HARVARD Kennedy School
JOHN F. KENNEDY SCHOOL OF GOVERNMENT

Self Control and Commitment: Can Decreasing the Liquidity of a Savings Account Increase Deposits?

Faculty Research Working Paper Series

John Beshears
Harvard Business School

James J. Choi
Yale School of Management

Christopher Harris
Cambridge University

David Laibson
Harvard University

Brigitte C. Madrian
Harvard Kennedy School

Jung Sakong
University of Chicago

August 2015

RWP15-048

Acknowledgements

This paper was formerly titled "Self Control and Liquidity: How to Design a Commitment Contract." This research was made possible by grants from the Russell Sage Foundation and the Sloan Foundation (joint grant 981011), the Pershing Square Fund for Research on the Foundations of Human Behavior, the National Institute on Aging of the National Institutes of Health (awards P01AG005842, P30AG034532, and R01AG021650), and the Social Security Administration (grant RRC08098400).

We have benefited from the comments of George-Marios Angeletos, B. Douglas Bernheim, Bruce Carlin, Stefano DellaVigna, Ulrike Malmendier, Sendhil Mullainathan, Christopher Parsons, and Ivan Werning. We thank Christopher Clayton, Peter Maxted, and Sean (Yixiang) Wang for excellent research assistance. The findings and conclusions expressed are solely those of the authors and do not represent the views of the Russell Sage Foundation, the Sloan Foundation, the Pershing Square Fund, the National Institutes of Health, the Social Security Administration, or any agency of the Federal Government. The views expressed herein are those of the authors and do not necessarily reflect the views of the Harvard Kennedy School or the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w21474.ack

Abstract

If individuals have self-control problems, they may take up commitment contracts that restrict their spending. We experimentally investigate how contract design affects the demand for commitment contracts. Each participant divides money between a liquid account, which permits unrestricted withdrawals, and a commitment account with withdrawal restrictions that are randomized across participants. When the two accounts pay the same interest rate, the most illiquid commitment account attracts more money than any of the other commitment accounts. We show theoretically that this pattern is consistent with the presence of sophisticated present-biased agents, who prefer more illiquid commitment accounts even if they are subject to uninsurable marginal utility shocks drawn from a broad class of distributions. When the commitment account pays a higher interest rate than the liquid account, the relationship between illiquidity and deposits is flat, suggesting that agents without present bias and/or naïve present-biased agents are also present in our sample.

In the first quarter of 2015, U.S. households held $14.4 trillion in employer-sponsored defined contribution savings plans and IRAs (Investment Company Institute, 2015). These retirement savings accounts are partially illiquid: withdrawals before age 59½ incur a 10% tax penalty.[1] There are at least two mutually compatible arguments for why early withdrawal penalties are socially desirable. First, the penalties may address moral hazard problems (discouraging mid-life spending reduces the social burden of supporting retirees). Second, the penalties may help agents with self-control problems commit not to prematurely spend their savings.[2]

If early withdrawal penalties are potentially beneficial, an open question is how high they should be. While higher penalties may reduce premature withdrawals, they may also discourage deposits, defeating the goal of raising net savings. On the other hand, if savers recognize that penalties help them overcome self-control problems, they may welcome higher penalties and make *more* deposits in response. In this paper, we use experiments to study how the demand for commitment savings accounts is affected by the strength of the commitment—that is, how illiquid the accounts are—and analyze theoretically whether present-biased agents would indeed make more commitment deposits in response to higher penalties.

The 1,045 participants in our two experiments are drawn from the American Life Panel, which is selected to be broadly representative of the U.S. adult population. Each participant is given $50, $100, or $500 and asked to allocate this money between a liquid account, which does not limit withdrawals, and either one or two commitment accounts. All participants have access to the same kind of liquid account, but the characteristics of

---

[1] However, it is often possible to access 401(k) account balances by taking a penalty-free loan. In addition, the penalty on withdrawals is sometimes waived. For example, for IRA accounts no penalty is charged when the account holder (i) is permanently or totally disabled; (ii) has medical expenses exceeding 7.5% of her adjusted gross income; (iii) uses the withdrawal to buy, build, or rebuild a home if the withdrawal is no more than $10,000 and she has not owned a home in the previous two years; (iv) uses the withdrawal to pay higher education costs; (v) uses the withdrawal to make a back tax payment to the IRS as the result of an IRS levy; (vi) uses the withdrawal to pay health insurance premiums (if unemployed for more than 12 weeks); (vii) receives distributions in the form of an annuity; (viii) uses the withdrawal to make a distribution to an alternate payee under a QDRO (Qualified Domestic Relation Order); or (ix) has been affected by certain natural disasters (e.g., Hurricanes Katrina and Sandy).

[2] There are of course other reasons for government intervention in retirement savings systems, such as adverse selection (Finkelstein and Poterba, 2004; Einav, Finkelstein, and Schrimpf, 2010).

the commitment account(s) vary randomly across participants. The commitment accounts either impose financial penalties on or completely prohibit withdrawals prior to a commitment date chosen by the participant at the start of the experiment (the commitment date can be as far as twelve months in the future). The commitment account interest rate also varies across participants.

We find that commitment is appealing. When only one commitment account is offered that pays the same interest rate as the liquid account, classical economic theory predicts that nothing should be deposited in the commitment account. Instead, participants allocate nearly half their money to the commitment account. Even when the commitment account's interest rate is *less* than that of the liquid account, the commitment account receives one-quarter of the participants' endowments. The chosen commitment dates are widely dispersed over the duration of the experiments, and only a handful are set at the earliest possible date.

More importantly, when the commitment account and the liquid account have the same interest rate, allocations to the commitment account *increase* with the degree of account illiquidity. Among participants offered only one liquid account and one commitment account, those whose commitment account has an early withdrawal penalty equal to 10% of the withdrawal allocate less to it than those whose commitment account has a 20% early withdrawal penalty, who in turn allocate less than those whose commitment account completely prohibits early withdrawals. When a participant has simultaneous access to a liquid account and two types of commitment accounts, the commitment account that prohibits early withdrawals receives twice as much money as the commitment account with the 10% penalty.

We extend the theoretical results of Amador, Werning, and Angeletos (2006) to show that the commitment deposits of sophisticated present-biased agents—agents who are aware of their self-control problems—increase with the commitment accounts' illiquidity under a wide range of assumptions, even if agents are subject to stochastic, uninsurable taste shocks that affect future marginal utility and create a motive to provide spending flexibility to the future self.[3] Therefore, both the positive overall demand for commitment and the increasing relationship between commitment deposits and the commitment account's illiquidity are evidence in favor of the presence of sophisticated present-biased individuals in the U.S. adult population.

However, our experimental results also indicate that making 401(k) and IRA early withdrawal penalties higher would not necessarily increase deposits to these accounts. We find that when the commitment account pays a higher interest rate than the liquid account—as is the case with 401(k) and IRA accounts—the empirical relationship between the commitment account's illiquidity and the deposits it attracts is flat. This result suggests that the U.S. adult population contains not only sophisticated present-biased individuals, but also naïve present-biased individuals (those who are unaware of their present bias) and/or individuals without present bias. When the commitment account pays an interest rate premium, these latter two groups make commitment deposits that are positive but diminishing with the commitment account's illiquidity. This decrease offsets the increase in commitment deposits by sophisticated present-biased individuals as illiquidity rises. Therefore, the aggregate relationship between commitment deposits and illiquidity can take any sign, including the flat relationship we observe in our data. In contrast, when the commitment account pays the same interest rate as the liquid account, individuals without present bias and naïve present-biased individuals deposit nothing to the commitment account regardless of its illiquidity, since individuals without present bias have no need for commitment and naïve present-biased individuals believe that they have no need for commitment. Only sophisticated present-biased individuals make commitment deposits in this case, causing the aggregate relationship between illiquidity and commitment deposits to be unambiguously positive, as we find empirically.

Demand for commitment devices has been documented in many different

[3] Our analysis is an example of the approach used in the new dynamic public finance literature (e.g., Golosov, Tsyvinski, and Werning, 2007). Our setting is not constructed to allow us to explore alternative mechanisms that could in principle mitigate self-control problems, such as trigger strategies that enable the agent to behave patiently, incentivized by the threat of a breakdown in the cooperative regime (Friedman 1971; Aumann and Shapley, 1976; Ainslie, 1991) and reputation strategies that encourage the agent to avoid lapses today because such lapses undermine credibility in the future (Kreps et al., 1982; Sobel, 1985; Benabou and Tirole, 2004).

domains of behavior: completing homework assignments for university courses (Ariely and Wertenbroch, 2002), cigarette smoking cessation (Gine, Karlan, and Zinman, 2010), avoiding distractions in a computer-based task (Houser et al., 2010), achieving workplace goals (Kaur, Kremer, and Mullainathan, forthcoming), reducing time spent playing online games (Chow, 2011), performing an unpleasant task (Augenblick, Niederle, and Sprenger, forthcoming), and going to the gym (Milkman, Minson, and Volpp, 2013; Royer, Stehr, and Sydnor, 2015). Our paper is most closely related to the work of Ashraf, Karlan, and Yin (2006), who offered Filipino households a savings account that did not allow withdrawals until a certain date had passed or a goal amount had been deposited. This illiquid account was taken up by 28% of households and increased savings among households that were offered the account.[4] Our key empirical contribution is to vary the characteristics of the commitment savings accounts within the same setting to investigate which contract design features are attractive to consumers.

Dupas and Robinson (2013) conduct a related field experiment in Kenya. Participants in one treatment group received a lockbox, where they could deposit money through a slit at the top, and the key to the lockbox. Participants in another treatment group received a physically identical lockbox, but a research officer held the key, which reduced the temptation to prematurely withdraw from the lockbox. The more illiquid second lockbox also had the additional restriction that savings in the lockbox could only be used to purchase preventative health products. Take-up of the lockbox where the participant held the key is higher than take-up of the lockbox where the research officer held the key. In contrast, when participants were offered an illiquid Rotating Savings and Credit Association (ROSCA) account that could only be spent on emergency health treatments, take-up is much higher than for the liquid lockbox.

Karlan and Linden (2014) report the results of a field experiment in Uganda. Students were offered a commitment savings account from which balances could only be withdrawn after the end of the school trimester, and 44% of students took up accounts.

[4] Kast, Meier, and Pomeranz (2014) and Brune et al. (forthcoming) also study take-up of commitment savings accounts and find similar results. For a caveat to the results of Ashraf, Karlan, and Yin (2006), see the related field experiment conducted by John (2015).

Although these students appear to value the lock-up period, they do not seem to value a restriction on what the savings in the account can be spent on. More students used the commitment savings account when withdrawals were available in the form of cash compared to when withdrawals were only available in the form of vouchers for the purchase of school supplies.

Finally, Burke, Luoto, and Perez-Arce (2014) build on our basic experimental design to compare a commitment savings account that prohibits early withdrawals to a savings account that creates a psychological commitment to savings but does not prohibit early withdrawals. They find that the latter account attracts greater balances on average, perhaps because the participants who were offered that account were encouraged to think about the importance of achieving their savings goals. This inducement may have increased their overall desire for savings, regardless of whether or not the savings account prohibits early withdrawals.

The paper proceeds as follows. We start with the empirical results obtained from two related experiments. Section I describes our experimental participant recruitment. Section II discusses the design of our first experiment, and Section III presents the first experiment's results. Sections IV and V respectively describe the design and results of our second experiment. Section VI presents our extension of the Amador, Werning, and Angeletos (2006) modeling framework and discusses the relationship between our theoretical results and our experimental findings on the demand for illiquid financial accounts. Section VII concludes.

I. Participant recruitment

We conducted our two experiments using participants from the RAND American Life Panel (ALP), a panel of respondents at least 18 years old who are selected to be broadly representative of the U.S. adult population. ALP respondents participate in approximately two half-hour surveys per month over the Internet, and respondents who do not have their own Internet access have it provided to them by RAND.

For the first experiment, RAND sent an email in early 2010 to 750 ALP members inviting them to participate in a year-long experiment on financial decision-making that would provide at least $40 in compensation. 495 members consented to participate, and all of them completed the study. Forty-one participants in the first experiment are in the same household as at least one other participant in the first experiment.[5]

The recruitment procedure for the second experiment mirrored the procedure for the first experiment. In early 2011, RAND emailed 737 ALP members inviting them to participate in an experiment that would provide approximately $100 in compensation. 550 of the invited members completed the study. There is no overlap between the participants in the first experiment and the participants in the second experiment. Furthermore, no participant in the second experiment is in the same household as another participant in the second experiment, although 21 participants in the second experiment are in the same household as a participant in the first experiment.[6]

The demographic characteristics of the participants, which were collected by RAND in other surveys, are summarized in Table 1. In both experiments, 43% of the participants are male, and their ages are distributed fairly evenly across six ten-year age categories. Nearly two-thirds have at least some college education. Less than 10% of participants have annual household income below $15,000, while 17% of participants have annual household income of at least $100,000. Two-thirds are married, and more than 60% are currently working. Approximately 80% are White/Caucasian, and approximately 10% are Black/African American. Finally, the median participant has one other member in his or her household.

