American Economic Review*, **96** (5), 1821–1834.

—, Looney, A. and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, **99** (4), 1145–1177.

Conley, T. and Udry, C. (2010). Learning about a new technology: Pineapple in ghana. *American Economic Review*, **100** (1), 35–69.

Cook, P. J. and Moore, M. J. (2000). Alcohol. *Handbook of Health Economics*, **1(30)**, 1629–1673.

— and Tauchen, G. (1982). The effect of liquor taxes on heavy drinking. *The Bell Journal of Economics*, **13** (2).

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, **47** (2), 315–372.

— and Pollet, J. (2009). Investor inattention and friday earnings announcements. *Journal of Finance*, **64**, 709–749.

Duflo, E., Kremer, M. and Robinson, J. (2008). How high are the returns to fertilizer use? *American Economic Review Papers and Proceedings*, **98** (2), 482–488.

—, — and — (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from kenya. *American Economic Review*, **101** (6), 2350–2390.

Dupas, P. (2014). Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. *Econometrica*, **82** (1), 197–228.

— and Robinson, J. (2013a). Savings constraints and microenterprise development: Evidence from a field experiment in kenya. *American Economic Journal: Applied Economics*, **5** (1), 163–192.

— and — (2013b). Why don't the poor save more? evidence from health savings experiments. *American Economic Review*, **103** (4), 1138–1171.

Ebrahim, I., Shapiro, C., Williams, A. and Fenwick, P. (2013). Alcohol and sleep i: Effects on nocturnal sleep. *Alcoholism: Clinical and Experimental Research*, **37** (4), 539–549.

Evenson, R. and Gollin, D. (2003). Assessing the impact of the green revolution, 1960 to 2000. *Science*, **300**, 758–762.

Fisher, I. (1926). *Prohibition at Its Worst*. New York: MacMillan.

— (1930). *The Theory of Interest*. London: MacMillan.

Foster, A. and Rosenzweig, M. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, **103** (6), 1176–1209.

— and — (2010). Microeconomics of technology adoption. *Annual Review of Economics*, **2**, 395–424.

Frederick, S., Loewenstein, G. and O'Donoghue, T. (2002). Time discounting and time preference: A critical review. **40** (351–401).

Fudenberg, D. and Levine, D. (2006). A dual-self model of impulse control. *American Economic Review*, **96** (3), 694–719.

Fuller, R. K., Branchey, L. and Brightwell, D. R. (1986). Disulfiram treatment of alcoholism. a veterans administration cooperative study. *JAMA*, **256**, 1449–1455.

Giancola, P. R., Josephs, R. A., Parrott, D. and Duke, A. A. (2010). Alcohol myopia revisited: Clarifying aggression and other acts of disinhibition through a distorted lens. *Perspectives on Psychological Science*, **5** (3), 265–278.

Gine, X., Karlan, D. and Zinman, J. (2010). Put your money where your butt is: A commitment contract for smoking cessation. *American Economic Journal: Applied Economics*, **2**, 213–235.

Glazer, J. and Weiss, A. (2007). A model of dysfunctional urges and addiction with an application to cigarette smoking. *The B.E. Journal of Economic Analysis & Policy*, **7** (1), Article 3.

Gruber, J. and Kőszegi, B. (2001). Is addiction rational? theory and evidence. *Quarterly Journal of Economics*, **116** (4), 1261–1303.

— and Kőszegi, B. (2004). Tax Incidence when Individuals are Time-inconsistent: the Case of Cigarette Excise Taxes. *Journal of Public Economics*, **88**, 1959–1987.

— and Mullainathan, S. (2005). Do cigarette taxes make smokers happier? *Advances in Economic Analysis and Policy*, **5** (1), Article 4.

Gul, F. and Pesendorfer, W. (2001). Temptation and self-control. *Econometrica*, **69** (6), 1403–1435.

Haerdle, W. (1990). *Applied Nonparametric Regression*. Cambridge University Press.

Harris, C. and Laibson, D. (2001). Dynamic choices of hyperbolic consumers. *Econometrica*, **69** (4), 935–957.

Hartzmark, S. M. (2015). The worst, the best, ignoring all the rest: The rank effect and trading behavior. *Review of Financial Studies*, **28** (4), 1024–1059.

Haushofer, J. and Fehr, E. (2014). On the psychology of poverty. *Science*, **344**, 862–867.

Heath, C., Huddart, S. and Lang, M. (1999). Psychological factors and stock option exercise. *The Quarterly Journal of Economics*, **114** (2), 601–627.

Hens, T. and Vlcek, M. (2011). Does prospect theory explain the disposition effect? *Journal of Behavioral Finance*, **12** (3), 141–157.

Hey, J. D. and Lee, J. (2005). Do subjects separate (or are they sophisticated)? *Experimental Economics*, **8**, 233–265.

Higgins, S. and Petry, N. M. (1999). Contingency management – incentives for sobriety. *Alcohol Research & Health*, pp. 122–127.

Higgins, S. T., Silverman, K., Simon, S. C. and Naito, N. A. (2012). Incentives in health – an introduction. *Preventive Medicine*, **55**, S2–6.

Hinnosaar, M. (2012). Time inconsistency and alcohol sales restrictions. *mimeo*.

Holt, C. (1986). Preference reversals and the independence axiom. *American Economic Review*, **76**, 508–513.

Houser, D., Schunk, D., Winter, J. and Xiao, E. (2010). Temptation and commitment in the laboratory. *mimeo*.

Hull, J. G. and Slone, L. B. (2004). Alcohol and self-regulation. *Handbook of Self-regulation: Research, and Applications*, pp. 466–491.

IIPS and International, M. (2007). *National Family Health Survey (NFHS-3), India, 2005-06: India (Volume I)*. Mumbai: International Institute for Population Sciences.

— and — (2008). *National Family Health Survey (NFHS-3), India, 2005-06: Tamil Nadu*. Mumbai: International Institute for Population Sciences.

Ivkovic, Z., Poterba, J. and Weissbenner, S. (2005). Tax-motivated trading by individual investors. *American Economic Review*, **95(5)**, 1605–1630.

John, A. (2014). When commitment fails – evidence from a regular saver product in the philippines. *LSE EOPP Discussion Paper 55*.

Jovanovic, B. and Nyarko, Y. (1996). Learning by doing and the choice of technology. *Econometrica*, **64** (6), 1299–1310.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, **47**, 263–291.

Karlan, D., McConnell, M., Mullainathan, S. and Zinman, J. (2014a). Getting to the top of mind: How reminders increase saving. *Management Science (forthcoming)*.

—, Ratan, A. L. and Zinman, J. (2014b). Savings by and for the poor: A research review and agenda. *Review of Income and Wealth*, **60** (1), 36–78.

Kast, F., Meier, S. and Pomeranz, D. (2014). Under-savers anonymous: Evidence on self-help groups and peer pressure as a savings commitment device. *HBS Working Paper 12-060*.

Kaur, S., Kremer, M. and Mullainathan, S. (2014). Self control at work. *Journal of Political Economy (forthcoming)*.

