



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

Citation

Beshears, John, James J. Choi, Christopher Harris, David Laibson, Brigitte C. Madrian, and Jung Sakong. "Which Early Withdrawal Penalty Attracts the Most Deposits to a Commitment Savings Account?" Art. 104144. Journal of Public Economics 183 (March 2020).

Published version

<https://doi.org/10.1016/j.jpubeco.2020.104144>

Link

<https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37364485>

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles (OAP), as set forth at

<https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFIOWI2M2EwYTg>

Accessibility

<https://accessibility.huit.harvard.edu/digital-accessibility-policy>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#)

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, Brigham Young University and NBER
Jung Sakong, Federal Reserve Bank of Chicago

November 22, 2019

Abstract: Previous research has shown that some people voluntarily use commitment contracts that restrict their own choice sets. We study how people divide money between two accounts: a liquid account that permits unrestricted withdrawals and a commitment account that is randomly assigned in a between-subject design to have either a 10% early withdrawal penalty, or a 20% early withdrawal penalty, or not to allow early withdrawals at all (i.e., an infinite penalty). When the liquid account and the commitment account pay the same interest rate, higher early-withdrawal penalties attract *more* commitment account deposits. This pattern is predicted by the hypothesis that some participants are partially- or fully-sophisticated present-biased agents. Such agents perceive that higher penalties generate greater scope for commitment by disincentivizing (penalized) early withdrawals. The experiment also shows that when the commitment account pays a *higher* interest rate than the liquid account, the positive empirical slope relating penalties and commitment deposits is flattened, suggesting that naïve present-biased agents or agents with standard exponential discounting are also in our sample. Across all of our experimental treatments, higher early withdrawal penalties on the commitment account sometimes increase and never reduce allocations to the commitment account.

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 greatly benefited from the comments of our editor, Johannes Spinnewijn and three anonymous referees, as well as George-Marios Angeletos, B. Douglas Bernheim, Bruce Carlin, Stefano DellaVigna, Luigi Guiso, Ulrike Malmendier, Sendhil Mullainathan, Christopher Parsons, and Ivan Werning. We deeply grateful to Christopher Clayton, Harry Kosowsky, Omeed Maghzi, Peter Maxted, Kartik Vira, and Sean (Yixiang) Wang for exceptional 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.

* Corresponding author: Harvard Business School, Soldiers Field, Boston, MA 02163, jbeshears@hbs.edu

In 2018, U.S. households held \$16.3 trillion in employer-sponsored defined contribution savings plans and IRAs (Investment Company Institute, 2019). 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 and other tax inducements to let balances accumulate in these accounts, 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 and \$0.21 of early withdrawals for which the penalty is waived (Argento, Bryant, and Sabelhaus, 2015).³ 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.

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).

³ See footnote 1 for instances in which the penalty is waived.

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 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. Moreover, if they were to allocate balances to a commitment account due only to an experimenter demand effect, there is no reason to anticipate that they would have higher commitment account allocations in treatment arms with higher early withdrawal penalties (in our between-subject design). Economic agents with exponential discounting or naïve agents with present-bias 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 become necessary.

Partially or fully sophisticated present-biased agents (agents who are at least somewhat aware of their self-control problems), perceive both costs and benefits of illiquidity. Not having access to assets when a legitimate liquidity need might occur is a cost of illiquidity. On the other hand, stronger commitment is afforded by higher early withdrawal penalties (Laibson, 1997). Indeed, in the absence of uncertainty, or under particular regularity conditions that we provide in the on-line appendix, sophisticated present-biased agents will allocate more assets to illiquid accounts the higher the early withdrawal penalties associated with those accounts. When there is no uncertainty (e.g., no taste shocks) the logic for this effect is easy to summarize. Sophisticated agents will not allocate funds to accounts where they expect to withdraw those funds and pay a penalty. So higher penalties enable sophisticated agents to more intensively use illiquid accounts. The higher the penalty, the more wealth early selves can store in the illiquid asset without generating gratuitous penalties from early withdrawals. The higher penalty is protective with respect to early withdrawals. It turns out that this logic for the case with no uncertainty extends to a wide range of leading cases with stochastic taste shocks (see appendices B and C).⁴

Thus, our empirically observed *increase* in commitment account deposits in treatment

⁴ 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.

arms that have higher early withdrawal penalties suggests the presence of fully or partially sophisticated present-biased agents.⁵ Importantly, in this analysis we identify the presence of sophistication by the *slope* of take-up with respect to commitment penalties, not just the level of take-up. While the level of take-up might partially reflect experimenter demand effects or participant indifference, the change in take-up as the commitment penalty grows suggests that increased contributions are driven by the increase in commitment penalties themselves.

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 is consistent with the U.S. adult population containing 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 any sign, including the roughly 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 (Giné, 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

⁵ Our results are also consistent with models of costly self-control (Gul and Pesendorfer, 2001), which imply demand for commitment among *time-consistent* agents. For experimental support for these models, see Sadoff, Samek, and Sprenger (2015) and Toussaert (2018).

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.⁶ While Ashraf et. al (2006) relied on participants to make future deposits, Brune Giné, Goldberg, and Yang (2016) offered Malawian tobacco crop farmers the opportunity to allocate their *existing* harvest proceeds into both a liquid savings account and a commitment savings account. They find that participants offered both accounts saved more than either control group participants or participants offered only the liquid savings account. 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.⁷ 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 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 is consistent with the presence of fully or partially sophisticated present-biased individuals who recognize the commitment benefits of higher early withdrawal penalties, the data also points to heterogeneity in sophistication/naiveté. Our results therefore

⁶ Kast, Meier, and Pomeranz (2018) also study take-up of commitment savings accounts and finds similar results.

⁷ 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.

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, while Amador, Werning, and Angeletos (2006), Galperti (2015), Beshears et al. (2019), 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 and discusses policy implications.

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.⁸

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

⁸ The following paragraph from the RAND website contains information on how the ALP forms its sample: "ALP members have been recruited from multiple sources over the years. Many ALP members were recruited from other completed surveys. The original ALP cohort, for example, was initially recruited for a RAND-University of Michigan collaboration on the Health and Retirement Survey. Since then, ALP members have been recruited from several other surveys and directly for the panel using multiple modes (in-person/face-to-face, telephone, and mail) and probability-based sampling methods, including address-based samples and telephone (random-digit dial) samples." - <https://www.rand.org/research/data/alp/panel/recruitment.html>

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.

The Harvard University Institutional Review Board approved both experiments, and informed consent was obtained from all participants in both experiments.

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-subject 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 instead used a between-subjects experimental design to make the decision task simple for participants and to avoid the potential experimenter demand effects associated with a within-subject design. Thus, each participant saw only one version of the commitment account.

The illiquidity of the commitment account varied across conditions. In all of these conditions, early withdrawals from the commitment account are defined as withdrawals

requested prior to a commitment date chosen (and permanently fixed) by the participant at the beginning of the experiment. Withdrawals from the commitment account made before this commitment date were penalized in different ways in the treatment arms. Early withdrawals were subject to a 10% penalty, a 20% penalty, or disallowed altogether. We asked participants to choose their own commitment dates to allow for heterogeneity in the horizons over which individuals wished to generate spending. 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).⁹ No version of the commitment account permitted withdrawals during the first week of the experiment. (For balance, the liquid account also did not permit 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-

⁹ We are cautious in generalizing our results due to important differences between our experiment and real-world 401(k) plans. First, our interest rates are much higher than market interest rates and our experimental endowments are small compared to actual 401(k) balances. Second, our experiment studies windfalls and not “earned” income, which may have different mental frames. Third, our experiment and the associated theoretical framework (see appendix B) require individuals to allocate a portion of wealth from a given endowment, whereas actual 401(k) plans require individuals to make regular deposits at each pay cycle. In a typical 401(k) setting, however, individuals set up automatic contributions for their future selves (rather than manually making each 401(k) contribution), and since individuals are unlikely to cancel their contributions (due to inertia/switching costs), the initial 401(k) allocation decision may serve as a form of partial commitment, generating some limited similarity with the once-and-for-all allocation decision in our experiment. On the other hand, 401(k) contributions that result from a default option (such as 401(k) contributions induced by automatic enrollment) contrast with the “active choice” allocation decision in our experiment.

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.¹⁰ 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

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

¹⁰ 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.

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. We report results from the exit questionnaire in Appendix Table A8.

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.¹¹ Table 3 shows the mean fraction allocated to the commitment account by experimental condition. We have three main results.¹²

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

¹¹ 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.

¹² 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.

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%. 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. Appendix Figure A12 additionally shows the distribution of days until commitment date by treatment arm. 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.¹³

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 taste shock distributions). Rising allocations to the commitment account is the pattern we *empirically* observe in the arms of the study in which the liquid account and the commitment account pay the *same* interest rate (i.e., 22%). The weaker relationship between allocations to the commitment account and commitment account illiquidity when the commitment account pays a *higher* interest rate than the liquid account (23% for the commitment account vs. 22% for the liquid

¹³ 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.

account) is theoretically predicted if there are also agents with standard exponential discounting and/or naïve present-biased agents among our experimental participants. When the commitment account has an interest rate premium, it attracts some deposits from these two groups.¹⁴

However, since they have no desire for commitment, their commitment account allocations decrease as the account becomes more illiquid, partially offsetting the rising allocations to the commitment account by sophisticated present-biased agents. This offset affect implies that the *slope* of allocations with respect to rising illiquidity is predicted to be lower in the arms of the study in which the commitment account has a 23% rate of interest than it is in the arms of the study in which the commitment account has a 22% rate of interest. (Recall that the liquid account has a 22% rate of interest in all arms of the study.) When the commitment account pays the same interest rate as the liquid account (i.e., the 22% arm for commitment account), the model predicts that both agents with standard exponential discounting and agents that have naïve present-bias will allocate no money to the commitment account regardless of its strictness. Therefore, the theoretically predicted relationship between rising withdrawal penalties and rising commitment account balances is driven by the sophisticated present-biased agents in the arms of the study in which the commitment account has the same interest rate as the liquid account.

We linked the data from our experiment with other participant data available from the RAND American Life Panel and 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. Appendix Table A7 shows a sample of these correlations.

B. Withdrawals

What happens to account balances after the initial allocation? For each participant and each 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

¹⁴ 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.

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.¹⁵ In order to facilitate the relevant comparisons, we present subsets of the seven conditions in each of the three graphs in Appendix Figure A6.

In all conditions, most of the experimental endowment 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 appears to show that withdrawals take place earlier in the experiment in the treatment arms in which 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 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.¹⁶ 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% 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.¹⁷

¹⁵ 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.

¹⁶ We display various withdrawal statistics in Appendix Table A4. Appendix Table A6 also shows statistics related to incurred penalties.

¹⁷ 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 we include the analysis because it allows us to examine withdrawal decisions around the date that a participant deems most relevant for commitment. We find that holding fixed the commitment account interest rate, participants who were not allowed to withdraw early have the highest balance

The net effect of commitment account illiquidity on balance ratios is complicated by the competing channels through which illiquidity may affect withdrawal behavior. On the one hand, commitment account illiquidity *positively* impacts balance ratios because individuals in the more illiquid treatment groups allocate more to their commitment accounts. On the other hand, commitment account illiquidity *negatively* impacts balance ratios because individuals in the more illiquid treatments set earlier commitment dates. Additionally, incurred penalties may lower balance ratios for individuals in the 10% or 20% early withdrawal penalty conditions as compared to individuals that cannot incur penalties in the no early withdrawals condition.¹⁸ These competing effects generate muddy predictions and lower our power to detect differences. In addition, (continuous) withdrawal decisions depend on many realizations that occur during the one-year duration of the experiment (e.g., liquidity shocks and taste shocks) while the ex-ante commitment decision depends on expectations about these events. Accordingly, withdrawal decisions are noisier than ex-ante commitment decisions, further challenging our power to make inferences about withdrawal effects by treatment arm.

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

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 conduct a similar analysis that adjusts for the fact that the mean commitment date differs across arms (see Appendix Discussion A1). While we find suggestive evidence that stronger commitment raises balance ratios, we again find no statistically significant differences between the averages.

¹⁸ It is possible that the savings goals set during the initial allocation decision impact later withdrawal behavior. If either the goals themselves or the withdrawal behavior originating from the goals differ by experimental condition, we might expect average balance ratios to differ as well.

(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 partially capture 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;¹⁹ 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,²⁰ 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

¹⁹ 401(k) hardship withdrawals differ from our safety-valve treatment in three key ways. First, some hardship withdrawals are still penalized with the 10% penalty. Second, the financial circumstance necessitating a hardship withdrawal must correspond with an IRS-listed financial hardship; in our safety-valve treatment, we did not specify qualifying financial hardships. Third, until recently employers had to verify that the requesting employee was indeed experiencing a financial hardship. An IRS memo distributed in 2017, however, changed hardship withdrawal rules to allow employers to offer self-substantiation for financial hardships, along the same lines as our safety-valve treatment. The Bipartisan Budget Act of 2018 additionally repealed the 6-month suspension of elective deferrals following the hardship withdrawal and removed the mandate that required individuals to take out a 401(k) loan prior to a hardship withdrawal. See <https://www.irs.gov/retirement-plans/retirement-plans-faqs-regarding-hardship-distributions> and <https://www.irs.gov/pub/foia/ig/spder/tege-04-0217-0008.pdf> for more information.

²⁰ Like the first experiment, the second experiment permitted withdrawals no sooner than one week after the initial allocation decision.

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.²¹ 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 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

²¹ 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.

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.²²

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. Appendix Figure A13 shows the distribution of days until 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 ($p = 0.137$).²³ 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 A10 shows the balance ratios over time for the four experimental conditions, breaking apart participants by whether they received their endowments allocated according to

²² 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.

²³ These two p -values are 0.101 and 0.099, respectively, when we control for participant characteristics.

their choices or 100% in the liquid account.²⁴ 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. Appendix Table A5 shows additional withdrawal statistics.

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

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

²⁴ 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.

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 is consistent with the presence of some partially- or fully-sophisticated present-biased agents. However, when the commitment account pays a higher interest rate than the liquid account, we find that this positive relationship is flattened: commitment account allocations do not robustly rise with the commitment account's degree of illiquidity. This flattening is consistent with the hypothesis that 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 (sophisticated, present-biased) 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 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—might not object to or even welcome increasing the illiquidity of retirement accounts. 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 and naïve present-biased individuals, as well as individuals with no present bias at all.

The results from the experiments reported in this paper raise the possibility that voluntary commitment accounts with modest financial incentives could improve the lifecycle welfare of both sophisticated agents (who understand the benefits of the penalties/illiquidity) and naïve agents (who invest in those commitment accounts for the excess return, despite, and not because of, the illiquidity). Our empirical results suggest that many households might be tolerant of highly illiquid retirement savings accounts if those accounts had a modest sweetener (e.g., a higher return than alternative liquid investments). Across all of our experimental treatments, higher early withdrawal penalties on the commitment account sometimes increase and never reduce allocations to the commitment account.

Highly illiquid accounts are socially optimal in economies with populations that have heterogeneous levels of present bias (e.g., see Moser and Olea de Souza e Silva, 2017; and Beshears et al., 2019). To a first approximation, socially optimal illiquidity is obtained when early withdrawal penalties are equal to the degree of present bias (in a two-period problem): i.e., the early withdrawal penalty should equal $1 - \beta$, where β is the present bias parameter. To see why, note that the planner would like equilibrium allocations to be characterized by the classical Euler Equation:

$$u'(c_t) = R\delta u'(c_{t+1}),$$

where u is a stationary utility function, c is consumption (with a time subscript), R is the gross rate of return, and δ is the exponential discount factor. If an agent has present bias, and the planner introduces an early withdrawal penalty, p , then the agent's actual (two-period²⁵) Euler Equation will be

$$u'(c_t) = \beta(1 + p)R\delta u'(c_{t+1}).$$

The planner's socially optimal intertemporal consumption allocation is obtained if $p \approx 1 - \beta$.²⁶

In populations with *heterogeneous* present bias where present-bias screening is either difficult because agents try to pool or challenging for political reasons (e.g., the government needs to treat everyone equally), society's retirement savings regime should be disproportionately targeted at the households with relatively extreme levels of present bias (i.e., those with lower values of β). These are the households at most risk of radically deviating from their optimal consumption path.²⁷ Accordingly, high penalties in universal retirement accounts will be (second-best) socially optimal. Despite numerous significant reservations about external validity, the results of our experiments hold out the possibility that long-run savings accounts with large early withdrawal penalties (or even complete illiquidity, as is the norm in social security systems or defined benefit pension systems) may be broadly popular, particularly if the commitment accounts are sugar-coated so they also appeal to agents who are naïve.

²⁵ In a problem with an arbitrary horizon the Euler Equation is characterized in Harris and Laibson (2001).

²⁶ The relationship $p = 1 - \beta$, is exact if the penalty is paid out of withdrawals, so that the Euler Equation is $(1 - p)u'(c_t) = \beta R\delta u'(c_{t+1})$.

²⁷ See Beshears et al (2019) for details of this argument.

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 Philippines." *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, 2019. "Optimal Illiquidity." Working paper.
- Brune, Lasse, Xavier Giné, 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.
- Giné, 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.
- Gul, Faruk, and Wolfgang Pesendorfer, 2001. "Temptation and Self-Control." *Econometrica* 69.6, pp. 1403-1435.

- Harris, Christopher, and David Laibson. "Dynamic choices of hyperbolic consumers." *Econometrica* 69, no. 4 (2001): 935-957.
- 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.
- Toussaert, Séverine, 2018. "Eliciting temptation and self-control through menu choices: a lab experiment." *Econometrica* 86.5, pp. 859-889.

Table 1. Participant Characteristics

Demographic characteristics for participants in the first experiment ($n = 495$) and the second experiment ($n = 550$). We additionally include two columns with US statistics from the CPS (among individuals 18+). Experiment 1 took place in 2010 and experiment 2 took place in 2011.

	Expt. 1	2010 CPS	Expt. 2	2011 CPS
Percent male	43%	48%	43%	48%
Age				
≤ 25	8%	15%	8%	15%
26-35	17%	18%	19%	17%
36-45	21%	18%	18%	17%
46-55	22%	19%	22%	19%
56-65	16%	15%	15%	15%
≥ 66	16%	16%	17%	16%
Education				
No high school diploma	3%	14%	5%	13%
High school graduate	32%	31%	29%	30%
Some college	20%	19%	23%	20%
Associate's degree	7%	9%	12%	9%
Bachelor's degree	24%	18%	19%	18%
Graduate degree	13%	9%	12%	10%
Annual Household Income				
< \$15,000	6%	9%	9%	10%
\$15,000 - \$34,999	19%	20%	20%	20%
\$35,000 - \$49,999	16%	14%	16%	13%
\$50,000 - \$74,999	27%	19%	22%	19%
\$75,000 - \$99,999	15%	13%	16%	13%
≥ \$100,000	17%	25%	17%	25%
Marital Status				
Married	68%	54%	66%	54%
Separated/divorced	11%	13%	14%	13%
Widowed	5%	6%	5%	6%
Never married	16%	27%	15%	27%
Race				
White/Caucasian	80%	81%	81%	80%
Black/African American	8%	12%	10%	12%
Amer. Indian or Alaskan Native	1%	1%	1%	1%
Asian or Pacific Islander	4%	5%	2%	5%
Other	6%	2%	5%	2%

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. Importantly, the interest rate on the *liquid* account is 22% percent in all experimental conditions.

Withdrawal restrictions on commitment account prior to commitment date	Commitment account interest rate			<i>p</i> -value of equality of means	
	21%	22%	23%	21% vs. 22%	22% vs. 23%
10% early withdrawal penalty	27.6 (2.8)	38.9 (3.4)	58.2 (3.4)	0.011	0.000
20% early withdrawal penalty	--	44.8 (3.4)	61.1 (3.4)	--	0.001
No early withdrawals	--	56.0 (4.1)	59.9 (3.6)	--	0.469
<i>p</i> -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 (12.0)	209.0 (13.4)	227.6 (12.3)	0.165	0.306
20% early withdrawal penalty	--	207.4 (12.5)	202.1 (13.7)	--	0.775
No early withdrawals	--	214.3 (14.1)	186.0 (12.6)	--	0.136
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 (7.3)	81.8 (9.1)	129.6 (10.6)	0.136	0.001
20% early withdrawal penalty	--	100.5 (10.9)	127.0 (12.3)	--	0.108
No early withdrawals	--	131.8 (13.9)	117.8 (11.2)	--	0.436
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).

Withdrawal restrictions on commitment account prior to commitment date	Endowment allocation		Total
	According to participant's choice	All in liquid account	
Safety valve (withdrawals only in financial emergencies)	85	65	150
10% early withdrawal penalty	54	46	100
No early withdrawals	60	90	150
Two commitment accounts: 10% early withdrawal penalty and no early withdrawals	70	80	150

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 (0.033)	-0.004 (0.047)	0.002 (0.059)	0.022 (0.066)	-0.027 (0.071)
10% early withdrawal penalty	0.120* (0.060)	0.121 (0.071)	0.156 (0.082)	0.197* (0.087)	0.143 (0.090)
No early withdrawals	0.070* (0.034)	0.149** (0.047)	0.127* (0.057)	0.092 (0.070)	0.114 (0.073)
Two commitment accounts	-0.038 (0.031)	0.029 (0.046)	0.026 (0.053)	0.035 (0.057)	0.064 (0.061)
Combined	0.044* (0.019)	0.069** (0.026)	0.069* (0.031)	0.078* (0.034)	0.067 (0.036)
Combined (excluding safety valve)	0.039 (0.023)	0.094** (0.031)	0.093** (0.036)	0.097** (0.040)	0.103* (0.042)

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

Online Appendix A to “Which Early Withdrawal Penalty Attracts the Most Deposits to a Commitment Savings Account?” Supplementary Tables, Figures, and Discussions

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

November 22, 2019

Appendix Table A1. Percent of Decisions Allocating Strictly Positive Amount to Commitment Account: Experiment 1
Appendix Table A2. Regression Analysis of Percent of Endowment Allocated to Commitment Account and Dollar-Weighted Days to Commitment Date: Experiment 1
Appendix Table A3. Percent of Participants Allocating Strictly Positive Amount to Commitment Account: Experiment 2
Appendix Table A4. Withdrawal Statistics: Experiment 1
Appendix Table A5. Withdrawal Statistics: Experiment 2
Appendix Table A6. Incurred Penalties: Experiment 1
Appendix Table A7. Regression Analysis of Percentage Allocated to the Commitment Account on Various ALP Survey Variables
Appendix Table A8. Exit Questionnaire: Experiment 1
Appendix Figure A1. Description of the Liquid Account
Appendix Figure A2. Description of the 22% Interest Rate, 10% Early Withdrawal Penalty Commitment Account
Appendix Figure A3. Example Allocation Page
Appendix Figure A4. Sample Weekly Email to Participant
Appendix Figure A5. Withdrawal Interface
Appendix Figure A6. Balance Ratios by Experimental Condition: Experiment 1
Appendix Figure A7. Balance Ratios by Experimental Condition: Experiment 1
Appendix Figure A8. Description of Two Commitment Accounts Offered Simultaneously

Appendix Figure A9. Description of the Safety Valve Commitment Account
Appendix Figure A10. Withdrawal Patterns for Own versus All Liquid Allocation: Experiment 2
Appendix Figure A11. Withdrawal Patterns for Own versus All Liquid Allocation: Experiment 2
Appendix Figure A12. Distribution of Days to Commitment Date: Experiment 1
Appendix Figure A13. Distribution of Days to Commitment Date: Experiment 2
Appendix Discussion A1. Balance Ratios Adjusted for Mean Commitment Dates Across Arms

Appendix Table A1. Percent of Decisions Allocating Strictly Positive Amount to Commitment Account: Experiment 1

For each experimental condition, this table reports the percent of decisions that allocate a strictly positive amount 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 comparing pairs of conditions, as indicated.

Withdrawal restrictions on commitment account prior to commitment date	Commitment account interest rate			p -value for test of equality	
	21%	22%	23%	21% vs. 22%	22% vs. 23%
10% early withdrawal penalty	0.681 (0.053)	0.722 (0.050)	0.902 (0.033)	0.566	0.003
20% early withdrawal penalty	--	0.789 (0.043)	0.926 (0.029)	--	0.008
No early withdrawals	--	0.823 (0.046)	0.892 (0.035)	--	0.232
p -value for test of equality					
10% penalty vs. 20% penalty	--	0.310	0.570		
10% penalty vs. no early w/d	--	0.142	0.841		
20% penalty vs. no early w/d	--	0.590	0.445		

Appendix Table A2. Regression Analysis of Percent of Endowment Allocated to Commitment Account and Dollar-Weighted Days to Commitment Date: Experiment 1

This table reports the results of ordinary least squares regressions that use the sample of all allocation decisions in the first experiment. There are three observations for every participant: one observation for each possible endowment amount. In the first column, the outcome variable is the percent of endowment allocated to the commitment account. In the second column, the outcome variable is the 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. The explanatory variables are indicator variables for different interest rates, indicator variables for different withdrawal restrictions on the commitment account prior to the commitment date, and the interactions of those indicator variables. The omitted category is the condition featuring a 22% interest rate and a 10% early withdrawal penalty for the commitment account. Standard errors clustered at the participant level are in parentheses.

Explanatory variable	% of endowment allocated to commitment account	Dollar-weighted days to commitment date
21% interest rate	-11.3* (4.4)	-17.4 (11.6)
23% interest rate	19.3** (4.8)	47.8** (13.9)
20% early withdrawal penalty	5.9 (4.8)	18.7 (14.1)
No early withdrawals	17.1** (5.3)	50.0** (16.5)
23% interest rate × 20% early withdrawal penalty	-2.9 (6.8)	-21.3 (21.5)
23% interest rate × no early withdrawals	-15.4* (7.2)	-61.7** (22.6)
Constant (22% interest rate, 10% early withdrawal penalty is omitted category)	38.9** (3.4)	81.8** (9.0)
R^2	0.128	0.060
N	1,485	1,485

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

**Appendix Table A3. Percent of Participants Allocating Strictly Positive Amount to
Commitment Account: Experiment 2**

For each experimental condition, this table reports the percent of participants allocating a strictly positive amount to a commitment account. For the condition offering two commitment accounts, the table also reports the percent of participants allocating a strictly positive amount to each individual commitment account. Standard errors are in parentheses.

Withdrawal restrictions on commitment account prior to commitment date	% of participants using commitment account
Safety valve (withdrawals only in financial emergencies)	75.3 (3.5)
10% early withdrawal penalty	83.0 (3.8)
No early withdrawals	90.7 (2.4)
Two commitment accounts: strictly positive allocation to either 10% early withdrawal penalty account or no early withdrawals account	80.7 (3.2)
10% early withdrawal penalty account	56.0 (4.1)
No early withdrawals account	75.3 (3.5)

Appendix Table A4. Withdrawal Statistics: Experiment 1

This table reports various withdrawal statistics for individuals in experiment 1. To note, withdrawal statistic comparisons by treatment condition are confounded by initial allocations into the commitment accounts.

What fraction of participants ever withdrew?

By treatment, we calculate the percentage of individuals that *ever* withdrew from the commitment account before the last day of the study conditional on allocating to the commitment account. We do not differentiate between pre and post commitment date withdrawals.

Withdrawal restrictions on commitment account prior to commitment date	Commitment account interest rate		
	21%	22%	23%
10% early withdrawal penalty	0.361	0.379	0.346
20% early withdrawal penalty		0.367	0.412
No early withdrawals		0.344	0.338

How many withdrawals did participants make?

By treatment, we calculate the mean number of withdrawals participants made before the last day of the study conditional on allocating to the commitment account. We do not differentiate between pre and post commitment date withdrawals.

Withdrawal restrictions on commitment account prior to commitment date	Commitment account interest rate		
	21%	22%	23%
10% early withdrawal penalty	0.639	0.621	0.628
20% early withdrawal penalty		0.633	0.706
No early withdrawals		0.531	0.544

How many dollars did participants earn?

For each individual in the study, we divide their end earnings (amount withdrawn + balance at the end) over their initial endowment, and we average this ratio across individuals in each treatment.

Withdrawal restrictions on commitment account prior to commitment date	Commitment account interest rate		
	21%	22%	23%
10% early withdrawal penalty	1.170	1.178	1.188
20% early withdrawal penalty		1.175	1.176
No early withdrawals		1.178	1.189

Appendix Table A5. Withdrawal Statistics: Experiment 2

This table reports various withdrawal statistics for individuals in experiment 2. To note, withdrawal statistic comparisons by treatment condition are confounded by initial allocations into the commitment accounts.

What fraction of participants ever withdrew?

By treatment, we calculate the percentage of individuals that *ever* withdrew from the commitment account before the last day of the study conditional on allocating to the commitment account. We do not differentiate between pre and post commitment date withdrawals.

Withdrawal restrictions on commitment account prior to commitment date	Endowment allocation	
	According to participant's choice	All in liquid account
Safety valve (withdrawals only in financial emergencies)	0.376	0.308
10% early withdrawal penalty	0.296	0.413
No early withdrawals	0.383	0.449
Two commitment accounts: 10% early withdrawal penalty and no early withdrawals	0.300	0.338

How many withdrawals did participants make?

By treatment, we calculate the mean number of withdrawals participants made before the last day of the study conditional on allocating to the commitment account. We do not differentiate between pre and post commitment date withdrawals.

Withdrawal restrictions on commitment account prior to commitment date	Endowment allocation	
	According to participant's choice	All in liquid account
Safety valve (withdrawals only in financial emergencies)	0.659	0.400
10% early withdrawal penalty	0.537	0.609
No early withdrawals	0.717	0.528
Two commitment accounts: 10% early withdrawal penalty and no early withdrawals	0.571	0.400

How many dollars did participants earn?

For each individual in the study, we divide their end amount (amount withdrawn + balance at the end) over their initial endowment, and we average this ratio across individuals in each treatment.

Withdrawal restrictions on commitment account prior to commitment date	Endowment allocation	
	According to participant's choice	All in liquid account
Safety valve (withdrawals only in financial emergencies)	1.096	1.096
10% early withdrawal penalty	1.096	1.082
No early withdrawals	1.099	1.086
Two commitment accounts: 10% early withdrawal penalty and no early withdrawals	1.098	1.095

Appendix Table A6. Incurred Penalties: Experiment 1
This table reports on incurred penalties for participants in experiment 1.

How many participants incurred penalties?

By treatment, conditional on allocating to the commitment account, we calculate the percentage of individuals that incurred a penalty (total number of individuals that incurred a penalty in parentheses).

Withdrawal restrictions on commitment account prior to commitment date	Commitment account interest rate		
	21%	22%	23%
10% early withdrawal penalty	0.106 (5)	0.045 (2)	0.057 (4)
20% early withdrawal penalty		0.033 (2)	0.049 (3)
No early withdrawals			

How do incurred penalties compare to initial endowments?

For each individual that incurred a penalty, we calculate their penalties divided by their initial endowments, and we average by treatment (total number of individuals that incurred a penalty in parentheses).

Withdrawal restrictions on commitment account prior to commitment date	Commitment account interest rate		
	21%	22%	23%
10% early withdrawal penalty	0.093 (5)	0.059 (2)	0.091 (4)
20% early withdrawal penalty		0.188 (2)	0.177 (3)
No early withdrawals			

Appendix Table A7. Regression Analysis of Percentage Allocated to the Commitment Account on Various ALP Survey Variables

This table reports a regression of the percentage allocated to the commitment account across all treatments on various dummy variables created from additional surveys administered by RAND. Below we display the regression results for select dummy variables. Though two of the correlations yield some level of statistical significance, neither of the variables are statistically significant after a Bonferroni correction.

Dummy Variable	Coefficient (T-statistic in parentheses)	Dummy Construction
Income in Quartile 2	0.018 (0.40)	= 1 if family income is greater than \$35,000 and less than \$60,000
Income in Quartile 3	0.041 (0.93)	= 1 if family income is greater than \$60,000 and less than \$100,000
Income in Quartile 4	0.045 (1.04)	= 1 if family income is greater than \$100,000
Net Wealth in Quartile 2	0.057 (1.32)	= 1 if net wealth is greater than \$2,500 and less than \$63,250
Net Wealth in Quartile 3	0.090* (2.03)	= 1 if net wealth is greater than \$63,250 and less than \$235,000
Net Wealth in Quartile 4	0.048 (1.15)	= 1 if net wealth is greater than \$235,000
Overweight	-0.001 (0.04)	= 1 if BMI is greater than 25
No Exercise	-0.028 (1.08)	= 1 if participants reported that they hardly ever or never engage in physical activity
Smokes Now	-0.060 (1.47)	= 1 if participants reported being a current cigarette smoker
Good at Math	0.036 (1.25)	= 1 if participants reported strongly agreeing or somewhat agreeing that they are good at math
Financial Confidence	-0.014 (0.34)	= 1 if participants reported strongly agreeing or somewhat agreeing that they are financially confident
Financial Assessment	-0.083** (2.81)	= 1 if participants reported strongly agreeing or somewhat agreeing that they are able to assess financial services
Emergency Savings	0.020 (0.64)	= 1 if participants reported strongly agreeing or somewhat agreeing that they could come up with \$2,000 if an unexpected need arose
Present Bias	0.056 (1.32)	= 1 if individuals are present biased ¹

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

¹ We calculated whether participants are present biased based on survey questions that measured preferences over financial prizes at different time periods. Specifically, participants answered two questions, one that asked whether participants would prefer \$1,000 today or \$1,250 next year, and another that asked individuals whether they would prefer \$1,000 next year or \$1,250 in two years. The survey also asked these same questions for preferences over \$1,000 and \$1,650. If individuals answered that they preferred the \$1,000 in the former question but the larger dollar amount in the latter question (for either the \$1,250 or the \$1,650 case), we recorded that the individuals were present-biased.

Appendix Table A8. Exit Questionnaire: Experiment 1

This table reports the results from exit questionnaire question 1. The question asked participants whether they would have in hindsight changed their allocation decision. Below we report the results by treatment condition.

Treatment Condition	More to liquid account	More to commitment account	Same allocation as before	Missing survey response	Number of observations
21% commitment account interest rate – 10% withdrawal penalty	6	7	48	11	72
22% commitment account interest rate – 10% withdrawal penalty	4	9	40	13	66
23% commitment account interest rate – 10% withdrawal penalty	5	10	39	24	78
22% commitment account interest rate – 20% withdrawal penalty	5	14	51	9	79
23% commitment account interest rate – 20% withdrawal penalty	3	10	37	18	68
22% commitment account interest rate – no withdrawals	3	4	38	19	64
22% commitment account interest rate – no withdrawals	4	7	40	17	68
Total	30^a	61^a	293	111	495

^a A hypothesis test on the equality of these two proportions yields a p-value of 0.0012.

Appendix Figure A1. Description of the Liquid Account

The **Freedom Account** is designed to let you access your money whenever you want. You can withdraw money from this account any time over the next year, starting one week from today.

Money in the Freedom Account will grow at an interest rate of 22% per year until you withdraw it. When you withdraw money from the Freedom 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.

<<Back Next>>



Appendix Figure A2. 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>>](#)



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>>](#)



Appendix Figure A3. 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.

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





Appendix Figure A4. 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](tel:866.591.2909)

Thanks!

www.rand.org/alp

Appendix Figure A5. Withdrawal Interface

The screenshot shows a web browser window titled "RAND | American Life Panel - Mozilla Firefox". The address bar displays "rand.org" and the URL "https://mmic.rand.org/panel/". The page content is titled "My accounts" with a link "back to my personal ALP page". Below this, a message states: "Please enter an amount you would like to withdraw in the appropriate box and click 'withdraw'".


The interface lists two accounts:

- Freedom Account**
remaining balance: \$100.70
Withdrawal input: \$
* withdraw
- Goal Account**
remaining balance: \$105.47
goal date: July 20th, 2010
Withdrawal input: \$
* withdraw

A footnote at the bottom states: "* If you make a withdrawal, a check will be mailed to you within the next three business days."


Appendix Figure A6. 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.



Appendix Figure A7. 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 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. 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.



Appendix Figure A8. 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>>](#)



Appendix Figure A9. Description of the Safety Valve Commitment Account, Withdrawal Screen for 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.




My accounts

[back to my personal ALP page](#)

You had told us that you wanted to save for:

a black dog

You requested an emergency withdrawal of \$25.00 from your Goal Account, but your goal date (April 25th, 2011) has not passed yet.


If you are experiencing a financial emergency, you can withdraw your money.

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'.

I attest that I have a financial emergency


Appendix Figure A10. 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.




Appendix Figure A11. 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 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 before the participant's earliest commitment date in the case of participants who had more than one), three days after the participant's commitment date (three days after the participant's latest commitment date in the case of participants who had more than one), and three days before remaining account balances were automatically disbursed to the participant. For participants who did not allocate funds to a commitment account, we use the actual balance and the hypothetical total balance on the initial deposit date when calculating the withdrawal measure for three days before the commitment date, and we use the actual balance and the hypothetical total balance three days after the initial deposit date when calculating the withdrawal measure for three days after the commitment date.




Appendix Figure A12. Distribution of Days to Commitment Date: Experiment 1

For each experimental condition, we calculate the cumulative distribution function for days to commitment date. In other words, for each day until commitment date, we plot the fraction of participants in that treatment that set days to commitment date equal or prior to that number of days.



Appendix Figure A13. Distribution of Days to Commitment Date: Experiment 2

For each experimental condition, we calculate the cumulative distribution function for days to commitment date. In other words, for each day until commitment date, we plot the fraction of participants in that treatment that set days to commitment date equal or prior to that number of days.



Appendix Discussion A1. Balance Ratios Adjusted for Mean Commitment Dates Across Arms

When calculating balance ratios, we can 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.

**Online Appendix B to “Which Early Withdrawal Penalty Attracts the
Most Deposits to a Commitment Savings Account?”**

**A Theory of the Commitment-Account Allocations of Sophisticated
Present-Biased Agents**

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

November 25, 2019

1. INTRODUCTION

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

$$D(\tau) = \begin{cases} 1 & \text{if } \tau = 0 \\ \beta & \text{if } \tau \geq 1 \end{cases},$$

where $0 < \beta < 1$ (Phelps and Pollak, 1968; Laibson, 1997).¹ 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 commitment.

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 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.

¹The analysis that follows would be nearly identical if we were to use the more general quasi-hyperbolic discount function given 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 $\delta = 1$.

ONLINE APPENDIX B TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE
MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

We briefly describe the key properties of the model below. Sections 1-16 of Online Appendix C provide a complete exposition and analysis of the model.

2. 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.

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 $\theta \in \Theta = [\underline{\theta}, \bar{\theta}]$ 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 3 below describes the set of commitment mechanisms available to self 0, and Section 4 sets out our assumptions on the distribution of θ . Note that the three-period structure maps directly onto our experimental setup, with period 0 corresponding to the initial allocation decision, period 1 corresponding to the time between the allocation decision and the commitment date (which was tailored in the experiment by each participant according to the time horizon over which the temptation to overspend was relevant), and period 2 corresponding to the time after the commitment date.

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:

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

ONLINE APPENDIX B TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE
MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

Here U_t is the utility function at time t . We assume that: $U_t : [0, \infty) \rightarrow [-\infty, \infty)$; $U'_t > 0$ and $U''_t < 0$ on $(0, \infty)$; and $U'_t(0+) = \infty$.²

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

3. 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 time 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. To map to our experimental design, y can be interpreted as the agent's total wealth if the participants integrate the experimental windfall with their other wealth, or it can be interpreted as only the windfall itself if the participants psychologically code the windfall as part of a separate mental account. The main implications of the model apply in both cases.

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 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 .³

²For 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 $U_t(0) = -\infty$.

³The precise statement of Constraint 2 runs as follows. 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) $c_2 - \tilde{c}_2 \leq (1 + \pi)(\tilde{c}_1 - c_1)$. In other words, self 1 can increase her consumption by any amount $\tilde{c}_1 - c_1$ between 0 and $\frac{1}{1+\pi} c_2$. Moreover, she can do this in such a way that the associated reduction $c_2 - \tilde{c}_2$ in consumption in period 2 is at most $\tilde{c}_1 - c_1$ plus the maximum penalty that can be placed on a withdrawal of $\tilde{c}_1 - c_1$, namely $\pi(\tilde{c}_1 - c_1)$. Notice that this penalty is paid out of period 2 consumption.

ONLINE APPENDIX B TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”




Figure B1: A budget set illustrating Constraints 1 and 2 of the model. This 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 $\pi = 0.1$).

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 .⁴ The only restriction on shape is that, starting from any consumption pair $(c_1, c_2) \in B$ such that $c_2 > 0$ and any $\Delta \in [0, \frac{c_2}{1+\pi}]$, self 1 must be able to transfer Δ units 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 the reduction resulting from the transfer Δ itself. However, this penalty will never be greater than $\pi \Delta$.⁵

A wide variety of budget sets satisfy Constraints 1 and 2. For example, the budget set shown in Figure B1 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 set at the encircled point is -1.3 . This is greater than -1.5 (the most negative slope that is permissible

⁴This is Constraint 1.

⁵This is Constraint 2.

ONLINE APPENDIX B TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”




Figure B2: 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.

when $\pi = 0.5$), but less than -1.1 (the most negative slope that is permissible when $\pi = 0.1$).⁶

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 B2 consists of just such a budget curve, together with all of the points of A that lie below or to the left of it.

⁶Notwithstanding 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 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.)

ONLINE APPENDIX B TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE
MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

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 .⁷ Then self 1 will face a two-part budget set.

4. 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 $(0, \infty)$.⁸

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

A3 Put $G(\theta) = (1 - \beta) \theta F'(\theta) + F(\theta)$. Then there exists $\theta_M \in [\underline{\theta}, \bar{\theta}]$ such that: (i) $G' \geq 0$ on $(0, \theta_M)$; and (ii) $G' \leq 0$ on (θ_M, ∞) .⁹

We now comment on these assumptions. A function $f : (0, \infty) \rightarrow \mathbb{R}$ is of bounded variation if and only if it is the difference of two bounded and non-decreasing functions $f_1, f_2 : (0, \infty) \rightarrow \mathbb{R}$. 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 $(0, \infty) \setminus [\underline{\theta}, \bar{\theta}]$. Notice that F' need not be continuous. In particular, it can jump up at $\underline{\theta}$ and down at $\bar{\theta}$. A3 means that G is first increasing and then decreasing. It implies that the support of F' is connected. It is preserved under truncation:

⁷By 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 $(1 + p) \Delta$.

⁸A 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 $\bar{\theta}$. Indeed, a good way of generating examples is to take a standard distribution and truncate it at suitable points $\underline{\theta}$ and $\bar{\theta}$. This procedure typically results in discontinuities in F' at $\underline{\theta}$ and $\bar{\theta}$.

⁹A3 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.

if a distribution function F satisfies A3, then so too does the distribution function obtained by truncating F at $\underline{\theta}$ and $\bar{\theta}$.

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 $\bar{\theta}$). More precisely, we have:

Remark 1. *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 $0 < \underline{\theta} < \bar{\theta} < \infty$, the distribution function F obtained by truncating D at $\underline{\theta}$ and $\bar{\theta}$ satisfies Assumptions A1-A3.¹⁰*

The four exceptions are the Beta, Cauchy, Log-Gamma and Pareto distributions. In the form in which it occurs in both Rice and Hogg, McKean, and Craig, the Beta 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.¹¹ 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.¹² Next, the Log-Gamma distribution occurs only in Hogg, McKean, and Craig. This distribution may or may not satisfy A3, depending on the parameters.¹³ Finally, Rice and Hogg, McKean and Craig each consider a single special case of the Pareto distribution. Both of these special cases satisfy

¹⁰Notice that 5 of these 14 distributions (namely the Burr, Chi-squared, F, Gamma and Weibull distributions) are unbounded at 0 for some parameter values. However, the truncated distributions all satisfy A1 because $\underline{\theta} > 0$.

¹¹The 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 $\zeta < 1$.

¹²The 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 $\frac{\mu}{\sigma}$ is large and positive.

¹³The 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 $\eta > 1 - \beta$.

A1-A3. However, in general, the Pareto distribution fails A3 for some choices of the parameter values.¹⁴ For additional discussion of these exceptional cases, see Section 17 of Online Appendix C.

5. 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 any withdrawals in period 1 but places no penalty on withdrawals in period 2.^{15,16} 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 2.¹⁷*

¹⁴For example, the density of the Pareto type II distribution is proportional to $(1 + \frac{x-\mu}{\sigma})^{-\zeta-1}$ 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 $\frac{\mu}{\sigma}$ is large and positive.

¹⁵There 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 $y - y_{\text{liquid}}$ 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 $(y_{\text{liquid}}, y - y_{\text{liquid}})$. On the other hand, if she places y_{liquid} in a fully liquid account and $y - y_{\text{liquid}}$ 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 $(y_{\text{liquid}}, y - y_{\text{liquid}})$. (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 $(y_{\text{liquid}}, y - y_{\text{liquid}})$. We shall not therefore distinguish between them in what follows.

¹⁶Ambrus and Egorov (2013) provide additional analysis of AWA’s model.

¹⁷As 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

ONLINE APPENDIX B TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

See Sections 1 through 10 of Online Appendix C.

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.

Moving on from Theorem 1, let us continue to suppose that $U_1 = U_2 = U$, and that U has constant relative risk aversion $\rho > 0$. But let us suppose now 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. Finally, let us denote the optimal allocations to these two accounts by y_{liquid} and y_{penalty} .

How will the allocation y_{penalty} to the partially illiquid account depend on p ? As with many questions in comparative statics, this question is easier to answer when y_{penalty} is unique. We therefore begin by introducing an additional assumption that, when taken in conjunction with our existing Assumptions A1-A3, ensures this:

A4 G is strictly increasing on $[\underline{\theta}, \theta_M)$.¹⁸

A4 strengthens Part (i) of A3 – which effectively requires that G is *weakly* increasing on $(0, \theta_M)$ – by requiring that G is *strictly* increasing on a part of this interval.¹⁹

Under Assumptions A1-A4, we obtain an explicit expression for the derivative of y_{penalty} with respect to p .²⁰ In the case where the maximum-penalty constraint is binding, in the sense that some high- θ types will choose to pay the penalty and consume out of the partially illiquid account, this expression can be decomposed into two contributions.²¹ The first of these is always positive. The second can be either positive or negative. In the case where the maximum-penalty constraint is slack, in

think rather of her as choosing an incentive-compatible consumption curve $c = (c_1, c_2) : [\underline{\theta}, \bar{\theta}] \rightarrow A$.

¹⁸More precisely, the right-continuous version of G is strictly increasing on $[\underline{\theta}, \theta_M)$.

¹⁹We do not need to strengthen Part (ii) of A3, which effectively requires that G is weakly decreasing on (θ_M, ∞) , e.g. by requiring that G is strictly decreasing on $(\theta_M, \underline{\theta}]$.

²⁰See Propositions 32, 40, 42 and 43 of Appendix C.

²¹See Propositions 32 and 40 of Appendix C.

the sense that even the highest- θ type will strictly prefer not to consume out of the partially illiquid account, this expression vanishes.²²

Hence, in order to find sufficient conditions under which y_{penalty} is non-decreasing in p , it suffices to find sufficient conditions under which the second contribution to the derivative of y_{penalty} with respect to p is non-negative. The required conditions depend on whether $\rho < 1$, $\rho = 1$ or $\rho > 1$.

Theorem 2. *Suppose that:*

1. *Assumption A4 is satisfied;*
2. *U has constant relative risk aversion $\rho < 1$;*
3. *$\theta_M = \underline{\theta}$.*

Then there exists $\pi_1 \in (0, \infty)$ such that y_{penalty} is strictly increasing on $(0, \pi_1]$ and constant on $[\pi_1, \infty)$.

In other words, if $\rho < 1$ and G is weakly decreasing on the whole of $(\underline{\theta}, \infty)$, then y_{penalty} is monotonic in p .²³

Theorem 3. *Suppose that:*

1. *Assumption A4 is satisfied;*
2. *U has constant relative risk aversion $\rho = 1$.*

Then there exists $\pi_1 \in (0, \infty)$ such that y_{penalty} is strictly increasing on $(0, \pi_1]$ and constant on $[\pi_1, \infty)$.

In other words, if $\rho = 1$ then, under no additional assumptions on G beyond A4 itself, y_{penalty} is monotonic in p .

Theorem 4. *Suppose that:*

²²See Propositions 42 and 43 of Appendix C.

²³Notice that, if G is weakly decreasing on the whole of $(\underline{\theta}, \infty)$, then G must be strictly increasing at $\underline{\theta}$. That is, we must have $\Delta G(\underline{\theta}) > 0$.

ONLINE APPENDIX B TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

1. *Assumption A4 is satisfied;*
2. *U has constant relative risk aversion $\rho > 1$;*
3. *$\theta_M = \bar{\theta}$.*

Then there exists $\pi_1 \in (0, \infty)$ such that y_{penalty} is strictly increasing on $(0, \pi_1]$ and constant on $[\pi_1, \infty)$.

In other words, if $\rho > 1$ and G is weakly increasing on the whole of $(0, \bar{\theta})$, then y_{penalty} is again monotonic in p .²⁴

Notice that the strategy of proof used to obtain Theorems 2 and 4 is quite extreme: there are two contributions to the derivative of y_{penalty} with respect to p , and we already know that the first of these is positive. Hence this contribution could easily outweigh the second contribution, even if the latter were negative. Nonetheless, we impose the extreme conditions $\theta_M = \underline{\theta}$ and $\theta_M = \bar{\theta}$ respectively to ensure that the second contribution weakly reinforces the first.²⁵ This suggests that the sufficient conditions for monotonicity contained in these theorems might be some way from being necessary. We have obtained limited confirmation for this suggestion: monotonicity does seem to hold in all of the simple examples that we have investigated numerically, and in most of which we do not have $\theta_M \in \{\underline{\theta}, \bar{\theta}\}$. However, we also have an analytical “counterexample” to each theorem. Hence, ultimately, it is a quantitative question whether monotonicity does or does not hold for any given calibration: the second contribution can certainly go the wrong way, and it can even outweigh the first contribution.

Theorem 3, by contrast, does not impose any extreme conditions on G : the fact that $\rho = 1$ ensures that the second contribution to the derivative of y_{penalty} with respect to p vanishes. We are therefore left with the first contribution, which is unambiguous.

²⁴Notice that, if G is weakly increasing on the whole of $(0, \bar{\theta})$, then G must be strictly decreasing at $\bar{\theta}$. That is, we must have $\Delta G(\bar{\theta}) < 0$.

²⁵These conditions are extreme in the sense that they put θ_M at the extremes of the support of F' .

Remark 2. *An interesting example is provided by the uniform distribution on $[a, b]$. For this distribution, we have*

$$G'(\theta) = \frac{2 - \beta}{b - a} > 0.$$

*Hence Theorems 3 and 4 apply, and we can be sure that the desired conclusion holds for all $\rho \geq 1$. On the other hand, if $\rho < 1$, then the second contribution to the derivative of y_{penalty} with respect to p is negative.*²⁶

Remark 3. *If we solve the differential equation $G'(\theta) = 0$, then we obtain (up to a multiplicative constant)*

$$F'(\theta) = \theta^{-\frac{2-\beta}{1-\beta}}.$$

This is a special case of the Pareto distribution. For this distribution, Theorems 2, 3 and 4 all apply, and we can be sure that the desired conclusion holds for all $\rho > 0$.

Finally, it is helpful to provide some intuition as to why y_{penalty} is weakly increasing in p when there is no uncertainty (i.e. θ is fixed). In equilibrium, self 0 uses the illiquid account to store all wealth that will be consumed in period 2. Moreover, self 0 will not store any wealth in the illiquid account that will end up being consumed in period 1, because such a strategy is strictly dominated by the alternative strategy that takes both the amount that is withdrawn from the illiquid account in period 1 and the amount that is paid in early-withdrawal penalty and reallocates that sum into the fully liquid account. Accordingly, self 0 will allocate resources to the illiquid account either (i) up to the point where self 0 has achieved its first-best optimum under commitment (which will be true with a high enough penalty),²⁷ or (ii) up to the point where self 1 is indifferent between consuming a marginal unit of consumption in period 1 or consuming $1 + p$ marginal units of consumption in period 2. In the first case, y_{penalty} will not change with further increases in p . In the second $c_1 = y_{\text{liquid}}$, $c_2 = y_{\text{penalty}}$, $y_{\text{liquid}} + y_{\text{penalty}} = y$ and

$$\theta U_1'(y_{\text{liquid}}) = (1 + p) \beta U_2'(y_{\text{penalty}}).$$

²⁶Our simulations suggest that, in the case of the uniform distribution, the desired conclusion holds even when $\rho < 1$.

²⁷Specifically, when θ is fixed, self 0 can achieve its first-best optimum if and only if $\beta \geq \frac{1}{1+p}$.

ONLINE APPENDIX B TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE
MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

For this case, strict concavity of the utility functions implies that y_{penalty} is strictly increasing in p . Intuitively, the higher the penalty, the more wealth self 0 can store in the illiquid asset without generating gratuitous penalties from early withdrawals.

Remark 4. *For more details on Theorems 2, 3 and 4, see Sections 10 through 14 of Online Appendix C.*

Remark 5. *Versions of Theorems 2, 3 and 4 all hold even in the absence of Assumption A4. See Section 15 of Online Appendix C.*

Moving on again, let us continue to suppose that $U_1 = U_2 = U$, that U has constant relative risk aversion $\rho > 0$ and that Assumption A4 is satisfied. Suppose further that self 0 must now 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} .

Theorem 5. *The liquid allocation y_{liquid} is unique and independent of p . By the same token, the total illiquid allocation $y_{\text{penalty}} + y_{\text{illiquid}}$ is unique and independent of p . Furthermore, self 0 weakly prefers the fully illiquid account to the partially illiquid account. Specifically, there exists $\pi_1 \in (0, \infty)$ such that:*

1. *For all $p \in (0, \pi_1)$, self 0 strictly prefers the fully illiquid account to the partially illiquid account. More precisely: self 0 places her total illiquid allocation $y_{\text{penalty}} + y_{\text{illiquid}}$ in the fully illiquid account; $y_{\text{penalty}} = 0$; and y_{illiquid} is unique and independent of p .*
2. *For all $p \in [\pi_1, \infty)$, 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 $y_{\text{penalty}} + y_{\text{illiquid}}$ is divided between the partially illiquid account and the fully illiquid account.*

ONLINE APPENDIX B TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

The logic behind Theorem 5 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 $p = \infty$ 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_{illiquid} , and consider two scenarios. In the first scenario, self 0 deposits y_{illiquid} in a fully illiquid account. In that case, there will be a θ_1 such that: (i) any self 1 of type $\theta \leq \theta_1$ will choose freely from the line segment joining the two points $(0, y)$ and $(y - y_{\text{illiquid}}, y_{\text{illiquid}})$; and (ii) any self of type $\theta \geq \theta_1$ will choose the point $(y - y_{\text{illiquid}}, y_{\text{illiquid}})$. In the second scenario, self 0 deposits y_{illiquid} in a partially illiquid account with penalty p . In that case, the behaviour of self 1 will be effectively unchanged from the first scenario if and only if $p \geq \pi_1$, where π_1 is the minimum penalty necessary to deter the $\bar{\theta}$ type of self 1 from increasing consumption above $y_{\text{liquid}} = y - y_{\text{illiquid}}$.

Hence, if $p \geq \pi_1$, 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 two 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 y_{penalty} 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 subjects allocated their endowments across three accounts: a liquid account, a partially illiquid account with a 10% penalty, and a fully illiquid

ONLINE APPENDIX B TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

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. The decision to allocate money to the partially illiquid account is consistent with the model. Theorem 5 predicts that subjects 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 5. 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 subjects 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, subjects 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.

Finally, it is important to emphasize that our theoretical analysis considers a sophisticated agent, who in period 0 fully anticipates the difference between the current self’s preferences and preferences as of period 1. There is evidence that many individuals in the population are only partially sophisticated—they understand that there is a divergence between current and future preferences but underappreciate the full

ONLINE APPENDIX B TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE
MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

extent of that divergence (John, 2018). From a descriptive perspective, our theoretical analysis of commitment account allocations also applies to the case of a partially sophisticated agent. However, the welfare implications may not apply. In particular, because a partially sophisticated agent makes commitment account allocation decisions in period 0 based on an incorrect forecast of consumption decisions in period 1, the period 0 decisions may not be optimal in the sense that welfare from the period 0 perspective may be improved by selecting different commitment account allocations. As the current analysis focuses on descriptive issues, we leave welfare analysis to other work (see Galperti, 2015; Beshears et al., 2017; and Moser and Olea de Souza e Silva, 2017).

ONLINE APPENDIX B TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE
MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

REFERENCES

- 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.
- Beshears, John, James J. Choi, Christopher Clayton, Christopher Harris, David Laibson, and Brigitte C. Madrian, 2017. “Optimal Illiquidity.” Working paper.
- Galperti, