Which Early Withdrawal Penalty Attracts the Most Deposits to a Commitment Savings Account?

John Beshears, Harvard University and NBER
James J. Choi, Yale University and NBER
Christopher Harris, University of Cambridge
David Laibson, Harvard University and NBER
Brigitte C. Madrian, Harvard University and NBER
Jung Sakong, University of Chicago

July 10, 2018

Abstract: Previous research has shown that some people voluntarily use commitment contracts that restrict their own choice sets. We use an experiment to study how people divide money between two accounts: a liquid account that permits unrestricted withdrawals and a commitment account that is randomly assigned between subjects to have a 10% early withdrawal penalty, a 20% early withdrawal penalty, or not allow early withdrawals (which is like an infinite penalty). When the two accounts pay the same interest rate, higher penalties attract more commitment account deposits, suggesting that some participants are sophisticated present-biased agents who understand the commitment benefits of higher penalties. However, the experiment also shows that when the commitment account pays a *higher* interest rate than the liquid account, the empirical relationship between penalties and commitment deposits is flat, suggesting that naïve present-biased agents or agents without present bias are also in our sample.

Keywords: quasi-hyperbolic discounting, present bias, sophistication, naiveté, commitment, flexibility, savings, contract design, defined contribution retirement plan, 401(k), IRA

This paper was formerly titled "Self Control and Liquidity: How to Design a Commitment Contract" and "Self Control and Commitment: Can Decreasing the Liquidity of a Savings Account Increase Deposits?" 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 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, Luigi Guiso, Ulrike Malmendier, Sendhil Mullainathan, Christopher Parsons, and Ivan Werning. We thank Christopher Clayton, Omeed Maghzian, 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, any agency of the Federal Government, or the NBER.

In 2017, U.S. households held $16.9 trillion in employer-sponsored defined contribution savings plans and IRAs (Investment Company Institute, 2018). These retirement savings accounts are partially illiquid: withdrawals before age 59½ incur an early withdrawal penalty equal to 10% of the withdrawal (in *addition* to any income taxes that are owed).[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] Despite the 10% penalty, however, early withdrawals from retirement accounts are substantial. For every dollar that households younger than age 55 in the U.S. contributed to retirement accounts in 2010, those same households had $0.20 of penalized early withdrawals (Argento, Bryant, and Sabelhaus, 2015). Retirement savings plan managers assert that this "leakage" is socially sub-optimal (Steyer, 2011). One potential solution to this perceived problem is to increase the penalty on early withdrawals to make retirement savings accounts more illiquid, as they are in several other developed countries (Beshears et al., 2015). How would households respond if the early withdrawal penalty in the U.S. were higher than 10%?

The answer to this question is unclear from a theoretical perspective. Although higher penalties will reduce early withdrawals, higher penalties will also discourage initial deposits for neoclassical economic agents who prefer liquidity, undermining the goal of raising net savings. On the other hand, some savers may believe that penalties help them partially overcome self-control problems. These households will perceive that higher penalties have both costs and benefits, so the impact of higher early withdrawal penalties on their deposits is ambiguous.

It is challenging to identify natural experiments that would permit an analysis of behavioral responses to variation in early withdrawal penalties, so in this paper, we use an

[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, no penalty is charged for IRA accounts 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). Finally, Roth IRAs have low (or even zero) penalties for withdrawals.

[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).

experimental approach to shed light on the issue. The results of our experiments cannot be applied directly to predict how individuals would respond to a change in U.S. policy regarding early withdrawal penalties, but the primary contribution of this paper is to use the control available in an experimental setting to study the underlying economic forces at play. In our experiments, a higher early withdrawal penalty does not discourage average deposits to an illiquid account. Indeed, under some conditions, a higher early withdrawal penalty *increases* deposits to the illiquid account, suggesting that sophisticated present-biased individuals are present in the population. However, we also find empirical evidence of heterogeneity in present bias, implying that policy makers must take multiple subpopulations into account when designing an optimal savings system.