Kaustia, M. (2010). Prospect theory and the disposition effect. *Journal of Financial and Quantitative Analysis*, **45** (03), 791–812.

Keren, G. and Roelofsma, P. (1995). Immediacy and certainty in intertemporal choice. *Organizational Behavior and Human Decision Processes*, **63** (3), 287–297.

Koszegi, B. and Rabin, M. (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics*, **121(4)**, 1133–1166.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, **112** (2), 443–478.

— (2001). A cue-theory of consumption. *Quarterly Journal of Economics*, **116** (1), 81–119.

— (2015). Why don't present-biased agents make commitments. *American Economic Review – Paper & Proceedings (forthcoming)*.

Lester, R., Ritvo, P., Mills, E. J., Kariri, A., Karanja, S., Chung, M. H., Jack, W., Habyarimana, J., Sadatsafavi, M., Najafzadeh, M., Marra, C. A., Estambale, B., Ngugi, E., Ball, T. B., Thabane, L., Gelmon, L. J., Kimani, J., Ackers, M. and Plummer, F. A. (2010). Effects of a mobile phone short message service on antiretroviral treatment adherence in kenya (weltel kenya1): a randomised trial. *The Lancet*, **376** (9755), 1838–1845.

MacKillopp, J., Amlung, M. T., Few, L. R., Ray, L. A., Sweet, L. H. and Munafò, M. R. (2011). Delayed reward discounting and addictive behavior: A meta-analysis. *Psychopharmacology*, **216**, 305–321.

Mani, A., Mullainathan, S., Shafir, E. and Zhao, J. (2013). Poverty impedes cognitive function. *Science*, **341**, 976–980.

Manning, W. G., Blumberg, L. and Moulton, L. H. (1995). The demand for alcohol: The differential response to price. *Journal of Health Economics*, **14**, 123–148.

Meng, J. (2013). Can prospect theory explain the disposition effect? a new perspective on reference points, mimeo.

Miles, W. R. (1924). *Alcohol and Human Efficiency: Experiments with Moderate Quantities and Dilute Solutions of Ethyl Alcohol on Human Subjects*. Carnegie Institute of Washington.

Milkman, K., Minson, J. and Volpp, K. (2014). Holding the Hunger Games Hostage at the Gym: An Evaluation of Temptation Bundling. *Management Science*, **60** (2), 283–299.

Mullainathan, S. and Shafir, E. (2013). Scarcity: Why having too little means so much. *Times Books*.

Munshi, K. (2004). Social learning in a heterogeneous population: Technology diffusion in the indian green revolution. *Journal of Development Economics*, **73**, 185–213.

Neufeld, K., Peters, D., Ranic, M., Bonuc, S. and Broonera, R. (2005). Regular use of alcohol and tobacco in india and its association with age, gender, and poverty. *Drug and Alcohol Dependence*, **77**, 283–291.

O'Daire, A. (2009). *Blood Alcohol, Breath Alcohol, Impairment and the Law – A Manual for Law Enforcement*. Bloomington: AuthorHouse.

Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance*, **53**, 1775–1798.

O'Donoghue, T. and Rabin, M. (1999). Doing it now or later. *American Economic Review*, **89** (1), 103–124.

— and — (2006). Optimal sin taxes. *Journal of Public Economics*, **90** (10), 1825–1849.

of the European Alcohol, S. G. and Forum, H. (2011). *Alcohol, Work, and Productivity*. European Commission.

O'Farrell, T. J., Allen, J. P. and Litten, R. Z. (1995). *Disulfiram (Antabuse) Contracts in Treatment of Alcoholism*. in: Integrating Behavioral Therapies With Medications in the Treatment of Drug Dependence, National Institute on Drug Abuse.

Ortner, C. N. M., MacDonald, T. K. and Olmstead, M. C. (2003). Alcohol intoxication reduces impulsivity in the delay-discounting paradigm. *Alcohol and Alcoholism*, **38** (2), 151–156.

Patel, N. P., Grandner, M. A., Xie, D., Branas, C. C. and Gooneratne, N. (2010). 'sleep disparity' in the population: poor sleep quality is strongly associated with poverty and ethnicity. *BMC Public Health*, **10**, 1–11.

Patel, V. (2007). Alcohol use and mental health in developing countries. *Annals of Epidemiology*, **17** (5), S87–S92.

Perry, J. L. and Carroll, M. E. (2008). The role of impulsive behavior in drug abuse. *Psychopharmacology*, **200**, 1–26.

Petry, N. M., Martin, B., Cooney, J. L. and Kranzler, H. R. (2000). Give them prizes, and they will come: Contingency management for treatment of alcohol dependence. *Journal of Consulting and Clinical Psychology*, **68** (2), 250–257.

Poleshuk, E. L. and Green, C. R. (2008). Socioeconomic disadvantage and pain. *Pain*, **136**, 235–238.

Prendergast, M., Modus, D., Finney, J., Greenwell, L. and Roll, J. (2006). Contingency management for treatment of substance use disorders – a meta-analysis. *Addiction*, **101**, 1546–1560.

Prina, S. (2014). Banking the poor via savings accounts: Evidence from a field experiment. *mimeo*.

Rahman, L. (2003). Alcohol prohibition and addictive consumption in india. *mimeo*.

Richards, J. B., Zhang, L., Mitchell, S. H. and de Wit, H. (1999). Delay or probability discounting in a model of impulsive behavior: Effect of alcohol. *Journal of the Experimental Analysis of Behavior*, **71**, 121–143.

Room, R. (2004). Smoking and drinking as complementary behaviours. *Biomedicine and Pharmacotherapy*, **58**, 111–115.

Royer, H., Stehr, M. and Sydnor, J. (2014). Incentives, commitments and habit formation in exercise: Evidence from a field experiment with workers at a fortune-500 company. *AEJ Applied (forthcoming)*.

Schaner, S. (2014). The persistent power of behavioral change: Long-run impacts of temporary savings subsidies for the poor. *mimeo*.

Schofield, H. (2014). The economic costs of low caloric intake: Evidence from india. *mimeo*.

Schuckit, M., Tipp, J. E., Reich, T., Hesselbrock, V. M. and Buchholz, K. K. (1995). The histories of withdrawal convulsions and delirium tremens in 1648 alcohol dependent subjects. *Addiction*, **90**, 1335–1347.

Shefrin, H. and Statman, M. (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance*, **40**, 777–790.

Starmer, C. and Sugden, R. (1991). Does the random-lottery incentive system elicit true preferences? *American Economic Review*, **81**, 971–978.

Steele, C. M., Critchlow, B. and Liu, T. J. (1985). Alcohol and social behavior ii: The helpful drunkard. *Journal of Personality and Social Psychology*, **48** (1), 34–46.

— and Josephs, R. A. (1988). Drinking your troubles away ii: An attention-allocation model of alcohol's effect on psychological stress. *Journal of Abnormal Psychology*, **97** (2), 196–205.

— and — (1990). Alcohol myopia – its prized and dangerous effects. *American Psychologist*, **45** (8), 921–933.

