Center for Financial Security

UNIVERSITY OF WISCONSIN-MADISON

Time-Inconsistency and Savings:

Experimental Evidence from Low-Income Tax Filers

Damon Jones*, *University of Chicago*
Aprajit Mahajan, *Stanford University*

Center for Financial Security
Working Paper 2011-CFS.6

October 2011

*Contact Author
damonjones@uchicago.edu

Time-Inconsistency and Savings: Experimental Evidence from Low-Income Tax Filers

Damon Jones* Aprajit Mahajan †

August 2011

Abstract

Preliminary – Please do not cite.

We describe a forthcoming field experiment designed to test theories of time-inconsistency, namely a $\beta - \delta$ model of quasi-hyperbolic discounting. The experiment takes place within the context of a savings decision made by low-income tax filers. These tax filers are often offered the option to place their income tax refund in an illiquid savings vehicle at the time that taxes are filed. In addition to testing for time-inconsistency, the results of our study can be used to improve the design of savings incentives for this population, as well as help evaluate the welfare effects of these savings programs. We present evidence from a small-scale, pilot version of the experiment. The pilot study results in point estimates for β and δ of 0.34 and 1.08 over an 8-month period, respectively. The small scale of the pilot limits the precision of these estimates. We discuss refinements to the experimental study in light of the pilot experience.

JEL Classification: D14, H24, K34

*University of Chicago, The Harris School and NBER, damonjones@uchicago.edu.

†Stanford University, amahajan@stanford.edu.

1 Introduction

A large body of economic research focuses on the determinants of household savings. One important concern in this literature is whether many households are saving adequately. For instance, in a survey administered by Choi et al. (2004) two-thirds of respondents view their current savings rate as being "too low." In the short-term, inadequate savings renders households vulnerable to negative income shocks (e.g. unforeseen health expenses). In the long-run, households may not be able to achieve a desired standard of living in retirement. At the macro level, a low national savings rate translates into lower levels of investment and thus lower rates of economic growth.

This project will use a set of specifically designed experimental interventions in conjunction with survey data collection to examine some of the leading explanations for household saving behavior. In particular, it will examine the relevance of self-control related problems in explaining observed saving behavior. The project will develop formal tests for the presence of self-control problems (such as impatience) by offering respondents monetary prizes for achieving saving goals. The main intervention offers respondents one of two types of financial rewards, immediate or delayed, if they save or commit to saving their income tax refund for a specific length of time. This will be combined with a second intervention, a soft-commitment device that offers households a financial reward if their actual saving decision at the time they receive a refund corresponds to their stated decision three months earlier. Together these interventions will allow the researcher to identify the parameters of a quasi-hyperbolic discounting model (which nests the standard model of time-consistent preferences).

Using the interventions and survey data the project can convincingly answer two key questions using field experimental evidence and shed light on a third, related policy question:

1. To what extent does impatience play a role in saving behavior for low-income households?
2. Can a better understanding of impatience help improve the design of policies aimed at encouraging savings among low-income households?
3. Do the policies under consideration, which encourage savings among low-income households at tax time, appear to be "nudging" agents in an optimal direction?

The answers to these questions have broad implications for understanding how U.S. households maintain buffer stock savings in the short term and save for old age. The answers will also inform the design of policies aimed at increasing their welfare, especially in the case of the lower-income households that will be a focus of the study. Concretely, the results will shed light on the design of financial tools that help lower income households smooth consumption over lumpy annual payments, such as the annual receipt of the Earned Income Tax Credit (EITC), and to reach desired savings goals.

Below we outline the design of a forthcoming field experiment. We also provide results from the pilot version of this study. Our preferred point estimates for β and δ over an 8-month horizon from a quasi-hyperbolic discounting model are 0.34 and 1.08 respectively. These translate into a one-year discount factor of 0.38, or equivalently, a one-year discount rate of 164%. Our results are limited by the sample size and other idiosyncratic factors associated with the pilot study and are therefore only suggestive. Our aim is to replicate the pilot study with larger sample size, generating more precise estimates.

2 Related Literature and Ongoing Work

Previous research has demonstrated that among tax filers there exists a demand for savings options (Beverly et al. (2006)), a positive effect of a match on savings (Duflo et al. (2006)) and a demand for illiquid savings options, such as a savings bond (Tufano (2008)). The PIs plan to expand on these studies by applying lessons from a second set of studies that demonstrate a positive effect of pre-commitment on savings. Commitment devices have been studied in the context of defined contribution retirement accounts (Bernatzi and Thaler (2004)), among consumers in the Philippines (Ashraf et al. (2006)) and for health interventions (Tarozzi et al. (2009), Mahajan and Tarozzi (2010)). The project's contribution includes examining the effects of commitment mechanisms in the context of income tax refund-based saving in the US. Furthermore, the experiment is designed to test theories of self-control and savings found in models of quasi-hyperbolic discounting (Laibson (1997), O'Donoghue and Rabin (1999)). Though their exists laboratory evidence in support of these models, ours will be one of very few field-experimental studies that yields direct empirical tests of these models. In particular, we seek quantitative estimates of β and δ, which compliment previous field studies that focus on qualitative evidence of time-inconsistency.

In ongoing work, Gine et al. (2010) implement a field experiment in which quantitative measures of time preference and time-inconsistency are correlated with the likelihood of revising one's savings plans. Whereas their study applies to rural households in a developing economy, Malawi, our study focuses on low-income households in an urban setting in the US. Meier and Sprenger (2010) elicit time-preference parameters from similar sample as ours: low-income tax filers in the US. These parameters are then correlated with credit scores and debt levels. In these studies time-preferences are elicited by incentivized choice experiments: either the "convex time budget" or "multiple price list" method.[1] They may be regarded as "artefactual field experiments" in the parlance of Harrison and List (2004). Our study contributes to this line of work by measuring time-preferences in the course of decisions that are routine to the agents, i.e. the decision to save one's income tax refund. Thus, we aim for a context that is akin to a "natural field experiment" (Harrison and List (2004)).[2]

The current project builds upon previous work by the PIs that explore the effects of commitment mechanisms on compliance with health-protecting behavior (Tarozzi et al. (2009)) and current work that develops econometric tools to identify the parameters of a quasi-hyperbolic discounting model in this context (Mahajan and Tarozzi (2010)). Furthermore, it directly relates to current work by the PIs which shows that default EITC rules generate lumpy pay schedules (Jones (Forthcoming)) and that this phenomenon may have implications for recipients' savings, borrowing and consumption smoothing (Jones (2010a)). The current project is at the nexus of these previous studies and is a natural extension of this work. Finally, while the full-scale implementation of this project will take place between 2011 and 2012, its design is informed by a smaller scale pilot study that the PIs have implemented between late 2010 and early 2011.

3 Study Design

This section will outline the study design which is then analyzed formally in Section 4 and implemented in a pilot study in Section 5. To fix ideas, consider an individual who files her taxes in February and chooses between receiving her income tax refund immediately or placing it in an illiquid savings account until October. Suppose further that the individual is offered a savings match as an incentive. Finally, suppose her match is increased by a small amount, and this additional incentive can either be received immediately (February) or is delayed until money is withdrawn from the savings account (October). Conversations with tax planners suggest that such individuals focus heavily on the prospect of receiving the refund as soon as possible and pay far less attention to the need to put some of the refund aside for future events or the opportunity to earn a match on these savings, which suggests they may have present-biased preferences. As compared to the baseline take-up rate of this account, the relative effect of the

[1] See Andreoni and Sprenger (2010) for a discussion and comparison of the two methods and see Frederick et al. (2002) for an extensive review of the time-preference literature.

[2] Hausman (1979) and Wanrner and Pleeter (2001) estimate time preference parameters in "natural" decisionmaking contexts, but do not explicitly test for time-inconsistent preferences. Harrison et al. (2005) test for dynamic inconsistency using an "artefactual" field experiment.

additional incentives, immediate or delayed, can give one a sense of the time discount between the two periods.[3]

Alternatively, suppose that the individual were asked at some earlier date (e.g. in December) to choose between receiving the refund in February or saving it until some future date (e.g. October) and thereby earning the match on the savings. Again, suppose that the incentive to save is increased by an amount, either delivered early in February or late in October. Standard theories of present-biased preferences predict that when asked at an earlier date, the respondent may be more inclined to choose the savings option. In addition, from this pre-tax season vantage point, the relative effect of the early and late incentives will be less pronounced. Intuitively, when making a decision in December, the tax filer already has to wait two months to receive the refund. Waiting an additional couple of months to receive a larger refund may not seem so bad. The individual is more aware of the benefits of saving vis-a-vis the cost of putting the money aside when both occur far enough in the future. The model also suggests that if the tax filer can make a binding, early decision in December to save in February, she may do so as a means of self-control. However, if the commitment device not completely binding, a sophisticated tax filer may avoid it, realistically believing that commitment will ultimately be undone. The use of a partial commitment mechanism will be key in allowing the researcher to infer preferences in December over outcomes in February and October.

The experimental design follows the outline sketched above. First, using information from the prior year's tax return, participants will be randomized into two treatment groups (the "commitment" and the "non-commitment" group). Within each group, clients will be further randomized into three sub-groups based on the timing of a monetary incentive (either baseline, early/immediate incentives or late/delayed incentives).[4] In the December preceding the tax season, the PIs will elicit a soft-commitment savings decision from members of the commitment treatment groups. After individuals have made a choice, they will be reminded of their prior decision when the actual tax season arrives and will receive a monetary reward if they honor their prior commitment.[5] Members of the non-commitment group will not be offered the soft-commitment device. However, at the tax site, members of both treatment groups will have the option of saving their income tax refund, and will receive a match provided that the savings remains in an account for a designated period of time. Finally, there are additional incentives to pre-commit or save, which are early/immediate or late/delayed, i.e. delivered in February or October, respectively. The study results in six different groups of participants: those with commitment and non-commitment, and within each of those two groups, those with an early/immediate incentive, those with a late/delayed incentive and a set of baseline group members. A theoretical model below will show how the different groups help in identifying key parameters in the model.

3.1 Sample

The project will proceed in two phases. In the first currently on-going phase the PIs are running a pilot study to help guide the final research design, which will be implemented in 2011-2012. For the pilot the sampling frame is a pool of previous clients from a Volunteer Income Tax Association (VITA) organization[6] – The Financial Clinic in New York City. We have chosen to sample from clients with tax refunds greater than $300.[7] In our pilot study, we reached out to 833 clients, who were randomly assigned to one of six treatment groups. Given the results of our pilot, we estimate a sample size of 1,238 for our full-scale study. We will likewise draw this sample from prior clients

[3] The use of monetary incentives to identify discount factors in part relies on an assumption of borrowing constraints. This is perhaps reasonable for the low-income EITC recipients in our study and can additionally be confirmed using survey data among participants.

[4] To emphasize the differences between treatment arms, we deem the February incentive the "early" incentive and the October incentive the "late" incentive when referring to the commitment group. However, within the non-commitment group, we use the terms "immediate" and "delayed" to describe the timing of incentives.

[5] This allows us to distinguish the soft-commitment from "cheap talk."

[6] Third-party tax preparation is quite common among low-income households, particularly those eligible for various refundable credits (which are the object of many saving promotion policies). For instance, IRS public use data suggests that 70% of EITC recipients in 2006 used a tax preparer

[7] Analysis of 2008 and 2009 tax data at this tax preparer indicated that past refunds are highly predictive of future refunds. Note that over 80% of US tax filers receive income-tax refunds with the share being even higher for low-income households (IRS public use data). In our pilot, the average refund size was about $2,095 for Tax Year (TY) 2009.

and walk-ins at the tax site.[8]

The population we are sampling from, while not representative of the US population, is of direct interest to policy makers. There are various state sponsored initiatives aimed at encouraging low-income tax filers to save[9] and our intervention is well suited to examine issues surrounding the take-up of at least one of these programs. In addition, we are not concerned that our estimates are biased due to the refund restrictions we have made. In particular, recent research by the PI (Jones (Forthcoming)) demonstrates that the income tax refund is the result of largely passive behavior on behalf of the tax filer, especially in the case of low-income tax filers that receive the EITC.[10]

3.2 Experimental Arms

We will assign each client to one of six potential groups. In particular, the client will belong to either the commitment or non-commitment treatment, and within each treatment to either the early/immediate match, the late/delayed match or the baseline group.

1. Tax filers in the non-commitment group will be contacted in the December preceding the tax season.[11] They will be invited to participate in the first wave of a longitudinal survey.

    (a) The participant will also be notified of a savings opportunity: the SaveUp Account. When they return to file their taxes, they will be offered the incentive of a savings match. In order to keep the match, the tax filer must keep her savings amount in an illiquid account for 8 months. In return for saving an amount of \$300 or greater, the tax filer will receive a 50% match for each dollar saved between \$300 and \$1,050.[12] This is the baseline group (Treatment Group 6). In addition, there will be two additional sub-groups who are offered either an "immediate" incentive of \$50 (Treatment Group 4), given at the time of filing taxes, or a "delayed" incentive of \$50 received after the 8 month saving period (Treatment Group 5), in return for saving.[13] A treatment member of the non-commitment group has the opportunity to receive up to \$425 in savings incentives. A baseline member of the non-commitment group can receive up to \$375 in savings incentives.

    (b) Survey data will be collected during the baseline interview, at the time of tax filing and 8 months after taxes are filed. The survey will measure outcomes that include level of savings, expenditures, income and

[8] Please refer to http://home.uchicago.edu/~j1s/Jones_Mahajan_Appendix_NSF.pdf for a detailed power calculation. One possible concern is that the total amount of monetary transfers, and by extension our total budget, depends on actual take-up and is therefore not certain ex ante. Fortunately, we can monitor take-up and therefore the maximum possible transfers in real time. To the extent that we experience higher than expected take-up, we can remain within our budget by scaling back the sample size. Moreover, our power calculations make clear that our required sample size is inversely related to the treatment effect. Thus, this reduction in sample size in response to underestimated treatment effects does not necessarily come at the expense of statistical power.

[9] Examples of this include: numerous changes to the tax filing process including the ability to have refunds deposited directly into savings accounts or invested in savings bonds, a national pilot to offer matched savings to low-income filers (Save USA) and potentially the Savers Credit, though this last program targets a broader range of incomes.

[10] In an appendix of that paper, the author shows that the high level of overwithholding cannot be fully explained by a preference for forced savings. Other work by the PI (Jones (2010b)) suggests that it's safe to assume that the Advance EITC and its low take-up can be safely ignored due to dramatically low awareness this program.

[11] For the pilot, we are also enrolling a set of non-commitment clients directly in February at tax-preparation time. Comparing outcomes for this group with outcomes for the non-commitment group members contacted in December should help us explore an alternative, low cost means of sample recruitment (we thank a previous reviewer for this suggestion).

[12] The match is generous, but in fact less generous than pre-existing programs offered by e.g. New York state. The Save NYC account offers a 50% match on every dollar, until \$1,000 (see http://www.nyc.gov/html/ofe/html/poverty/save.shtml). Furthermore, this program is currently being piloted in several cities, demonstrating scalability. To date, however, take-up results have been quite disappointing – less than 5%. Our contribution would be to (1) identify plausible causes for extremely low take-up for apparently very generous accounts and (2) identify, conditional on offering such an attractive incentive, ways of encouraging savings at lower costs. That is, can one generate an equal amount of savings with a lower match by exploiting time-inconsistency? Importantly, the Save NYC account is not simultaneously offered to our study participants.

[13] Even though the "immediate" match is given up front, it is forfeited by deducting the match amount from the savings account in the event of an early withdrawal.

measures of credit constraints, income volatility and relevant demographic information.[14]

2. Tax filers in the commitment group will be contacted in the December preceding the tax season and invited to participate in a longitudinal survey as well. In addition, they will be offered the soft-commitment option, which will be presented as follows:

   (a) In December, the individual will be asked to make a soft-commitment regarding a savings account: the SaveUpFront Account. They have the option of pre-committing to save or pre-committing to not save. The savings account is similar to before, with a minimum deposit of $300, and a 50% match on every additional dollar deposited up until $950.[15] The commitment is "soft" in that the ultimate savings decision can deviate from the commitment. However, the pre-commitment still matters for future incentives. If the tax filer pre-commits, the savings account will include an additional $100 in savings matches for saving at least $300. Alternatively, if the tax filer pre-commits to not saving, she still has the option to save at the tax site, but now will receive a $75 payment in October should she not save. Thus, the commitment reinforces decisions in either direction, and therefore can be distinguished from "cheap talk."[16] This is the baseline commitment group (Treatment Group 3).

   (b) Some commitment group members will additionally be offered an incentive for pre-committing to save. For pre-committing to save, the individual will receive a $50 incentive. This incentive will either be an "early" incentive, given at the time of filing taxes (Treatment Group 1), or a "late" one received 8 months after the tax season (Treatment Group 2). Furthermore, the incentive is kept regardless of the final savings outcome.[17] A treatment member of the commitment group has the opportunity to receive up to $475 in savings and commitment incentives. A baseline member of the commitment group may receive up to $425 in savings.[18]

   (c) At the time of tax preparation the individual will be reminded of their prior soft-commitment and the monetary rewards available should they honor their prior soft-commitment.

   (d) Survey data similar to what is collected from members of the non-commitment group will be collected from commitment group members as well.

Due to space limitations, we have not included visual depictions of the incentives. However, extensive-form representations of the incentives for each of the six groups are available in the appendix.[19] We strongly encourage the reader to view these visuals to gain a better understanding of the incentives.

3.3 Time Line

December 2010

Both commitment and non-commitment group members in the pilot study have been contacted and notified of the

[14] Measuring savings and debt overtime will allow us see how realistic an assumption of credit constraints is. We can also see if savings or debt are endogenously changed between December and tax season in response to anticipation of the savings incentive.

[15] We choose different upper limits for the two groups since commitment group members receive additional incentives for following through with commitments. In particular, this allows our scripts across treatments to be roughly equivalent when we state, you will have the opportunity to receive as much as "$X" in savings incentives. Furthermore, we see that the upper limit is typically not binding in our pilot study.

[16] The pre-commitment decision is symmetric in the sense that individuals can commit to either saving or not saving. This is because the pre-commitment decision is based both on one's preferred outcome in February and what they believe to be the most likely outcome. For individuals with time-inconsistent preferences, the two may not be the same. In particular, there may be very low-probability savers for whom committing not to save is optimal.

[17] As will be explained below, making this incentive independent of the final savings decision is key for separately identifying the discount parameters in a quasi-hyperbolic discounting model.

[18] We should reiterate that take-up of these accounts is relatively low, so that realized payments will tend to be far below maximum potential payments.

[19] http://home.uchicago.edu/~j1s/Jones_Mahajan_Appendix_NSF.pdf

savings incentives. Additionally, commitment group members are offered the soft-commitment.

February 2011-april 2011

Participants in the pilot study make an actual decision regarding the savings option while filing their tax return. Commitment group members are additionally reminded of their prior, soft-commitment. Surveys are conducted among pilot study participants to help inform the final design of the study. In addition, participants are recruited for participation in the next round of the study.

October - December 2011

Both commitment and non-commitment group members in the full-scale study are contacted and notified of the savings incentives. Additionally, commitment group members are offered the soft-commitment. Savings matches from the pilot study are deposited into accounts.

February 2012-April 2012

Participants in the full-scale study make a final decision regarding the savings option while filing their tax return. Commitment group members are additionally reminded of their prior, soft-commitment. Surveys of all participants are also administered in order to provide additional, qualitative data.

October - December 2012

Savings matches from the full-scale study are deposited into accounts.

3.4 Discussion

As described above, the contributions of the project are both an empirical test of economic theory and the evaluation of financial policy tools used among low-income tax filers. First, we join a relatively recent literature that uses field experiments to test theories of time-inconsistency. We add to this literature by obtaining additional quantitative estimates of the parameters of a quasi-hyperbolic discounting model and elicit these preferences from choices that take place in a "natural" decision setting. Our study also takes place in a policy-relevant setting. Policies and research surrounding the financial decisions made upon receiving one's income tax refund have developed into a virtual cottage industry. In this field, comprised of practitioners and researchers, it is commonly thought a prudent idea to save some or all of the income tax refund, hence the prevalence of savings incentive programs. In our study, we remain agnostic as to whether the savings is the optimal decision, but nonetheless aim to contribute to this policy question.

First, our perturbation of existing savings incentives allows us to compare alternative methods of encouraging savings. In particular, our pilot study below suggests that varying the timing of savings incentives can result in a more cost-effective program design. A larger question is whether this savings results in a welfare improvement, at the very least for the participants. On the one hand, income tax refunds are comprised of overwithholdings and lagged transfers, which suggests that they are prime candidates for spending or debt reduction. On the other hand, these lumpy payments may provide a buffer of savings moving forward that is otherwise hard to build up in the presence of self-control problems. Even so, placing the buffer in an illiquid account limits the insurance possibilities. However, research shows that even at the monthly frequency, lumpy benefits appear to be drawn down too quickly.[20] Thus, it remains an empirical question whether placing a portion of the refund in an illiquid account can aid in a more even, intra-annual allocation of the income tax refund and help provide a buffer against late-year shocks by stemming the draw-down of the income tax refund earlier in the year.[21] With our experimental variation in savings

[20] Shapiro (2005) shows that during the course of a month, food stamp benefits tend to be exhausted at rate that is consistent with impatience, i.e. time-inconsistent preferences.

[21] For this reason, we shortened the duration of the savings requirement to 8 months, relative to the 12 months that is typically offered in the field. In the latter case, money is held until about the time another lumpy income tax refund is received, while in the former case, the money becomes available at a more intermediate point between lumpy tax transfers. As is discussed below in our model, this informs our decision to allow the commitment device to be symmetric, as saving may not be the optimal decision for everyone in the study.

account use and our longitudinal survey, we can speak to this policy-relevant question by estimating the causal effect of the illiquid savings account on consumption smoothing.

4 Model

This section outlines a basic version of the model that informs the experimental design and empirical methodology. We model the savings account as an "investment good" as in DellaVigna and Malmendier (2004) (see also DellaVigna and Malmendier (2002) and DellaVigna (2009) or "immediate cost" activities in O'Donoghue and Rabin (1999)).[22] We begin by specifying preferences and beliefs and then by outlining for each period the state space, action space and per-period payoffs.[23] Before moving forward, it will be useful to summarize the key features of the model:

1. Our identification of time preferences is driven by variation in savings incentives at different points in time. In particular, we measure the effect of an "immediate" incentive relative to that of a "delayed" incentive and then measure the effect of a two future incentives, an "early" and "late" incentive. Under the null hypothesis, these comparisons should yield similar results, while time-inconsistency implies a discrepancy between the two.

2. A "soft-commitment" decision features prominently in our model, as it makes decisions prior to tax season incentive compatible. This decision is by design non-binding, reflecting both institutional limitations and the fact that exact refund amounts and liquidity concerns are not known with certainty ex ante. It is important to note, however, that our test of time-consistency is not based on detecting a demand for commitment, as individuals do not choose between being in the commitment or non-commitment groups. Moreover, a subsequent reversal of the commitment is not equivalent to time-inconsistency in our model. As such, intra-personal comparisons of the commitment decision and actual savings outcome do not constitute our formal test.[24]

3. Along the same lines, a divergence in the levels of aggregate pre-commitment and aggregate savings probabilities is not equivalent to time-inconsistency in our model. Rather, it is the marginal effect of savings incentives relative to these baseline probabilities that is used as described below.

4. Our test of time-consistency is not based on the direct effect of the "soft-commitment" on savings probabilities. Nonetheless, we can measure this effect in the event that it informs the effective design of savings incentives.

4.1 Preferences and Beliefs

Individuals maximize utility functions, separable in time, with potentially quasi-hyperbolic discounting, i.e. "$\beta - \delta$" preferences (see Strotz (1955), Phelps and Pollak (1985), Laibson (1997) and O'Donoghue and Rabin (1999)). That is, in period t, the present discounted value of a consumption stream is

$$U_t = u(c_t) + \beta \sum_{\tau=t+1}^{T} \delta^{\tau-t} \cdot u(c_\tau). \tag{1}$$

[22] In mapping the model to the actual field experiment, the discrete choice modeled here is the decision to save at least the minimum amount listed above. Individuals have an additional continuous choice of the amount to save beyond the minimum savings threshold. But note, the immediate and delayed incentives are conditional on the extensive margin decision, and thus map into a discrete choice model. We discuss below how to potentially use the additional continuous choice to estimate time-preferences as in Andreoni and Sprenger (2010).

[23] In an appendix, we extend this model to one of dynamic discrete choice with potentially time-inconsistent agents (see Fang and Wang (2008) and Mahajan and Tarozzi (2010) for models along similar lines).

[24] Though we do not make explicit use of revision behavior, this does not mean that it is uninformative regarding time consistency. We can observe revision behavior based on participant decisions and/or early withdrawal of savings. In addition, we allow for shocks to information regarding the refund level. We can explore the importance of this uncertainty by collecting expectations regarding the refund level in our survey prior to the tax season.

where $\beta < 1$ generates time-inconsistent preferences. In addition, the agent holds beliefs about future objective functions, parameterized by $\hat{\beta}$. That is, in period t, the agent believes that the following objective function will be maximized in k periods:

$$U_{t+k} = u(c_{t+k}) + \hat{\beta} \sum_{\tau=t+k+1}^{T} \delta^{\tau-t} \cdot u(c_\tau). \tag{2}$$

where $\hat{\beta} < \beta$ is referred to as partial naïveté (see O'Donoghue and Rabin (2001)). Our null hypothesis is that individuals are exponential discounters: $\beta = \hat{\beta} = 1$. Note that, in period t, an individual who is time-inconsistent discounts utility between period $t+k$ and $t+k+j$ by a factor of δ^j, but believes that when period $t+k$ arrives, utility will be discounted by a factor of $\hat{\beta}\delta^j$. When period $t+k$ actually arrives, utility is discounted by $\beta\delta^j$.

4.2 State Spaces, Action Spaces and Payoffs

PERIOD 1

State Space: $x_1 \in \mathcal{X}_1$ where x_1 are pre-intervention observables that potentially affect the agent's utility in period 1.

Action Space: $a_1 \in \{0, 1\}$[25] where $a_1 = 1$ indicates that the agent chooses in period 1 to the "pre-commitment" decision to save in period 2. $a_1 = 0$ if the agent chooses the "pre-commitment" decision to not save in period 2. For simplicity in exposition the model does not consider an alternative third action in which the agent rejects the first two actions.[26]

Utility Flow in Period 1: To focus only on essential details it is assumed that the individual receives no direct utility flow from actions taken in period 1. All incentives will be derived from the effect on utility flows in periods 2 and 3 as a result of actions taken in period 1.

PERIOD 2

State Space: $x_2 \in \mathcal{X}_2 \equiv \{0, 1\}$ where $x_2 = a_1$ just records the first period decision. It is straightforward to incorporate additional state variables using standard approaches but those are eliminated here in the interest of brevity. Note that members of the non-commitment option group do not have an opportunity to make a period 1 decision and therefore, the state variable is irrelevant for their period 2 decision.

Action Space: $a_2 \in \{0, 1\}$ where $a_2 = 1$ indicates that the agent chooses to save and $a_2 = 0$ if the agent decides not to save.

Utility Flow in Period 2: Payoffs in the second period vary by treatment group. In the non-commitment group (NC) agents receive a payment of i if they decide to save, representing the "immediate" savings incentive. This is the additional \$50 incentive for saving, which is contrasted with an equivalent "delayed" incentive received in period 3. In the commitment group (C), if the agent had chosen to pre-commit to saving $a_1 = 1$ she receives a payoff e (*irrespective* of her actual savings decision). This is referred to as the "early" incentive for pre-committing, which contrasts with the "late" incentive in period 3. In addition, agents in both groups are assumed to pay a cost c (i.e. the cost of saving c dollars, including a binding borrowing constraint) if they decide to save.

The experimental design yields the following values for the $u(x_2, a_2)$ for the different treatment groups. Utility is denoted as a function f of the payoffs since in future work the model will be expanded to explore what can be learned about $f(\cdot)$ experimentally. However, for this proposal the model will deal mostly with the case where f is

[25] Note that in the model, we only model actions related to the project choices. In principle, agents could make other decisions (e.g. change their saving and/or consumption behavior) in response to this news. If, for instance, individuals can easily borrow, then they can take other actions that would render the findings from our study uninteresting. However, we assume here that individuals face credit constraints (and we attempt to quantify them via the survey). If individuals are credit constrained, then the savings decision at tax-time is an inter-temporal utility trade-off rather than a mere arbitrage opportunity. Individuals can still re-optimize, however, by delaying debt-repayment. We hope to capture this via our survey instruments.

[26] It can be shown that this third action is weakly dominated.

the identity function, i.e. utility is quasi-linear. This is a strong assumption and discussion of relaxing it can be found below.[27] Payoffs for period 2 are given by

Table 1: Period 2 Payoffs by Treatment Group

(x_2, a_2)	$u_{NC}(x_2, a_2)$	$u_C(x_2, a_2)$
(1,1)	$f(i - c)$	$f(e - c)$
(1,0)	0	$f(e)$
(0,1)	$f(i - c)$	$f(-c)$
(0,0)	0	0

Note in the first column that $u_{NC}(x_2, a_2)$ is invariant to the state variable, since there is no pre-commitment for the non-commitment group. Looking at the first row of the last column, one can see for example that when a pre-commitment is made, i.e. $x_2 = 1$, the early savings incentive e is received. Furthermore, a cost c of saving is incurred for all agents that save in either group, i.e. when $a_2 = 1$. For each group, a normalization of utility in all states with respect to a reference action is needed. In particular, one can normalize $u(x_2, 0)$ for all $x_2 \in \{0, 1\}$. For the non-commitment group, this normalization is $u(x_2, 0) \equiv 0$. For the commitment group, the normalization assumes $f(e)$ is known.

<u>Period 3</u>

State Space: $x_3 \in \mathcal{X}_3 \equiv \{0, 1\} \times \{0, 1\}$ where $x_3 = (x_2, a_2)$.

Agents take no action in this period and payoffs are a function of group status and state variables. In the non-commitment group, if the agent had not saved in period 2, then she receives a payoff of 0. If she has saved, she receives a savings incentive of d. This payoff d is referred to as the "delayed" payoff. In the commitment group, the agent receives a payoff of p if the pre-commitment and actual savings decision coincide, i.e. $x_2 = a_2$. That is, if she had pre-committed to not saving and fulfilled that commitment in period 2 or if the agent had pre-committed to saving in period 1 and she followed through with that commitment, she again receives the a payoff p. Finally, for members of the commitment group, a payoff l is received *irrespective* of her period 2 saving decision. This is the "late" incentive for pre-committing to save. Finally, agents in either group who chose to save in period 2 receive a payoff of b, which is in effect the value of the amount now available to be withdrawn from the savings account.[28]

Table 2: Period 3 Payoffs by Treatment Group

(x_3)	$u_{NC}(x_3)$	$u_C(x_3)$
(1,1)	$f(b + d)$	$f(b + l + p)$
(1,0)	0	$f(l)$
(0,1)	$f(b + d)$	$f(b)$
(0,0)	0	$f(p)$

<u>State Transitions</u>: State transitions $\mathrm{dF}(x_{t+1}|x_t, a_t)$ are extremely simple in this basic framework since they are completely deterministic functions of previous period actions.

<u>Discussion</u>: Looking at Tables 1 and 2 one can get a general sense of the source of identification. First, looking at the columns for $u_{NC}(\cdot)$ one will notice that the incentives (i, d) vary depending on the decision in Period 2, a_2 ($\equiv x_3$) and are by construction invariant to the state variable x_2 ($\equiv a_1$) since members of the NC group do not make

[27] This type of quasi-linearity is found in other models that analyze present-biased preferences, e.g. DellaVigna and Malmendier (2002).

[28] This amount b is inclusive of the variable portion of the savings match, i.e. the 50% on each dollar above $300.

pre-commitments. Thus, by observing the response of a_2 to experimental variation in (i, d), one learns about Period 2 preferences over utility in Period 2 (i) relative to utility in Period 3 (d). In contrast, one can see in the columns for $u_C(\cdot)$, (e, l) vary depending on the decision in Period 1, a_1 ($\equiv x_2$), but are invariant to Period 2 decisions, a_2 ($\equiv x_3$). Thus, by observing the response of a_1 to experimental variation in (e, l) one learns about Period 1 preferences over utility in Period 2 (e) relative to utility in Period 3 (l). Finally, the comparison of the preferences in Period 1 and Period 2 provide the grounds for testing for time consistency. Note also that in Table 2 the commitment reward p depends on both the Period 1 decision and the Period 2 decision. All things equal, pre-committing to save makes saving more likely to occur, and vice versa. In this sense, it is the mechanism by which Period 1 decisions can alter Period 2 decisions, and is necessary for making Period 1 decisions nontrivial (without $p \geq l$ all commitment group agents would choose $a_1 = 1$).[29]

4.3 Solving the Model

To solve the model, we make the following econometric assumptions. Agents are indexed by the variable n. To capture uncertainty over time, we assume that individuals do not know in period 1 the precise values of the cost and benefits of savings (c, b), but rather they know the joint distribution of these parameters with c.d.f. $G_n(c, b)$. We model the savings the decision as the consumption of an investment good, as in DellaVigna and Malmendier (2004). In our particular case, you can imagine that individuals do not know exactly what their income tax refund will be, which generates uncertainty over c and b.[30] The subscript n allows for ex-ante cross-sectional heterogeneity in preferences. Because c and b will be unobservable to the researcher, we will make use of the average c.d.f. of savings preferences $G(c, b) = \int G_n(c, b)\,\mathrm{dF}(n)$. In period 2, c and b are revealed to the agent, but are still unobservable to the researcher. We will also make a structural assumption that the utility function f is the identity function, i.e. preferences are quasi-linear. Using these assumptions, we solve the model by backward induction. Since no actions are made in period 3, we begin with period 2.

Period 2

Table 3: Net value of saving in Period 2 by Treatment Group

x_2	Non-Commitment Group	Commitment Group
1	$f(i-c) + \beta\delta f(b+d)$	$f(e-c) - f(e) + \beta\delta(f(b+l+p) - f(l))$
0	$f(i-c) + \beta\delta f(b+d)$	$f(-c) + \beta\delta\,(f(b) - f(p))$

The net value of saving for agents in commitment and non-commitment groups are shown in Table 3. Under the assumption that $f(x) = x$, we see that the agent in the non-commitment group will save ($a_2 = 1$) if:

$$i - c + \beta\delta\,(b+d) \geq 0 \tag{3}$$

For individuals in the commitment group, the decision depends on the whether or not one has pre-committed to save

[29] Again, note that the discrete decisions a_1 and a_2 are the decisions to pre-commit or to actually save above the minimum savings threshold or not. The matches (e, l) and (i, d) are the incentives received for pre-committing or actually saving more than these thresholds. The variable part of the savings match is proportional to the amount saved and is for convenience collapsed into the (c, b) variables in the model.