The 1,045 participants in our two online experiments are drawn from the American Life Panel, a sample of U.S. adults who regularly take part in online research studies. Each participant is given $50, $100, or $500. Participants are asked to allocate this endowment between a liquid account, which does not limit withdrawals in any way, and one or more commitment accounts. All participants have access to the same type of liquid account (in particular, every participant receives the same interest rate from the liquid account), but the characteristics of the commitment accounts vary across participants. Each commitment account has a commitment date that is selected by the participant at the start of the experiment and may be up to one year in the future. The commitment account either penalizes withdrawals before the commitment date or prohibits such early withdrawals altogether; these penalties/prohibitions are randomly assigned in the experiment. The interest rates on the commitment accounts also vary randomly across participants.

When we offer participants only one commitment account and set its interest rate equal to the interest rate on the liquid account, allocations to the commitment account increase as its early withdrawal penalty rises (across subjects) from 10% to 20% to not allowing any early withdrawals (which is like an infinite penalty). In another arm of the study, we give participants simultaneous access to a liquid account and two types of commitment accounts, one with a 10% early withdrawal penalty and one that does not allow early withdrawals. The commitment account with the 10% early withdrawal penalty receives half as much money as the commitment account that prohibits early withdrawals.

These experimental results are consistent with the presence of fully or partially sophisticated present-biased agents in the sample. Individuals without present bias and naïve present-biased individuals (those who are present-biased but do not anticipate their present bias) would not allocate balances to a commitment account in the first place or, if they were to allocate balances to a commitment account due to experimenter demand effects, would likely not increase their commitment account allocations in response to higher early withdrawal penalties. These agents do not perceive a benefit from higher penalties, as they believe that they have no need for commitment. They only perceive the cost of greater financial losses if early withdrawals are necessary. Fully or partially sophisticated present-biased agents (agents who are at least somewhat aware of their self-control problems), on the other hand, perceive the same cost but also perceive the benefit of the stronger commitment afforded by higher early withdrawal penalties (Laibson, 1997). Their commitment account allocations may increase or decrease with higher penalties.[3] Thus, the empirically observed aggregate increase in commitment account deposits in response to higher penalties suggests the presence of fully or partially sophisticated present-biased agents.

However, in our experiments, higher early withdrawal penalties do not always increase deposits to commitment accounts. We find that when we offer participants only one commitment account and set its interest rate to be slightly higher than the interest rate on the liquid account—as is the case with 401(k) accounts and IRAs, which both have tax-preferred status—deposits to the commitment account essentially do not respond to rising early withdrawal penalties. This result suggests that the U.S. adult population contains not only sophisticated present-biased individuals, but also individuals without present bias or naïve present-biased individuals. When the commitment account pays an interest rate premium, these latter two groups make deposits to commitment accounts that are positive but diminishing with the commitment account's early withdrawal penalty. This decrease offsets the increase in commitment account deposits by sophisticated present-biased individuals as the early withdrawal penalty rises. Therefore, the aggregate relationship between commitment deposits and the early withdrawal penalty can take

[3] In fact, in Online Appendix B, we extend the theoretical analysis of Amador, Werning, and Angeletos (2006) and show that the benefit of the stronger commitment afforded by higher early withdrawal penalties tends to outweigh the cost when it comes to determining the relationship between higher penalties and commitment account allocations. In the model, fully or partially sophisticated present-biased agents are subject to stochastic, uninsurable taste shocks drawn from a broad class of distributions that affect future marginal utility and create a motive to provide spending flexibility to the future self. We provide conditions under which the desire for commitment outweighs the desire for flexibility in the sense that commitment account deposits increase with the commitment accounts' early withdrawal penalty.

any sign, including the flat relationship we observe in our data.