— and Southwick, L. (1985). Alcohol and social behavior i: The psychology of drunken excess. *Journal of Personality and Social Psychology*, **48** (1), 18–34.

Subramanian, S. V., Nandy, S., Irving, M., Gordon, D. and Smith, G. D. (2005). Role of Socioeconomic Markers and State Prohibition Policy in Predicting Alcohol Consumption among Men and Women in India: a Multilevel Statistical Analysis. *Bulletin of the World Health Organization*, **83** (11), 829–836.

Suh, J. J., Pettinati, H. M., Kampman, K. M. and O'Brien, C. P. (2006). The status of disulfiram – a half of a century later. *Journal of Clinical Psychopharmacology*, **2**, 290–302.

Thaler, R. and Benartzi, S. (2004). Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving. *Journal of Political Economy*, **112** (S1), S164–S187.

— and Shefrin, H. M. (1981). An Economic Theory of Self-Control. *Journal of Political Economy*, **89**, 392–406.

Thornton, M. (1991). *The Economics of Prohibition*. University of Utah Press.

USAID (2003). *Poverty Reduction in Uganda: A Background Paper*. Bureau for Policy and Program Coordination Working Paper No. 3.

Volpp, K. G., John, L. K., Troxel, A., Norton, A., Fassbender, L. and Loewenstein, G. (2008). Financial incentive-based approaches for weight-loss: a randomized trial. *JAMA*, **300**, 2631–2637.

Vuchinich, R. E. and Simpson, C. A. (1999). *Delayed-Reward Discounting in Alcohol Abuse*. in: The Economic Analysis of Substance Use and Abuse: an Integration of Econometrics and Behavioral Economics Research.

Wagenaar, A. C., Salois, M. J. and Komro, K. A. (2009). Effect of beverage alcohol price and tax levels on drinking: a meta-analysis of 1003 estimates from 112 studies. *Addiction*, **104**, 179–190.

Wakker, P. P. (2007). Message to referees who want to embark on yet another discussion of the random-lottery incentive system for individual choice. *mimeo*.

Weber, B. J. and Chapman, G. B. (2005). The Combined Effects of Risk and Time on Choice: Does Uncertainty Eliminate the Immediacy Effect? Does Delay Eliminate the Certainty Effect? *Organizational Behavior and Human Decision Processes*, **96**, 104–118.

Weber, M. and Camerer, C. F. (1998). The disposition effect in securities trading: an experimental analysis. *Journal of Economic Behavior and Organization*, **33** (2), 167 – 184.

Wetterling, T., Kanitz, R.-D., Veltrup, C. and Driessen, M. (1994). Clinical predictors of alcohol withdrawal delirium. *Alcoholism: Clinical and Experimental Research*, **18** (5), 1100–1102.

WHO (2001). *The Alcohol Use Disorder Identification Test: Guidelines for Use in Primary Care*. Geneva: World Health Organization.

— (2014). *Global Status Report on Alcohol and Health 2014*. Geneva: World Health Organization.

Woodrow, K. M. and Eltherington, L. G. (1988). Feeling no pain – alcohol as an analgesic. *Pain*, **32**, 159–163.

Chapter 5

Appendix

5.1 Supplementary Figure and Tables

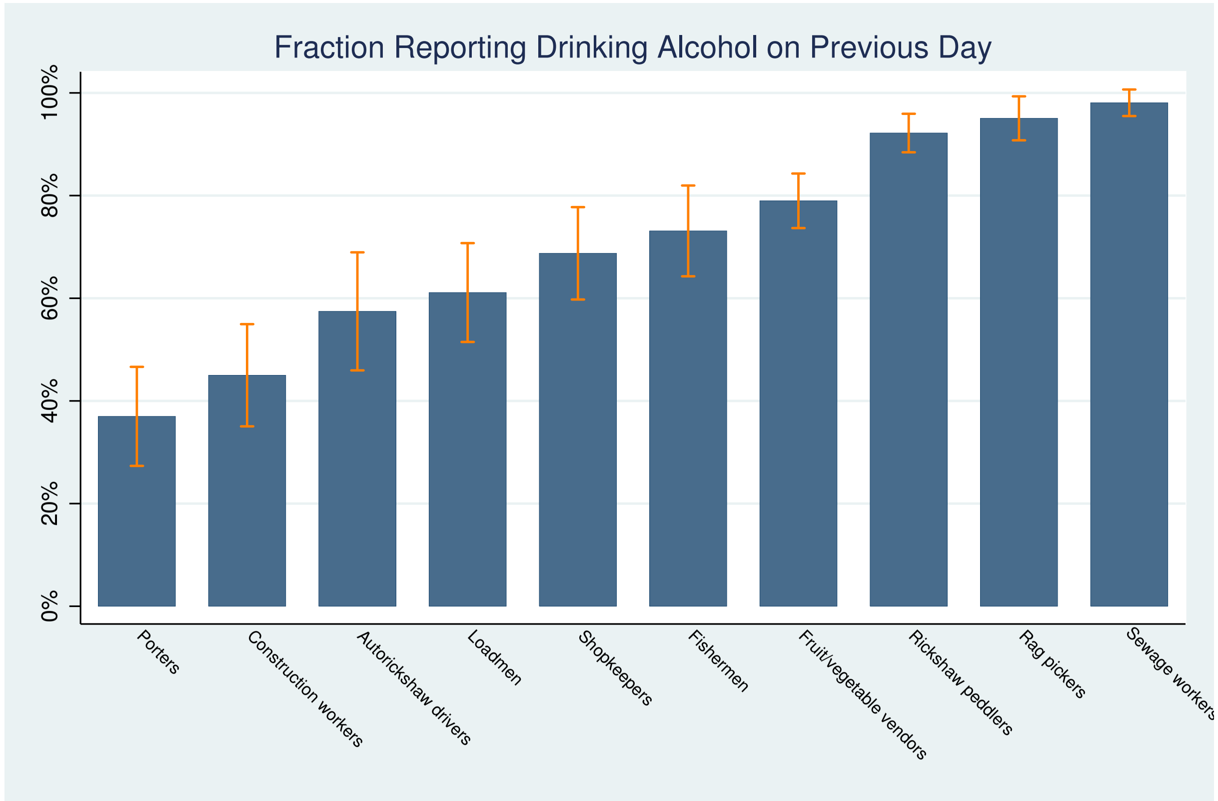

Notes: This figure depicts the prevalence of alcohol consumption among males in ten different low-income professions in Chennai, India, as measured by the fraction of individuals who reported consuming alcohol on the previous day. The underlying data from these figures are from a short survey conducted with a total sample size of 1,227 individuals. The number of individuals surveyed in each profession varies from 75 (auto rickshaw drivers) to 230 (fruit and vegetable vendors). Error bars show 95 percent confidence intervals.