II. Design of Experiment 1

A. Experimental conditions

Participants in our first experiment allocated an experimental endowment between a liquid account and a commitment account. We randomly assigned each participant to one of seven experimental conditions. The features of the liquid account were constant across conditions, but the features of the commitment account varied. Withdrawals from the liquid account were allowed without restriction or penalty at any time starting one week from the participant's initial participation in the experiment. The illiquidity of the

[5] Our results remain unchanged if these 41 participants are dropped from the sample.
[6] Our results remain unchanged if these 23 participants are dropped from the sample.

commitment account varied across conditions: early withdrawals, defined as withdrawals requested prior to a commitment date chosen by the participant at the outset of the experiment, were subject to a penalty equal to 10% of the withdrawal, subject to a penalty equal to 20% of the withdrawal, or disallowed altogether. The 10% penalty condition was chosen to mirror the existing penalty levied on non-hardship pre-retirement 401(k) and IRA withdrawals in the U.S. The no early withdrawal condition mirrors the complete lack of pre-retirement liquidity in some defined contribution retirement savings systems in others countries (Beshears et al., 2015). No version of the commitment account permitted withdrawals during the first week of the experiment. Balances in the liquid account earned a 22% annual interest rate, while balances in the commitment account earned a 21%, 22%, or 23% annual interest rate. The account interest rates were chosen to be higher than typical credit card interest rates so that most participants would not find it advantageous to allocate money to the liquid account in order to withdraw it immediately to pay down credit card debt.

Table 2 summarizes the experimental design and gives the number of participants in each condition.[7] Instead of having a full 3 × 3 factorial design involving nine types of commitment accounts (all three interest rates and all three degrees of illiquidity), the experiment omitted the two arms where the commitment account has a 21% interest rate and (i) imposes a 20% early withdrawal penalty, or (ii) prohibits early withdrawals. We anticipated that commitment accounts with a 21% interest rate would not attract large allocations, so we did not want to devote much of our sample to those conditions. However, we did want to compare commitment account allocations when the commitment account interest rate was lower than, equal to, or higher than the liquid account interest rate. Therefore, we included one condition where the commitment account paid a 21% interest rate.

B. Initial allocation task

Participants clicked through a series of screens describing the details of the experiment. They would receive $50, $100, or $500, depending on a random number

[7] The number of participants is not perfectly balanced across cells because the ALP's random assignment algorithm made the cell sizes equal only in expectation; the realized cell sizes could differ from each other.

drawn in the next national Powerball lottery. Their task was to make three allocation decisions: divide each of the possible monetary endowments between a liquid account and a commitment account. They would receive weekly emails that displayed their account balances and a link to the webpage where they could request withdrawals (including partial withdrawals). They could also log into the study website at any time to view their balances and request withdrawals. Transfers between the two accounts would be impossible after the initial allocation, and withdrawal requests would result in a check being mailed to the participant within three business days.

Throughout the experiment, the liquid account was labeled the “Freedom Account,” and the commitment account was labeled the “Goal Account.” These labels were intended to help participants remember each account’s rules and understand their purposes. The description of the liquid account emphasized that it permitted flexibility. The description of the commitment account emphasized that it could help participants reach their savings goals. Participants using the commitment account would have to select a commitment date (labeled the “goal date”) no later than one year from the current date, and this date might be associated with a gift purchase, a vacation, another special event, or no particular purpose. Appendix Figures 1 and 2 show the screens explaining the accounts.

All participants allocated the $50 endowment first, the $100 allocation second, and the $500 allocation third. Whenever participants allocated any money to the commitment account, they were invited (but not required) to associate a goal with the commitment account. Appendix Figure 3 shows an example allocation page.

Finally, participants chose four Powerball numbers. In the twice-weekly Powerball lottery, six integers from 1 to 39 are randomly drawn without replacement, and one of these numbers is designated as the “Powerball.” All numbers have an equal likelihood of being the Powerball. If the Powerball in the next drawing was the first or second number chosen by the participant, she received a $500 endowment in the experiment; if the third or fourth number chosen by the participant, she received $100; and if any other number, she received $50. The money was then allocated between the two accounts according to the participant's stated wishes for the given monetary amount. After the Powerball drawing, participants received emails indicating the dollar amount they were given and reminding them of the allocation they had chosen for that amount. All participants chose their allocations between February 1, 2010, and February 11, 2010.

C. Withdrawals

Appendix Figure 4 shows an example of the weekly email sent to participants, and Appendix Figure 5 shows the summary webpage participants saw when they logged into the experimental website. When a participant requested a withdrawal, a message asked the participant to confirm the withdrawal amount and the amount by which the account balance would be reduced.

If participants withdrew all the money from their accounts before a year had elapsed, they were asked to complete an exit questionnaire asking whether any parts of the study were confusing and whether they would have changed any of their decisions in the experiment with the benefit of hindsight. If participants still had money in their accounts one year after their initial allocation decision, their remaining balances were automatically disbursed to them, and they were asked to complete the same exit questionnaire.

III. Results of Experiment 1

A. Initial allocations

We first examine the initial allocation decisions of participants. We treat each participant's three allocation decisions as three separate observations, and we perform statistical inference using standard errors clustered at the participant level.[8] Table 3 shows the mean fraction allocated to the commitment account by experimental condition. We have three main results.[9]

First, about half of initial balances are allocated to the commitment account when it has the same interest rate as the liquid account (22% column), and about one-quarter of initial balances are allocated to the commitment account when it has a 1% lower interest

[8] Commitment account allocations generally increase as the initial endowment amount increases, but our results are qualitatively similar if we separately examine $50 allocation decisions, $100 allocation decisions, or $500 allocation decisions.

[9] Our results are nearly identical if we control for participant characteristics using regressions.

rate than the liquid account (21% column). Thus, it seems that participants place some value on commitment.

Second, when the commitment account and the liquid account have the same interest rate (22% column), stricter commitment accounts are more attractive. As we move from a 10% early withdrawal penalty to a 20% early withdrawal penalty to a complete prohibition on early withdrawals, the fraction allocated to the commitment account rises from 39% to 45% to 56%. The first and second percentages are not statistically significantly distinguishable from each other, but the first and third are, as well as the second and third. This result gives us some confidence that the value participants place on commitment is not purely due to experimenter demand effects. Although demand effects could explain why a positive amount is deposited to commitment accounts, it is not obvious that demand effects would become stronger as the commitment account becomes more illiquid. Variation in illiquidity occurred exclusively between participants, and participants were not aware that illiquidity varied across participants.

The effect of increasing the commitment account's illiquidity can be benchmarked against the effect of increasing the commitment account's interest rate. Comparing across conditions with a 10% early withdrawal penalty, as the commitment account's interest rate rises from 21% to 22% to 23%, the fraction allocated to it rises from 28% to 39% to 58%. The differences across these three conditions are statistically significant. Thus, starting with a 10% penalty commitment account with a 22% interest rate, moving to a prohibition on early withdrawals has approximately the same effect on commitment account usage as increasing the interest rate to 23%.

Third, when the interest rate on the commitment account is higher than the interest rate on the liquid account, the relationship between commitment account allocations and illiquidity disappears. Commitment accounts with a 23% interest rate attract approximately 60% of the endowment regardless of their early withdrawal policy.

When participants allocate money to a commitment account, they are required to specify a commitment date before which early withdrawal restrictions apply. Table 4 shows the mean number of days between the participant's initial allocation date and his commitment date. This mean varies between 186 days and 234 days across conditions. A comprehensive measure of commitment must take into account both the amount of money committed and the time until the commitment date. Thus, for each allocation decision, we calculate the dollar-weighted days to commitment date, which is the fraction of balances allocated to the commitment account multiplied by the number of days between the allocation decision date and the commitment date.

Table 5 displays the mean dollar-weighted days to commitment date by experimental condition. The results are similar to what we found for percentage allocations to the commitment account, but slightly weaker statistically. When the commitment account pays a 22% interest rate, the mean dollar-weighted days to commitment date increases from 82 to 101 to 132 as we move from a 10% early withdrawal penalty to a 20% early withdrawal penalty to a prohibition on early withdrawals. When the commitment account has a 10% penalty on early withdrawals, the mean dollar-weighted days to commitment date increases from 64 to 82 to 130 as the interest rate increases from 21% to 22% to 23%. When the commitment account pays a 23% interest rate, the mean dollar-weighted days to commitment date has no relationship with illiquidity.[10]

In Section VI, we show theoretically that sophisticated present-biased agents will allocate more to the commitment account as its illiquidity rises under a wide range of assumptions, which is exactly the pattern we observe when both the liquid account and the commitment account pay the same interest rate. The lack of a relationship between allocations and commitment account illiquidity when the commitment account pays a higher interest rate than the liquid account can be explained if there are agents without present bias and/or naïve present-biased agents also present among our participants. Because of the commitment account's interest rate premium, it attracts some deposits

[10] A participant who is offered a commitment account with a 23% interest rate might allocate the entire endowment to the commitment account but choose the earliest possible commitment date in order to earn the higher interest rate while avoiding commitment. We see little evidence of this behavior. Of the 214 participants who had access to the 23% interest rate commitment account, only four participants selected goal dates within the first two weeks after the initial allocation decision.

from these two groups.[11] However, since they have no desire for commitment, their commitment account allocations decrease as the account becomes more illiquid, offsetting the rising allocations to the commitment account by sophisticated present-biased agents. The result is little aggregate relationship between allocations and commitment strictness. On the other hand, when the commitment account pays the same interest rate as the liquid account, agents without present bias and naïve present-biased agents allocate no money to the commitment account regardless of its strictness. Therefore, the aggregate relationship between allocations and withdrawal penalties is driven entirely by the sophisticated present-biased agents in this case.

We linked the data from our experiment with other participant data available from the RAND American Life Panel, and in untabulated analysis, we examined correlations between commitment account allocations in the experiment and variables such as credit card usage. We did not identify any correlations that survive correction for multiple testing.

B. Withdrawals

What happens to account balances after the initial allocation? For each participant and day during the year-long experiment, we calculate the sum of the liquid account and commitment account balances that the participant would have had if no withdrawals had been requested. This hypothetical total balance uses the allocation decision for the one endowment amount that the participant ended up receiving ($50, $100, or $500). We then calculate the ratio of the participant's actual balance to the hypothetical total balance on each day, and plot the mean of this ratio against the number of days since the endowment was received.[12] In order to facilitate the relevant comparisons, we present subsets of the seven conditions in each of the three graphs in Appendix Figure 6, with some conditions appearing in more than one graph.

[11] In theory, agents who believe themselves to be time-consistent should choose the earliest possible commitment date for their commitment account. The absence of such behavior may be due to an experimenter demand effect, where participants feel that they are "misbehaving" if they game the system by allocating money to the commitment account while creating negligible commitment.

[12] Recall that there was a gap between when the allocation decision was made and when the endowment was received because we needed to wait for the next Powerball lottery drawing to determine how large the participant's endowment would be.

In all conditions, most of the money stays in the accounts until the very end of the experiment. The lowest ending mean balance ratio is 0.626, and the highest is 0.723. The top graph in Appendix Figure 6 shows that withdrawals take place earlier in the experiment when the interest rate on the commitment account is lower. Holding fixed the withdrawal penalty at 10%, the average balance ratio across all the days after endowment receipt is 0.814 when the commitment account interest rate is 21%, 0.831 when the commitment account interest rate is 22%, and 0.869 when the commitment account interest rate is 23%. However, with a standard error on each average of about 0.03, we do not have the statistical power to reject their equality.

The next two graphs indicate that withdrawal patterns do not vary strongly with the commitment account's degree of illiquidity. When both the commitment account and the liquid account have the same interest rate, the average balance ratio across all days is 0.831 with a 10% early withdrawal penalty, 0.837 with a 20% early withdrawal penalty, and 0.827 with no early withdrawals allowed. When the commitment account has a 1% higher interest rate, the average balance ratio across days is 0.869 with a 10% early withdrawal penalty, 0.829 with a 20% early withdrawal penalty, and 0.857 with no early withdrawals allowed. We cannot reject the hypothesis that the average balance ratio does not change as illiquidity varies while holding fixed the commitment account interest rate.

The above comparisons are imperfect measures of how commitment affects withdrawals because the averages include days subsequent to each individual's commitment date, a period during which the commitment account is fully liquid. If we instead examine the balance ratio averaged only over days on which the commitment account is illiquid, we find some suggestive evidence that access to the fully illiquid account raises balances. In order to adjust for the fact that the mean commitment date differs across arms, we use the following procedure. Let the "adjustment factor" for participant i be the difference between the mean commitment date (measured in days since endowment receipt) in i's experimental arm and the earliest mean commitment date among the arms being compared. Let i's "adjusted commitment date" be the larger of zero and i's goal date minus the adjustment factor. If there were no censoring at zero, this adjustment would equalize the mean commitment date across the arms being compared. We then compute commitment period balance ratios for each participant by averaging that participant's daily balance ratios from the endowment receipt date to the adjusted commitment date. If a participant allocated zero dollars to the commitment account or had an adjusted commitment date of zero, we classify the participant as having made no withdrawals during the commitment period, and we therefore assign that participant a commitment period balance ratio of one.

When the commitment account and liquid account have the same interest rate, the average commitment period balance ratio is 0.967 with a 10% penalty, 0.961 with a 20% penalty, and 0.982 with no early withdrawals allowed. When the commitment account has a 1% higher interest rate, the averages are 0.932, 0.950, and 0.967, respectively. However, holding fixed the commitment account interest rate, there are no statistically significant differences among these averages, as the standard errors of the averages range from 0.009 to 0.022.