[30] For example, in period 1, an individual who plans to save the entire refund may know that the cost will be $c = c_n + \varepsilon$ and benefit will be $b = Rc$, where ε is a mean-zero shock to the refund level and R is the gross return of the savings account. Indeed, we remind the study participants during our phone interview of their prior year income tax refund, which is predictive of future refunds.

(i.e. $x_2 \equiv a_1 = 1$). The agent's decision can be summarized as follows:

$$a_2 = \begin{cases} 1 & \text{if } a_1 = 1 \text{ and} - c + \beta\delta b \geq -\beta\delta p \\ 1 & \text{if } a_1 = 0 \text{ and } - c + \beta\delta b \geq \beta\delta p \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

<u>Period 1</u>

Only members of the commitment group make decisions in period 1. The expected value of pre-committing to save is

$$V_{a_1=1} = \iint\limits_{-c+\hat{\beta}\delta b \geq -\hat{\beta}\delta p} \left[e - c + \delta\left(b + l + p\right)\right] \mathrm{dG}_n\left(c, b\right) + \iint\limits_{-c+\hat{\beta}\delta b < -\hat{\beta}\delta p} \left[e + \delta l\right] \mathrm{dG}_n\left(c, b\right) \tag{5}$$

where we have again used the fact that $f(x) = x$. The first component is the payoff in the event that the agent actually follows through with the saving pre-commitment and the second component is the payoff should the agent fail to follow through. Note that the argument inside the integral includes a discount factor of δ, as it reflects the agent's preferences in period 1 over payoffs in period 2 and 3. However, the limits of integration feature a discount factor of $\hat{\beta}\delta$, as the likelihood of following through with the commitment is based on the period 1 agent's beliefs about period 2 preferences over payoffs in period 2 and 3. It is this sense in which agents take into account the implications of their pre-commitment decisions on future outcomes. Similarly, the expected value of pre-committing to not save is

$$V_{a_1=0} = \iint\limits_{-c+\hat{\beta}\delta b \geq \hat{\beta}\delta p} \left[-c + \delta b\right] \mathrm{dG}_n\left(c, b\right) + \iint\limits_{-c+\hat{\beta}\delta b < \hat{\beta}\delta p} \left[p\right] \mathrm{dG}_n\left(c, b\right) \tag{6}$$

The agent will pre-commit to saving if $V_{a_1=1} \geq V_{a_1=0}$.

4.4 Estimation

We now outline our estimation strategy for β and δ. First, using variation in (i, d), we estimate $\beta\delta$ from non-commitment group members in period 2. Note from (3) that

$$\begin{aligned} \mathbb{E}\left[a_2 | NC\right] &= \Pr\left(a_2 = 1 | NC\right) \\ &= \iint\limits_{-c+\beta\delta b \geq -i-\beta\delta d} \mathrm{dG}\left(c, b\right) \end{aligned}$$

It follows that

$$\frac{\partial \mathbb{E}\left[a_2 | NC\right] / \partial d}{\partial \mathbb{E}\left[a_2 | NC\right] / \partial i} = \beta\delta$$

Next, using variation in (e, l), we estimate δ. Using (5) and (6) we see that:

$$\begin{aligned} \mathbb{E}\left[a_1 | C\right] &= \Pr\left(a_1 = 1 | C\right) \\ &= \Pr\left(V_{a_1=1} \geq V_{a_1=0}\right) \\ &= \Pr\left(e - \delta l \geq \iint\limits_{-c+\hat{\beta}\delta b \geq \hat{\beta}\delta p} \left[-\delta p\right] \mathrm{dG}_n\left(c, b\right) + \iint\limits_{\hat{\beta}\delta p > -c+\hat{\beta}\delta b \geq -\hat{\beta}\delta p} \left[-c + \delta b\right] \mathrm{dG}_n\left(c, b\right) \right. \\ &\qquad \left. + \iint\limits_{-\hat{\beta}\delta p > -c+\hat{\beta}\delta b} \left[\delta p\right] \mathrm{dG}_n\left(c, b\right)\right) \end{aligned}$$

Again, we can show that

$$\frac{\partial \mathbb{E}\left[a_1|\, C\right]/\, \partial l}{\partial \mathbb{E}\left[a_1|\, C\right]/\, \partial e} = \delta$$

Thus, we use four reduced form parameters to estimate time preferences. In particular, we estimate the following linear probability model among the commitment group members:

$$a_1 = b_e T_1 + b_l T_2 + \varepsilon \tag{7}$$

and a similar model is estimated among non-commitment group members:

$$a_2 = b_i T_4 + b_d T_5 + \varepsilon \tag{8}$$

Here T_j are treatment group indicators and the b_m measure the early, late, immediate and delayed incentive effects. The omitted group in (7) is Treatment Group 3 while the omitted group in (8) is Treatment Group 6.

5 Pilot Study Results

Below we present preliminary results from our pilot version of the field experiment. The results will be restricted to estimation of time preference parameters. Evaluation of the savings incentive program and its effect on consumption smoothing are omitted, as necessary survey data are currently in the process of being collected.

5.1 Data

We use a combination of administrative and survey data to test for and measure the degree of time-inconsistency in the sample. We have access to administrative tax return data from IRS Form 1040. This includes filing status, age, adjusted gross income (AGI), unemployment insurance, number of dependents, earned income tax credit (EITC), child tax credit, withholdings, tax liability, federal and state income tax refund, and whether or not direct deposit was used to receive a refund. We have this data for both tax year (TY) 2009 and TY 2010, which are collected between February and April (tax-filing season) in 2010 and 2011 respectively. Additional client intake data for each tax-filing season are collected by the collaborating non-profit. This data include gender, race/ethnicity, primary language, marital status, highest level of education, zip code, tax filing date, tax preparation site location and type of bank accounts owned. We will refer to TY 2009 and tax filing season 2010 as the "prior year" and TY 2010/tax filing season 2011 data as the "current year". We have prior year data for all baseline participants and only have current year data for those participants that returned to file their taxes with the non-profit.

During our initial phone interviews, which happen in the winter prior to the current tax filing season, we collect data on savings level and debt levels. Following Lusardi et al. (2011), we measure credit constraints by asking participants the likelihood of unexpectedly owing $2,000, the ability to come up with this amount of cash and the potential source of this cash. In the spirit of Ashraf et al. (2006), we ask a series of questions aimed at detecting time-inconsistency. Finally, after explaining the SaveUp account to participants, we quiz them on the parameters of the savings vehicle, to measure comprehension. During the initial phone survey, the pre-commitment decision is also recorded. In our follow up survey, which takes place during the current tax filing season we repeat questions regarding savings, debt, credit constraints and comprehension. The savings decision and savings amount is also recorded on site.

5.2 Baseline Descriptive Statistics

Table 4 provides descriptive statistics for the six treatment groups in the pilot study. All of the characteristics presented have been established prior to the experimental intervention. As can be seen, there is treatment groups are appropriately balanced at the onset. The average age is about 41, and the sample is two-thirds female. The overrepresentation of females is in part due to income restrictions on free tax preparation clients and the concentration of EITC benefits among single female-headed households. For this same reason, average income is low, at $17,600 and the share married is only 11 percent. The average federal refund is $2,500 and state refund is $666, which is well above the maximum allowable SaveUp deposit. About one-half of the sample has a high school equivalent degree, but only 4 percent have a college degree or higher. Nearly half of the sample is African-American and another 30 percent are Hispanic and only 5 percent are white. The low share of Asian filers, less than 3 percent, is in part due to the sample being restricted to English and Spanish speaking tax filers. Finally, three-fourths of the sample reports some type of bank account.

5.3 Sample Selection and Attrition

One challenge to conducting the pilot study is a high rate of sample attrition. Table 5 shows the likelihood of making to each key stage of the study. The baseline group was selected from the pool of prior year tax clients at the collaborating non-profit. We restricted the study to tax filers who had a 2009 federal income tax refund above $300 and who spoke either English or Spanish. This resulted in an initial pool of 833 potential study participants. The first key stage is the pre-tax season phone call during which consent was obtained and pre-commitment decisions were recorded. As can be seen, the majority of attrition is due to the inability to reach participants by phone. Less than one-third is reached by phone. Once on the phone, an additional 13 percent are lost due to failure to obtain consent for the study. Finally, another 6 percent are loss between the pre-tax-filing season phone call and appearance at the tax site, where final savings decisions are made.

We suspect that the primary reason for attrition is that we are attempting to reach low-income households. In addition, the contact information we had for these households was collected in the prior tax season and may have a high likelihood of changing among this group. Furthermore, our pilot was originally designed to contact all households within a two-week period. In fact, we extended outreach to about one month and could have used more time to reach participants. Note that attrition is not as dramatic if you define the baseline sample as the group successfully reached by phone. From this perspective, the consent rate rises from 15 percent to 51 percent and the share that appears at the tax site rises from 9 percent to 28 percent. In the results below we report intent-to-treat (ITT) effects using the entire sample as well as treatment effects conditional on surviving to each stage of the study.

Though attrition itself is of concern, an even larger issue is differential attrition across treatment groups. Table 5 compares attrition between groups. The first thing to note is that the non-commitment groups, Groups 4, 5 and 5, have a significantly lower probability of being reached by phone and consenting by phone. Recall that prior to being called, participants received treatment specific informational mailings. It is possible that differential attrition arises because commitment groups can on average receive larger financial incentives than non-commitment groups. The second thing to note is that within commitment treatment arms, attrition does not significantly differ by treatment group. That is, Groups 1, 2 and 3 do not differ from each other, and Groups 4, 5 and 6 do not differ from each other. This is important for our estimates below. For a majority of the analysis below, we will be conducting comparisons within commitment treatment arm. We address nonrandom attrition below.

Table 4: Baseline Descriptive Statistics

	(1)	(2)	(3)	(4)	(5)	(6)
	Commitment Groups			Non-Commitment Groups		
	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6
Age	41.44	42.80	41.84	40.84	38.47*	40.83
	[1.18]	[1.13]	[1.27]	[1.14]	[1.21]	[1.24]
Female	0.66	0.66	0.72	0.65	0.57	0.69
	[0.04]	[0.04]	[0.04]	[0.04]	[0.04]	[0.04]
AGI (2009)	18,234	18,681	17,986	17,459	17,479	15,813*
	[971]	[1,059]	[1,110]	[948]	[1,048]	[883]
Refund (2009)	2,214	2,222	2,132	1,990	1,858	2,157
	[178]	[193]	[175]	[156]	[145]	[178]
NY Refund (2009)	721	642	713	641	554*	730
	[83]	[80]	[70]	[69]	[57]	[73]
Dependents	0.71	0.60	0.65	0.54	0.55	0.65
	[0.09]	[0.08]	[0.07]	[0.07]	[0.07]	[0.08]
Married	0.09	0.11	0.14	0.12	0.14	0.10
	[0.02]	[0.03]	[0.03]	[0.03]	[0.03]	[0.03]
High School	0.54	0.48	0.51	0.53	0.54	0.52
	[0.04]	[0.04]	[0.04]	[0.04]	[0.04]	[0.04]
College+	0.04	0.04	0.04	0.04	0.04	0.05
	[0.02]	[0.02]	[0.02]	[0.02]	[0.02]	[0.02]
African-American	0.49	0.50	0.50	0.53	0.55	0.47
	[0.04]	[0.04]	[0.04]	[0.04]	[0.04]	[0.04]
Asian	0.04	0.03	0.03	0.01	0.01	0.04
	[0.02]	[0.01]	[0.01]	[0.01]	[0.01]	[0.02]
Hispanic	0.30	0.31	0.33	0.31	0.29	0.31
	[0.04]	[0.04]	[0.04]	[0.04]	[0.04]	[0.04]
White	0.04	0.07	0.04	0.06	0.06	0.07
	[0.02]	[0.02]	[0.02]	[0.02]	[0.02]	[0.02]
Bank Account	0.79	0.72	0.75	0.73	0.81	0.77
	[0.03]	[0.04]	[0.04]	[0.04]	[0.03]	[0.04]
N	140	140	137	139	140	137

Note: Descriptive statistics for 6 treatment groups are tax year 2009 and tax-filing season 2010 variables established prior to the intervention. Robust standard errors are reported in brackets. One, two and three stars denote statistically significant difference from Treatment Group 1 at the 10, 5 and 1 percent levels respectively.

Table 5: Treatment Group Survival Rates

	(1)	(2)	(3)	(4)	(5)	(6)
	Commitment Groups			Non-Commitment Groups		
	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6
Reached by Phone	0.39	0.38	0.31	0.15	0.23	0.20
	[0.04]	[0.04]	[0.04]	[0.03]	[0.04]	[0.03]
p-value between	.	0.90	0.16	0.00***	0.00***	0.00***
p-value within	0.16	0.21	.	0.25	0.62	.
Consented to Study on Phone	0.21	0.20	0.14	0.09	0.13	0.11
	[0.03]	[0.03]	[0.03]	[0.02]	[0.03]	[0.03]
p-value between	.	0.88	0.13	0.00***	0.08*	0.02**
p-value within	0.13	0.17	.	0.52	0.62	.
Appeared On Site	0.08	0.13	0.09	0.04	0.08	0.10
	[0.02]	[0.03]	[0.03]	[0.02]	[0.02]	[0.03]
p-value between	.	0.17	0.63	0.21	1.00	0.49
p-value within	0.63	0.37	.	0.06*	0.49	.
Consented to Study on Phone	0.54	0.51	0.43	0.52	0.56	0.54
(Conditional on Phone Contact)	[0.07]	[0.07]	[0.08]	[0.11]	[0.09]	[0.10]
p-value between	.	0.78	0.29	0.92	0.82	0.99
p-value within	0.29	0.43	.	0.93	0.84	.
Appeared On Site	0.20	0.30	0.28	0.23	0.30	0.38
(Conditional on Phone Contact)	[0.05]	[0.06]	[0.07]	[0.08]	[0.08]	[0.08]
p-value between	.	0.21	0.34	0.76	0.29	0.06*
p-value within	0.34	0.85	.	0.20	0.46	.
N	140	140	137	139	140	137

Note: Sample survival rates are the probability of remaining in the study at each stage of the experiment. Two sets of p-values are reported. The "between" p-value measures compares each treatment group to Treatment Group 1, while the "within" p-value compares either Treatment Groups 1 and 2 to Treatment Group 3 or Treatment Groups 4 and 5 to Treatment Group 6. One, two and three stars denote statistically significant differences at the 10, 5 and 1 percent level respectively.

Table 6: ITT Estimates - Precommitment Decision

	(1)	(2)	(3)	(4)
	Treatment Group 1 (early incentive)		Treatment Group 2 (late incentive)	
Treatment Effect	0.092	0.087	0.092	0.088
	[0.035]***	[0.036]**	[0.035]***	[0.035]**
Mean Outcome	0.051	0.056	0.051	0.056
	[0.019]***	[0.020]***	[0.019]***	[0.020]***
Treatment Bounds				
Upper Bound	0.885	0.879	0.892	0.887
	[0.028]***	[0.030]***	[0.027]***	[0.028]***
Lower Bound	−0.770	−0.767	−0.770	−0.764
	[0.038]***	[0.040]***	[0.038]***	[0.039]***
N	417	401	417	401
Controls	No	Yes	No	Yes

Note: Treatment effects for Groups 1 and 2 are relative to Treatment Group 3. Upper and lower bounds are calculated using methods outlined by Horowitz and Manksi (2000). One, two and three stars denote statistical significance at the 10, 5 and 1 percent level respectively.

5.4 Intent to Treat Effects

Recall from our theoretical model that we would like to estimate 4 treatment effects and use them to recover time preference parameters. First, comparing pre-commitment probabilities between Treatment Group 1 and Treatment Group 3 gives us the "early" incentive effect, while comparing Treatment Group 2 to Treatment Group 3 gives us the "delayed" incentive effect. We can then use these to estimate δ. Similarly, comparing savings probabilities between Treatment Groups 4 and 5 relative to Treatment Group 6 gives us the "immediate" and "delayed" incentive effects. These are in turn used to estimate $\beta\delta$. As mentioned above, one challenge will be addressing sample attrition. Our first approach is to use all the original study participants and to code those who dropped out of the study as a zero outcome, giving us an "intent-to-treat" (ITT) effect. The ITT effect on pre-commitment is shown in the first row of Table 6. Treatment effects are relative to Treatment Group 3. The early and delayed incentives increase the likelihood of pre-committing by about 9 percentage points, from a baseline of 5 percentage points. The effect is identical across the early and delayed treatment groups, which is consistent with a value of $\delta = 1$. This result is not sensitive to introducing control variables.

The ITT effect on savings is shown in Table 7. Treatment effects here are relative to Treatment Group 6. When not controlling for observables, we actually have small negative effect for Treatment Group 4, as compared to a slight positive effect for Treatment Group 6. Our null hypothesis predicts similar, nonnegative effects for these two groups. However, our estimates are confounding the treatment effects with attrition. Thus, our negative estimate for Group 4 is caused by relatively high attrition in Group 4. Note, that after controlling for observables, the results are now nonnegative and slightly larger for Treatment Group 4. We suspect that the control variables may help account for low attrition in Group 4. Nonetheless, the overall small effect, 0.002 - 0.003 relative to a baseline of 0.044 is in part due to the large drop-off in sample size once we condition on tax site appearance. We discuss ways to account for attrition below.

Table 7: ITT Estimates - Savings Decision

	(1)	(2)	(3)	(4)
	Treatment Group 4 (immediate incentive)		Treatment Group 5 (delayed incentive)	
Treatment Effect	−0.001	0.003	0.006	0.002
	[0.025]	[0.026]	[0.025]	[0.026]
Mean Outcome	0.044	0.044	0.044	0.044
	[0.018]**	[0.018]**	[0.018]**	[0.018]**
Treatment Bounds				
Upper Bound	0.956	0.956	0.928	0.922
	[0.018]***	[0.018]***	[0.023]***	[0.023]***
Lower Bound	−0.898	−0.887	−0.892	−0.879
	[0.027]***	[0.028]***	[0.027]***	[0.030]***
N	416	394	416	394
Controls	No	Yes	No	Yes

Note: Treatment effects for Groups 4 and 5 are relative to Treatment Group 6. Upper and lower bounds are calculated using methods outlined by Horowitz and Manksi (2000). One, two and three stars denote statistical significance at the 10, 5 and 1 percent level respectively.

5.5 Treatment Effects Conditional on Participation

The ITT effects above reflect both attrition and take up. If we assume that attrition across treatment groups is not correlated with the outcome variables, then we can potentially estimate treatment effects that are more relevant to our theoretical exercise. The first row of Table 8 provides treatment effects on pre-commitment conditional on successfully obtaining consent on the phone. That is, the group of individuals who either said "yes" or "no" to the pre-commitment question. We now observe treatment effects of an order of magnitude larger, roughly a 0.2 - 0.3 probability increase relative to a baseline of about 0.4. However, the treatment effect still does not vary much when comparing the early to the delayed incentive. In Table 9 we also see larger treatment effects for savings. The effect for Treatment Group 4 is a 0.48 to 0.57 increase in savings probability relative to a baseline of 0.45. Furthermore, the immediate incentive (Treatment Group 4) is between 2 - 3 times as large as the delayed incentive, suggestive of impatience in the context of our theoretical model.

5.6 Treatment Bounds

One concern is that attrition is confounding our estimates. When estimating the ITT effect, we impute a zero for pre-commitment and savings outcomes for those that drop from the sample. While this is technically correct, we may alternatively view these outcomes as missing, since the treatment was not implemented for these participants. Viewed this way, we cannot obtain a point estimate for our treatment effects, but we may nonetheless bound these estimates using a method outlined by Horowitz and Manksi (2000). Unfortunately, in our current setting, attrition is such that the bounds are only marginally informative. Given a binary response, we know the treatment effects are bounded by [-1,1]. As seen in the bottom rows of Tables 6 and 7 our bounds are not much tighter. We now turn to an alternative method of achieving tighter bounds.

Table 8: ITT Estimates Conditional on Phone Consent - Precommitment Decision

	(1)	(2)	(3)	(4)
	Treatment Group 1 (early incentive)		Treatment Group 2 (late incentive)	
Treatment Effect	0.321	0.220	0.346	0.274
	[0.143]**	[0.140]	[0.143]**	[0.137]**
Mean Outcome	0.368	0.430	0.368	0.430
	[0.113]***	[0.111]***	[0.113]***	[0.111]***
Treament Bounds				
Upper Bound	0.443	0.347	0.459	0.388
	[0.132]***	[0.130]***	[0.134]***	[0.128]***
Lower Bound	0.113	0.051	0.152	0.122
	[0.161]	[0.149]	[0.166]	[0.157]
Participation Rate - Phone Consent				
Treatment	0.207	0.210	0.200	0.201
	[0.034]***	[0.034]***	[0.034]***	[0.034]***
Control	0.139	0.148	0.139	0.148
	[0.030]***	[0.032]***	[0.030]***	[0.032]***
N	417	401	417	401
Controls	No	Yes	No	Yes

Note: Treatment effects for Groups 1 and 2 are relative to Treatment Group 3 and conditional on participating in the initial phone interview. Upper and lower bounds are calculated using methods outlined by Behaghel et al. (2009). Participation rates are in absolute terms, where "Control" refers to Treatment Group 3. One, two and three stars denote statistical significance at the 10, 5 and 1 percent level respectively.

Table 9: ITT Estimates Conditional on Site Appearance - Savings Decision

	(1)	(2)	(3)	(4)
	Treatment Group 4 (immediate incentive)		Treatment Group 5 (delayed incentive)	
Treatment Effect	0.571	0.476	0.208	0.145
	[0.139]***	[0.253]*	[0.207]	[0.246]
Mean Outcome	0.429	0.457	0.429	0.457
	[0.139]***	[0.126]***	[0.139]***	[0.126]***
Participation Rate - Site Appearance				
Treatment	0.043	0.047	0.079	0.072
	[0.017]**	[0.019]**	[0.023]***	[0.023]***
Control	0.102	0.110	0.102	0.110
	[0.026]***	[0.028]***	[0.026]***	[0.028]***
N	416	394	416	394
Controls	No	Yes	No	Yes

Note: Treatment effects for Groups 4 and 5 are relative to Treatment Group 6 and conditional on appearing at the tax site. Upper and lower bounds are calculated using methods outlined by Behaghel et al. (2009). Participation rates are in absolute terms, where "Control" refers to Treatment Group 6. One, two and three stars denote statistical significance at the 10, 5 and 1 percent level respectively.

Differential attrition is potentially a larger concern for our estimates that condition on participation. We can use a method proposed by Behaghel et al. (2009), which allows us to bound the ITT effect for the subsample that participate when offered the early/immediate or late/delayed incentive. This effect is akin to the treatment on the treated. The method allows for a differential effect of the treatment on participation, but relies on an assumption of monotonicity. That is, these bounds assume that members who receive additional incentives are no less likely to participate than when given the baseline pre-commitment or savings treatment. For our pre-commitment estimates, we show evidence consistent with this assumption at the bottom of Table 8. Treatment Groups 1 and 2 have a higher participation rate than Treatment Group 3. Using this monotonicity, we bound the conditional treatment effect in Table 8. As can be seen, the bounds are much tighter than those for the unconditional ITT, and are similar for the early and delayed effects.

We cannot, unfortunately, implement the same bounding for the technique for the savings decisions. This is because the probability of participation is lower for Treatment Groups 4 and 5 as compared to Treatment Group 6. This is shown at the bottom of Table 9. Therefore, we do not calculate the analogous bounds for these treatment effects. We discuss this issue further below.[31]

5.7 Estimating β and δ

Using the methods outlined above, we can use the four estimated treatment effects to recover time preference parameters. These estimates vary depending on the specification chose. Table 10 summarizes the estimates using the unconditional ITT results from Tables 6 and 7. In Columns (1) and (2) we have estimates of δ very close to 1, albeit with relatively large confidence intervals. Column (3) has a sign and magnitude far out of line with our theoretical

[31] In our appendix, we outline the bounding methods described here.

Table 10: Beta Delta Estimates - ITT

	(1)	(2)	(3)	(4)
	δ		β	
Point Estimates	1.000	1.016	-9.846	0.867
	[0.456]**	[0.476]**	[405.169]	[8.744]
Upper Bound	9.718	10.218	$-1.42e{+}04$	$3,329.063$
	[3.617]***	[4.052]**	[$5.53e{+}05$]	[30531.388]
Lower Bound	0.104	0.100	0.001	0.000
	[0.039]***	[0.038]***	[0.003]	[0.003]
N	417	401	833	795
Controls	No	Yes	No	Yes

Note: Estimates for δ and β are calculated using the results from Tables (6) and (7). One, two and three stars denote statistical significance at the 10, 5 and 1 percent level respectively.

predictions. As mentioned above, this is most likely attributable to high rates of attrition in Treatment Group 4. In Column (4), we have more reasonable estimates after controlling for observable characteristics. In both cases, the estimates of β are dramatically imprecise. In the bottom two rows of Table 10, we see that the bounding method alluded to in Section 5.6 is uninformative, as expected.[32] Below, we explore an alternative method that results in tighter bounds.

Table 11: Beta Delta Estimates - Conditional on Participation

	(1)	(2)	(3)	(4)
	δ		β	
Point Estimates	1.077	1.244	0.338	0.245
	[0.395]***	[0.588]**	[0.301]	[0.318]
Upper Bound	4.078	7.656	1.058	0.870
	[5.914]	[22.881]	[1.459]	[1.578]
Lower Bound	0.344	0.351	0.089	0.040
	[0.384]	[0.476]	[0.148]	[0.128]
N	417	401	448	431
Controls	No	Yes	No	Yes

Note: Estimates for δ and β are calculated using the results from Tables (8) and (9). One, two and three stars denote statistical significance at the 10, 5 and 1 percent level respectively.

In Table 11 we present estimates for β and δ based on the results in Tables 8 and 9. These results are conditional on participation at either the phone interview or tax site stage. Compared to the unconditional ITT estimates, these estimates are much more stable and precisely estimated. Nonetheless, the small sample size results in wide

[32] Because our lower bounds on the treatment effects in Tables 6 and 7 are negative, we do not use them in these bounds. We use the upper bounds as follows. Recall that the estimates for β and δ are ratios of treatment effects. To get upper bounds, we substitute upper bounds for the treatment effects that show up in the numerator, and to get lower bounds we do the same for treatment effects in the denominator.

confidence intervals. The estimates for δ are slightly larger than 1. The estimates of β are now smaller in magnitude, between 0.25 and 0.34. In both cases, we can rule out a value of $\beta = 1$ with our 95 percent confidence interval. By conditioning on participation, we are implicitly assuming that attrition was not differentially affected by treatment group. It is the case that attrition is not statistically different within commitment arms, as can be seen in Table 5. Alternatively, if differential attrition is a valid concern, our results will be biased. However, we can still bound the estimates as discussed in Section 5.6.[33] The bottom two rows of Table 11 present our upper and lower bounds for β and δ. In this case, it is only when we use control variables that we can rule out a value of $\beta = 1$ with our upper bound.

5.8 Discussion

Overall, we find evidence that suggests that our study participants are have time-inconsistent preferences. Our preferred point estimates for β and δ over an 8-month horizon are 0.34 and 1.08 respectively. These translate into a one-year discount factor of 0.38, or equivalently, a one-year discount rate of 164%. As compared to prior structural estimates, our discount rate is higher than that implied in the life-cycle model estimated by Laibson et al. (2007) (49%), similar to the rate among low-wage workers estimated by DellaVigna and Paserman (2005) (153%) and lower than the rate among single-women with children estimated by Fang and Silverman (2009) (238%). One challenge to estimation in our context is sample attrition. In particular, differential savings rates across treatment groups may be attributed to sample selection bias. We attempt to account for this by estimating bounds on our time-preferences, but lose much of our precision in the process. As detailed below, a major refinement will involve reducing the extent of sample attrition in our full-scale study. In addition, our results suggest that more immediate incentives for savings can have as much as $2 - 3$ times as much an effect on the likelihood of saving. Finally, we are still in the process of collecting longitudinal survey data that can speak to the effect of the illiquid savings account on consumption smoothing.

6 Refinements and Extensions to the Study Design

The pilot study has generated several key insights for the design of our full-scale study. First, we underestimated the level of attrition that would occur. However, we have identified what we believe to be the major factors contributing to this phenomenon. We plan to allocate a much larger window of time, on the order of $1 - 2$ months, and a greater number of research assistants to reaching participants by phone. In addition, we will primarily focus our phone outreach on members of the commitment group and will rely more heavily on tax site recruitment for members of the non-commitment group.[34] On another note, our treatment effects are significantly higher than expected. This tends to reduce the required sample size for a given level of significance and power. Our projected sample size has reduced from about 3,000 to approximately 1,200 participants. This smaller sampling frame allows us to exclusively operate with our already established non-profit collaborator.

Another limitation of our study is that we impose the structural restriction of quasi-linearity. Given feedback from reviewers, we propose two strategies for relaxing this assumption. First, we can attempt to elicit risk preferences within our survey using methods similar to Harrison et al. (2007). These measures of risk preferences can then be used to impose an isoelastic functional form on our $f(\cdot)$ function. In an appendix, we describe the extension of our model

[33] As before, we can only calculate a partial upper bound, owing to the fact that the monotonicity assumption does not hold in the case of savings. Thus, we can only bound δ correctly. To bound β, we use our previous estimate for β in combination with the upper or lower bounds for δ.

[34] Attrition was also high due to an idiosyncratic factor. Many participants declined our study because the prior clients in our sample were in part drawn from a different tax site than the one used in our experiment. This misstep can be completely avoided in our full-scale study. In addition, we plan for our phone calls to be sufficiently spaced from major holidays.

to a standard dynamic discrete choice framework, which can accommodate the additional structure.[35] Alternatively, we may use the "convex time budget" (Andreoni and Sprenger (2010)) method of simultaneously eliciting risk and time preferences using the additional information imbedded in agents' continuous savings decision.[36] Along the same lines, we plan to add additional survey questions in order to gauge the level of uncertainty with regard to one's income tax refund. This information may facilitate an analysis of sophistication with regard to time preferences.

7 Research Partners

The PIs will require participation by a tax preparation organization in order to implement the field study. We have identified the major collaborator, The Financial Clinic, based in New York City that is working with us on the pilot and will provide the sample for the final study. The Financial Clinic is a nonprofit financial development organization that helps low-income families to achieve financial stability by providing legal support and financial counseling. This organization is particularly well suited for the study in that they maintain contact with their clientele on a year-round basis. Thus, they are able to accommodate the longitudinal structure of the project. In addition, they have previously offered a savings vehicle to tax preparation clients, the Save NYC account. Importantly, the Save NYC account is not simultaneously offered to our study participants. Prior users of the account may continue to the use the account, but this only applied to 5 people within our pilot sample.[37]

8 Research Output and Broader Impact

In terms of promotion of teaching and training, the PIs plan to include graduate students in economics as research assistants (RA), who will help with the on-the-ground implementation of the project. The budget includes allotments for one RA during the pilot study and two RAs during the full-scale intervention. These graduate students will gain invaluable experience with field research in economics.

The project requires collaboration with a non-profit organizations, such as The Financial Clinic. The PI's interaction with these groups represents significant enhancement in the infrastructure for research, dissemination of scientific understanding and benefit to society at large. Such organizations can benefit greatly from the research expertise of the PIs in evaluating and improving upon their policies and financial tools for lower-income tax filers. Furthermore, researchers in general will benefit from improved lines of communication between academic and the nonprofit sector, which will foster future research projects similar to the current one. Finally, the clients of these organizations will benefit from rigorous evaluation and redesign of the financial services provided by these public sector organizations.

The project will also help broaden the participation of underrepresented groups. First, one of the PIs is an African-American, junior faculty member in economics, an underrepresented group within the profession. Second, one of the PIs is a past participant in the American Economics Association's (AEA) Committee on the Status of Minority Groups in the Economics Profession (CSMGEP) Pipeline Mentoring Program, and has presented research and participated as a panelist at CSMGEP's Pipeline Conferences every year for the past 4 years (Jones (2006), Jones (2009)). CSMGEP's Pipeline Conference is an event that brings minority economists and students (graduate and undergraduate) together for workshops that cover a number of areas in economics, as well as topics such as conducting research, success in graduate schools, the economics job market, and the experience of minorities in

[35] http://home.uchicago.edu/~j1s/Jones_Mahajan_Appendix_NSF.pdf

[36] Note, in our pilot, a continuous savings decision is only collected during the actual savings decision. To implement the convex time budget method in our full-scale study, we would require a projected, continuous savings level, along with our pre-commitment decision.

[37] In addition, we don't believe the Saver's Credit to pose a significant threat to our study, as there is low take-up of this credit in our study.

the field of economics. The PIs plan to present the results of both the pilot study and full-scale project at future Pipeline Conferences in 2011 and 2012, providing underrepresented graduate students greater exposure to research and methods in the field. Finally, the PIs plan to advertise the graduate student RA positions among CSMGEP Pipeline graduate students who may be available in the New York region.

In terms of dissemination, the PIs plan to eventually publish the results in a peer-reviewed academic journal. In addition, the PIs will make use of outlets that will distribute their work beyond academia. As a faculty research fellow of the National Bureau of Economic Research (NBER), a PI can share the results of the research as an NBER working paper, which is distributed to numerous journalists and NBER subscribers. In addition, a version of the results will be written for a non-scientific audience as a policy brief at the Stanford Institute for Economic Policy Research (SIEPR).[38] The PIs will also disseminate the results to the broader non-profit tax-preparation community in partnership with the Financial Clinic. In fact, one of the PIs participated in the 2011 National Community Tax Coalition conference in Chicago, presenting preliminary results of the pilot study and also sharing lessons from the researcher/practitioner collaborative process (Jones and Panday (2011)). In addition, the PIs presented preliminary results at the Financial Literacy Research Consortium Spring Workshop at the University of Wisconsin, Madison to an audience of academics and practitioners. Finally, all the papers from the project will be freely available on the PI's websites.

[38] The PIs have published previous work in policy brief format, which was circulated among the many, non-academic, SIEPR subscribers (Jones (2010c)).

References

Aguirregabiria, Victor and Pedro Mira (2010), "Dynamic discrete choice structural models: A survey." *Journal of Econometrics*, 156, 38–67, URL http://linkinghub.elsevier.com/retrieve/pii/S0304407609001985.

Andreoni, James and Charles Sprenger (2010), "Estimating time preferences from convex budgets." NBER Working Paper 16347, URL http://www.nber.org/paers/w16347.