Figure 5.1: *Prevalence of Alcohol Consumption among Low-Income Males in Chennai*

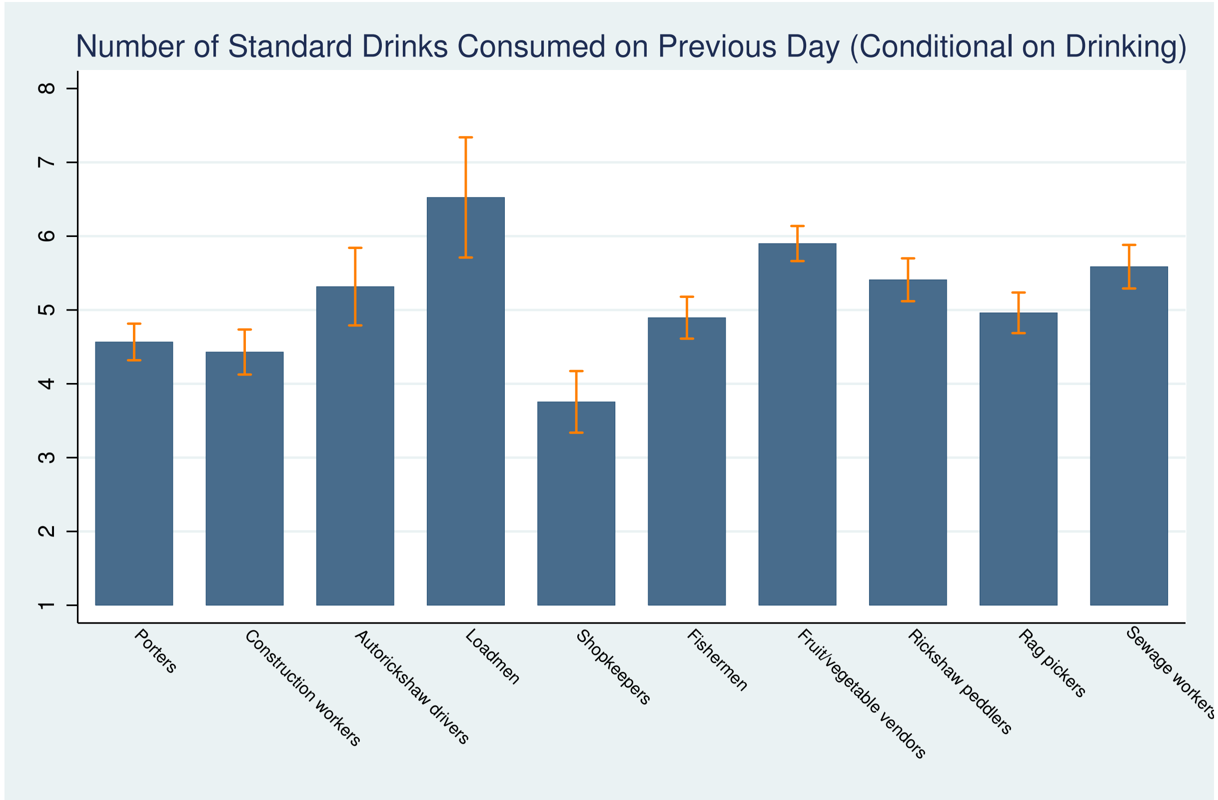

Notes: This figure shows the number of standard drinks consumed on the previous day, conditional on reporting any alcohol consumption on the previous day as described in Figure 5.1. Reported consumption levels are converted into standard drinks according to WHO (2001). A small bottle of beer (330 ml at 5% alcohol), a glass of wine (140 ml at 12% alcohol), or a shot of hard liquor (40 ml at 40% alcohol) each contain about one standard drink. Error bars measure 95 percent confidence intervals.

Figure 5.2: *Prevalence of Alcohol Consumption among Low-Income Males in Chennai*

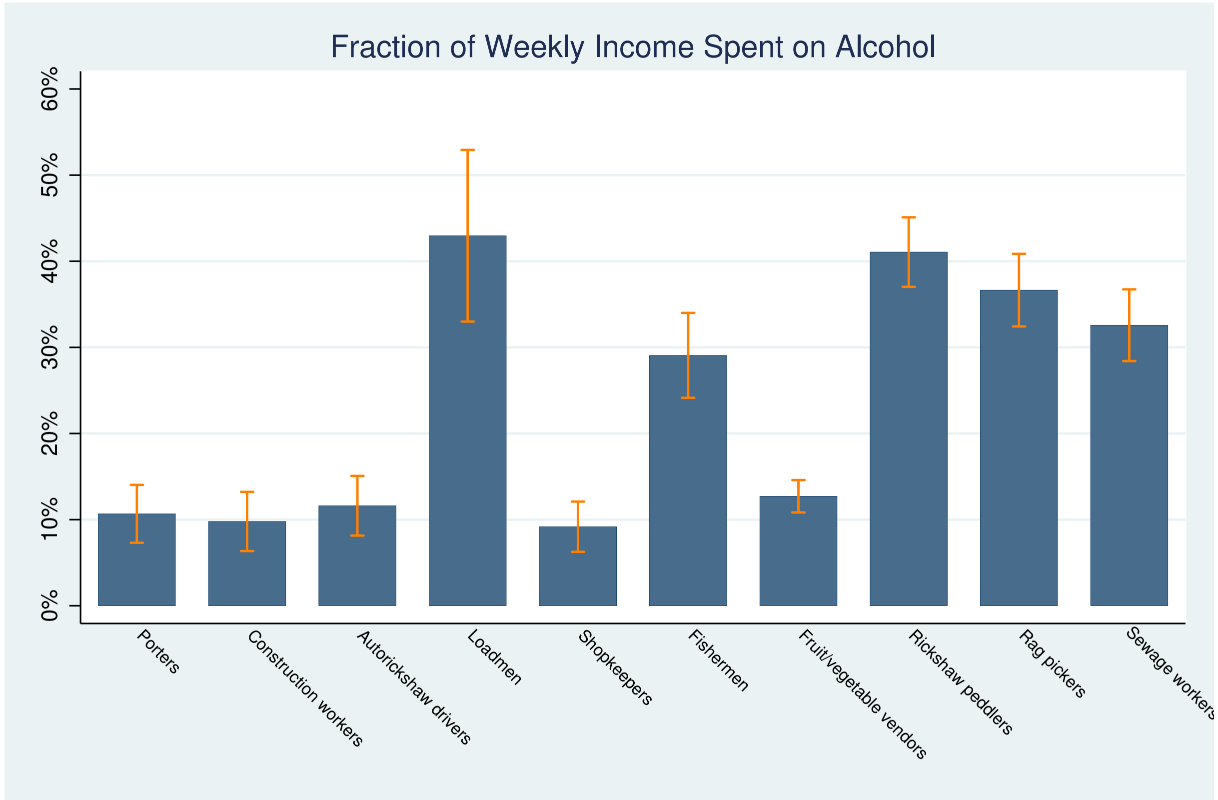

Notes: This figure shows the fraction of weekly income spent on alcohol for the sample described in Figure 5.1. For each individual, the fraction spent on alcohol is calculated by dividing reported weekly alcohol expenditures by reported weekly earnings. Weekly alcohol expenditures are calculated by multiplying the number of days the individual reported consuming alcohol in the previous week times the amount spent on alcohol per drinking day. Weekly earnings are calculated by the number of days worked during the previous week times the amount earned per working day. Error bars measure 95 percent confidence intervals.