Our failure to find a significant effect of illiquidity on withdrawal behavior may be due to the fact that illiquidity did not vary sufficiently across our arms for us to be able to detect its effect with our sample sizes. Even a 10% penalty was quite a deterrent to early withdrawals. Only 3.0% of participants with a 22% commitment account interest rate and 5.1% of participants with a 23% commitment account interest rate made an early withdrawal when facing a 10% penalty, even though 10.6% of these participants with a 22% commitment account interest rate and 26.9% of these participants with a 23% commitment account interest rate allocated money to their commitment account and had less than $1 left in their liquid account on the commitment date. This means that much of the variation in illiquidity across arms came from the amount allocated to the liquid account. However, when the commitment account interest rate was 22% or 23%, the amount allocated to the liquid account was no more than 17 percentage points higher in the 10% or 20% penalty arms relative to the no-early-withdrawal arms. Furthermore, 69% of participants in the no-early-withdrawal arms allocated no money to the commitment account or had at least $1 left in their liquid account when their commitment date arrived. This means that the extra allocations to the liquid account in the less restrictive arms would affect the withdrawals of only 31% of participants. Therefore, we would expect any average differences across arms in balances prior to the commitment date to be relatively small.

IV. Design of Experiment 2

Our second experiment investigates several questions motivated by the first experiment. First, do voluntary commitment accounts discourage withdrawals? To address this, we introduce greater exogenous variation in the strength of commitment in order to be able to detect withdrawal effects more reliably. Second, given some participants' preference for more illiquid commitment accounts, why are such commitment products rarely observed in the market? We test one hypothesis: a highly illiquid commitment account is attractive when compared only to a fully liquid account, but unattractive when a less illiquid commitment account is added to the choice set, since the latter makes the highly illiquid account seem like an extreme option (Simonson, 1989). Furthermore, the complexity of choosing from a set with multiple commitment accounts may make individuals favor the simple liquid account (Redelmeier and Shafir, 1995). Finally, strict commitment has the advantage of preventing overspending but does not allow participants to access their funds in a financial emergency. Is a commitment account that offers early liquidity only in the event of an emergency more attractive to participants than a commitment account that prohibits all early withdrawals?

A. Experimental conditions

Participants in our second experiment were randomized into four treatment conditions. In all conditions (and consistent with the first experiment), participants had access to a liquid account that paid a 22% interest rate and allowed penalty-free withdrawals. In contrast to the first experiment, the commitment accounts in this experiment always paid a 22% interest rate and varied across conditions only in their illiquidity characteristics. Two conditions mimicked conditions in the first experiment for the purposes of replication. In the first arm (for replication), participants allocated their endowment between the liquid account and a commitment account that imposed a 10%

penalty on withdrawals before the participant's chosen commitment date. In the second arm (for replication), participants allocated their endowment between the liquid account and a commitment account that prohibited withdrawals before the participant's self-selected commitment date. In the third arm, participants allocated their endowment among the liquid account and two different commitment accounts, one that imposed a 10% penalty on early withdrawals and the other that prohibited early withdrawals (mirroring the different goal accounts available to participants in the first two arms of the experiment). Participants in this third arm could pick any convex combination across the three accounts, and each commitment account could be assigned its own commitment date if both were used. In the fourth and final arm, participants allocated their endowment between a liquid account and a new type of commitment account with a "safety valve" feature that prohibited early withdrawals unless a participant indicated that the funds were needed for a financial emergency. Financial emergencies would not be verified, but participants were asked to indicate honestly whether or not they were experiencing a financial emergency. The safety valve commitment account attempts to impose a psychological cost of lying only on participants who make an early withdrawal when they are not experiencing a financial emergency, creating a state-contingent early withdrawal penalty. This account was chosen to mirror the provisions that exist in 401(k) and IRA accounts that allow for penalty-free pre-retirement withdrawals in the case of certain financial hardships; some other countries with defined contribution retirement savings systems also allow for pre-retirement withdrawals only in the case of certain financial hardships (Beshears et al., 2015).

After participants indicated their desired allocations, they were randomly assigned to receive either $100 allocated according to their wishes or $100 allocated entirely to the liquid account. Table 6 shows the number of participants assigned to each experimental condition, broken out into the number who received allocations according to their wishes and the number who received all of their funds in the liquid account. We did not stratify by experimental condition when randomly assigning participants to receive their chosen allocations or the 100% liquid account allocation, so the distribution of participants within each experimental condition is unbalanced.

B. Initial allocation task

Participants were told that they would receive $100 to allocate between the accounts offered in their condition. The liquid account was again labeled the “Freedom Account,” and the commitment accounts were again labeled “Goal Accounts.” The experimental website would display balances and allow withdrawal requests at any time,[13] and weekly emails would also display balances and a link to the withdrawal webpage. Transfers between the accounts would not be allowed, and checks would be mailed within three business days of a withdrawal request.

The descriptions of the liquid account, the 10% penalty commitment account, and the no-early-withdrawal commitment account were the same as the descriptions used in the first experiment. When the 10% penalty account and the no-early-withdrawal account were offered simultaneously, they were labeled “Goal Account A” and “Goal Account B,” respectively (see Appendix Figure 7). Participants learned that the two commitment accounts could be assigned distinct commitment dates (again labeled “goal dates”). In the case of the safety valve account, participants were informed that early withdrawals were possible only when a financial emergency occurred. Participants would be the sole judges of whether or not an emergency was actually occurring (see Appendix Figure 8).

Participants were then told that they would receive their chosen allocation with 50% probability and an allocation selected by the experimenters with 50% probability. They did not know that the allocation selected by the experimenters would place all of the money in the liquid account. A computer rather than a public randomizing device was used for this randomization procedure.

Finally, participants made their allocation and commitment date choices. Participants were then informed whether they were receiving their chosen allocation or the 100% liquid account allocation.

[13] Like the first experiment, the second experiment permitted withdrawals no sooner than one week after the initial allocation decision.

Participants completed this initial phase of the experiment between February 14, 2011, and March 2, 2011. The experiment ended for all participants on September 1, 2011. Therefore, unlike the one-year duration of the first experiment, the second experiment's duration was only about half a year.

C. Withdrawals

All participants who requested withdrawals were asked to confirm their requests. In addition, participants who wished to make early withdrawals from the safety valve account were shown the following text:

> We are relying on you to be *honest* in judging whether you have a financial emergency. If you are sure you want to make a withdrawal, please type the sentence below, then click "Next." Otherwise, click "Cancel my withdrawal."

The sentence that these participants were asked to type was, "I attest that I have a financial emergency." However, the website accepted any entered text.

The second experiment gave an exit questionnaire to participants who withdrew all of their money before September 1, 2011. Participants who had remaining balances on September 1, 2011, automatically received checks for their balances and received emails with links to the same exit questionnaire. The exit questionnaire gave participants the opportunity to identify confusing aspects of the experiment.[14] Also, whenever participants in the second experiment made any withdrawals (including partial withdrawals) before September 1, 2011, they were given the option to provide the reasons for the withdrawal.

V. Results of Experiment 2

A. Initial allocations

Table 7 shows the mean fraction of the endowment allocated to a commitment account in each experimental condition. When participants are offered only the liquid account and the 10% penalty account, the commitment account receives 46% of the endowment. When participants are offered only the liquid account and the no-early-

[14] In contrast to the first experiment, participants in the second experiment were not asked to explain anything that they would have done differently in retrospect.

withdrawal account, the mean commitment account allocation is 54%, which is significantly higher ($p = 0.034$) than the 46% allocation in the former condition. Thus, we replicate the findings from the first experiment that commitment is desirable, and stronger commitment is more attractive when the commitment and liquid accounts pay the same interest rate.

The no-early-withdrawal account is appealing even when it is offered in the same choice set as the 10% penalty account. In this arm, the no-early-withdrawal account attracts 34% of the endowment, while the 10% penalty account attracts only 16%, a difference that is highly significant ($p < 0.001$). We therefore find no evidence that the lack of strict commitment accounts in the marketplace is due to the simultaneous presence of partially illiquid accounts.

Surprisingly, total allocations to commitment accounts are not higher when two commitment accounts are available rather than one. With two commitment accounts, the commitment accounts receive 50% of the endowment in total. This is halfway between the 46% allocation when the 10% penalty account is the only commitment account and the 54% allocation when the no-early-withdrawal account is the only commitment account. It is possible that the availability of two commitment accounts makes the allocation decision more complex, leading participants to view the simple and distinct liquid account as more desirable (Redelmeier and Shafir, 1995). Intuitively, if a participant has a hard time choosing between two similar commitment accounts, the participant may take the exit strategy of adopting a conflict-avoiding alternative (i.e., the liquid account). This is an instance of "reason-based choice" (Shafir, Simonson, and Tversky, 1993).

Our attempt to create a commitment account that is more appealing than the no-early-withdrawal account was unsuccessful. The safety valve account receives a mean allocation of 45%. This is statistically indistinguishable from the 46% allocation to the 10% penalty account when it is the only commitment account available, and significantly less ($p = 0.018$) than the 54% allocation to the no-early-withdrawal account when it is the only commitment account available. It may be that the psychological cost of lying about a financial emergency in order to make a withdrawal is too low for the safety valve commitment account to be a strong commitment device.[15]

Table 8 displays the mean days between the initial allocation date and the commitment date, and Table 9 shows the mean dollar-weighted days to commitment date. The results in Table 9 are in line with the initial commitment account allocations in Table 7. Mean dollar-weighted days to commitment date rises from 62 to 64 to 75 in the single commitment account conditions as the commitment account changes from safety valve to 10% penalty to no early withdrawals. The difference between the safety valve and no early withdrawal conditions is significant ($p = 0.046$), but not the difference between the 10% penalty and no early withdrawal conditions ($p = 0.137$).[16] When two commitment accounts are available, the mean dollar-weighted days to commitment date of 71 lies between the values in the arms where only one commitment account is available and the commitment account either imposes a 10% penalty or does not allow early withdrawals.

B. Withdrawals

Because we randomly assigned half of participants to receive all of their endowment in the liquid account, we have greater exogenous variation in liquidity than in the first experiment, which we can use to identify whether the commitment accounts help participants save more. Appendix Figure 9 shows the balance ratios over time for the four experimental conditions, breaking apart participants by whether they received their endowments allocated according to their choices or 100% in the liquid account.[17] Because participants made initial allocation decisions on different dates but completed the experiment on the same date (September 1, 2011), some participants participated in the experiment for slightly longer periods of time than others. The figure displays only

[15] All of the allocation results are qualitatively unchanged if we adjust for participant characteristics using regressions, except that the difference between the safety valve account allocation and the no early withdrawal account allocation when only one commitment account is offered is significant at only the 10% level.

[16] These two *p*-values are 0.101 and 0.099, respectively, when we control for participant characteristics.

[17] For one participant in the no early withdrawal condition, we have conflicting records as to whether the participant was randomly assigned to receive the chosen commitment account allocation or was randomly assigned to receive the entire endowment in the liquid account. We drop this participant from the data set when analyzing withdrawal patterns, but the results do not change materially if we assume that the participant was randomly assigned to one group or the other.

the first 183 days since endowment receipt, so that the sample remains constant within each graph.

Consistent with the safety valve account being a weak commitment device, the balance ratios for those in the safety valve condition do not markedly differ when participants receive all of their endowment in the liquid account instead of according to their chosen allocation. In contrast, balance ratios are substantially lower in the 10% penalty and no early withdrawal conditions with only one commitment account if all of the endowment was deposited into the liquid account. The same pattern emerges when there are two commitment accounts, although the gap is much smaller. In Table 10, we report the difference in balance ratio means within condition at selected points in time during the experiment, as well as for the four experimental conditions pooled. The results for the pooled sample suggest that the commitment accounts do significantly reduce withdrawals.

Withdrawals among participants who receive their entire endowment in the liquid account provide a measure of self-control problems. It would be interesting if there were a correlation between these withdrawals and chosen commitment account allocations, but we do not find such a correlation in our sample. However, it is difficult to draw strong conclusions from this analysis because withdrawals also reflect liquidity shocks and are therefore at best a noisy measure of self-control problems.

VI. A theory of the commitment account allocations of sophisticated present-biased agents

To study the tradeoff between commitment and flexibility in a consumption/savings context, Amador, Werning, and Angeletos (2006; hereafter AWA) use a model with three conceptual ingredients.

First, AWA assume dynamically inconsistent preferences generated by the present-biased discount function

$$D(\tau) = \left\{ \begin{array}{ccc} 1 & \text{if} & \tau = 0 \\ \beta & \text{if} & \tau \geq 1 \end{array} \right\},$$

where $0 < \beta < 1$ (Phelps and Pollak, 1968; Laibson, 1997).[18] This discount function implies that, from the perspective of period 0, the agent is more patient about tradeoffs between periods 1 and 2 than she will be when period 1 actually arrives:

$$\frac{D(1)}{D(2)} = \frac{\beta}{\beta} < \frac{1}{\beta} = \frac{D(0)}{D(1)}.$$

Dynamically inconsistent preferences generate a motivation for precommitment.

Second, they assume that the agent experiences transitory taste shocks that are not observable in advance and are not contractable. Such taste shocks generate a motivation to give future selves flexibility in choosing the consumption path.

Third, they assume that the agent has a very general commitment technology. Specifically, she can manipulate the choice sets of future selves, trading off the benefits of commitment (preventing later selves from overconsuming) and the costs of commitment (preventing later selves from responding flexibly to the taste shocks).

We enrich AWA's analysis by placing a bound on the strength of the commitment technology. We show that, in this more general setting, the agent can still achieve the (second-best) optimum using a simple commitment mechanism. Furthermore, we vary the bound and explore the implications for the choice of commitment mechanism. These comparative statics enable us to compare the model's predictions with the behavior of our experimental participants.

We briefly describe the key properties of the model below. Appendices A-P provide a complete exposition and analysis of the model.

A. Timing and preferences

The simplest model that elicits a tradeoff between commitment and flexibility has three periods: an initial period in which some degree of commitment is created with respect to future decisions; a following period in which a consumption/savings choice is made with immediate utility consequences; and a final period in which residual wealth is consumed.