Ashraf, Nava, Dean Karlan, and Wesley Yin (2006), "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines." *Quarterly Journal of Economics*, 121, 635–672, URL http://www.mitpressjournals.org/doi/abs/10.1162/qjec.2006.121.2.635.

Behaghel, Luc, Bruno Crépon, Marc Gurgand, and Thomas Le Barbanchon (2009), "Sample attrition bias in randomized experiments: A tale of two surveys." Paris School of Economics Working Paper No. 2009-15, URL http://www.pse.ens.fr/document/wp200915.pdf.

Bernatzi, Shlomo and Richard Thaler (2004), "Save more tomorrow ^{TM}: using behavioral economics to increase employee saving." *Journal of Political Economy*, 112, S164–S187.

Beverly, Sondra, Daniel Schneider, and Peter Tufano (2006), "Splitting tax refunds and building savings: An empirical test." In *NBER Tax Policy and the Economy*, volume 20, 111–161, Cambridge: The MIT Press.

Choi, J., D. Laibson, B. Madrian, and A. Metrick (2004), "For better or for worse: Default effects and 401(k) savings behavior." In *Perspectives on the Economics of Aging* (David A. Wise, ed.), 81–121, Chicago: University of Chicago Press.

DellaVigna, S (2009), "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*, 47, 315 – 372.

DellaVigna, Stefano and Ulrike Malmendier (2002), "Overestimating self-control: Evidence from the health club industry." UC Berkeley Working Paper, URL http://www.econ.berkeley.edu/~sdellavi/wp/gymemp02-10-30.pdf%.

DellaVigna, Stefano and Ulrike Malmendier (2004), "Contract design and self-control: Theory and evidence." *Quarterly Journal of Economics*, 119, p353–402.

DellaVigna, Stefano and Daniele Paserman (2005), "Job search and impatience." *Journal of Labor Economics*, 23, 527 – 588.

Duflo, Esther, William Gale, Jeffrey Liebman, Peter Orzag, and Emmanuel Saez (2006), "Saving incentives for low- and middle-income families: Evidence from a field experiment with h&r block." *Quarterly Journal of Economics*, 121, 1311–1346.

Fang, Hanming and Dan Silverman (2009), "Time-inconsistency and welfare program participation: Evidence from the nlsy." *International Economic Review*, 50, 1043 – 1077.

Fang, Hanming and Y. Wang (2008), "Estimating dynamic discrete choice models with hyperbolic discounting, with an application to mammography decisions." URL http://www.hss.caltech.edu/{~}mshum/gradio/hyp%erbolic6.pdf.

Frederick, Shane, George Loewenstein, and Ted O'donoghue (2002), "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 40, 351–401, URL http://www.atypon-link.com/AEAP/doi/abs/10.1257/002205102320161311.

Gine, Xavier, Jennifer Goldberg, Dan Silverman, and Dean Yang (2010), "Revising commitments: Time preference and time-inconsistency in the field." Mimeo: University of Michigan, URL http://www-personal.umich.edu/~dansilv/research.htm.

Gul, F. and W. Pesendorfer (2004), "Self-control and the theory of consumption." *Econometrica*, 72, 119–158, URL http://www.jstor.org/stable/3598852.

Gul, Faruk and Wolfgang Pesendorfer (2001), "Temptation and Self-Control." *Econometrica*, 69, 1403–1435, URL http://www.blackwell-synergy.com/doi/abs/10.1111/1468-0262.00252.

Harrison, Gleen, Morten Igel Lau, and E. Elisabet Rutström (2005), "Dynamic consistency in denmark: A longitudinal field experiment." Mimeo, Georgia State University. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=698483.

Harrison, Glenn W., Morten I. Lau, and E. Elisabet Rutström (2007), "Estimating Risk Attitudes in Denmark: A Field Experiment." *Scandinavian Journal of Economics*, 109, 341–368, URL http://www.blackwell-synergy.com/doi/abs/10.1111/j.1467-9442.2007.00496.x.

Harrison, Glenn W. and John A. List (2004), "Field experiments." *Journal of Economic Literature*, 122, 1009–1055.

Hausman, Jerry (1979), "Individual discount rates and the purchase and utilization of enery-using durables." *The Bell Journal of Economics*, 10, 33–54.

Horowitz, Joel and Charles Manksi (2000), "Nonparametric analysis of randomized experiments with missing covariate and outcome data." *Journal of the American Statistical Association*, 95, 77–84.

Jones, Damon (2006), "The advance earned income tax credit: Why is participation so low?" URL http://www.vanderbilt.edu/AEA/CSMGEP/pipeline/archives/conference2006.html. CSMGEP Pipeline Conference Presentation.

Jones, Damon (2009), "CSMGEP pipeline conference panel: Choosing and thriving in graduate school." URL http://econ.ucsb.edu/aeastp/2009pipeline.html.

Jones, Damon (2010a), "Consumption smoothing and the earned income tax credit." URL http://home.uchicago.edu/~j1s. Working Paper, University of Chicago.

Jones, Damon (2010b), "Information, preferences and public benefit participation: Experimental evidence form the advance eitc and 401(k) savings." *American Economic Journal: Applied Economics*, 2, 147 – 163.

Jones, Damon (2010c), "Withholdings, salience and tax policy." SIEPR Policy Brief, URL http://siepr/publicationsprofile/2101.

Jones, Damon (Forthcoming), "Inertia and overwithholding: Explaining the prevalence of income tax refunds." *American Economic Journal: Economic Policy*, URL http://home.uchicago.edu/~j1s/Jones_Overwithholding.pdf.

Jones, Damon and Ambika Panday (2011), "Collaborate, investigate & evaluate: Maximizing vita site partnerships." URL http://www.cvent.com/events/mapping-the-future-charting-a-ne%w-direction/agenda-c2669277d5664d459e457f4d6fc92459.aspx. Presentation at the National Community Tax Coalition Conference in Chicago, IL.

Laibson, David (1997), "Golden eggs and hyperbolic discounting." *Quarterly Journal of Economics*, 63, 443–477.

Laibson, David, Andrea Repetto, and Jeremy Tobacman (2007), "Estimating discount functions with consumption choices over the lifecycle." NBER Working Paper No. 13314, URL http://www.nber.org/papers/w133141.

Lee, David S. (2002), "Trimming bounds on treatment effects with missing outcomes." UC Berkeley Center For Labor Economics Working Paper No. 51, URL http://www.princeton.edu/~davidlee/wp/WP51Lee.pdf.

Lusardi, Annamaria, Daniel J. Schneider, and Peter Tufano (2011), "Financially fragile households: Evidence and implications." NBER Working Paper No. 17072, URL http://www.nber.org/papers/w17072.

Mahajan, Aprajit and Alessandro Tarozzi (2010), "Time Inconsistency , Expectations and Technology Adoption: The case of Insecticide Treated Nets.", URL http://www.stanford.edu/{~}axl/. Working Paper, Stanford University.

Meier, Stephen and Charles Sprenger (2010), "Present-biased preferences and credit card borrowing." *American Economic Journal: Applied Economics*, 2, 193–210.

O'Donoghue, T. and M. Rabin (1999), "Doing it now or later." *American Economic Review*, 89, 103–124, URL http://www.jstor.org/stable/116981.

O'Donoghue, Ted and Matthew Rabin (2001), "Choice and procrastination." *Quarterly Journal of Economics*, 116, 121–160.

Phelps, Edmund S. and Robert A. Pollak (1985), "On second-best national saving and game-equilibrium growth." *American*, 75, 1071–7082.

Shapiro, Jesse M. (2005), "Is there a daily discount rate? evidence from the food stamp nutrition cycle." *Journal of Public Economics*, 89, 303–325.

Strotz, R. H. (1955), "Myopia and Inconsistency in Dynamic Utility Maximization." *The Review of Economic Studies*, 23, 165, URL http://www.jstor.org/stable/2295722?origin=crossref.

Tarozzi, A., A. Mahajan, J. Yoong, and B. Blackburn (2009), "Commitment mechanisms and compliance with health-protecting behavior: Preliminary evidence from Orissa, India." *American Economic Review (Papers and Proceedings)*, 99, 231–235, URL http://www.atypon-link.com/AEAP/doi/abs/10.1257/aer.99.2.231.

Tufano, Peter (2008), "Just keep my money! supporting tax-time savings with us savings bonds." Harvard Business School Finance Working Paper No. 09-059.

Wanrner, John T. and Saul Pleeter (2001), "The personal discount rate: Evidence from military downsizing programs." *American Economic Review*, 91, 33–53.

A Appendix

A.1 Power Calculations

In a standard context, one tests the null hypothesis of no treatment effect. In such a case, we typically estimate a treatment effect from the following linear regression:

$$a = bT + \Gamma \mathbf{X} + \varepsilon$$

where T is a dummy variable indicating (randomly assigned) treatment group membership. In our example T may indicate that a tax filer was offered a savings match and a indicates whether the individual used the savings account, and $\mathbf{X}$ is a vector of baseline characteristics and includes an intercept, which help increase the precision of the estimates but are independent of T. For testing the hypothesis, $H_0 : b = 0$, it is straightforward to calculate the required sample size for a test at a 5% significance level and 80% power:

$$N = \frac{\pi (1 - \pi) \left(1 - R^2\right)}{\bar{T} \left(1 - \bar{T}\right)} \times \left(\frac{2.49}{d}\right)^2$$

where π is the share of participants saving, $\bar{T}$ is the share of participants in the treatment group, R^2 is taken from the regression above, the 2.49 follows from the size and power of the test and d is the minimum detectable effect, or the smallest b that we can rule out. Our case is slightly more complex. First, we are estimating four treatment effects, two in each of two groups:

$$a_c = b_e T_1 + b_l T_2 + \Gamma_c \mathbf{X} + \varepsilon_c \tag{A.1}$$

$$a_n = b_i T_4 + b_d T_5 + \Gamma_n \mathbf{X} + \varepsilon_n \tag{A.2}$$

where a_c is the commitment decision for the commitment group and a_n is the savings decision for the non-commitment group. The four treatment indicators, $T_j,\ j \in \{1, 2, 4, 5\}$, are the early, late, immediate and delayed treatments, respectively. The commitment group in each regression is the baseline sub-group (i.e. Treatment Groups 3 and 6). For example, the coefficient b_i measures how much more likely an individual with an immediate incentive will save relative to an individual in the non-commitment group that received the baseline savings match. Ultimately, we will be comparing the ratio of the early and late treatments in the commitment context to the ratio of the immediate and delayed treatments in the non-commitment context. We would like to test the following nonlinear hypothesis:

$$H_0 : \triangle = \frac{b_l}{b_e} - \frac{b_d}{b_i} = 0$$

against the alternative hypothesis:

$$H_A : \triangle = \frac{b_l}{b_e} - \frac{b_d}{b_i} > 0$$

Note that given our model, we can rewrite the statistic $\triangle$ as:

$$\triangle = \frac{b_l}{b_e} - \frac{b_d}{b_i} = \delta - \beta\delta = (1 - \beta)\,\delta$$

So that our we are testing whether or not $\beta = 1$. Now we will derive the necessary sample size for this test. Using standard results, we have the variance-covariance matrix for the four reduce form parameters:

$$V\left(b_e, b_l, b_i, b_d\right) =$$

$$\begin{bmatrix} \frac{\sigma_c^2\left(1-R_c^2\right)}{\sigma_{T_1}^2\left(1-R_{T_1}^2\right)N_c} & -\frac{\sigma_c^2\left(1-R_c^2\right)\gamma_{T_2}}{\sigma_{T_2}^2\left(1-R_{T_2}^2\right)N_c} & \multicolumn{2}{c}{\underset{2\times 2}{\mathbf{0}}} \\ -\frac{\sigma_c^2\left(1-R_c^2\right)\gamma_{T_1}}{\sigma_{T_1}^2\left(1-R_{T_1}^2\right)N_c} & \frac{\sigma_c^2\left(1-R_c^2\right)}{\sigma_{T_2}^2\left(1-R_{T_2}^2\right)N_c} & & \\ \multicolumn{2}{c}{\underset{2\times 2}{\mathbf{0}}} & \frac{\sigma_n^2\left(1-R_n^2\right)}{\sigma_{T_4}^2\left(1-R_{T4}^2\right)N_n} & -\frac{\sigma_n^2\left(1-R_c^2\right)\gamma_{T_5}}{\sigma_{T_5}^2\left(1-R_{T_5}^2\right)N_n} \\ & & -\frac{\sigma_n^2\left(1-R_n^2\right)\gamma_{T_4}}{\sigma_{T_4}^2\left(1-R_{T4}^2\right)N_n} & \frac{\sigma_n^2\left(1-R_n^2\right)}{\sigma_{T_5}^2\left(1-R_{T_5}^2\right)N_n} \end{bmatrix}$$

where N_c and N_n are the sample sizes, σ_c^2 and σ_n^2 are the variance of the dependent variable and R_c^2 and R_n^2 are the R^2 statistics from (A.1) and (A.2) above. The γ_{T_j}'s are defined from the following regressions:

$$\begin{aligned} T_1 &= \gamma_{T_1} T_2 + \Gamma_1 \mathbf{X} + u_1 \\ T_2 &= \gamma_{T_2} T_1 + \Gamma_2 \mathbf{X} + u_2 \\ T_4 &= \gamma_{T_4} T_5 + \Gamma_4 \mathbf{X} + u_4 \\ T_5 &= \gamma_{T_5} T_4 + \Gamma_5 \mathbf{X} + u_5 \end{aligned}$$

and the $\sigma_{T_j}^2$ and $R_{T_j}^2$ are likewise the variance of the treatment variables and R^2 from these ancillary regressions. Note that due to random assignment the $\Gamma_j = 0$ for all j. We know exactly what the coefficients, γ_{T_j}, are from the T_j regressions as well as the $\sigma_{T_j}^2$ and $R_{T_j}^2$, since they are dictated by random assignment. So the covariance can be reduced to:

$$V\left(b_e, b_l, b_i, b_d\right) =$$

$$\begin{bmatrix} \frac{\sigma_c^2\left(1-R_c^2\right)\left(1-\bar{T}_2\right)}{\bar{T}_1\left(1-\bar{T}_1-\bar{T}_2\right)N_c} & \frac{\sigma_c^2\left(1-R_c^2\right)}{\left(1-\bar{T}_1-\bar{T}_2\right)N_c} & \multicolumn{2}{c}{\underset{2\times 2}{\mathbf{0}}} \\ \frac{\sigma_c^2\left(1-R_c^2\right)}{\left(1-\bar{T}_1-\bar{T}_2\right)N_c} & \frac{\sigma_c^2\left(1-R_c^2\right)\left(1-\bar{T}_1\right)}{\bar{T}_2\left(1-\bar{T}_1-\bar{T}_2\right)N_c} & & \\ \multicolumn{2}{c}{\underset{2\times 2}{\mathbf{0}}} & \frac{\sigma_n^2\left(1-R_n^2\right)\left(1-\bar{T}_5\right)}{\bar{T}_4\left(1-\bar{T}_4-\bar{T}_5\right)N_n} & \frac{\sigma_n^2\left(1-R_n^2\right)}{\left(1-\bar{T}_4-\bar{T}_5\right)N_n} \\ & & \frac{\sigma_n^2\left(1-R_n^2\right)}{\left(1-\bar{T}_4-\bar{T}_5\right)N_n} & \frac{\sigma_n^2\left(1-R_n^2\right)\left(1-\bar{T}_4\right)}{\bar{T}_5\left(1-\bar{T}_4-\bar{T}_5\right)N_n} \end{bmatrix}$$

Where $\bar{T}_j$ is the share of individuals in a particular regression, (A.1) or (A.2), that are in treatment group j. Note, because we have an equal probability of assignment to treatment groups, $\bar{T}_j = \bar{T} = 1/3$. Next, we define the Jacobian of the test statistic:

$$J\left(\triangle\right) = \left[\frac{-b_l}{b_e^2}, \frac{1}{b_e}, \frac{b_d}{b_i^2}, \frac{-1}{b_i}\right]$$

and therefore, by the delta method, we have:

$$V\left(\triangle\right) = J\left(\triangle\right) V\left(b\right) J\left(\triangle\right)'$$

where $b = (b_1, b_2, b_3, b_4)$. Furthermore, under the null, we have:

$$\begin{aligned}\frac{b_l}{b_e} &= \delta \\ \frac{b_d}{b_i} &= \beta\delta\end{aligned}$$

Thus, the variance of the test statistic reduces to:

$$V(\triangle) = \left[\frac{1+\delta^2-\delta}{b_e^2}\cdot\frac{\bar{a}_c(1-\bar{a}_c)\left(1-R_c^2\right)}{N_c} + \frac{1+\beta^2\delta^2-\beta\delta}{b_i^2}\cdot\frac{\bar{a}_n(1-\bar{a}_n)\left(1-R_n^2\right)}{N_n}\right]\times\frac{\left(1-\bar{T}\right)}{\bar{T}\left(1-2\bar{T}\right)}$$

Where $\bar{a}_c$ is the share of individuals in the commitment group that make pre-commitments and $\bar{a}_n$ is the share of individuals in the non-commitment group that save. Now we can derive an expression for our required sample sizes, at a 5% significance level and 80% power[39]:

$$\begin{aligned}N_n &= \left[\frac{1+\delta^2-\delta}{b_e^2}\cdot\frac{\bar{a}_c(1-\bar{a}_c)\left(1-R_c^2\right)}{\theta_c} + \frac{1+\beta^2\delta^2-\beta\delta}{b_i^2}\cdot\bar{a}_n(1-\bar{a}_n)\left(1-R_n^2\right)\right]\times\frac{\left(1-\bar{T}\right)}{\bar{T}\left(1-2\bar{T}\right)}\times\left(\frac{2.49}{\delta-\beta\delta}\right)^2 \\ N_c &= \theta_c N_n\end{aligned}$$

As can be seen, we have one degree of freedom, in that we can adjust the relative samples in the commitment and non-commitment groups while maintaining significance and power. This can be useful if one outcome, say commitment, generates more precise coefficients. Or if one type of treatment is more expensive than the other. Our total sample size for the study will be $N^* = N_c + N_n$.