Figure 5.3: *Fraction of Weekly Income Spent on Alcohol*

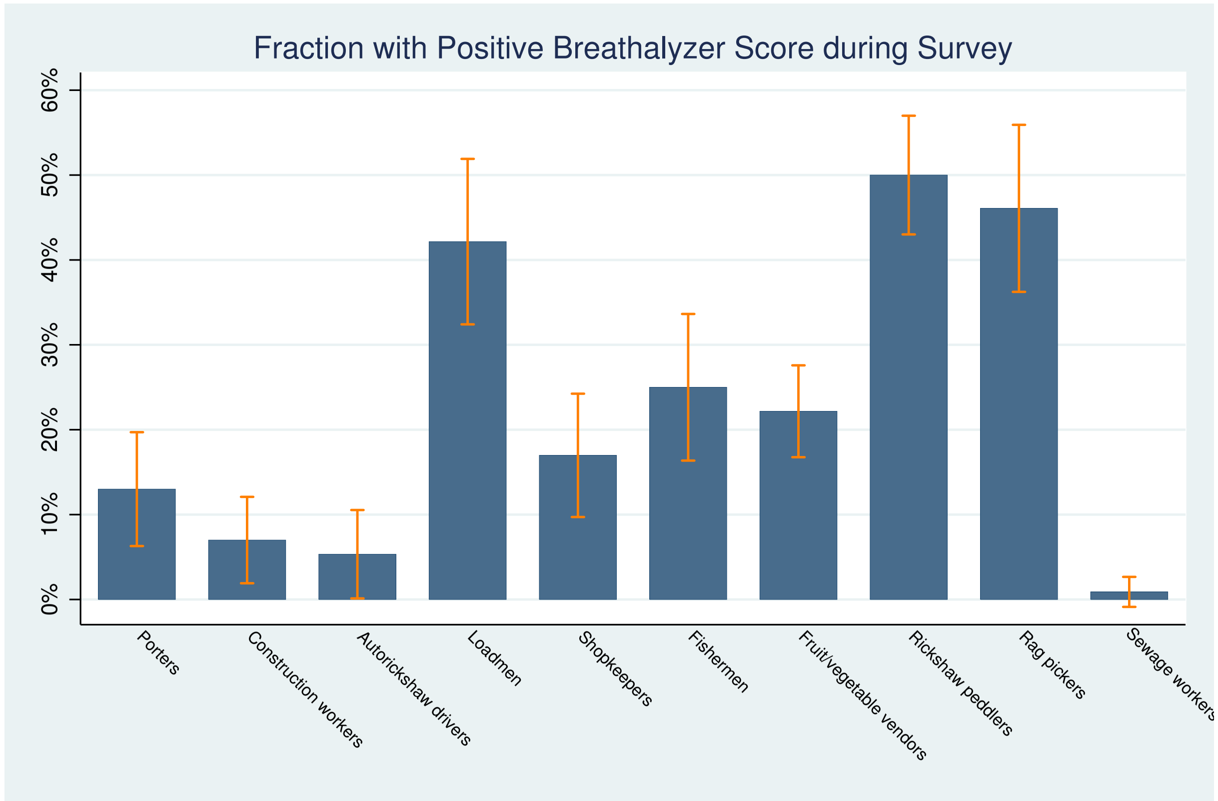

Notes: This figure shows the fraction of individuals who were inebriated at the time of the survey, as measured by having a positive blood alcohol content in a breathalyzer test (BAC > 0). The sample is the same as described in Figure 5.1. All surveys were conducted during the day, i.e. between 8 am and 6 pm. Error bars measure 95 percent confidence intervals.

Figure 5.4: *Fraction with Positive Breathalyzer Score*

(a) *All individuals*

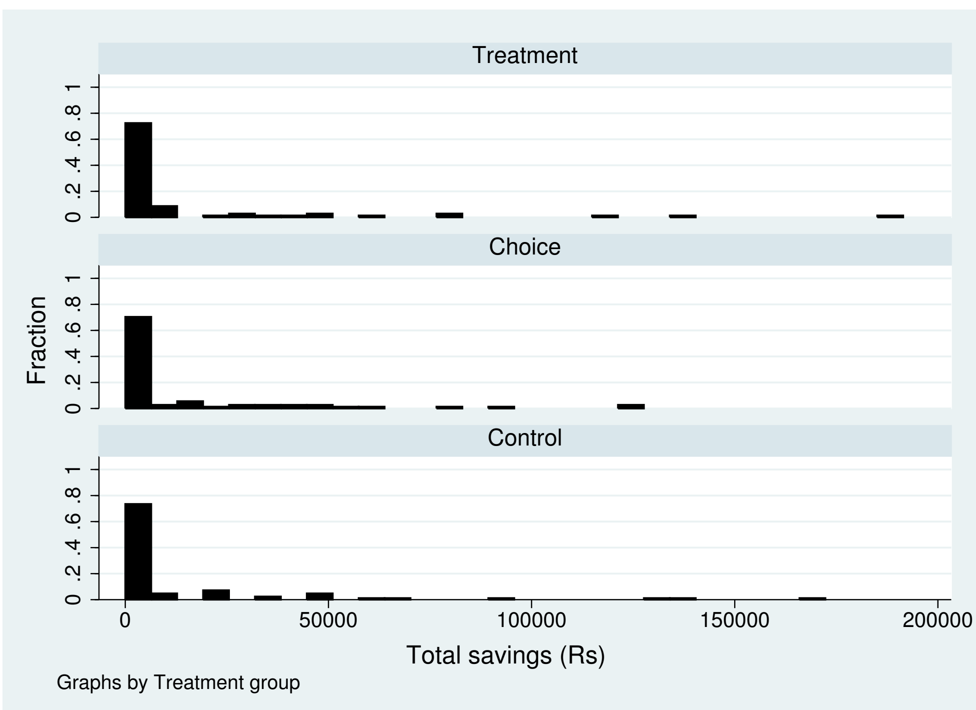

(b) *Only individuals with savings below Rs. 200,000*

Figure 5.5: *Reported Sum of Total Savings by Incentive Treatment Group at Baseline*

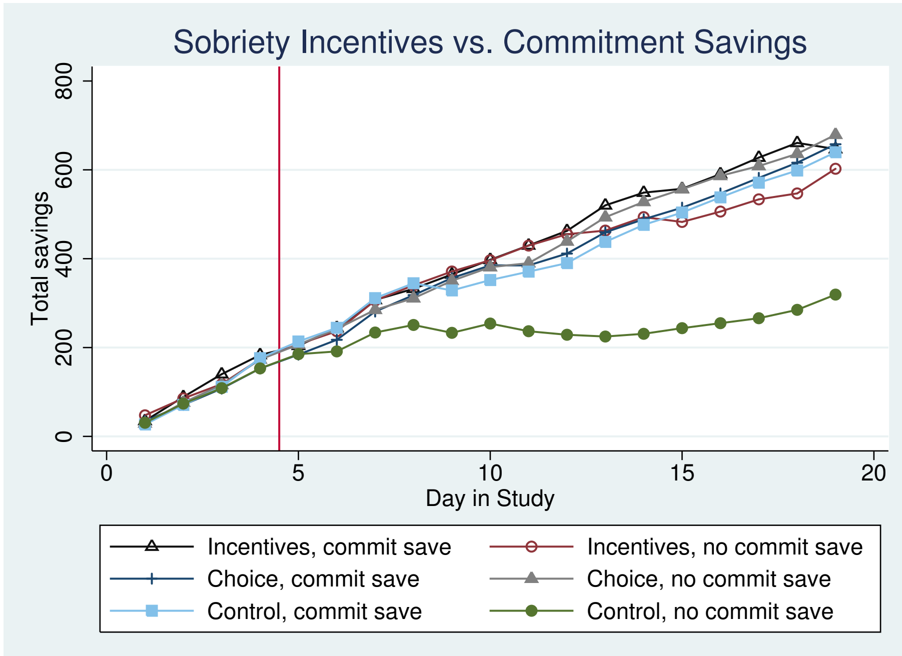

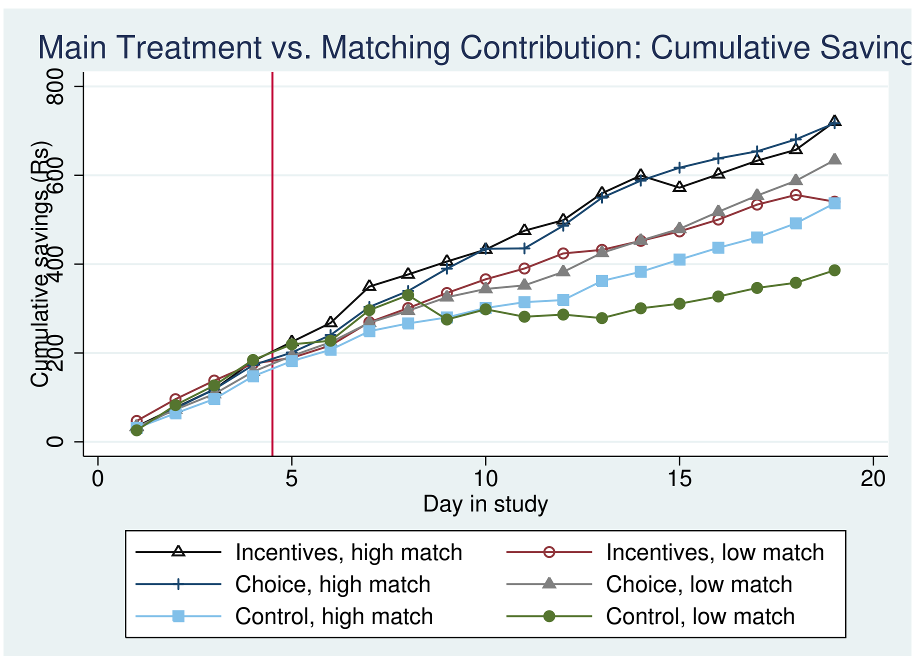

Notes: This figure shows the interaction between the cross-randomized sobriety incentives and savings treatments. The figure is the same as Figure **??**, except for the fact that the two sobriety incentive treatment groups are shown separately rather than pooled (as in Figure **??**).

Figure 5.6: *Interaction between Sobriety Incentives (not pooled) and Savings Treatments*

Table 5.1: *Balance Table for Main Demographs*

	Treatment groups			p value for test of:		
	Control	Incentives	Choice	1=2	1=3	$1 = (2 \cup 3)$
	(1)	(2)	(3)	(4)	(5)	(6)
Age	36.54	35.27	35.08	0.43	0.29	0.30
	(9.96)	(9.92)	(7.40)			
Married	0.82	0.80	0.81	0.80	0.92	0.84
	(0.39)	(0.40)	(0.39)			
Number of children	1.80	1.77	1.80	0.93	0.98	0.97
	(1.19)	(1.55)	(1.19)			
Lives with wife in Chennai	0.73	0.72	0.73	0.82	0.98	0.88
	(0.44)	(0.45)	(0.45)			
Wife earned income during past month	0.24	0.17	0.28	0.27	0.58	0.80
	(0.43)	(0.38)	(0.45)			
Years of education	4.89	5.45	5.49	0.38	0.34	0.28
	(3.93)	(3.95)	(3.92)			
Able to read the newspaper	0.63	0.62	0.63	0.93	1.00	0.96
	(0.49)	(0.49)	(0.49)			
Added 7 plus 9 correctly	0.86	0.77	0.77	0.20	0.19	0.12
	(0.35)	(0.42)	(0.42)			
Multiplied 5 times 7 correctly	0.48	0.41	0.47	0.36	0.85	0.53
	(0.50)	(0.50)	(0.50)			
Distance of home from office (km)	2.64	2.30	2.65	0.20	0.99	0.54
	(2.15)	(1.06)	(1.72)			
Years lived in Chennai	31.57	27.77	29.16	0.04**	0.17	0.05*
	(12.19)	(11.10)	(9.81)			
Reports having ration card	0.65	0.52	0.61	0.11	0.63	0.22
	(0.48)	(0.50)	(0.49)			
Has electricity	0.81	0.68	0.75	0.07*	0.37	0.10
	(0.40)	(0.47)	(0.44)			
Owns TV	0.76	0.59	0.68	0.03**	0.27	0.05**
	(0.43)	(0.50)	(0.47)			
Happiness ladder score (0 to 10)	5.73	5.46	5.76	0.43	0.94	0.68
	(2.14)	(2.08)	(2.11)			

Notes: This table shows balance checks for main demographics across the incentive treatment groups. Columns 1 through 3 show sample means for individuals in the Control Group (1), Incentive Group (2), and the Choice Group (3), respectively. Standard deviations are in parentheses. Columns 4 through 6 show p-values of OLS regressions of each variable on dummies for each treatment group. Columns 4 and 5 shows p-values of tests for equality of means between the Incentive and Choice Groups compared to the Control Group, respectively. Column 6 shows the corresponding p-values for comparisons between the Control Group and the Incentive and Choice Groups combined.