Demand for commitment devices has been documented in many different 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., 2018), reducing time spent playing online games (Acland and Chow, 2018), going to the gym (Milkman, Minson, and Volpp, 2013; Royer, Stehr, and Sydnor, 2015), performing an unpleasant task (Augenblick, Niederle, and Sprenger, 2015), achieving workplace goals (Kaur, Kremer, and Mullainathan, 2015), selecting food items (Sadoff, Samek, and Sprenger, 2015), reducing alcohol consumption (Schilbach, 2018), and repaying debt (Cho and Rust, 2017). Our paper is most closely related to previous work on commitment savings accounts. Ashraf, Karlan, and Yin (2006) offered Filipino households a savings account that did not allow withdrawals until a certain date had passed or a certain 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] Further research on this topic has examined how deposits to commitment savings accounts vary according to the features of those accounts, including the presence of restrictions on the types of items that can be purchased with the money in the accounts (Dupas and Robinson, 2013; Karlan and Linden, 2014), the existence of physical barriers to accessing account balances, such as lockboxes for which a third party and not the saver has the key (Dupas and Robinson, 2013), and the imposition of psychological barriers to early withdrawals (Burke, Luoto, and Perez-Arce, forthcoming).

Our paper is distinct from these prior studies because we take inspiration from the structure of 401(k) accounts and IRAs and focus on the effect of varying the *financial* penalty for early withdrawals, conditional on offering a commitment savings account in the first place.[5] Financial penalties may have effects that are different from the effects of the other barriers to early withdrawals studied previously because, for example, people value commitment but dislike restrictions on the types of items they can purchase when they make withdrawals. Indeed, we

[4] Kast, Meier, and Pomeranz (2018) and Brune et al. (2016) also study take-up of commitment savings accounts and find similar results.

[5] Our second experiment does have one treatment arm that imposes a psychological barrier to early withdrawals. Participants must declare that they have a financial emergency if they wish to make early withdrawals from this account. If there is a psychological cost to lying, this account imposes a psychological penalty on early withdrawals that are not triggered by an emergency. We are primarily interested in this arm because it mimics the fact that IRAs and many 401(k) plans permit penalty-free withdrawals when the account holder is facing a financial hardship.

find that increasing the early withdrawal penalty can lead to higher commitment savings account deposits, while other researchers have found that imposing restrictions on the items that can be purchased using account balances can reduce deposits (Dupas and Robinson, 2013; Karlan and Linden, 2014).

While our evidence indicates the presence of fully or partially sophisticated present-biased individuals who recognize the commitment benefits of higher early withdrawal penalties, the data also suggest that there is heterogeneity in present bias in the sample. Our results therefore accord with previous work documenting present bias heterogeneity (Augenblick, Niederle, and Sprenger, 2015), and a contribution of our paper is to draw out the implications of this heterogeneity for the relationship between commitment account deposits and the level of early withdrawal penalties. In a complementary experiment, John (2018) allows individuals to select their own financial penalties for failing to follow through on their savings plans, and more than half of the participants end up paying the self-chosen penalty. Her results suggest that many participants in the experiment are partially but not fully sophisticated regarding their self-control problems. Thus, the welfare implications of increasing early withdrawal penalties for commitment savings accounts are far from clear. The current paper focuses on the descriptive question of how individuals respond to higher early withdrawal penalties, but Amador, Werning, and Angeletos (2006), Galperti (2015), Beshears et al. (2017), and Moser and Olea de Souza e Silva (2017) analyze the question of optimal commitment account design from a social welfare perspective.

This paper proceeds as follows. 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 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 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.

Conducting the experiments through the ALP offers several advantages. First, because ALP members have an ongoing relationship with RAND, they are likely to trust that the experimental procedures described to them, especially regarding the detailed rules of the financial accounts, will be carried out as promised. Second, ALP members are accustomed to reading experimental instructions, so they are likely to understand the nature of the decisions that they are asked to make. Indeed, responses to our debriefing questionnaire suggest that participants did not find our instructions confusing. Third, the private nature of an ALP member's participation in the study over the Internet casts doubt on some alternative interpretations of the demand for commitment savings accounts. For example, some individuals may make deposits to commitment accounts not because they have self-control problems but instead because commitment accounts protect financial resources from family members' and friends' requests for money. It is unlikely that participants in our experiments would make deposits to our commitment accounts for this reason, as even the liquid account that we offer to participants is difficult for others to observe and hence largely protected from others' requests. A small number of individuals in our experiments are in the same household as other participants and may therefore have their experimental participation observed, but these individuals do not drive our results—our conclusions do not change if these individuals are dropped from the analysis.