[18] The analysis that follows would be nearly identical if we were to use the more general quasi-hyperbolic discount function by $D(0) = 1$ and $D(\tau) = \beta\delta^{\tau}$ for $\tau \geq 1$, where $0 < \beta < 1$ and $0 < \delta \leq 1$. For simplicity, we follow AWA and set $\delta = 1$.

Period 0. Self 0 chooses the commitment mechanism that will govern the choices of selves 1 and 2. (There is no consumption in period 0.)

Period 1. A taste shock $\theta \in \Theta = [\underline{\theta}, \overline{\theta}]$ is realized. Self 1 observes θ and makes a consumption/savings decision, subject to the constraints imposed by the commitment mechanism chosen by self 0.

Period 2. Self 2 consumes all remaining wealth.

Section VI.B below describes the set of commitment mechanisms available to self 0, and Section VI.C sets out our assumptions on the distribution of θ.

Let c_1 and c_2 denote the consumption levels of selves 1 and 2. Then underlying preferences at dates 0, 1, and 2 can be specified as follows:

$$\begin{aligned} \text{utility of self 0} &= \beta\theta U_1(c_1) + \beta U_2(c_2) \\ \text{utility of self 1} &= \theta U_1(c_1) + \beta U_2(c_2) \\ \text{utility of self 2} &= U_2(c_2) \end{aligned}$$

Here U_t is the utility function at time t. We assume that: $U_t : [0,\infty) \to [-\infty,\infty)$; $U_t' > 0$; $U_t'' < 0$ on $(0,\infty)$; and $U_t'(0+) = \infty$.[19]

We also assume that self 0 fully understands and anticipates the preferences of self 1. That is, we assume that the agent is sophisticated.

B. Commitment technology

A commitment mechanism is modeled as a budget set B chosen by self 0. This B is the set of consumption pairs (c_1, c_2) that can be chosen by self 1. Recall that the taste shock is not yet observable in period 0, and that it will only be privately observable in period 1, so B cannot be conditioned on the realization of the taste shock.

Let $y > 0$ be the agent's exogenous budget and, without loss of generality, let the gross interest rate be unity. Let the "ambient budget set" A be the set of all consumption pairs (c_1, c_2) such that $c_1, c_2 \geq 0$ and $c_1 + c_2 \leq y$. Fix a parameter $\pi \in [0,\infty)$, which will

[19] For example, it could be that U_t has constant relative risk aversion $\rho_t > 0$. In that case: if $\rho_t \in (0,1)$, then $U_t(0) > -\infty$; and if $\rho_t \in [1, \infty)$, then $U_t(0) = -\infty$. In particular, we do not require $U_t(0) = -\infty$.

bound the strength of the commitment mechanism. Then the budget set *B* chosen by self 0 must satisfy the following two constraints:

Constraint 1. *B* is a non-empty compact subset of *A*.

Constraint 2. The penalty for transferring consumption from period 2 to period 1 is no greater than π.[20]

In other words, self 0 can choose a budget set of almost any size and shape. The only restriction on size is that *B* must be small enough to fit inside *A*. The only restriction on shape is that, starting from any consumption pair $(c_1, c_2) \in B$ such that $c_2 > 0$, self 1 must always be able to transfer a small amount $\Delta > 0$ of consumption from period 2 to period 1. She may face a penalty for doing so, in the form of a reduction in consumption in period 2 (over and above that resulting from the transfer itself). However, this penalty will never be greater than $\pi\Delta$.[21]

A wide variety of budget sets satisfy Constraints 1 and 2. For example, the budget set shown in Figure 1 consists of: (i) a downward sloping budget curve that begins on the c_2 axis and ends on the c_1 axis; and (ii) all the points of *A* that lie below or to the left of the budget curve. It obviously satisfies Constraint 1. It satisfies Constraint 2 if $\pi = 0.5$, but not if $\pi = 0.1$. Indeed, the slope of the budget curve at the encircled point is –1.3. This is greater than –1.5 (the minimum slope that is permissible when $\pi = 0.5$), but less than –1.1 (the minimum slope that is permissible when $\pi = 0.1$).[22]

[20] To be precise, we require that, for all $(c_1, c_2) \in B$ and all $\tilde{c}_1 \in [c_1, c_1 + \frac{1}{1+\pi} c_2]$, there exists $\tilde{c}_2$ such that: (i) $(\tilde{c}_1, \tilde{c}_2) \in B$; and (ii) $\tilde{c}_2 \geq c_2 - (1+\pi)(\tilde{c}_1 - c_1)$. In other words, if we take any point (c_1, c_2) in *B*, and if we draw a line *L* of slope $-(1+\pi)$ through (c_1, c_2), then: associated with every point $(\tilde{c}_1, \hat{c}_2)$ on the segment of *L* that joins the points (c_1, c_2) and $\left(c_1 + \frac{1}{1+\pi} c_2, 0\right)$, there is a point $(\tilde{c}_1, \tilde{c}_2) \in B$ that either coincides with $(\tilde{c}_1, \hat{c}_2)$ or else lies vertically above it. (Note that $\left(c_1 + \frac{1}{1+\pi} c_2, 0\right)$ is the point at which *L* crosses the c_1 axis.)

[21] Perhaps more intuitively, this property can also be expressed in differential terms: starting from any consumption pair $(c_1, c_2) \in B$ such that $c_2 > 0$, it must be possible to draw a curve in *B* of slope at least $-(1+\pi)$, and to continue that curve up to the point at which it first hits the c_1 axis.

[22] Notwithstanding its obvious generality, this budget set is still special in a number of respects. We give four examples. First, there is nothing in Constraints 1 and 2 that requires that the budget curve be downward sloping. Indeed, these constraints place no upper bound at all on the slope of the budget curve. Second, there is no reason why the budget curve needs to begin on the c_2 axis. It could perfectly well begin

However, as we shall see, the optimum can be obtained using a particularly simple kind of budget set, namely a two-part budget set. Such budget sets consist of: (i) a budget curve that has slopes of –1 and –(1 + *p*) to the left and right of a kink at (c_1^*, c_2^*); and (ii) none, some, or all of the points of *A* that lie below or to the left of the budget curve. For example, the budget set shown in Figure 2 consists of just such a budget curve, together with all of the points of *A* that lie below or to the left of it.

Better still, two-part budget sets arise naturally in practical applications. Indeed, suppose that self 0 sets up two separate accounts: (i) a fully liquid account with balance c_1^*; and (ii) a partially illiquid account with balance c_2^* and an early withdrawal penalty *p*.[23] Then self 1 will face a two-part budget set.

C. Distribution of the taste shock

AWA show that their problem can be reduced to a problem in the class of optimization problems identified and analyzed by Luenberger (1969). We follow AWA's lead. We make the following assumptions on the distribution function *F* of the taste shock *θ*.

A1 Both *F* and *F′* are functions of bounded variation on (0, ∞).[24]

A2 The support of *F′* is contained in $[\underline{\theta}, \overline{\theta}]$, where $0 < \underline{\theta} < \overline{\theta} < \infty$.

---

at some $(c_1^0, c_2^0) \in A$ for which $c_1^0 > 0$. (Constraint 2 does, however, require that the budget curve end on the c_1 axis.) Third, there is no reason why points below or to the left of the budget curve need be included. Fourth, there is no reason why the budget curve need be connected. It could perfectly well consist of two or more components. For example, a first component might begin at some $(c_1^0, c_2^0) \in A$ for which $c_1^0 > 0$ and end at some $(c_1^1, c_2^1) \in A$ for which $c_2^1 = 0$. A second component might then begin at some $(c_1^2, c_2^2) \in A$ for which $c_1^2 > c_1^1$ and end at some $(c_1^3, c_2^3) \in A$ for which $c_2^3 = 0$. (Constraint 2 does, however, rule out budget curves consisting of a finite set of points, unless these points all lie on the c_1 axis.)

[23] By saying that self 1 pays a penalty *p* on early withdrawals from the second account, we mean that if she consumes Δ from the second account then that account is debited (1 + *p*)Δ.

[24] A sufficient condition for A1 is that: (i) *F′* and *F′′* both exist; and (ii) $\int_0^\infty |F'(\theta)|\, d\theta$ and $\int_0^\infty |F''(\theta)|\, d\theta$ are both finite. In other words, if one walks along the graph of *F′* or *F′′*, then the total vertical distance travelled (both up and down) is finite. We do not use this stronger condition because we want to allow for densities, like that of the uniform distribution, that have jumps at $\underline{\theta}$ and $\overline{\theta}$. Indeed, a good way of generating examples is to take a standard distribution and truncate it at suitable points $\underline{\theta}$ and $\overline{\theta}$. This procedure typically results in discontinuities in *F* and *F′* at $\underline{\theta}$ and $\overline{\theta}$.

A3 Put $G(\theta) = (1-\beta)\theta F'(\theta) + F(\theta)$. Then there exists $\theta_M \in [\underline{\theta}, \overline{\theta}]$ such that: (i) $G' \geq 0$ on $(0, \theta_M)$; and (ii) $G' \leq 0$ on (θ_M, ∞).[25]

We now comment on these assumptions. A function $f : (0, \infty) \to \mathbb{R}$ is of bounded variation if and only if it is the difference of two bounded and non-decreasing functions $f_1, f_2 : (0, \infty) \to \mathbb{R}$. Since *F* is a distribution function, it is automatically a function of bounded variation. The substance of A1 is therefore the requirement that *F* has a density *F′* that is a function of bounded variation. A2 means that *F′* = 0 on $(0, \infty) \setminus [\underline{\theta}, \overline{\theta}]$. Notice that *F′* need not be continuous. In particular, it can jump up at $\underline{\theta}$ and down at $\overline{\theta}$. A3 means that *G* is first increasing and then decreasing. It implies that the support of *F′* is connected. It is preserved under truncation: if a distribution function *F* satisfies A3, then so too does the distribution function obtained by truncating F at $\underline{\theta}$ and $\overline{\theta}$.

A3 is satisfied by many of the distributions that one encounters in practice. To illustrate this point, we have made a list of all the distributions occurring in either of two leading statistics textbooks: Rice (1995) and Hogg, McKean, and Craig (2005). This list contains 18 distributions. Of these, 14 satisfy A1-A3 for all parameter values (including $\underline{\theta}$ and $\overline{\theta}$). More precisely, we have:

Remark *Suppose that D is one of the Burr, Chi-squared, Exponential, Extreme Value, F, Gamma, Gompertz, Log-Normal, Maxwell, Normal, Rayleigh, t, Uniform, and Weibull distributions. Then, for any* $0 < \underline{\theta} < \overline{\theta} < \infty$, *the distribution function F obtained by truncating D at* $\underline{\theta}$ *and* $\overline{\theta}$ *satisfies Assumptions A1-A3.*[26]

The four exceptions are the Beta, Cauchy, Log-Gamma, and Pareto distributions. In the form in which it appears in both Rice and Hogg, McKean, and Craig, the Beta

---

[25] A3 is slightly stronger than the analogous assumption in AWA, namely their Assumption A. However: (i) it is not clear that our results for the model with $\pi < \infty$ actually hold under AWA's A; (ii) A3 is easier to state than AWA's A; and (iii) it is easier to check whether a given distribution satisfies A3 than to check whether it satisfies AWA's A.

[26] Notice that 5 of these 14 distributions (namely the Burr, Chi-squared, F, Gamma, and Weibull distributions) are unbounded at zero for some parameter values. However, the truncated distributions all satisfy A1 because $\underline{\theta} > 0$.

distribution does in fact satisfy A1-A3. However, for our purposes, it is more natural to consider a generalization of the Beta distribution for which the support is a compact interval contained in $(0,\infty)$. For this distribution, A3 is not always satisfied.[27] Similarly, the standard Cauchy distribution, which is the form of the Cauchy distribution considered in both Rice and Hogg, McKean, and Craig, satisfies A1-A3. However, in its general form, the Cauchy distribution fails A3 for some choices of the parameter values.[28] Next, the Log-Gamma distribution appears only in Hogg, McKean, and Craig. This distribution may or may not satisfy A3, depending on the parameters.[29] Finally, Rice and Hogg, McKean, and Craig each consider a special case of the Pareto distribution. Both of these special cases satisfy A1-A3. However, in general, the Pareto distribution fails A3 for some choices of the parameter values.[30] For additional discussion of these exceptional cases, see Appendix Q.

D. Theorems and relationship to experimental results

AWA show that, when there is no bound on the strength of the commitment technology, an optimal choice for self 0 is a minimum-savings rule. In our terminology, this can be expressed by saying that an optimal choice for self 0 is to divide her endowment y between two accounts: (i) a fully liquid account that places no penalty on withdrawals in either period 1 or period 2; and (ii) a fully illiquid account that disallows

[27] The density of the generalization of the Beta that we consider is proportional to $(x-a)^{\zeta-1}(b-x)^{\eta-1}$ on the interval (a,b), where $0 < a < b < \infty$ and $\zeta, \eta > 0$. Exceptions to A3 occur when $\zeta < 1$.

[28] The density of the general form of the Cauchy distribution is proportional to $\left(1+\left(\frac{x-\mu}{\sigma}\right)^2\right)^{-1}$ on $\mathbb{R}$, where $\mu \in \mathbb{R}$ is a location parameter and $\sigma > 0$ is a scale parameter. Exceptions to A3 occur when μ/σ is large and positive.

[29] The density of the Log-Gamma distribution is proportional to $x^{-\frac{\eta+1}{\eta}}(\log(x))^{\zeta-1}$ on $(1,\infty)$, where $\zeta, \eta > 0$. Exceptions to A3 occur when $\zeta < 1$ and $\eta > 1-\beta$.