Table 5.2: *Balance Table for Work and Savings*

	Treatment groups			p value for test of:		
	Control	Incentives	Choice	1=2	1=3	$1 = (2 \cup 3)$
	(1)	(2)	(3)	(4)	(5)	(6)
Years worked as a rickshaw puller	14.06	12.49	12.81	0.29	0.34	0.25
	(9.53)	(8.78)	(6.73)			
# of days worked last week	5.41	5.18	5.43	0.36	0.94	0.60
	(1.35)	(1.65)	(1.39)			
Has regular employment arrangement	0.47	0.52	0.47	0.53	0.97	0.74
	(0.50)	(0.50)	(0.50)			
Owns rickshaw	0.17	0.25	0.28	0.20	$0.10^{\star}$	$0.08^{\star}$
	(0.38)	(0.44)	(0.45)			
Says 'no money' reason for not owning rickshaw	0.61	0.65	0.59	0.67	0.72	0.98
	(0.49)	(0.48)	(0.50)			
Reported labor income in Phase 1 (Rs/day)	291.86	301.08	273.94	0.69	0.39	0.79
	(119.97)	(160.54)	(138.33)			
Total savings (Rs)	13261	23903	38184	0.22	0.13	$0.07^{\star}$
	(31197)	(67739)	(139224)			
Total borrowings (Rs)	11711	5648	7913	0.11	0.36	0.18
	(29606)	(15762)	(22253)			
Savings at study office in Phase 1 (Rs/day)	40.98	44.67	41.04	0.62	0.99	0.77
	(41.93)	(49.28)	(48.25)			

Notes: This table shows balance checks for work- and savings-related variables across the incentive treatment groups. Columns 1 through 3 show sample means for individuals in the Control Group (1), Incentive Group (2), and the Choice Group (3), respectively. Standard deviations are in parentheses. Columns 4 through 6 show p-values of OLS regressions of each variable on dummies for each treatment group. Columns 4 and 5 shows p-values of tests for equality of means between the Incentive and Choice Groups compared to the Control Group, respectively. Column 6 shows the corresponding p-values for comparisons between the Control Group and the Incentive and Choice Groups combined.

Table 5.3: *Balance Table for Alcohol Consumption*

	Treatment groups			p value for test of:		
	Control	Incentives	Choice	1=2	1=3	$1 = (2 \cup 3)$
	(1)	(2)	(3)	(4)	(5)	(6)
Years drinking alcohol	12.89	11.68	12.86	0.42	0.99	0.65
	(10.02)	(8.42)	(9.03)			
Number of drinking days per week	6.72	6.83	6.68	0.39	0.70	0.77
	(0.80)	(0.76)	(0.60)			
Drinks usually hard liquor ($\geq$ 40 % alcohol)	0.99	1.00	0.99	0.32	0.94	0.71
	(0.11)	(0.00)	(0.12)			
Alcohol expenditures in Phase 1 (Rs/day)	91.95	87.09	81.92	0.39	0.07*	0.12
	(37.03)	(32.48)	(32.98)			
# of standard drinks per day in Phase 1	6.17	5.71	5.80	0.21	0.31	0.19
	(2.29)	(2.17)	(2.18)			
# of std drinks during day in Phase 1	2.13	2.45	2.40	0.38	0.42	0.31
	(2.01)	(2.48)	(2.10)			
Baseline fraction sober	0.49	0.45	0.43	0.48	0.30	0.30
	(0.40)	(0.43)	(0.41)			
Alcohol Use Disorders Identification Test score	14.61	13.94	14.69	0.44	0.92	0.67
	(4.32)	(6.16)	(4.98)			
Drinks usually alone	0.87	0.82	0.85	0.40	0.80	0.51
	(0.34)	(0.39)	(0.36)			
Reports life would be better if liquor stores closed	0.84	0.80	0.77	0.52	0.27	0.29
	(0.37)	(0.40)	(0.42)			
In favor of prohibition	0.81	0.77	0.84	0.62	0.59	0.99
	(0.40)	(0.42)	(0.37)			
Would increase liquor prices	0.07	0.14	0.12	0.18	0.32	0.15
	(0.26)	(0.35)	(0.33)			

Notes: This table shows balance checks for alcohol-related variables across the incentive treatment groups. Columns 1 through 3 show sample means for individuals in the Control Group (1), Incentive Group (2), and the Choice Group (3), respectively. Standard deviations are in parentheses. Columns 4 through 6 show p-values of OLS regressions of each variable on dummies for each treatment group. Columns 4 and 5 shows p-values of tests for equality of means between the Incentive and Choice Groups compared to the Control Group, respectively. Column 6 shows the corresponding p-values for comparisons between the Control Group and the Incentive and Choice Groups combined.

Table 5.4: *Effect of Sobriety Incentives on Family Resources*

VARIABLES	(1) Wife	(2) Wife	(3) Wife	(4) Other	(5) Other	(6) Other	(7) Total	(8) Total	(9) Total
Incentives	19.89	10.93		-10.21	-9.90		9.68	1.03	
	(19.068)	(16.914)		(7.482)	(7.591)		(17.914)	(14.942)	
Choice	16.03	21.94		-9.85	-7.65		6.18	14.30	
	(20.117)	(16.590)		(8.441)	(8.379)		(19.924)	(15.585)	
Pooled alcohol treat			16.94			-8.67			8.28
			(13.969)			(7.115)			(12.699)
Observations	2,991	2,991	2,991	2,991	2,991	2,991	2,991	2,991	2,991
R-squared	0.002	0.127	0.126	0.006	0.082	0.082	0.000	0.144	0.143
Baseline survey controls	NO	YES	YES	NO	YES	YES	NO	YES	YES
Phase 1 controls	NO	YES	YES	NO	YES	YES	NO	YES	YES
Control group mean	148.7	148.7	148.7	25.13	25.13	25.13	173.9	173.9	173.9

Notes: This table shows the impact of the two sobriety incentive treatments on family resources.

1. All regressions use data from day 5 (the first day of sobriety incentives) through day 19 (the last day of sobriety incentives) of the study.

2. The outcome variables are (i) money given to the wife (Rs./day; always zero for unmarried individuals) (ii) other family expenses (the sum of money given to other family members and direct household expenses), and (iii) total family resources (i.e. the sum of (i) and (ii)).

3. The data used in the regressions is from retrospective surveys on the consecutive study days, during which individuals are asked about each of the above variables on the previous day. In addition, if individuals missed a day or two (and on Mondays), they were asked about the same outcomes two or three days ago, respectively.

4. Standard errors are in parentheses, clustered by individual. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively. Phase 1 and baseline survey controls are the same as in the above tables.

Table 5.5: *Expenses on Food, Coffee & Tea, and Tobacco & Paan*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
VARIABLES	Food	Food	Food	Cof/Tea	Cof/Tea	Cof/Tea	Tob/Paan	Tob/Paan	Tob/Paan
Incentives	3.03	5.83		0.02	0.38		2.13	2.58	
	(6.609)	(6.126)		(1.013)	(1.015)		(1.818)	(1.732)	
Choice	-3.45	3.05		-0.14	0.02		-2.95*	-2.35	
	(5.907)	(5.771)		(1.011)	(0.938)		(1.557)	(1.545)	
Pooled alcohol treat			4.34			0.18			-0.06
			(5.085)			(0.840)			(1.409)
Observations	1,034	1,034	1,034	1,047	1,047	1,047	1,047	1,047	1,047
R-squared	0.003	0.154	0.153	0.000	0.117	0.117	0.026	0.086	0.065
Baseline survey controls	NO	YES	YES	NO	YES	YES	NO	YES	YES
Phase 1 controls	NO	YES	YES	NO	YES	YES	NO	YES	YES
Control group mean	50.93	50.93	50.93	4.522	4.522	4.522	10.52	10.52	10.52

Notes: This table shows the impact of the two sobriety incentive treatments on other expenditures.