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.

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 23 participants in the second experiment are in the same household as a participant in the first experiment.

In both experiments, some ALP members who were invited to participate did not enroll in the study, so our experimental samples may not be representative of the U.S. adult population. However, while the lack of representativeness implies that the magnitudes of the effects observed in the experiments may not generalize to the U.S. adult population, it should not affect our main qualitative conclusions regarding the existence of individuals who, when asked to allocate resources between a liquid account and a commitment account with the same interest rate, respond to an increase in the early withdrawal penalty by increasing their commitment account deposits.

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. A within-subjects experimental design in which a given participant made allocation decisions for several different versions of the commitment account would have had the desirable property of eliciting individual-level demand for commitment account deposits as account features vary, but we used a between-subjects experimental design to make the decision task simple for participants and to avoid the potentially strong experimenter demand effects associated with a within-subjects design. Thus, each participant saw only one version of the commitment account.

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 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. We asked participants to choose their own commitment dates to allow for heterogeneity in the horizons over which individuals faced self-control problems. 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 other 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 just to withdraw it immediately to pay down credit card debt. Of course, savings accounts outside of our experiment have much lower interest rates, and the level of the experimental accounts' interest rates may affect the demand for commitment and how commitment account deposits respond to account liquidity. High interest rates may make illiquidity more attractive because it helps to lock in high returns, or high interest rates may make illiquidity less attractive because the high interest rates themselves serve as a deterrent to early withdrawals, rendering withdrawal restrictions superfluous. However, these issues do not pose a problem for our research design. Our conceptual arguments regarding fully sophisticated, partially sophisticated, and naïve present-biased agents and agents without present bias rely only on the liquid account and commitment account interest rates being equal, and our experiment is intended to produce generalizable insight into the qualitative impact of varying commitment account illiquidity, not the quantitative magnitude of the impact.

Table 2 summarizes the experimental design and gives the number of participants in each condition.[6] 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

[6] 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.

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

When individuals began participating in the experiment, they first saw a series of screens describing the details of the experiment. They would receive \$50, \$100, or \$500, depending on a random number 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 A1 and A2 show the screens explaining the accounts. Note that the experiment did not have a condition in which an account was labeled the “Goal Account” but was not associated with early withdrawal restrictions, so we cannot isolate the effect of account labeling. Instead, the labeling was held constant across all of the experimental conditions. Thus, while labeling was a relevant contextual factor, the design allows us to isolate the effect of varying the degree of commitment account liquidity, which is our primary research question.

All participants allocated the \$50 endowment first, the \$100 endowment second, and the \$500 endowment third. Whenever participants allocated any money to the commitment account, they were invited but not required to associate a goal with the commitment account (see Appendix Figure A3). The \$50, \$100, or \$500 endowment is a windfall, and participants' decisions when allocating a windfall between the liquid account and the commitment account may differ from the decisions they would make if they were allocating money they already had. Nonetheless, the relationship between commitment account allocations and account withdrawal restrictions in our experiment sheds light on how individuals think about the use of illiquid accounts.

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 Powerball was the third or fourth number chosen by the participant, she received \$100; and otherwise, 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 A4 shows an example of the weekly email sent to participants, and Appendix Figure A5 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.[7] Table 3 shows the mean fraction allocated to the commitment account by experimental condition. We have three main results.[8]

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 in Table 3, averaging across all penalty types), and about one-quarter of initial balances are allocated to the commitment account when it has a lower interest rate than the liquid account (21% column). Thus, it seems that some participants value commitment, as they are willing to use the commitment account despite earning no additional interest or even forgoing interest. Of course, positive demand for the commitment account could be due to experimenter demand effects, so we do not emphasize this result. We are primarily interested in how commitment account demand varies as the illiquidity of the account increases.

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%.

[7] Across all experimental conditions, 42% of participants allocate the same fraction of the endowment to the commitment account for all three allocation decisions. Among participants who do not choose the same allocation for all three decisions, 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. We speculate that changing the endowment amount changes the set of items that come to mind as temptation goods or consumption goals, sometimes leading to changes in the fraction of the endowment allocated to the commitment account.

[8] Our results are nearly identical if we control for participant characteristics using regressions. Appendix Table A1 shows that we see similar patterns when we examine the extensive margin of commitment account utilization, although the statistical significance of the differences is weaker.

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 why 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 (23% column). Commitment accounts with a 23% interest rate attract approximately 60% of the endowment regardless of their early withdrawal policy. Appendix Table A2 uses a regression framework to show that the negative interaction between the effect of the 23% interest rate (relative to the 22% interest rate) and the effect of complete illiquidity (relative to the 10% early withdrawal penalty) is statistically significant.

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 average varies between 186 days and 234 days across conditions. An alternative measure of commitment takes 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.[9]

In Online Appendix B, 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 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 also agents without present bias and/or naïve present-biased agents among our participants. Because of the commitment account's interest rate premium, it attracts some deposits from these two groups.[10] 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.

---

[9] 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.

[10] 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.

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 hypothesis 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 we plot the mean of this ratio against the number of days since the endowment was received.[11] In order to facilitate the relevant comparisons, we present subsets of the seven conditions in each of the three graphs in Appendix Figure A6, with some conditions appearing in more than one graph.

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 A6 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 in Appendix Figure A6 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 higher interest rate than the liquid account, the average balance ratio across days is 0.869 with a 10%

[11] 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.

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. To offer a different perspective on withdrawal decisions, Appendix Figure A7 shows average balance ratios for each experimental condition at four points in time: on the day of the initial deposit into participant accounts, three days before the commitment date, three days after the commitment date, and three days before remaining account balances were automatically disbursed. For participants who did not allocate any funds to a commitment account, we use the balance ratio on the initial deposit date as the balance ratio three days before the commitment date, and we use the balance ratio three days after the initial deposit date as the balance ratio three days after the commitment date.

This analysis of withdrawals is imperfect because the commitment date is an endogenous decision that is influenced by treatment assignment, but the perspective is nonetheless informative because it allows us to examine withdrawal decisions around the date that a participant deems most relevant for commitment. We view the results shown in Appendix Figure A7 as complementary to those shown in Appendix Figure A6. We find that holding fixed the commitment account interest rate, participants who were not allowed to withdraw early have the highest balance ratio three days before the commitment date. When the commitment account pays a 22% interest rate, the balance ratio is 0.939 for the 10% penalty condition, 0.926 for the 20% penalty condition, and 0.948 for the no-withdrawal condition. When the commitment account pays a 23% interest rate, the balance ratio is 0.903 for the 10% penalty condition, 0.894 for the 20% penalty condition, and 0.953 for the no-withdrawal condition. However, these differences within interest rate condition are not statistically significant.

We can also adjust for the fact that the mean commitment date differs across arms. 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 commitment 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.

Again, we find suggestive evidence that stronger commitment raises balance ratios. 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 higher interest rate than the liquid account, 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 participants with a 22% commitment account interest rate and 26.9% of 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 the second 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,[12] and weekly emails would also

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

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 A8). 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 A9).

Participants were 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.

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.[13] 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-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

[13] 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.

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.[14]

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 ($p = 0.137$).[15] 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.

[14] 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. Appendix Table A3 shows results for the extensive margin of commitment account utilization.

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

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 A10 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.[16] 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 the first 183 days since endowment receipt, so that the sample remains constant within each graph. To provide a complementary perspective, Appendix Figure A11 shows mean balance ratios in each of the experimental conditions, separately for participants who received their own allocation choices and those who received the entire endowment in the liquid account, at four points in time: the day of the initial deposit into the participant's accounts, three days before the participant's commitment date, three days after the participant's commitment date, and three days before remaining account balances were automatically disbursed to the participant.

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. Of course, we do not observe participants' other financial

[16] 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.

accounts, so higher balances in the experimental accounts may be offset by lower balances in accounts outside the experiment.

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. Conclusion

This paper studies the demand for commitment devices in the form of illiquid financial accounts, focusing on individuals' responses to variation in early withdrawal penalties. When we ask experimental participants to allocate an endowment between a liquid account and a commitment account with the same interest rate, we find that commitment account allocations are increasing in the commitment account's degree of illiquidity. This result indicates the presence of sophisticated present-biased agents in the sample. However, when the commitment account pays a higher interest rate than the liquid account, we find that commitment account allocations do not vary with the commitment account's degree of illiquidity, which can be explained if naïve present-biased individuals or individuals without 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 401(k) and IRA contributions despite the desire for strict commitment within a segment of the population.

Many U.S. 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). Our experimental results indicate that a fraction of the population—those present-biased individuals who are sophisticated about their present bias—would welcome increasing the illiquidity of retirement accounts, but future work should address the challenge of designing the liquidity features of an optimal retirement savings system that takes into account the presence of both sophisticated present-biased individuals and other types of individuals in the population.

References

Acland, Dan and Vinci Chow, 2018. "Self-Control and Demand for Commitment in Online Game Playing: Evidence from a Field Experiment." *Journal of the Economic Science Association* 4.1, pp. 46-62.

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

Argento, Robert, Victoria L. Bryant, and John Sabelhaus, 2015. "Early Withdrawals from Retirement Accounts during the Great Recession." *Contemporary Economic Policy* 33.1, pp. 1-16.

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, 2015. "Working Over Time: Dynamic Inconsistency in Real Effort Tasks." *Quarterly Journal of Economics* 130.3, pp. 1067-1115.

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.

Beshears, John, James J. Choi, Christopher Clayton, Christopher Harris, David Laibson, and Brigitte C. Madrian, 2017. "Optimal Illiquidity." Working paper.

Brune, Lasse, Xavier Gine, Jessica Goldberg, and Dean Yang, 2016. "Facilitating Savings for Agriculture: Field Experimental Evidence from Malawi." *Economic Development and Cultural Change* 64.2, pp. 187-220.

Burke, Jeremy, Jill E. Luoto, and Francisco Perez-Arce, forthcoming. "Soft versus Hard Commitments: A Test on Savings Behaviors." *Journal of Consumer Affairs*.

Cho, SungJin and John Rust, 2017. "Precommitments for Financial Self-Control? Micro Evidence from the 2003 Korean Credit Crisis." *Journal of Political Economy* 125.5, pp. 1413-1464.

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.

Galperti, Simone, 2015. "Commitment, Flexibility, and Optimal Screening of Time Inconsistency." *Econometrica* 83.4, pp. 1425-1465.

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.

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.

Houser, Daniel, Daniel Schunk, Joachim Winter, and Erte Xiao, 2018. "Temptation and Commitment in the Laboratory." *Games and Economic Behavior* 107, pp. 329-344.

Investment Company Institute, 2018. "The U.S. Retirement Market, First Quarter 2018." Washington, D.C.: Investment Company Institute. https://www.ici.org/info/ret_18_q1_data.xls accessed July 9, 2018.

John, Anett, 2018. "When Commitment Fails: Evidence from a Field Experiment." 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, 2018. "Saving More in Groups: Field Experimental Evidence from Chile." *Journal of Development Economics* 133, pp. 275-294.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan, 2015. "Self-Control at Work." *Journal of Political Economy* 123.6, pp. 1227-1277.

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

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.

Moser, Christian and Pedro Olea de Souza e Silva, 2017. "Optimal Paternalistic Savings Policies." Working paper.

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.

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.

Sadoff, Sally, Anya Samek, and Charles Sprenger, 2015. "Dynamic Inconsistency in Food Choice: Experimental Evidence from a Food Desert." Working paper.

Schilbach, Frank, 2018. "Alcohol and Self-Control: A Field Experiment in India." Working paper.

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.

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 errors 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 at a given number of days since the start of the experiment, we calculate the ratio of their actual balances in the experimental accounts to the hypothetical balances in the experimental accounts had the participant not made any withdrawals. The 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.