[30] For example, the density of the Pareto type II distribution is proportional to $\left(1+\frac{x-\mu}{\sigma}\right)^{-\zeta-1}$ on (μ,∞), where $\mu \in \mathbb{R}$ is a location parameter, $\sigma > 0$ is a scale parameter and $\zeta > 0$ is a shape parameter. Exceptions to A3 occur when ζ is small and μ/σ is large and positive.

any withdrawals in period 1 but places no penalty on withdrawals in period 2.[31,32] Our first result generalizes AWA's result to the case in which there is a bound π on the strength of the commitment mechanism.

Theorem 1 *Suppose that* $U_1 = U_2 = U$ *, and that U has constant relative risk aversion* $\rho > 0$*. Then an optimal choice for self 0 is to divide her endowment y between two accounts: (i) a fully liquid account with no penalty on withdrawals in either period 1 or period 2; and (ii) a partially illiquid account, with a penalty* $p = \pi$ *on withdrawals in period 1 and no penalty on withdrawals in period 2.*[33]

See Appendices B through I.

Theorem 1 implies that there is no advantage to self 0 in using more than two accounts, in using accounts with more complex conditions attached to them, in using accounts with a penalty $p < \pi$, or in using some commitment mechanism other than accounts.

Our next result gives a sufficient condition under which self 0's allocation to the partially illiquid account is monotonic in p.

[31] There is a small technical difference between a fully illiquid account and a partially illiquid account with penalty $p = \infty$. If self 0 places y_{liquid} in a fully liquid account and $y - y_{liquid}$ in a fully illiquid account, then she is effectively choosing a budget set that consists of the line segment joining the two points $(0, y)$ and $(y_{liquid}, y - y_{liquid})$. On the other hand, if she places y_{liquid} in a fully liquid account and $y - y_{liquid}$ in a partially illiquid account with penalty $p = \infty$, then she is effectively choosing a budget set that consists of all points on *or vertically below* the line segment joining the two points $(0, y)$ and $(y_{liquid}, y - y_{liquid})$. (In effect, an illiquid account with penalty $p = \infty$ gives self 1 the possibility of free disposal, whereas a fully illiquid account does not.) Of course, these two mechanisms are equivalent from the point of view of self 0, since self 1 will always choose from the line segment joining the two points $(0, y)$ and $(y_{liquid}, y - y_{liquid})$. Therefore, we shall not distinguish between them in what follows.

[32] Ambrus and Egorov (2013) provide additional analysis of AWA's model.

[33] As the wording of the Theorem implies, the optimal choice of self 0 is not unique. Indeed, as long as one thinks of self 0 as choosing a budget set B, her optimal choice is inherently non-unique. This is because, starting from any given B (optimal or not), one can make equivalent budget sets by adding or removing consumption pairs that would not be chosen by any type. This particular form of non-uniqueness can be eliminated if, instead of thinking of self 0 as choosing a budget set B, we think rather of her as choosing an incentive-compatible consumption curve $c = (c_1, c_2) : [\underline{\theta}, \overline{\theta}] \to A$.

Theorem 2 *Suppose that* $U_1 = U_2 = U$ *, and that U has constant relative risk aversion* $\rho = 1$ *(i.e., that U = log).*[34] *Suppose further that self 0 must divide her endowment y between two accounts: (i) a fully liquid account with no penalty on withdrawals in either period 1 or period 2; and (ii) a partially illiquid account with a penalty p on withdrawals in period 1 and no penalty on withdrawals in period 2. Denote the optimal allocations to the two accounts by* y_{liquid} *and* $y_{penalty}$*. Then the correspondence mapping p to the set of possible choices for* $y_{penalty}$ *is weakly increasing. Specifically:*[35]

1. *For small values of p, there may be a non-trivial interval* $[\underline{y}_{penalty}, \overline{y}_{penalty}]$ *of possible choices for* $y_{penalty}$ *.*[36] *However,* $\underline{y}_{penalty}$ *is strictly increasing in p, and* $\overline{y}_{penalty}$ *is independent of p.*
2. *For intermediate values of p, there is a unique choice for* $y_{penalty}$*.*[37,38] *This* $y_{penalty}$

[34] We have conducted a number of simulations for the case $\rho \neq 1$. These simulations suggest that, in practice, $y_{penalty}$ is non-decreasing in *p* for reasonable calibrations of the model. However, our analytic results reveal that the simulation results are not completely generalizable, and we can construct counterexamples to the claim that $y_{penalty}$ is non-decreasing in *p* when $\rho \neq 1$. We recover the result that $y_{penalty}$ is non-decreasing in *p* when $\rho \neq 1$ if we impose further restrictions on the distribution function *F* beyond Assumptions A1-A3 (see Appendix O).

[35] In the current case, there exist two cutoffs $\pi_0 \geq 0$ and $\pi_1 > \pi_0$. We say that *p* is small if $p \in (0, \pi_0)$, intermediate if $p \in (\pi_0, \pi_1)$, and large if $p \in (\pi_1, \infty)$. The significance of π_0 is that the optimal allocation to the partially illiquid account is non-unique when $p < \pi_0$ and unique when $p > \pi_0$. The significance of π_1 is that the maximum-penalty constraint is strictly binding when $p < \pi_1$ and strictly slack when $p > \pi_1$. We typically have $\pi_0 = 0$. In that case, non-uniqueness occurs only when *p* = 0, and then only for the trivial reason that the liquid account and the partially illiquid account are indistinguishable (and self 0 is therefore indifferent as to how she divides *y* between the two accounts). There is, however, a very specific scenario in which $\pi_0 > 0$. In that scenario, non-uniqueness occurs for all *p* in the non-trivial range $(0, \pi_0)$.

[36] The form of non-uniqueness discussed here arises even if (as in Footnote 33) we think of self 0 as choosing an incentive-compatible consumption curve. It can be eliminated, as we show in Appendix J, by introducing an additional assumption, namely: (A4) *G* is strictly increasing on $[\underline{\theta}, \theta_M)$. However: (i) A4 is a little out of keeping with the rest of our framework; (ii) we can obtain satisfactory comparative-statics results without A4; and (iii) the form of non-uniqueness described rarely arises even in the absence of A4. We do not therefore include A4 among our basic assumptions.

[37] In other words, if one thinks of self 0 as taking the two accounts (including the penalty *p* on the partially illiquid account) as given, then there is only one way in which she can allocate *y* between the two accounts.

[38] Self 0's choice is actually unique in a stronger sense than that given in Footnote 37. Suppose that we confine her to using at most two accounts, and that we place an upper bound π on the penalty that she can attach to either of these two accounts. Then she will attach a penalty of 0 to one account and the maximum

is strictly increasing in p.

3. *For large values of p, there is again a unique choice for* $y_{penalty}$.[39] *This* $y_{penalty}$ *is independent of p.*

See Appendices J through O.

Theorem 2 aligns with our empirical results when the liquid account and the illiquid account have the same interest rate (although log utility is a knife-edge case). Specifically, as *p* rises exogenously across these experimental arms, the allocation to the partially illiquid account rises. This is true in both the original experiment (where *p* rises from 10% to 20% to ∞) and the second experiment (where *p* rises from 10% to ∞).

Theorem 3 *Suppose that* $U_1 = U_2 = U$ *, and that U has constant relative risk aversion* $\rho > 0$. *Suppose further that self 0 must divide her endowment y among three accounts: (i) a fully liquid account with no penalty on withdrawals in either period 1 or period 2; (ii) a partially illiquid account with a penalty p on withdrawals in period 1 and no penalty on withdrawals in period 2; and (iii) a fully illiquid account that disallows any withdrawals in period 1 but places no penalty on withdrawals in period 2. Denote the optimal allocations to the three accounts by* y_{liquid}, $y_{penalty}$, *and* $y_{illiquid}$. *Then her liquid allocation* y_{liquid} *is unique and independent of p. By the same token, her total illiquid allocation* $y_{penalty} + y_{illiquid}$ *is unique and independent of p. Furthermore, self 0 weakly prefers the fully illiquid account to the partially illiquid account. Specifically:*[40]

1. *For small and intermediate values of p, self 0 strictly prefers the fully illiquid account to the partially illiquid account. More precisely: self 0 places her total*

---

penalty of π to the other. Moreover there is only one way in which she can allocate *y* between the two accounts.

[39] Self 0's choice is once again unique in a stronger sense than that given in Footnote 37. Suppose that, as in Footnote 38, we confine her to using at most two accounts. Then she will attach a penalty of 0 to one account and any penalty *p* in the range $[\pi_1, \pi]$ to the other. (The point here is that any penalty of at least π_1 is sufficient to deter self 1 from making any withdrawal from the partially illiquid account.) So her choice of accounts is no longer unique. However, there is only one way in which she can allocate *y* between the two accounts, and this allocation is independent of the penalty on the partially illiquid account.

[40] We say that *p* is small if $p \in (0, \pi_0)$, intermediate if $p \in (\pi_0, \pi_1)$ and large if $p \in (\pi_1, \infty)$. Cf. Theorem 2 and Footnote 35.

illiquid allocation $y_{penalty} + y_{illiquid}$ *in the fully illiquid account;* $y_{penalty} = 0$*; and* $y_{illiquid}$ *is unique and independent of p.*

2. *For large values of p, self 0 is indifferent between the fully illiquid account and the partially illiquid account. More precisely: self 0 does not care how her total illiquid allocation* $y_{penalty} + y_{illiquid}$ *is divided between the partially illiquid account and the fully illiquid account.*

See Appendices M and N.2.

The logic behind Theorem 3 runs as follows. First, it follows directly from the formulation of our problem that the maximum expected utility of self 0 is weakly increasing in *π*. Second, if $U_1 = U_2 = U$ and *U* has constant relative risk aversion, then self 0 can achieve this maximum using two accounts, namely a fully liquid account and a partially illiquid account with penalty $p = \pi$. Hence, if we restrict self 0 to dividing her endowment between a fully liquid account and a partially illiquid account with penalty *p*, then she will always weakly prefer a higher *p*. In particular, she will like *p* = ∞ best of all. In other words, she weakly prefers the fully illiquid account to the partially illiquid account.

Now suppose that the optimal allocation to a fully illiquid account is y_{liquid}. If self 0 deposits $y_{illiquid}$ in a *fully illiquid* account, then there will be a θ_1 such that: (i) any self 1 of type $\theta \le \theta_1$ will choose freely from the line segment joining the points (0, *y*) and $(y - y_{illiquid}, y_{illiquid})$; and (ii) any self 1 of type $\theta \ge \theta_1$ will end up choosing the endpoint $(y - y_{illiquid}, y_{illiquid})$. If, however, self 0 deposits $y_{illiquid}$ in a *partially illiquid* account with penalty *p*, then the behavior of self 1 will not change if and only if $p \ge \pi_1$, where π_1 is the minimum penalty necessary to deter the $\overline{\theta}$ type of self 1 from increasing consumption above $y_{liquid} = y - y_{illiquid}$.

Hence, if $p \ge \pi_1$, then self 0 can attain the maximum expected utility associated with a fully illiquid account by using a partially illiquid account with a penalty *p* instead.

She will therefore be indifferent between these accounts. On the other hand, if $p < \pi_1$, then a penalty of *p* is no longer sufficient to deter the $\bar{\theta}$ type of self 1 from increasing consumption above y_{liquid}. Hence the behavior of self 1 will change if self 0 deposits $y_{illiquid}$ in a partially illiquid account with penalty *p*. Furthermore, it can be shown that, even when self 0 makes the optimal allocation to the partially illiquid account, her expected utility will still be *strictly lower* than the expected utility that she can obtain from the fully illiquid account. She will therefore strictly prefer the fully illiquid account.

This prediction of an overall weak preference for the fully illiquid account over the partially illiquid account is consistent with our empirical results in the experimental arm in which participants allocated their endowments across three accounts: a liquid account, a partially illiquid account with a 10% penalty, and a fully illiquid account. Among the participants in this experimental arm, 37 allocated money to the fully illiquid account but not to the partially illiquid account, while only 8 allocated money to the partially illiquid account but not to the fully illiquid account (76 allocated money to both illiquid accounts, and 29 allocated money to neither). The average allocations to the accounts follow a similar pattern: the partially illiquid account attracts 16% of the endowment, while the fully illiquid account attracts 34% of the endowment. Furthermore, Theorem 3 predicts that participants who allocate money to the partially illiquid account do so because the 10% penalty is above π_1 and therefore sufficient to deter early withdrawals. There were 42 participants who allocated money to the partially illiquid account and were randomly assigned to receive their chosen allocation (instead of having all of their endowment placed in the liquid account), and out of those 42 participants, only one made a withdrawal from the partially illiquid account before the goal date.

Thus, the data tend to support Theorem 3. However, it would be necessary to extend the model to accommodate some of the nuances of the experimental design. Most importantly, participants in the study were allowed to set *different* goal dates for the fully illiquid account and the partially illiquid account, and 55 out of the 76 participants who allocated money to both accounts took advantage of this flexibility. Among the experimental participants who chose to allocate money to both the partially illiquid account and the fully illiquid account, the average goal horizon for the partially illiquid account was 116 days, and the average goal horizon for the fully illiquid account was 145 days, a difference that is statistically significant at the 1% level in a paired *t*-test. Hence, participants tended to use the partially illiquid account to create short-run commitments and the fully illiquid account to create long-run commitments. We do not know if participants would prefer to use the fully illiquid account to create commitments at all horizons (if they were given the option to do so).

VII. Conclusion

This paper studies the demand for commitment devices in the form of illiquid financial accounts. We experimentally identify the features of commitment accounts that make them attractive by asking participants to allocate an endowment between a liquid account and one or more commitment accounts, which place restrictions on early withdrawals.

We find that participants allocate meaningful fractions of their endowments to commitment accounts. When the commitment account and the liquid account have the same interest rate, commitment account allocations are increasing in the commitment account’s degree of illiquidity. We extend the theoretical framework of Amador, Werning, and Angeletos (2006) to show that this result is predicted by a model in which sophisticated present-biased agents face a tradeoff between commitment and flexibility, under a wide range of assumptions. However, when the commitment account pays a higher interest rate than the liquid account, commitment account allocations do not vary with the commitment account’s degree of illiquidity, which can be explained if individuals without present bias and/or present-biased individuals who are naïve about their present bias are also in our sample. Thus, increasing the illiquidity of 401(k) and IRA accounts, which yield higher after-tax returns than more liquid accounts, may not increase aggregate contributions to those accounts despite the desire for strict commitment within a segment of the population.

Many retirement savings accounts only weakly restrict pre-retirement spending. Withdrawals from 401(k) plans and IRAs before the age of 59½ generate only a 10% tax penalty, and there are many classes of withdrawals from these accounts that are penalty-free. It is estimated that 46% of workers with 401(k) accounts who leave their jobs receive their 401(k) balances as a lump-sum withdrawal (Hewitt Associates, 2009), and retirement savings plan managers assert that this "leakage" is socially sub-optimal (Steyer, 2011). Beshears et al. (2015) find that U.S. employer-based defined contribution accounts are unusually liquid in a comparison across six developed countries. Future work should explore the optimal amount of illiquidity that a retirement savings system should have.

References

Ainslie, George, 1991. "Derivation of 'Rational' Economic Behavior from Hyperbolic Discount Curves." *American Economic Review Papers and Proceedings* 81.2, pp. 334-340.

Amador, Manuel, Ivan Werning, and George-Marios Angeletos, 2006. "Commitment vs. Flexibility." *Econometrica* 74.2, pp. 365-396.

Ambrus, Attila and Georgy Egorov, 2013. "Comment on 'Commitment vs. Flexibility.'" *Econometrica* 81.5, pp. 2113-2124.

Ariely, Dan and Klaus Wertenbroch, 2002. "Procrastination, Deadlines, and Performance: Self-Control by Precommitment." *Psychological Science* 13.2, pp. 219-224.

Ashraf, Nava, Dean Karlan, and Wesley Yin, 2006. "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Phillipines." *Quarterly Journal of Economics* 121.2, pp. 635-672.

Augenblick, Ned, Muriel Niederle, and Charles Sprenger, forthcoming. "Working Over Time: Dynamic Inconsistency in Real Effort Tasks." *Quarterly Journal of Economics*.

Aumann, Robert and Lloyd Shapley, 1976. "Long Term Competition: A Game Theoretic Analysis." Mimeo, Hebrew University.

Benabou, Roland and Jean Tirole, 2004. "Willpower and Personal Rules." *Journal of Political Economy* 112.4, pp. 848-886.

Beshears, John, James J. Choi, Joshua Hurwitz, David Laibson, and Brigitte C. Madrian, 2015. "Liquidity in Retirement Savings Systems: An International Comparison." *American Economic Review Papers and Proceedings* 105.5, pp. 420-425.

Brune, Lasse, Xavier Gine, Jessica Goldberg, and Dean Yang, forthcoming. "Facilitating Savings for Agriculture: Field Experimental Evidence from Malawi." *Economic Development and Cultural Change*.

Burke, Jeremy, Jill E. Luoto, and Francisco Perez-Arce, 2014. "Soft versus Hard Commitments: A Test on Savings Behaviors." RAND Working Paper No. WR-1055.

Chow, Vinci Y. C., 2011. "Demand for a Commitment Device in Online Gaming." Working paper.

Dupas, Pascaline and Jonathan Robinson, 2013. "Why Don't the Poor Save More? Evidence from Health Savings Experiments." *American Economic Review* 103.4, pp. 1138-1171.

Einav, Liran, Amy Finkelstein, and Paul Schrimpf, 2010. "Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market." *Econometrica* 78.3, pp. 1031-1092.

Finkelstein, Amy and James Poterba, 2004. "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market." *Journal of Political Economy* 112.1, pp. 183-208.

Friedman, James W., 1971. "A Non-cooperative Equilibrium for Supergames." *Review of Economic Studies* 38.1, pp. 1-12.

Gine, Xavier, Dean Karlan, and Jonathan Zinman, 2010. "Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation." *American Economic Journal: Applied Economics* 2.4, pp. 213-235.

Golosov, Mikhail, Aleh Tsyvinski, and Ivan Werning, 2007. "New Dynamic Public Finance: A User's Guide." In Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds., *NBER Macroeconomics Annual 2006, Volume 21*, Cambridge: MIT Press, pp. 317-388.

Hewitt Associates, 2009. "The Erosion of Retirement Security from Cash-outs: Analysis and Recommendations." Available at http://www.aon.com/human-capital-consulting/thought-leadership/retirement/reports-pubs_retirement_cash_outs.jsp, accessed April 18, 2011.

Hogg, Robert V., Joseph W. McKean, and Allen T. Craig, 2005. *Introduction to Mathematical Statistics*. 6th Edition. Upper Saddle River: Prentice Hall.

Houser, Daniel, Daniel Schunk, Joachim Winter, and Erte Xiao, 2010. "Temptation and Commitment in the Laboratory." Institute for Empirical Research in Economics, University of Zurich Working Paper No. 488.

Investment Company Institute, 2015. "The U.S. Retirement Market, First Quarter 2015." Washington, D.C.: Investment Company Institute. http://www.ici.org/info/ret_15_q1_data.xls, accessed July 11, 2015.

John, Anett, 2015. "When Commitment Fails: Evidence from a Regular Saver Product in the Philippines." Working paper.

Karlan, Dean and Leigh L. Linden, 2014. "Loose Knots: Strong versus Weak Commitments to Save for Education in Uganda." NBER Working Paper No. 19863.

Kast, Felipe, Stephan Meier, and Dina Pomeranz, 2014. "Under-Savers Anonymous: Evidence on Self-Help Groups and Peer Pressure as a Savings Commitment Device." Harvard Business School Working Paper No. 12-060.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan, forthcoming. "Self-Control at Work." *Journal of Political Economy*.

Kreps, David M., Paul Milgrom, John Roberts, and Robert Wilson, 1982. "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory* 27.2, pp. 245-252.

Laibson, David, 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112.2, pp. 443-478.

Luenberger, David G., 1969. *Optimization by Vector Space Methods*. New York: John Wiley & Sons, Inc.

Milkman, Katherine L., Julia A. Minson, and Kevin G. M. Volpp, 2013. "Holding the Hunger Games Hostage at the Gym: An Evaluation of Temptation Bundling." *Management Science* 60.2, pp. 283-299.

Phelps, E. S. and R. A. Pollak, 1968. "On Second-Best National Saving and Game-Equilibrium Growth." *Review of Economic Studies* 35.2, pp. 185-199.

Redelmeier, Donald A. and Eldar Shafir, 1995. "Medical Decision Making in Situations That Offer Multiple Alternatives." *Journal of the American Medical Association* 273.4, pp. 302-305.

Rice, John A., 1995. *Mathematical Statistics and Data Analysis*. 2nd Edition. Belmont: Duxbury Press.

Royer, Heather, Mark Stehr, and Justin Sydnor, 2015. "Incentives, Commitments and Habit Formation in Exercise: Evidence from a Field Experiment with Workers at a Fortune-500 Company." *American Economic Journal: Applied Economics* 7.3, pp. 51-84.

Shafir, Eldar, Itamar Simonson, and Amos Tversky, 1993. "Reason-based Choice." *Cognition* 49.1-2, pp. 11-36.

Simonson, Itamar, 1989. "Choice Based on Reasons: The Case of Attraction and Compromise Effects." *Journal of Consumer Research* 16.2, pp. 158-174.

Sobel, Joel, 1985. "A Theory of Credibility." *Review of Economic Studies* 52.4, pp. 557-573.

Steyer, Robert, 2011. "DC Plan Leakage Problem Alarming, Solutions Evasive." *Pensions & Investments*, April 4, 2011. Available at http://www.pionline.com/apps/pbcs.dll/article?AID=/20110404/PRINTSUB/304049977&crit=leakage, accessed April 18, 2011.

Table 1. Participant Characteristics

This table summarizes participants' demographic characteristics in the first experiment (n = 495) and the second experiment (n = 550).

	Expt. 1	Expt. 2		Expt. 1	Expt. 2
Percent male	43%	43%	Marital status		
			Married	68%	66%
Age			Separated/divorced	11%	14%
≤ 25	8%	8%	Widowed	5%	5%
26-35	17%	19%	Never married	16%	15%
36-45	21%	18%			
46-55	22%	22%	Job status (overlapping categories)		
56-65	16%	15%	Working now	63%	60%
≥ 66	16%	17%	Unemployed	8%	9%
			Temporary layoff	1%	1%
Education			Disabled	4%	6%
No high school diploma	3%	5%	Retired	19%	19%
			Homemaker	10%	11%
High school graduate	32%	29%			
Some college	20%	23%	Race		
Associate's degree	7%	12%	White/Caucasian	80%	81%
Bachelor's degree	24%	19%	Black/African American	8%	10%
Graduate degree	13%	12%			
			American Indian or Alaskan Native	1%	1%
Annual household income					
< \$15,000	6%	9%	Asian or Pacific Islander	4%	2%
\$15,000 - \$34,999	19%	20%			
\$35,000 - \$49,999	16%	16%			
\$50,000 - \$74,999	27%	22%			
\$75,000 - \$99,999	15%	16%			
≥ \$100,000	17%	17%			

Table 2. Sample Size in Each Experimental Condition: Experiment 1

This table reports the number of participants who were assigned to each experimental condition in Experiment 1 (February 1, 2010, to February 13, 2011).

Withdrawal restrictions on commitment account prior to commitment date	Commitment account interest rate 21%	22%	23%
10% early withdrawal penalty	72	66	78
20% early withdrawal penalty	0	79	68
No early withdrawals	0	64	68

Table 3. Percent of Endowment Allocated to Commitment Account: Experiment 1

For each experimental condition, this table reports the mean percent of endowment allocated to the commitment account. There are three observations for every participant: one observation for each possible endowment amount. Standard errors clustered at the participant level are in parentheses. The table also gives *p*-values from tests of equality of means, as indicated.

Withdrawal restrictions on commitment account prior to commitment date	Commitment account interest rate			*p*-value of equality of means	
	21%	22%	23%	21% vs. 22%	22% vs. 23%
10% early withdrawal penalty	27.6	38.9	58.2	0.011	0.000
	(2.8)	(3.4)	(3.4)		
20% early withdrawal penalty	--	44.8	61.1	--	0.001
		(3.4)	(3.4)		
No early withdrawals	--	56.0	59.9	--	0.469
		(4.1)	(3.6)		
p-value of equality of means					
10% penalty vs. 20% penalty	--	0.220	0.539		
10% penalty vs. no early w/d	--	0.002	0.719		
20% penalty vs. no early w/d	--	0.035	0.809		

Table 4. Days to Commitment Date: Experiment 1

For each experimental condition, this table reports the mean days between the initial allocation decision date and the commitment date. There are up to three observations for every participant: one observation for each possible endowment amount. If a participant allocates no money to the commitment account for a given endowment amount, the days to commitment date for that participant and endowment amount is treated as missing. Standard errors clustered at the participant level are in parentheses. The table also gives *p*-values from tests of equality of means, as indicated.

Withdrawal restrictions on commitment account prior to commitment date	Commitment account interest rate			*p*-value of equality of means	
	21%	22%	23%	21% vs. 22%	22% vs. 23%
10% early withdrawal penalty	234.0	209.0	227.6	0.165	0.306
	(12.0)	(13.4)	(12.3)		
20% early withdrawal penalty	--	207.4	202.1	--	0.775
		(12.5)	(13.7)		
No early withdrawals	--	214.3	186.0	--	0.136
		(14.1)	(12.6)		
p-value of equality of means					
10% penalty vs. 20% penalty	--	0.931	0.167		
10% penalty vs. no early w/d	--	0.785	0.019		
20% penalty vs. no early w/d	--	0.716	0.384		

Table 5. Dollar-Weighted Days to Commitment Date: Experiment 1

For each experimental condition, this table reports the mean dollar-weighted days to commitment date, which is the fraction of the endowment initially allocated to the commitment account multiplied by the number of days separating the initial allocation decision date and the commitment date. There are three observations for every participant: one observation for each possible endowment amount. Standard errors clustered at the participant level are in parentheses. The table also gives *p*-values from tests of equality of means, as indicated.

Withdrawal restrictions on commitment account prior to commitment date	Commitment account interest rate			*p*-value of equality of means	
	21%	22%	23%	21% vs. 22%	22% vs. 23%
10% early withdrawal penalty	64.3	81.8	129.6	0.136	0.001
	(7.3)	(9.1)	(10.6)		
20% early withdrawal penalty	--	100.5	127.0	--	0.108
		(10.9)	(12.3)		
No early withdrawals	--	131.8	117.8	--	0.436
		(13.9)	(11.2)		
p-value of equality of means					
10% penalty vs. 20% penalty	--	0.188	0.872		
10% penalty vs. no early w/d	--	0.003	0.447		
20% penalty vs. no early w/d	--	0.078	0.584		

Table 6. Sample Size in Each Experimental Condition: Experiment 2

This table reports the number of participants who were assigned to each experimental condition in Experiment 2 (February 14, 2011, to September 1, 2011).

	Endowment allocation	
Withdrawal restrictions on commitment account prior to commitment date	According to participant's choice	All in liquid account
Safety valve (withdrawals only in financial emergencies)	85	65
10% early withdrawal penalty	54	46
No early withdrawals	60	90
Two commitment accounts: 10% early withdrawal penalty and no early withdrawals	70	80

Table 7. Percent of Endowment Allocated to Commitment Account: Experiment 2

For each experimental condition, this table reports the mean percent of endowment allocated to a commitment account. For the condition offering two commitment accounts, mean allocations are also reported for each individual commitment account. Standard are in parentheses.

Withdrawal restrictions on commitment account prior to commitment date	% allocated to commitment account
Safety valve (withdrawals only in financial emergencies)	45.3 (2.7)
10% early withdrawal penalty	45.8 (2.9)
No early withdrawals	53.7 (2.3)
Two commitment accounts: 10% early withdrawal penalty and no early withdrawals	50.1 (2.7)
Allocation to 10% early withdrawal penalty account	16.2 (1.4)
Allocation to no early withdrawals account	33.9 (2.4)

Table 8. Days to Commitment Date: Experiment 2

For each experimental condition, this table reports the mean days between the initial allocation decision date and the commitment date. If a participant allocates no money to a commitment account, the days to commitment date for that participant and commitment account is treated as missing. Standard errors are in parentheses. The table also gives *p*-values from tests of equality of means, as indicated.

Withdrawal restrictions on commitment account prior to commitment date	Days to commitment date	*p*-value of equality of means vs. no early withdrawals only
Safety valve (withdrawals only in financial emergencies)	135.4 (5.4)	0.923
10% early withdrawal penalty	135.6 (6.0)	0.900
No early withdrawals	134.7 (4.5)	--
Two commitment accounts		
10% early withdrawal penalty	116.3 (6.5)	0.020
No early withdrawals	148.7 (5.5)	0.050

Table 9. Dollar-Weighted Days to Commitment Date: Experiment 2

For each experimental condition, this table reports the mean dollar-weighted days to commitment date. When one commitment account is offered, dollar-weighted days to commitment date is defined as the fraction of the endowment initially allocated to the commitment account multiplied by the number of days separating the initial allocation date and the commitment date. When two commitment accounts are offered, dollar-weighted days to commitment date is obtained by calculating this product for each account and taking the sum. Standard errors are in parentheses. The table also gives *p*-values from tests of equality of means, as indicated.

Withdrawal restrictions on commitment account prior to commitment date	Dollar-weighted days to commitment date	*p*-value of equality of means vs. no early withdrawals only
Safety valve (withdrawals only in financial emergencies)	62.0 (4.6)	0.046
10% early withdrawal penalty	64.4 (5.5)	0.137
No early withdrawals	74.8 (4.4)	--
Two commitment accounts: 10% early withdrawal penalty and no early withdrawals	71.3 (4.8)	0.587

Table 10. Mean Withdrawal Measure for Own versus All Liquid Allocation: Experiment 2
For each participant and for a given number of days since the start of the experiment, we calculate the ratio of actual balances in experimental accounts to hypothetical balances in experimental accounts had the participant not made any withdrawals. For each experimental condition and for the experimental conditions combined, this table reports the mean difference between the balance ratio at various dates for participants who were randomly assigned to receive their chosen allocations versus participants who were randomly assigned to receive their entire endowment in the liquid account. Standard errors robust to heteroskedasticity are in parentheses.

Withdrawal restrictions on commitment account prior to commitment date	Own allocation vs. all in liquid account mean difference				
	Days since initial deposit into participant accounts				
	20	60	100	140	180
Safety valve (withdrawals only in financial emergencies)	0.049	-0.004	0.002	0.022	-0.027
	(0.033)	(0.047)	(0.059)	(0.066)	(0.071)
10% early withdrawal penalty	0.120*	0.121	0.156	0.197*	0.143
	(0.060)	(0.071)	(0.082)	(0.087)	(0.090)
No early withdrawals	0.070*	0.149**	0.127*	0.092	0.114
	(0.034)	(0.047)	(0.057)	(0.070)	(0.073)
Two commitment accounts	-0.038	0.029	0.026	0.035	0.064
	(0.031)	(0.046)	(0.053)	(0.057)	(0.061)
Combined	0.044*	0.069**	0.069*	0.078*	0.067
	(0.019)	(0.026)	(0.031)	(0.034)	(0.036)

* Significant at the 5% level. ** Significant at the 1% level.

Figure 1. Illustrative Budget Set

This figure shows a budget set that illustrates Constraints 1 and 2 of the model. The budget set is a non-empty compact subset of the ambient budget set and therefore satisfies Constraint 1. It satisfies Constraint 2 if $\pi = 0.5$, but not if $\pi = 0.1$. The slope at the encircled point is -1.3. This is greater than -1.5 (the minimum slope that is permissible when $\pi = 0.5$), but less than -1.1 (the minimum slope that is permissible when $\pi = 0.1$).

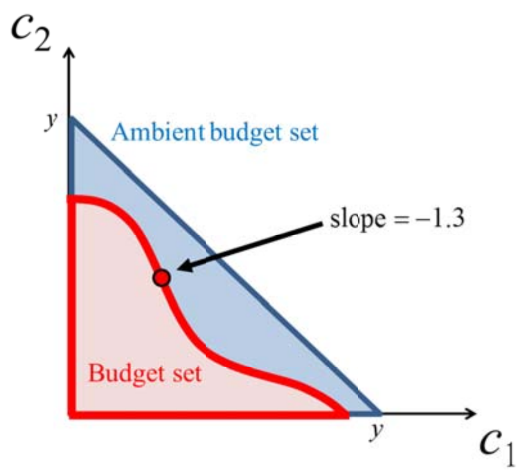

Figure 2. Two-Part Budget Set

This figure shows a two-part budget set. Such budget sets consist of: (i) a budget curve that has slopes of -1 and $-(1+p)$ to the left and right of a kink at (c_1^*, c_2^*); and (ii) none, some or all of the points of the ambient budget set that lie below or to the left of the budget curve.

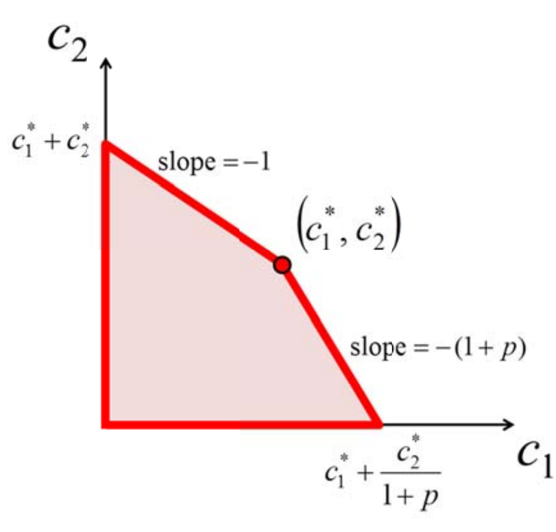

Appendix Figure 1. Description of the Liquid Account

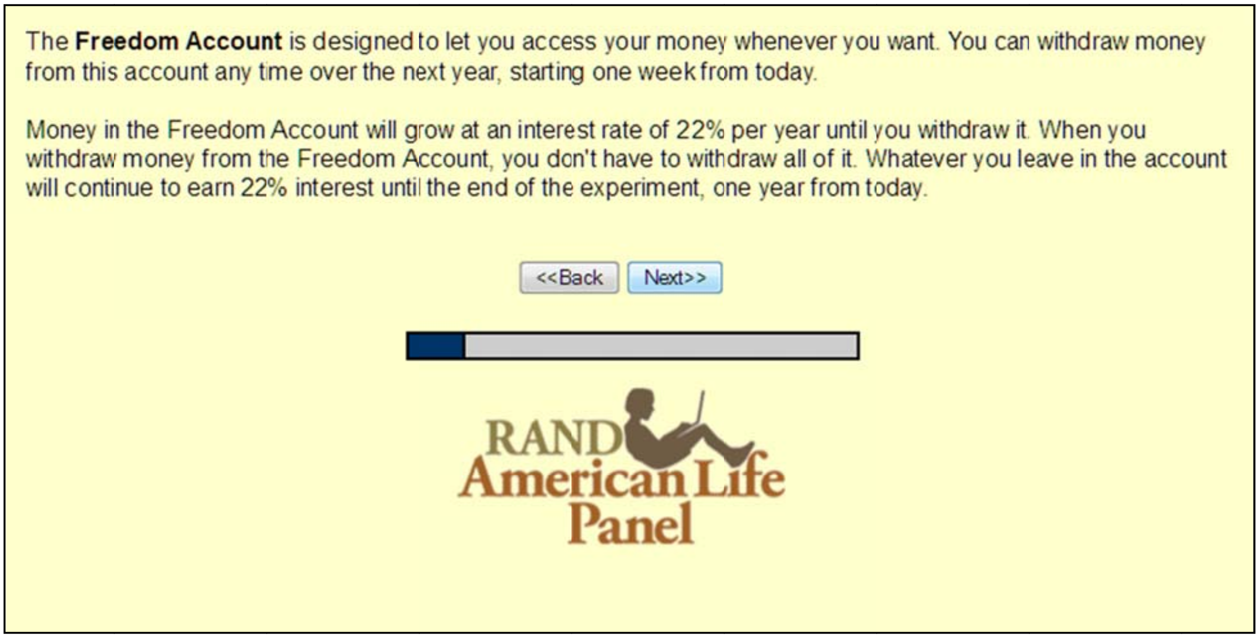

Appendix Figure 2. Description of the 22% Interest Rate, 10% Early Withdrawal Penalty Commitment Account

The **Goal Account** is designed to help you save. You can withdraw money from this account without penalty any time after a goal date that you pick. Setting a goal for yourself and picking the right goal date can help you avoid the temptation to spend your money too soon.

Money in the Goal Account will grow at an interest rate of 22% per year, both before and after the goal date, until you withdraw it. When you withdraw money from the Goal Account, you don't have to withdraw all of it. Whatever you leave in the account will continue to earn 22% interest until the end of the experiment, one year from today.

As explained earlier, if you withdraw money from the Goal Account before your goal date, you will incur a penalty equal to 10% of the amount you withdraw.

<<Back Next>>

RAND American Life Panel

To use the Goal Account, you will need to pick a goal date. You might want to pick a date based on something you want to save money for, like a birthday gift, holiday presents, vacation, or any other special purchase that you plan to make. You can also use the Goal Account as a way to help you save, even if you don't have a special purchase in mind.

<<Back Next>>

RAND American Life Panel

Appendix Figure 3. Example Allocation Page

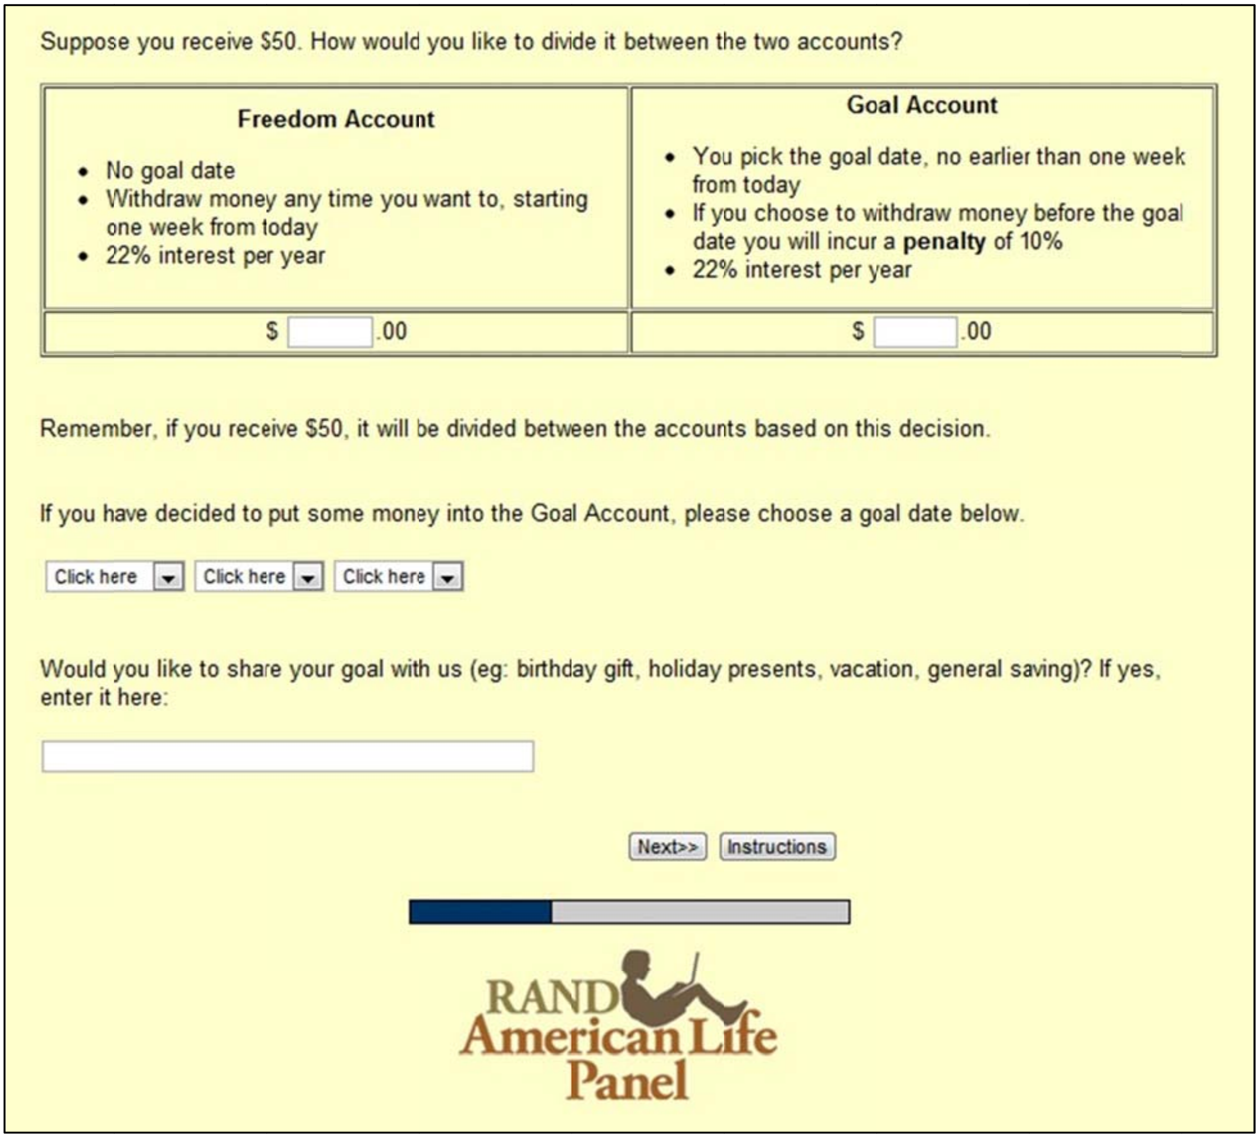

Suppose you receive $50. How would you like to divide it between the two accounts?

Freedom Account	Goal Account
• No goal date • Withdraw money any time you want to, starting one week from today • 22% interest per year	• You pick the goal date, no earlier than one week from today • If you choose to withdraw money before the goal date you will incur a **penalty** of 10% • 22% interest per year
$ .00	$ .00

Remember, if you receive $50, it will be divided between the accounts based on this decision.

If you have decided to put some money into the Goal Account, please choose a goal date below.

Click here Click here Click here

Would you like to share your goal with us (eg: birthday gift, holiday presents, vacation, general saving)? If yes, enter it here:

Next>> Instructions

RAND American Life Panel

Appendix Figure 4. Sample Weekly Email to Participant

Dear Participant,

This is a breakdown of your current balances:

Freedom Account: $24.25
Goal Account: $53.18
Goal Date: July 20th, 2010

If you wish to withdraw any money from your accounts, please go to your panel pages and click on the "Savings Game" button: https://mmic.rand.org/panel

If you have any questions about this game or your accounts, please feel free to contact us at webhelp@rand.org or 866.591.2909

Thanks!
www.rand.org/alp

Appendix Figure 5. Withdrawal Interface

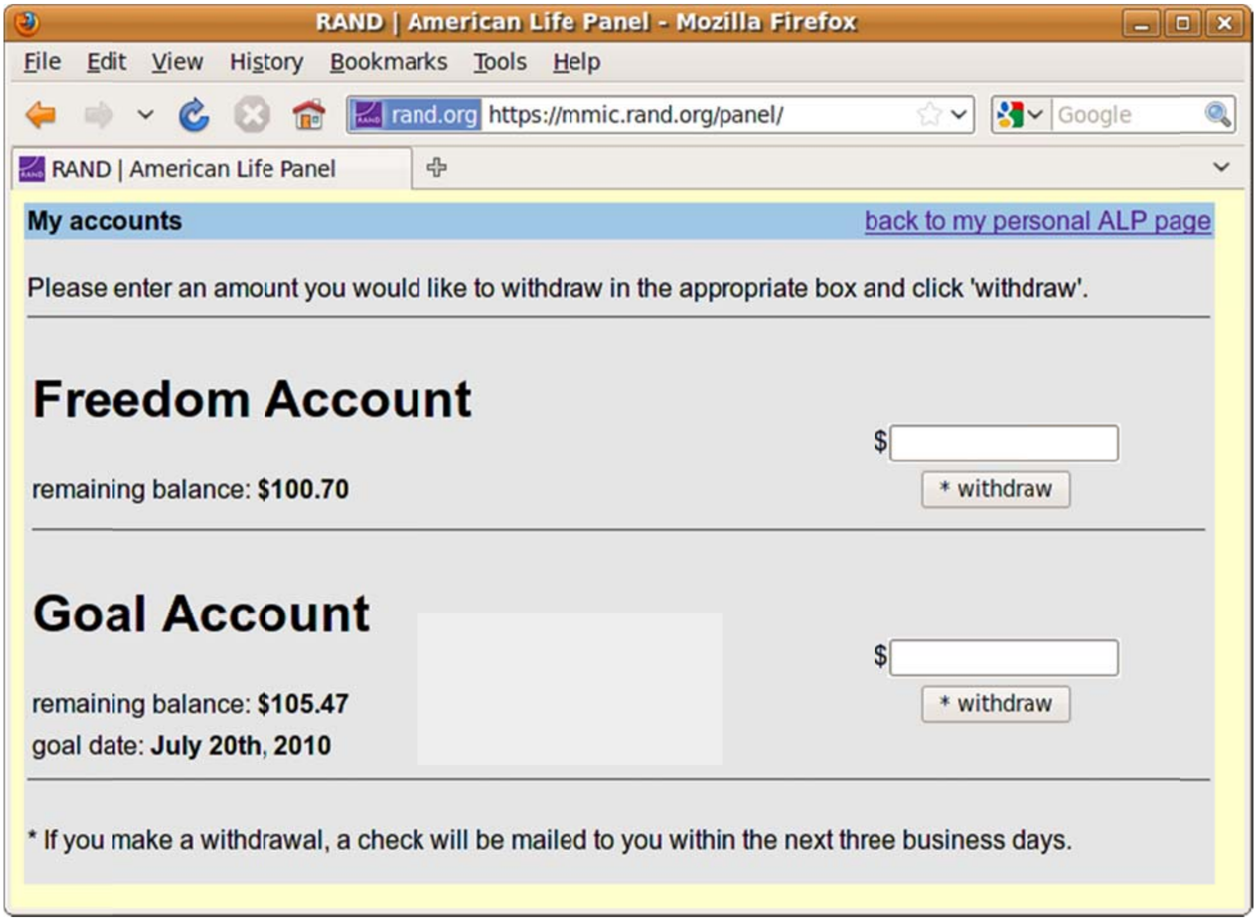

Appendix Figure 6. Balance Ratios by Experimental Condition: Experiment 1

For each experimental condition, these figures show withdrawal patterns over the course of the experiment. For each participant and for each day, we calculate the sum of the liquid account and commitment account balances that the participant would have had if no withdrawals had been requested. This hypothetical total balance takes as given the participant's initial allocation between the liquid account and the commitment account, and it uses the allocation decision that applies to the ex post realization of the endowment amount ($50, $100, or $500). We then calculate the ratio of the participant's actual balance to the hypothetical total balance, and we plot the mean of this ratio against the number of days since the initial deposit into the participant's accounts.

10% withdrawal penalty conditions

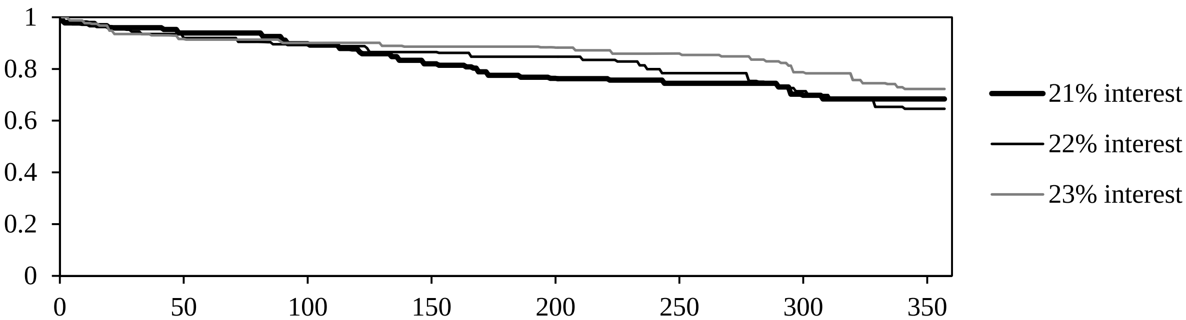

22% commitment account interest rate conditions

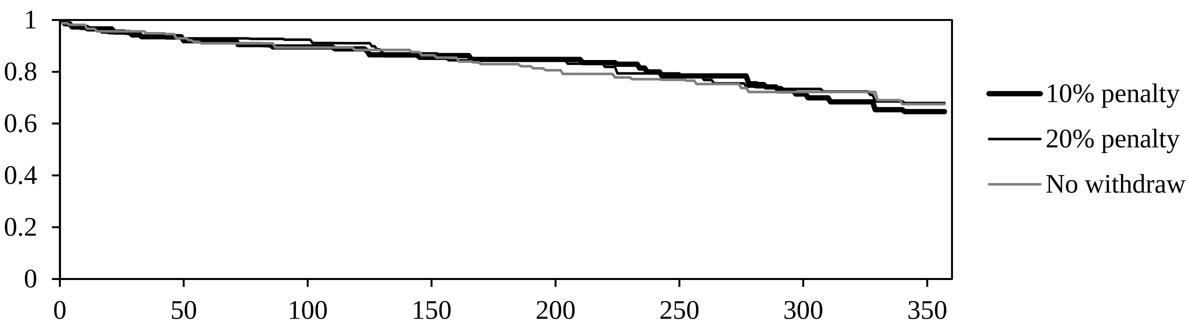

23% commitment account interest rate conditions

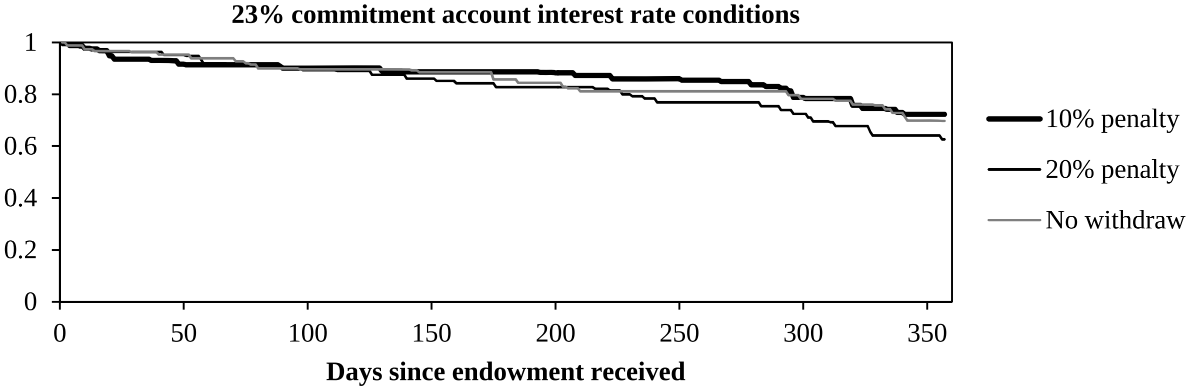

Appendix Figure 7. Description of Two Commitment Accounts Offered Simultaneously

The **Goal Accounts** are designed to help you save. You can withdraw money from these accounts any time **on or after** goal dates that you pick. Setting goals for yourself and picking the right goal dates can help you avoid the temptation to spend your money too soon.

There are two types of Goal Accounts:

- Goal Account A (10% Penalty) allows you to withdraw your money **before** its goal date, but you will be charged a 10% penalty on early withdrawals. For example, if you withdraw $10 before your goal date, your account balance will be reduced by $11.
- Goal Account B (No Withdrawal) does **not** allow withdrawals **before** its goal date.

If you choose to use both Goal Accounts, you can pick a different goal date for each Goal Account, or you can pick the same goal date.

Money in both Goal Accounts will grow at an interest rate of 22% per year, both before and after the goal date, until you withdraw it. When you withdraw money from a Goal Account, you don't have to withdraw all of it. Whatever you leave in the accounts will continue to earn 22% interest until the end of the experiment on September 1, 2011.

<<Back Next>>

RAND American Life Panel

Appendix Figure 8. Description of the Safety Valve Commitment Account

The **Goal Account** is designed to help you save. You can withdraw money from this account any time **on or after** a goal date that you pick. Setting a goal for yourself and picking the right goal date can help you avoid the temptation to spend your money too soon.

You cannot withdraw from this account **before** the goal date, except in the case of a financial emergency. If you have a financial emergency, you can make an early withdrawal. We are relying on you to be honest in judging whether you have a financial emergency.

Money in the Goal Account will grow at an interest rate of 22% per year, both before and after the goal date, until you withdraw it. When you withdraw money from the Goal Account, you don't have to withdraw all of it. Whatever you leave in the account will continue to earn 22% interest until the end of the experiment on September 1, 2011.

<<Back Next>>

RAND American Life Panel

Appendix Figure 9. Withdrawal Patterns for Own versus All Liquid Allocation: Experiment 2

For each experimental condition, these figures show withdrawal patterns over the course of the experiment for participants who were randomly assigned to receive their chosen allocations and for participants who were randomly assigned to receive their entire endowment in the liquid account. For each participant and for each day, we calculate the sum of the liquid account and commitment account balances that the participant would have had if no withdrawals had been requested. We then calculate the ratio of the participant's actual balance to this hypothetical total balance, and we plot the mean of this ratio against the number of days since the initial deposit into the participant's accounts.