1. All regressions use data from day 5 (the first day of sobriety incentives) through day 19 (the last day of sobriety incentives) of the study. Individuals were only asked about the below variables every third day (the timing was unannounced).

2. The outcome variables are (i) money given to the wife (Rs./day; always zero for unmarried individuals) (ii) other family expenses (the sum of money given to other family members and direct household expenses), and (iii) total family resources (i.e. the sum of (i) and (ii)).

3. Standard errors are in parentheses, clustered by individual. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively. Phase 1 and baseline survey controls are the same as in the above tables.

Table 5.6: *Attrition and Inconsistencies of Choices*

	Choice Group			Incentive Group	Control Group
	Week 1	Week 2	Week 3	Week 3	Week 3
Present & consistent (%)	88.0	89.3	88.0	90.1	86.7
Absent (%)	5.3	6.7	6.7	5.6	6.0
Inconsistent (%)	6.7	4.0	5.3	4.2	7.2

Notes: This table shows the fraction of individuals who were present and made consistent choices by treatment group and week of study. During a given choice session, an individual chose inconsistently if he chose Option B for the unconditional amount Y_1, but Option A for the unconditional amount Y_2 with $Y_2 > Y_1$. For instance, his choices are inconsistent if he preferred Option B in Choice 1, but not in Choice 3.

Table 5.7: *Summary of Choices in Choice Group Over Time*

	Option A		Option B	Percent choosing A		
Choice	BAC > 0	BAC $= 0$	regardless of BAC	Week 1	Week 2	Week 3
(1)	Rs. 60	Rs. 120	Rs. 90	60.0	62.7	57.3
(2)	Rs. 60	Rs. 120	Rs. 120	46.7	52.0	44.0
(3)	Rs. 60	Rs. 120	Rs. 150	30.7	33.3	40.0

Notes: This table shows the fraction of individuals among the Choice Group who preferred incentives over unconditional amounts for each of the choices by week of study. Individuals who were either absent or did not choose consistently are counted as *not* preferring incentives.

5.2 More Detailed Model Solution

5.2.1 Solution for the Case of Isoelastic Utility

This section provides the solution of the model described in section 5.1 for the commonly used case of isolelastic utility.

No commitment savings. Equations (1.7) and (1.9) become

$$c_2^{-\gamma} = \beta(1+M)c_3^{-\gamma} \tag{5.1}$$

$$c_1^{-\gamma} = \left[\beta\frac{dc_2}{dY_2} + \left(1 - \frac{dc_2}{dY_2}\right)\right] c_2^{-\gamma} \tag{5.2}$$

Using (1.8) and (5.1), we can solve for c_3 and c_2 as functions of Y_2:

$$c_3 = \left(\frac{1+M}{1+\theta}\right) Y_2 \quad \text{and} \quad c_2 = \left(\frac{\theta}{1+\theta}\right) Y_2. \tag{5.3}$$

where $\theta \equiv (\beta(1+M))^{\frac{-1}{\gamma}}(1+M)$. This implies $\frac{dc_2}{Y_2} = \frac{\theta}{1+\theta}$ and, using (5.2), we get

$$c_1 = \left(\frac{1+\beta\theta}{1+\theta}\right)^{\frac{-1}{\gamma}} c_2. \tag{5.4}$$

Using the budget constraint and rewriting (5.1) to $c_2 = \frac{\theta}{1+M}c_3$, this yields

$$c_3^{\mathrm{NC}} = \frac{Y(1+M)}{1+\theta+\theta\left[\frac{1+\beta\theta}{1+\theta}\right]^{\frac{-1}{\gamma}}}. \tag{5.5}$$

Commitment savings. Equations (1.10) and (1.11) become

$$c_2 = (1+M)^{\frac{-1}{\gamma}} c_3, \tag{5.6}$$

$$c_1 = \beta^{\frac{-1}{\gamma}} c_2 = \left(\frac{\theta}{1+M}\right) c_3. \tag{5.7}$$

Using the budget constraint (1.12), this implies

$$c_3^{\mathrm{C}} = \frac{Y(1+M)}{1+\theta+(1+M)^{1-\frac{1}{\gamma}}}. \tag{5.8}$$

5.2.2 A Special Case: Log Utility

This section considers a special case of log utility ($\gamma = 1$), i.e. $u(c_t) = \log(c_t)$.

No commitment savings. Equations (1.7) and (1.9) become

$$c_3 = \beta(1+M)c_2 \tag{5.9}$$
$$c_2 = \left[\beta\frac{dc_2}{dY_2} + \left(1 - \frac{dc_2}{dY_2}\right)\right]c_1 \tag{5.10}$$

Using $c_3 = (Y_2 - c_2)(1+M)$, we use (5.9) to solve for c_3 and c_2 as functions of Y_2:

$$c_2 = \frac{1}{1+\beta}Y_2 \quad \text{and} \quad c_3 = \frac{\beta(1+M)}{1+\beta}Y_2 \tag{5.11}$$

This implies $\frac{dc_2}{dY_2} = \frac{1}{1+\beta}$ and, hence $c_2 = \frac{2\beta}{1+\beta}c_1$ and $c_3 = (1+M)\frac{2\beta^2}{1+\beta}c_1$. Hence, we get

$$c_1 = Y - c_2 - \frac{c_3}{1+M} = Y - \frac{2\beta}{1+\beta}c_1 - \frac{2\beta^2}{1+\beta}c_1 = \frac{Y}{1 + \frac{2\beta}{1+\beta} + \frac{2\beta^2}{1+\beta}} \tag{5.12}$$

This implies $c_3^{\text{NC}} = \frac{2\beta^2}{1+3\beta+2\beta^2}Y(1+M)$.

Commitment savings. Consider now the solution for the commitment savings case. Equations (1.10) and (1.11) become

$$c_2 = \beta c_1 \qquad c_3 = (1+M)c_2 \tag{5.13}$$

Using the budget constraint (1.12), this yields

$$c_3^C = (Y - c_1 - c_2)(1+M) \tag{5.14}$$
$$= Y(1+M) - \frac{c_3}{\beta} - c_3 \tag{5.15}$$
$$= \frac{\beta}{1+2\beta}Y(1+M) \tag{5.16}$$

Comparing the two solutions yields

$$\Delta \equiv c_3^{\text{C}} - c_3^{\text{NC}} = \left[\frac{\beta(1-\beta)}{(1+2\beta)(1+\beta)}\right]Y(1+M) \tag{5.17}$$

Taking the derivative of the expression in brackets with respect to β yields

$$\frac{\partial[\cdot]}{\partial\beta} = \frac{1 - 2\beta - 5\beta^2}{(1 + 3\beta + 2\beta^2)^2} \tag{5.18}$$

This expression is positive for $0 \leq \beta \approx 0.29$ and negative for $0.29 \approx \beta \leq 1$.