There are many unknowns in this expression. In particular, the terms b_e^2 and b_i^2 enter into the denominator, which will tend to inflate our sample size needs. For time preference parameters, we will use estimates from the literature: $\beta = 0.7$ and $\delta = 0.97$. From our pilot study, we have $\pi_c \approx 0.58$, $b_e \approx 0.32$ and $R_c^2 \approx 0.33$, $\pi_n \approx 0.70$, $b_i \approx 0.57$ and $R_n^2 \approx 0.62$. With these values, the sample size minimizing value of θ is 2.82 and $N^* = 1,238$, with $N_n = 324$ and $N_c = 914$.

[39] Here, we have used the fact that at the given significance and power levels, $\delta - \beta\delta = 2.49\sqrt{V(\Delta)}$.

A.2 Extensive Form Representation of SaveUp and SaveUpFront Accounts

The examples below correspond to a tax filer who is deciding whether or not to save $600 of the income tax refund.

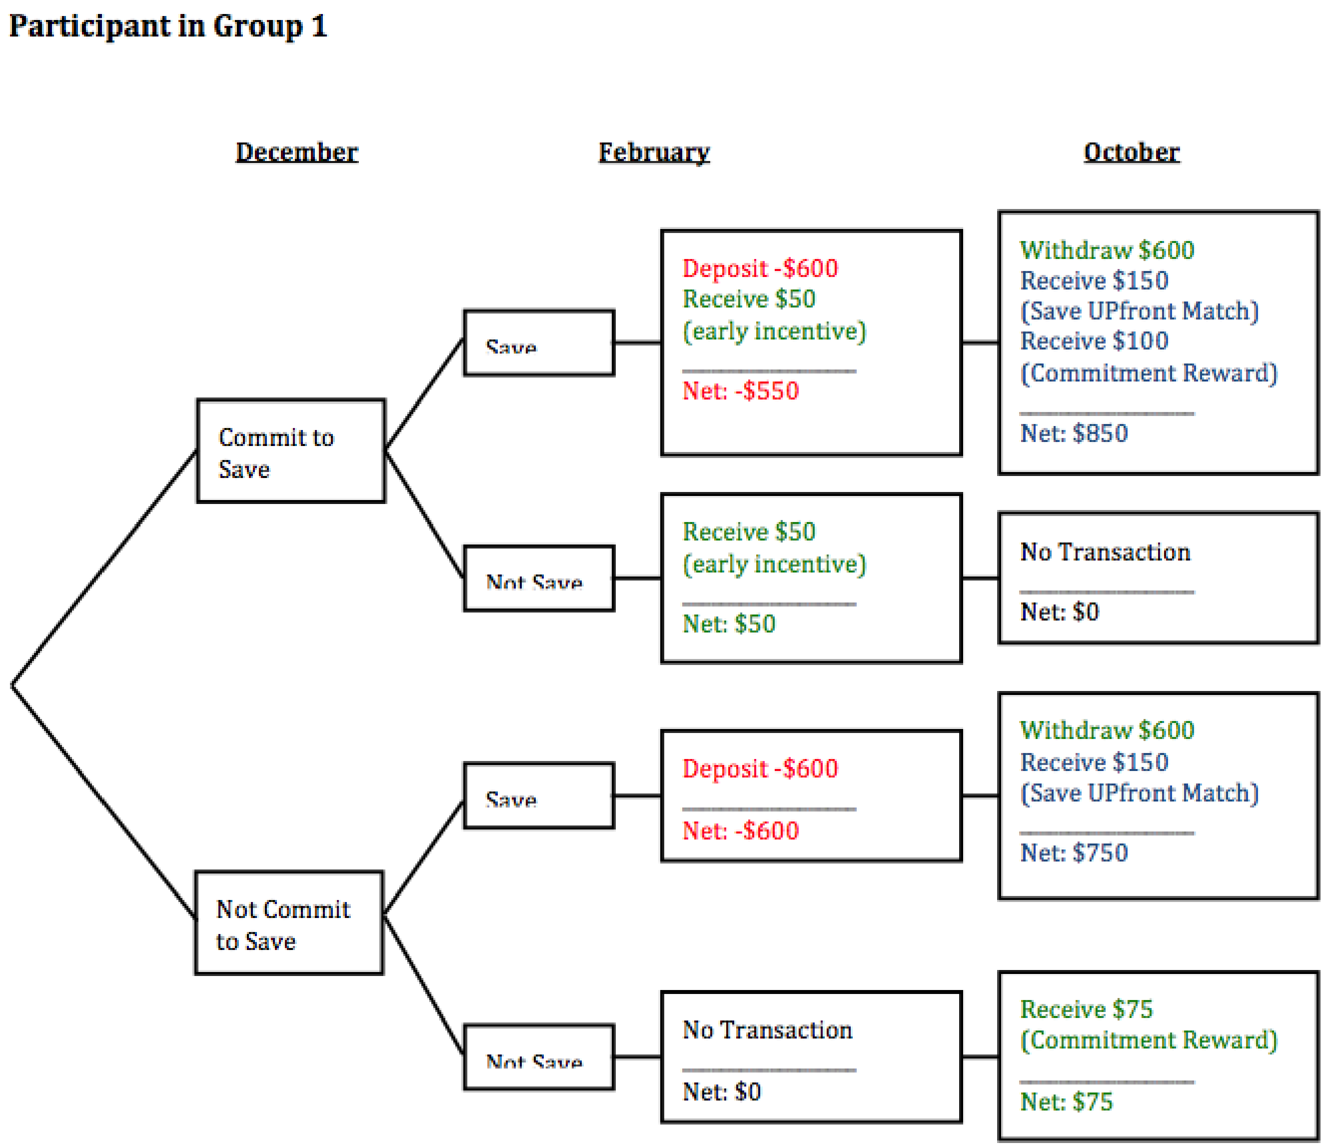

Participant in Group 2

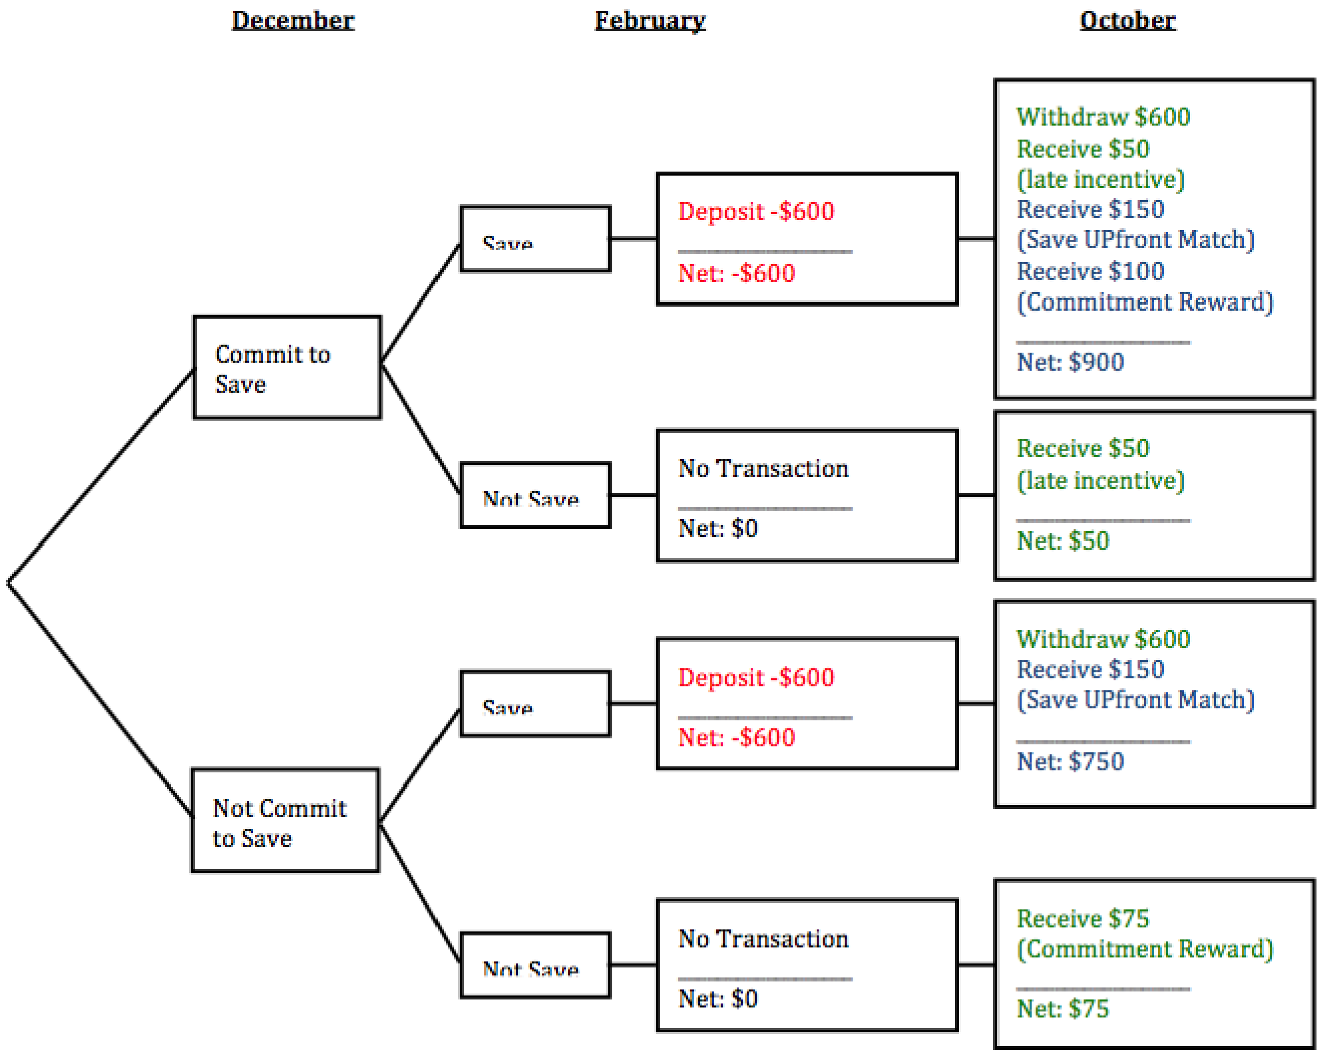

Participant in Group 3

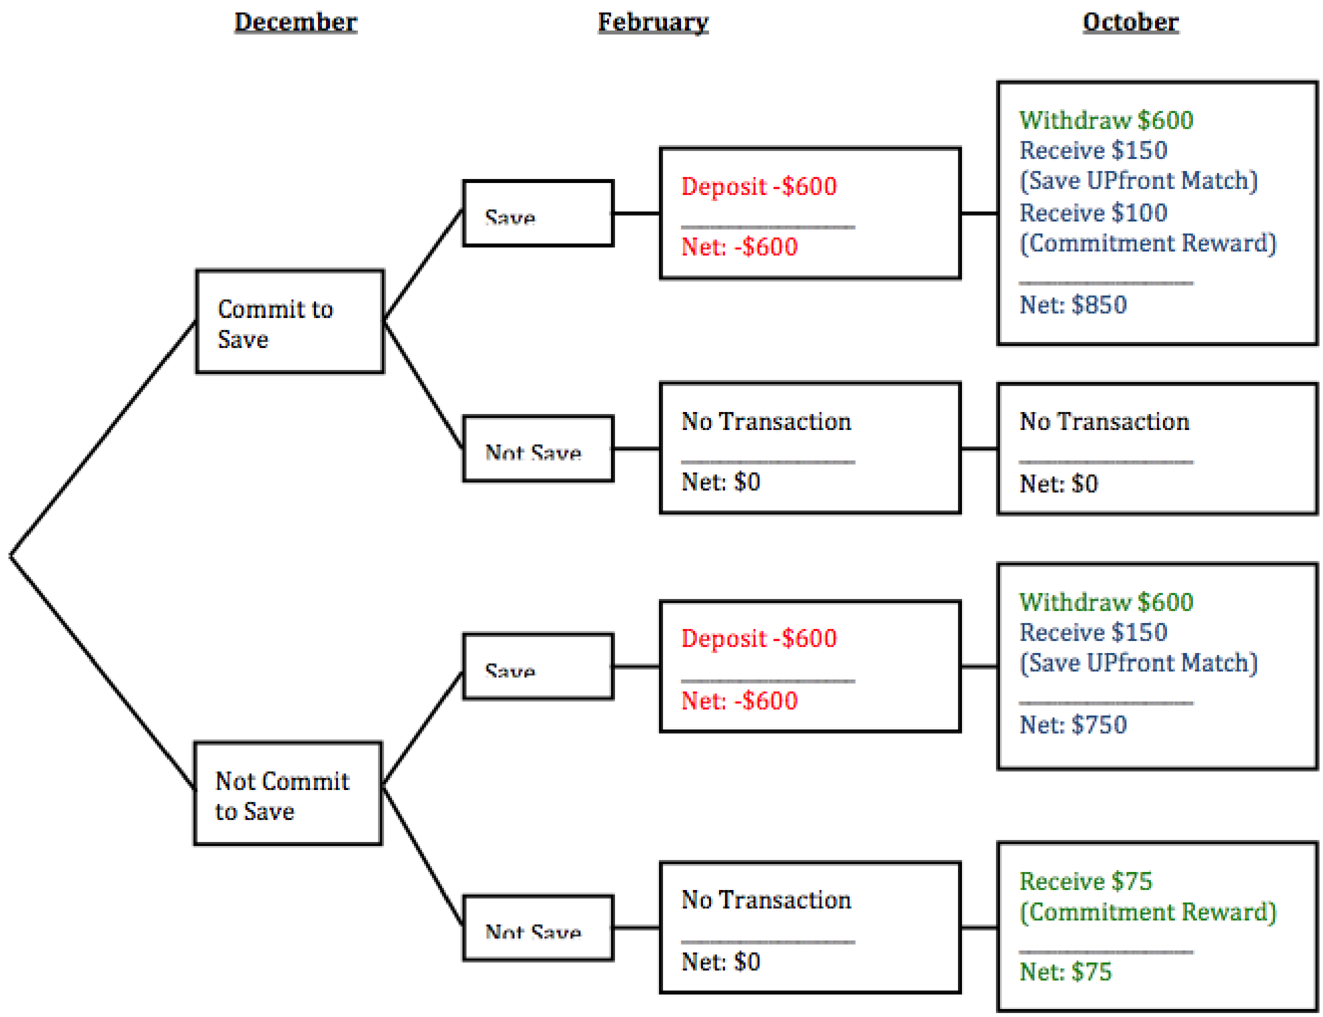

Participant in Group 4

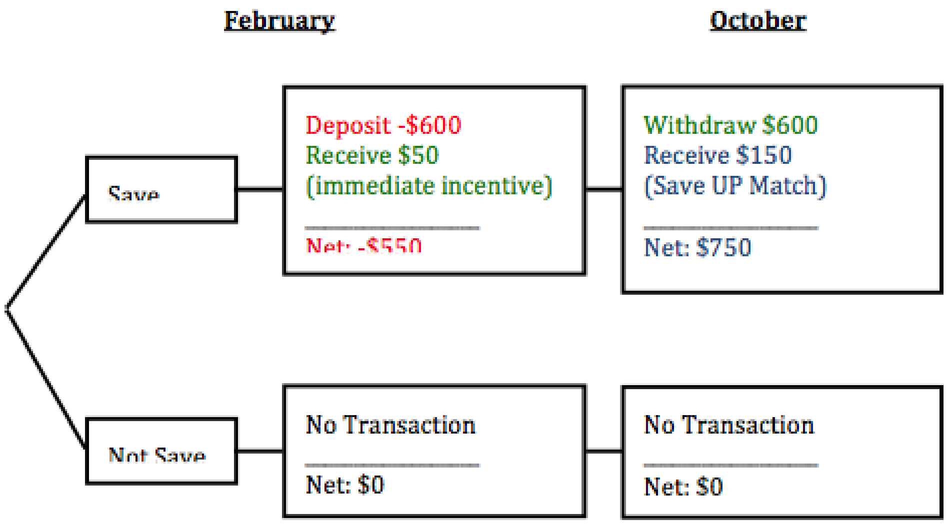

Participant in Group 5

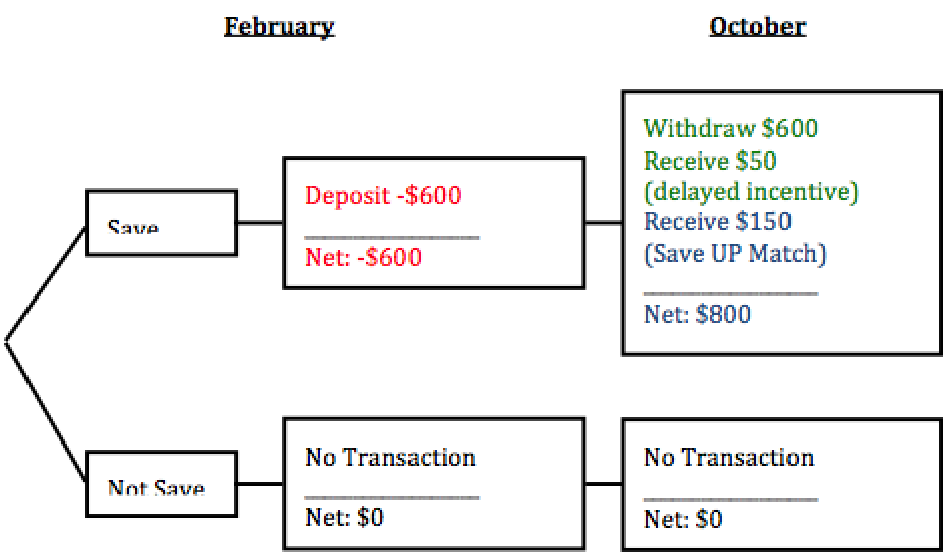

Participant in Group 6

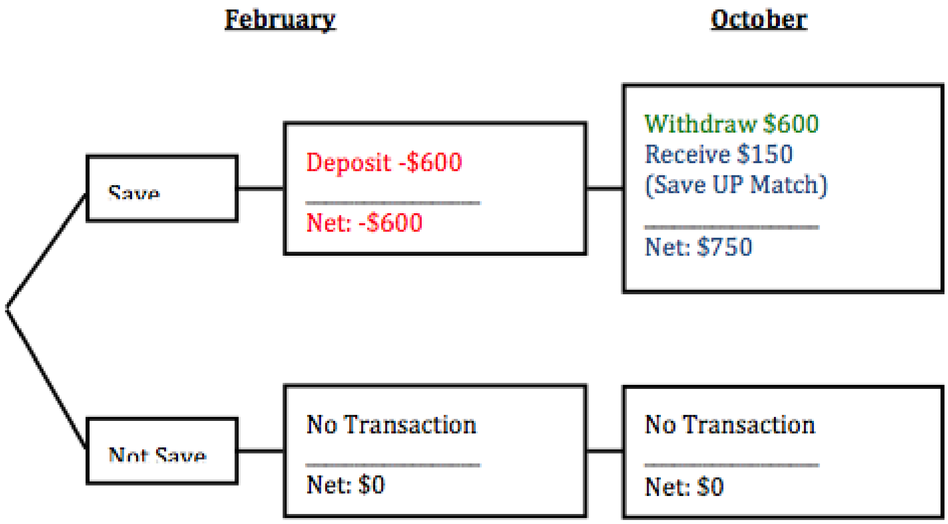

A.3 Treatment Bounding Method

A.3.1 Setup

- Define Z to be treatment assignment, which is binary valued and equal to 1 if the unit is assigned to the treatment group and 0 if the unit is assigned to the control group. (In our case, the control group is the baseline treatment group 3 or 6, and the treatment is one of the early/immediate or late/delayed incentive groups)

- Define C_z, a binary variable, which indicates participation in the study. This is either phone contact status or site appearance for a unit if her treatment is assigned to z. Define the observed random variable

$$C = C_1 Z + C_0 (1 - Z)$$

- Define Y_z to be the pre-commitment decision or savings decision if the unit's treatment assignment is set to z and define

$$Y = Y_1 Z + Y_0 (1 - Z)$$

- We assumed that $(Y_1, Y_0, C_1, C_0) \perp Z$

- The following distributions are identified

$$\Pr(Y, Z | C = 1)$$
$$\Pr(Z, C)$$

A.3.2 Bounding $\theta_1 \equiv \mathbb{E}(Y_1 - Y_0)$

Following Horowitz and Manksi (2000), we derive bounds on this ITT effect. Note, that for each z:

$$\begin{aligned}
\mathbb{E}(Y_z) &\equiv \mathbb{E}(Y|Z=z, C=1)\Pr(C=1|Z=z) + \mathbb{E}(Y|Z=z, C=0)\Pr(C=0|Z=z) \\
&= \mathbb{E}(Y|Z=z, C=1)\Pr(C=1|Z=z) + \mathbb{E}(Y_z|Z=z, C=0)\Pr(C=0|Z=z) \\
&= \mathbb{E}(Y|Z=z, C=1)\Pr(C=1|Z=z) + \mathbb{E}(Y_z|C=0)\Pr(C=0|Z=z)
\end{aligned}$$

Since, $Y_z \in \{0,1\}$ we can bound this by using the two extreme assumptions:

$$\begin{aligned}
C_z &= 0 \Longrightarrow Y_z = 1 \\
C_z &= 0 \Longrightarrow Y_z = 0
\end{aligned}$$

which in turn imply that:

$$\mathbb{E}(Y_z) \in \left[\mathbb{E}(Y|Z=z, C=1)\Pr(C=1|Z=z), \mathbb{E}(Y|Z=z, C=1)\Pr(C=1|Z=z) + \Pr(C=0|Z=z)\right]$$

Taking the lower bound for $\mathbb{E}(Y_1)$ and the upper bound for $\mathbb{E}(Y_0)$ to generate a lower bound on θ_1 and vice versa for an upper bound on θ_1, we have

$$\begin{aligned}
\theta_1 \in\ & [\mathbb{E}(Y|Z=1, C=1)\Pr(C=1|Z=1) - \mathbb{E}(Y|Z=0, C=1)\Pr(C=1|Z=0) - \Pr(C=0|Z=0), \\
& \mathbb{E}(Y|Z=1, C=1)\Pr(C=1|Z=1) + \Pr(C=0|Z=1) - \mathbb{E}(Y|Z=0, C=1)\Pr(C=1|Z=0)]
\end{aligned}$$

Thus, we have bounds on θ_1

$$\begin{aligned}
\theta_{1,ub} &= \mathbb{E}(YC|Z=1) - \mathbb{E}(YC|Z=0) + \Pr(C=0|Z=1) \\
\theta_{1,lb} &= \mathbb{E}(YC|Z=1) - \mathbb{E}(YC|Z=0) - \Pr(C=0|Z=0)
\end{aligned}$$

A.3.3 Bounding $\theta_2 \equiv \mathbb{E}(Y_1 - Y_0|C_1 = 1)$

We can potentially get sharper bounds on a similar treatment effect for a subgroup, namely those who are successfully contacted by phone or appearing at the tax site. Following Behaghel et al. (2009), we can see that:

$$\begin{aligned}
\theta_2 &\equiv \mathbb{E}(Y_1 - Y_0|C_1 = 1) \\
&= \frac{\mathbb{E}((Y_1 - Y_0)C_1)}{\Pr(C_1 = 1)} \\
&= \frac{\mathbb{E}(Y_1C_1 - Y_0C_0 - Y_0(C_1 - C_0))}{\Pr(C_1 = 1)} \\
&= \frac{\mathbb{E}(Y_1C_1) - \mathbb{E}(Y_0C_0) - \mathbb{E}(Y_0(C_1 - C_0))}{\Pr(C_1 = 1)} \\
&= \frac{\mathbb{E}(YC|Z=1) - \mathbb{E}(YC|Z=0) - \mathbb{E}(Y_0(C_1 - C_0))}{\Pr(C=1|Z=1)}
\end{aligned}$$

To further simplify, use an assumption of monotonicity:

Assumption 1 $C_1 \geq C_0$

Now, assuming $C_1 \geq C_0$ and using the fact that $Y_0 \in \{0, 1\}$

$$\begin{aligned} 0 &\leq \mathbb{E}\left(Y_0\left(C_1 - C_0\right)\right) \\ &\leq \mathbb{E}\left(C_1 - C_0\right) \\ &= \Pr\left(C_1 = 1\right) - \Pr\left(C_0 = 1\right) \\ &= \Pr\left(C = 1 | Z = 1\right) - \Pr\left(C = 1 | Z = 0\right) \end{aligned}$$

Thus, we have bounds on θ_2:

$$\begin{aligned} \theta_{2,ub} &= \frac{1}{\Pr\left(C = 1 | Z = 1\right)}\left(\mathbb{E}\left(YC | Z = 1\right) - \mathbb{E}\left(YC | Z = 0\right)\right) \\ \theta_{2,lb} &= \frac{1}{\Pr\left(C = 1 | Z = 1\right)}\left(\mathbb{E}\left(YC | Z = 1\right) - \mathbb{E}\left(YC | Z = 0\right) - \left(\Pr\left(C = 1 | Z = 1\right) - \Pr\left(C = 1 | Z = 0\right)\right)\right) \end{aligned}$$

A.3.4 Bounding $\theta_3 \equiv \mathbb{E}\left(Y_1 - Y_0 | C_0 = 1, C_1 = 1\right)$

Following Lee (2002), we can also estimate the treatment effect for a smaller subsample:

$$\begin{aligned} \theta_3 &\equiv \mathbb{E}\left(Y_1 - Y_0 | C_0 = 1, C_1 = 1\right) \\ &= \mathbb{E}\left(Y_1 | C_0 = 1, C_1 = 1\right) - \mathbb{E}\left(Y_0 | C_0 = 1, C_1 = 1\right) \end{aligned}$$

Again, assuming monotonicity:

$$\begin{aligned} C_1 &\geq C_0 \\ \Rightarrow \Pr\left(C_1 = 0 | C_0 = 1\right) &= 0 \\ \Rightarrow \Pr\left(C_1 = 1 | C_0 = 1\right) &= 1 \end{aligned}$$

It follows that:

$$\begin{aligned} \mathbb{E}\left(Y_0 | C_0 = 1, C_1 = 1\right) &= \mathbb{E}\left(Y_0 | C_0 = 1, C_1 = 1\right)\Pr\left(C_1 = 1 | C_0 = 1\right) \\ &\quad + \mathbb{E}\left(Y_0 | C_0 = 1, C_1 = 0\right)\Pr\left(C_1 = 0 | C_0 = 1\right) \\ &= \mathbb{E}\left(Y_0 | C_0 = 1\right) \\ &= \mathbb{E}\left(Y | C = 1, Z = 0\right) \end{aligned}$$

Thus, we have that:

$$\theta_3 = \mathbb{E}\left(Y_1 | C_0 = 1, C_1 = 1\right) - \mathbb{E}\left(Y | C = 1, Z = 0\right)$$

Also, we have that:

$$\mathbb{E}\left(Y_1 | C_1 = 1\right) = p\mathbb{E}\left(Y_1 | C_0 = 1, C_1 = 1\right) + (1 - p)\,\mathbb{E}\left(Y_1 | C_0 = 0, C_1 = 1\right)$$

where:

$$
\begin{aligned}
p &= \Pr\left(C_0 = 1 | C_1 = 1\right) \\
&= \frac{\Pr\left(C_0 = 1, C_1 = 1\right)}{\Pr\left(C_1 = 1\right)} \\
&= \frac{\Pr\left(C_0 = 1\right)}{\Pr\left(C_1 = 1\right)} \\
&= \frac{\Pr\left(C = 1 | Z = 0\right)}{\Pr\left(C = 1 | Z = 1\right)}
\end{aligned}
$$

The second to last line follows from monotonicity, and the last line follows from independence of (C_1, C_0) and Z. Now, we know that $C_1 = C_0 = 1$ for p of the individuals for whom $Z = 1$ and $C = 1$. Therefore, their mean Y_1 is bounded above by the top p of that group and bounded below by the bottom p of that group. Thus, we have the following:

$$
\mathbb{E}\left(Y | Z = 1, C = 1, Y \geq y_{1-p}\right) \geq \mathbb{E}\left(Y_1 | C_0 = 1, C_1 = 1\right) \geq \mathbb{E}\left(Y | Z = 1, C = 1, Y \leq y_p\right)
$$

where:

$$
y_q \equiv G^{-1}\left(q\right), \text{ with } G \text{ the c.d.f. of } Y \text{ conditional on } Z = 1,\ C = 1.
$$

Note, the definition of y_q is not well defined in the case of this discrete function, since G is not a one-to-one function. In any event, since $Y \in \{0, 1\}$, we know that:

$$
\begin{aligned}
\mathbb{E}\left(Y | Z = 1, C = 1, Y \geq y_{1-p}\right) &= \min\left(1, \frac{\mathbb{E}\left(Y | Z = 1, C = 1\right)}{p}\right) \\
\mathbb{E}\left(Y | Z = 1, C = 1, Y \leq y_p\right) &= \max\left(0, \frac{\mathbb{E}\left(Y | Z = 1, C = 1\right)}{p} - \frac{1-p}{p}\right)
\end{aligned}
$$

Looking at the upper bound, for example, if the share of those for whom $Z = 1$ and $C = 1$ with $Y = 1$ is larger than p, then $Y = 1$ for everyone in the top p of that group. Otherwise, everyone with $Y = 1$ is in the top p of that group. So, the mean Y of the top p is found by taking the fraction of everyone with a $Y = 1$ and normalizing it to the top p in the group. Finally, we have the bounds:

$$
\begin{aligned}
\theta_{3,ub} &= \min\left\{1, \frac{\mathbb{E}\left(Y | Z = 1, C = 1\right)}{p}\right\} - \mathbb{E}\left(Y | C = 1, Z = 0\right) \\
\theta_{3,lb} &= \max\left\{0, \frac{\mathbb{E}\left(Y | Z = 1, C = 1\right)}{p} - \frac{1-p}{p}\right\} - \mathbb{E}\left(Y | C = 1, Z = 0\right)
\end{aligned}
$$

A.3.5 Comparing the width of the bounds

First, we see that the width of the bounds on θ_1 is

$$
\theta_{1,ub} - \theta_{1,lb} = \Pr\left(C = 0 | Z = 1\right) + \Pr\left(C = 0 | Z = 0\right)
$$

We can also see from above that the width of the bounds on θ_2 is:

$$
\theta_{2,ub} - \theta_{2,lb} = \frac{\Pr\left(C = 1 | Z = 1\right) - \Pr\left(C = 1 | Z = 0\right)}{\Pr\left(C = 1 | Z = 1\right)}
$$

Comparing the first two sets of bounds, we have that

$$\begin{aligned}
\theta_{1,ub}-\theta_{1,lb} &= \Pr(C=0|Z=1)+\Pr(C=0|Z=0) \\
&\geq \Pr(C=0|Z=0) \\
&= 1-\Pr(C=1|Z=0) \\
&\geq 1-\frac{\Pr(C=1|Z=0)}{\Pr(C=1|Z=1)} \\
&= \frac{\Pr(C=1|Z=1)-\Pr(C=1|Z=0)}{\Pr(C=1|Z=1)} \\
&= \theta_{2,ub}-\theta_{2,lb}
\end{aligned}$$

The width of the bounds on θ_3 is:

$$\theta_{3,ub}-\theta_{3,lb}=\begin{cases} 1 & \text{if } \min\{\mathbb{E}(Y|Z=1,C=1),1-\mathbb{E}(Y|Z=1,C=1)\}\geq p \\ \frac{1-p}{p} & \text{if } p\geq\max\{\mathbb{E}(Y|Z=1,C=1),1-\mathbb{E}(Y|Z=1,C=1)\} \\ \frac{1-\mathbb{E}(Y|Z=1,C=1)}{p} & \text{if } \mathbb{E}(Y|Z=1,C=1)\geq p\geq 1-\mathbb{E}(Y|Z=1,C=1) \\ \frac{\mathbb{E}(Y|Z=1,C=1)}{p} & \text{if } 1-\mathbb{E}(Y|Z=1,C=1)\geq p\geq\mathbb{E}(Y|Z=1,C=1) \end{cases}$$

When comparing the second and third set of bounds, we have four cases. In the first case, when

$$\min\left(\mathbb{E}(Y|Z=1,C=1),1-\mathbb{E}(Y|Z=1,C=1)\right)\geq p$$

it's clear that:

$$\begin{aligned}
\theta_{2,ub}-\theta_{2,lb} &= \frac{\Pr(C=1|Z=1)-\Pr(C=1|Z=0)}{\Pr(C=1|Z=1)} \\
&= 1-p \\
&\leq 1 \\
&= \theta_{3,ub}-\theta_{3,lb}
\end{aligned}$$

In the second case, when $p\geq\max\left(\mathbb{E}(Y|Z=1,C=1),1-\mathbb{E}(Y|Z=1,C=1)\right)$, we can show that:

$$\begin{aligned}
\theta_{2,ub}-\theta_{2,lb} &= \frac{\Pr(C=1|Z=1)-\Pr(C=1|Z=0)}{\Pr(C=1|Z=1)} \\
&= 1-p \\
&\leq \frac{1-p}{p} \\
&= \theta_{3,ub}-\theta_{3,lb}
\end{aligned}$$

The third line follows from the monotonocity assumption, which implies that $p\leq 1$. In the last two cases, the relative size of the bound widths is ambiguous. First, when $\mathbb{E}(Y|Z=1,C=1)\geq p\geq 1-\mathbb{E}(Y|Z=1,C=1)$, then we have:

$$\begin{aligned}
p(1-p) &\leq 1-\mathbb{E}(Y|Z=1,C=1) \\
\Rightarrow \theta_{2,ub}-\theta_{2,lb} &\leq \theta_{3,ub}-\theta_{3,lb}
\end{aligned}$$

Second, when $1-\mathbb{E}(Y|Z=1,C=1)\geq p\geq \mathbb{E}(Y|Z=1,C=1)$ we have:

$$\begin{aligned} p(1-p) &\leq \mathbb{E}(Y|Z=1,C=1) \\ \Rightarrow \theta_{2,ub}-\theta_{2,lb} &\leq \theta_{3,ub}-\theta_{3,lb} \end{aligned}$$

Comparing the first and third set of bounds similarly depends on the values of p and $\mathbb{E}(Y|Z=1,C=1)$.

A.4 Extension of Model to Structural Dynamic Discrete Choice Framework

By imposing additional structural assumptions, we can recast our model in the standard dynamic discrete choice framework with potentially time-inconsistent agents (see Fang and Wang (2008) and Mahajan and Tarozzi (2010) for models along similar lines).

A.4.1 State Spaces, Action Spaces and Payoffs

<u>Period 1</u>

State Space: $x_1 \in \mathcal{X}_1$ where x_1 are pre-intervention observables that affect the agent's utility in period 1.

Action Space: $a_1 \in \{0,1\}$[40] where $a_1 = 1$ indicates that the agent chooses the soft "pre-commitment" decision to save in period 2. $a_1 = 0$ if the agent chooses the soft "pre-commitment" decision to not save in period 2. For simplicity in exposition the model does not consider an alternative third action in which the agent rejects the first two actions.[41]

Utility in Period 1: The model assumes standard per-period preferences. In particular, it requires that the unobserved component of utility is additively separable from and independent of the observed determinants of utility x_1. Specifically, $u(x_1, a_1, \epsilon_{a_1}) = \kappa(x_1, a_1) + \epsilon_{a_1}$. To focus only on essential details it is assumed that $\kappa(x_1, a_1) \equiv \kappa$ for some constant κ. (ϵ_0, ϵ_1) are unobserved i.i.d. random variables and $\epsilon_1 - \epsilon_0$ represents the net benefit of deciding to pre-commit to save against pre-committing to not saving.

<u>Period 2</u>

State Space: $x_2 \in \mathcal{X}_2 \equiv \{0,1\}$ where $x_2 = a_1$ just records the first period decision. It is straightforward to incorporate additional state variables using standard approaches but those are eliminated here in the interest of brevity.

Action Space: $a_2 \in \{0,1\}$ where $a_2 = 1$ indicates that the agent chooses to save and $a_2 = 0$ if the agent decides not to save.

Utility in Period 2: Payoffs in the second period vary by treatment group. In the non-commitment group (NC) agents receive a payment of i if they decide to save, representing the "immediate" savings incentive. This is the additional \$50 incentive for saving, which is contrasted with an equivalent "delayed" incentive received in period 3. In the commitment group (C), if the agent had chosen to pre-commit to saving $a_1 = 1$ she receives a payoff e (*irrespective* of her actual savings decision). This is referred to as the "early" incentive for pre-committing, which contrasts with the "late" incentive in period 3. In addition, agents in both groups are assumed to pay a cost c (i.e. the cost of saving c dollars, including a binding borrowing constraint) if they decide to save. As per usual, the model imposes additive separability between the unobserved determinants of choice: $u(x_2, a_2) = u_m(x_2, a_2) + \epsilon_{a_2}$, where $m \in \{NC, C\}$.

[40] Note that in the model, we only model actions related to the project choices. In principle, agents could make other decisions (e.g. change their saving and/or consumption behaviour) in response to this news. If for instance, individuals can easily borrow, then they can take other actions that would render the findings from our study uninteresting. However, we assume here that individuals face credit constraints (and we attempt to quantify them via the survey). If individuals are credit constrained, then the savings decision at tax-time is an inter-temporal utility trade-off rather than a mere arbitrage opportunity. Individuals can still re-optimize, however, by delaying debt-repayment. We hope to capture this via our survey instruments.

[41] It can be shown that this third action is weakly dominated.

The experimental design yields the following values for the $u_m(x_2, a_2)$ for the different treatment groups. Utility is denoted as a function f of the payoffs since in future work the model will be expanded to explore what can be learned about $f(\cdot)$ experimentally. However, for this proposal the model will deal mostly with the case where f is the identity function, i.e. utility is quasi-linear. This is a strong assumption and discussion of relaxing it can be found in our discussion.[42] Payoffs for period 2 are given by

Table A.1: Period 2 Payoffs by Treatment Group

(x_2, a_2)	$u_{NC}(x_2, a_2)$	$u_C(x_2, a_2)$
(1,1)	$f(i-c)$	$f(e-c)$
(1,0)	0	$f(e)$
(0,1)	$f(i-c)$	$f(-c)$
(0,0)	0	0

Note in the first column that $u_{NC}(x_2, a_2)$ is invariant to the state variable, since there is no pre-commitment for the non-commitment group. Looking at the first row of the last column, one can see for example that when a pre-commitment is made, i.e. $x_2 = 1$, the early savings incentive e is received. Furthermore, a cost c of saving is incurred for all agents that save in either group, i.e. when $a_2 = 1$. For each group, a normalization of utility in all states with respect to a reference action is needed. In particular, one can normalize $u(x_2, 0)$ for all $x_2 \in \{0, 1\}$. For the non-commitment group, this normalization is $u(x_2, 0) \equiv 0$. For the commitment group, the normalization assumes $f(e)$ is known.

<u>Period 3</u>

State Space: $x_3 \in \mathcal{X}_3 \equiv \{0, 1\} \times \{0, 1\}$ where $x_3 = (x_2, a_2)$.

Agents take no action in this period and payoffs are a function of group status and state variables. In the non-commitment group, if the agent had not saved in period 2, then she receives a payoff of 0. If she has saved, she receives a savings incentive of d. This payoff d is referred to as the "delayed" payoff. In the commitment group, the agent receives a payoff of p if the pre-commitment and actual savings decision coincide, i.e. $x_2 = a_2$. That is, if she had pre-committed to not saving and fulfilled that commitment in period 2 or if the agent had pre-committed to saving in period 1 and she followed through with that commitment, she again receives the a payoff p. Finally, for members of the commitment group, a payoff l is received *irrespective* of her period 2 saving decision. This is the "late" incentive for pre-committing to save. Finally, agents in either group who chose to save in period 2 receive a payoff of b, which is in effect the value of the amount now available to be withdrawn from the savings account.[43]

Table A.2: Period 3 Payoffs by Treatment Group

(x_3)	$u_{NC}(x_3)$	$u_C(x_3)$
(1,1)	$f(b+d)$	$f(b+l+p)$
(1,0)	0	$f(l)$
(0,1)	$f(b+d)$	$f(b)$
(0,0)	0	$f(p)$

[42] This type of quasi-linearity is found in other models that analyze present-biased preferences, e.g. DellaVigna and Malmendier (2002).
[43] This amount b is inclusive of the variable portion of the savings match, i.e. the 50% on each dollar above \$300.

STATE TRANSITIONS: State transitions $\mathrm{dF}(x_{t+1}|x_t, a_t)$ are extremely simple in this basic framework since they are completely deterministic functions of previous period actions.

DISCUSSION: Looking at Tables A.1 and A.2 one can get a general sense of the source of identification. First, looking at the columns for $u_{NC}(\cdot)$ one will notice that the incentives (i, d) vary depending on the decision in Period 2, a_2 ($\equiv x_3$) and are by construction invariant to the state variable x_2 ($\equiv a_1$) since members of the NC group do not make pre-commitments. Thus, by observing the response of a_2 to experimental variation in (i, d), one learns about Period 2 preferences over utility in Period 2 (i) relative to utility in Period 3 (d). In contrast, one can see in the columns for $u_C(\cdot)$, (e, l) vary depending on the decision in Period 1, a_1 ($\equiv x_2$), but are invariant to Period 2 decisions, a_2 ($\equiv x_3$). Thus, by observing the response of a_1 to experimental variation in (e, l) one learns about Period 1 preferences over utility in Period 2 (e) relative to utility in Period 3 (l). Finally, the comparison of the preferences in Period 1 and Period 2 provide the grounds for testing for time consistency. Note also that in Table A.2 the commitment reward p depends on both the Period 1 decision and the Period 2 decision. All things equal, pre-committing to save makes saving more likely to occur, and vice versa. In this sense, it is the mechanism by which Period 1 decisions can alter Period 2 decisions, and is necessary for making Period 1 decisions nontrivial (without $p \geq l$ all commitment group agents would choose $a_1 = 1$).[44]

A.4.2 Solving the Model

PERIOD 2

Since the model is a finite horizon problem, it is solved in a straightforward manner by backward induction. The first departure from standard dynamic discrete choice modeling is that the model allows for agents to be potentially present-biased. One can parameterize this time inconsistency using the tractable (β, δ) formulation described in Strotz (1955).[45] In particular, utility between now and the next period is discounted by a factor of $\beta\delta$, while utility between any two adjacent periods in the future is discounted only by a factor of δ. In period 2, an agent in group G will choose to save if

$$\underbrace{u_G(x_2, 1) + \beta\delta \int u(x_3)\,\mathrm{dF}(x_3|a_2 = 1, x_2)}_{v_G(x_2,1,\beta)} + \epsilon_1 > \underbrace{u_G(x_2, 0) + \beta\delta \int u(x_3)\,\mathrm{dF}(x_3|a_2 = 0, x_2)}_{v_G(x_2,0,\beta)} + \epsilon_0$$

where $v_G(x_2, a, \beta)$ is the choice a specific value function for treatment group G. There is an emphasis on the dependence of this function on the hyperbolic parameter β since it will play a key role in the subsequent analysis.

For the following, it will be assumed that the errors (ϵ_0, ϵ_1) are i.i.d. generalized extreme value (GEV). This is done for brevity in exposition, but note that the GEV assumption is not important for the identification results. In particular the results stated below will continue to hold as long as the errors are independent and have support over the entire real line. Using the GEV assumption, one can relate the observed choice probability $\mathbb{P}_G(a_2|x_2)$ for treatment group G to the model as

$$\mathbb{P}(a_2 = 1|x_2) = \mathrm{Logit}(v_G(x_2, 1, \beta) - v_G(x_2, 0, \beta))$$

where $\mathrm{Logit}(x) = \frac{\exp(x)}{1+\exp(x)}$. Since the differences in choice specific value functions will be useful later, the following table records the values here for the basic model.

[44] Again, note that the discrete decisions a_1 and a_2 are the decisions to pre-commit or to actually save above the minimum savings threshold or not. The matches (e, l) and (i, d) are the incentives received for pre-committing or actually saving more than these thresholds. The variable part of the savings match is proportional to the amount saved and is for convenience collapsed into the (c, b) variables in the model.

[45] Although this is the dominant model used in empirical investigations, this is certainly not the only possible formulation (see for instance Gul and Pesendorfer (2001), Gul and Pesendorfer (2004)).

Table A.3: Net value of saving in Period 2 by Treatment Group

x_2	$v_{NC}(x_2,1,\beta)-v_{NC}(x_2,0,\beta)$	$v_C(x_2,1,\beta)-v_C(x_2,0,\beta)$
1	$f(i-c)+\beta\delta f(b+d)$	$f(e-c)-f(e)+\beta\delta(f(b+l+p)-f(l))$
0	$f(i-c)+\beta\delta f(b+d)$	$f(-c)+\beta\delta\left(f(b)-f\left(p\right)\right)$

First, note that experimental variation in (i,d) as well as knowledge of the payoff functions $f(\cdot)$ identifies the product $\beta\delta$ in two rows of the first column of the previous table (recall that (c,b) are assumed known). Identification of $(\beta\delta)$ in the case where $f(\cdot)$ is unknown is considerably more complex and will depend upon the details of the class of functions to which $f(\cdot)$ belongs.

<u>Period 1</u>

For period 1, again assuming the error terms are i.i.d. GEV

$$\mathbb{P}(a_1=1|x_1)=\text{Logit}\left(v_C(x_1,1,\beta)-v_C(x_1,0,\beta)\right)$$

where as before, $v_C(\cdot)$ denotes the choice specific value function for members of the commitment group. Choice in period 1 is critical for identifying the time preference parameters. In particular, one can use the structure of the experimental payoffs (as well as the linearity assumption) to separately identify both the hyperbolic parameter β and the exponential parameter δ.

The exercise begins by defining the first period choice specific value functions

$$v_C(x_1,a,\beta)=u(x_1)+\beta\delta\int v_C^*(x_2)\,\mathrm{dF}(x_2|a_1=a)$$

where it is assumed that first period (pre-intervention) utility levels $u(x_1)$ are identical across treatment groups due to the randomization. Next, one can define the Emax function

$$v_C^*(x_2)=\int\sum_{a=0}^{1}\left(v_C(x_2,a,1)+\epsilon_a\right)\mathbb{I}\left(v_C(x_2,a,\hat{\beta})+\epsilon_a>v_C(x_2,a',\hat{\beta})+\epsilon_{a'}\ ,\ a\neq a'\right)\mathrm{dF}(\epsilon_0,\epsilon_1),$$

where $\mathbb{I}\left(\cdot\right)$ is the indicator function. The agent (while in period 1) and for given values of (ϵ_0,ϵ_1) thinks that she will choose to save in period 2 if and only if

$$v_C(x_2,1,\hat{\beta})+\epsilon_1>v_C(x_2,0,\hat{\beta})+\epsilon_0$$

so that in period 1, the agent assumes that her extent of present-bias in period 2 will be $\hat{\beta}$. For time consistent agents $\hat{\beta}=1$. In addition, naive time-inconsistent agents will also set $\hat{\beta}=1$ even though when they actually are in period 2 their extent of present-bias will be $\beta<1$. Finally, sophisticated time-inconsistent agents realize that their period 2 self is present-biased and set $\hat{\beta}=\beta$.

Under the assumption that (ϵ_0,ϵ_1) are i.i.d. GEV random variables it is easy to show that

$$v_C^*(x_2)=\sum_{s=0}^{1}\mathbb{P}_C(\tilde{a}_2=s|x_2)\left(v_C(x_2,s,1)-v_C(x_2,s,\hat{\beta})\right)+\gamma_{\texttt{euler}}+\log\left(\sum_{j=0}^{1}\exp(v_C(x_2,j,\hat{\beta}))\right)$$

where $\gamma_{\texttt{euler}}$ is Euler's constant and

$$\mathbb{P}_C(\tilde{a}_2 = s|x_2) = \text{Logit}\left(v_C(x_2, s, \hat{\beta}) - v_C(x_2, s', \hat{\beta});\;\; s \neq s'\right)$$

The term $v_C(x_2, s, 1) - v_C(x_2, s, \hat{\beta})$ is non-standard and captures the difference in the period 2 choice specific value function at two different values of the hyperbolic parameter. In the standard case $\hat{\beta} = 1$ and so this term disappears and what is left is the standard expression (see e.g. Aguirregabiria and Mira (2010)). It can be shown that

$$\begin{aligned}
v_C(x_1, 1, \beta) &= u(x_1) + \beta\delta v^*(1) \\
&= u(x_1) + \beta\delta\bigg\{\mathbf{A}\left((1-\hat{\beta})\delta f(b+d)\right) + (1-\mathbf{A})(1-\hat{\beta})\delta f(l) \\
&\quad + \log\left(e^{f(e)+\hat{\beta}f(l)} + \exp\left(f(e-c) + \hat{\beta}\delta f(b+l+p)\right)\right)\bigg\}
\end{aligned}$$

where

$$\mathbf{A} = \text{Logit}\left(f(e-c) - f(e) + \hat{\beta}\delta(f(b+l+p) - f(l))\right)$$

In what follows the further assumption that $f(i-c) - f(i) = f(-c)$ and $f(b+l+p) - f(l) = f(b+p)$ will be utilized. This linearity assumption is critical for the identification of the time preference parameters. The PIs are currently working on adding different experimental manipulations (increasing the number of treatments) to relax this condition. In particular, the PIs have shown that with the addition of another intervention it is possible to at least test (an implication of) this restriction. Assuming linearity, the value function for the soft-commitment decision to save is

$$v_C(x_1, 1, \beta) = u(x_1) + \beta\delta\bigg(f(e) + \delta f(l) + \log\left(1 + \exp(f(-c) + \hat{\beta}f(b))\right)\bigg)$$

The value function for the soft-commitment decision to not save is given by

$$\begin{aligned}
&v_C(x_1, 0, \beta) = u(x_1) + \beta\delta v^*(0) \\
&= u(x_1) + \beta\delta\bigg\{\text{Logit}\left(f(-c) + \hat{\beta}\delta\left(f(b) - f\left(p\right)\right)\right)\left((1-\hat{\beta})\delta\left(f(b) - f\left(p\right)\right)\right) \\
&+ \log\left(1 + \exp\left(f(-c) + \hat{\beta}\delta\left(f(b) - f\left(p\right)\right)\right)\right)\bigg\}
\end{aligned}$$

<u>IDENTIFYING TIME PREFERENCE PARAMETERS:</u> $(\delta, \beta, \hat{\beta})$

First, note that $\beta\delta$ is identified from the first step. Next, one can show that as long as the intervention payoffs (i, d) and (e, l) have at least two points of support conditional on (b, c) then δ (and therefore β) as well as $\hat{\beta}$ are identified. The identification argument is done by contradiction and is available upon request.

Time-Inconsistency and Savings:
Experimental Evidence from Low-Income Tax Filers

Center for Financial Security Working Paper 2011-CFS.6

Center for Financial Security
School of Human Ecology
University of Wisconsin-Madison
Sterling Hall Mailroom B605
475 North Charter Street
Madison, Wisconsin 53706
608.262.6766
cfs@mailplus.wisc.edu

RESEARCH · PRACTICE · POLICY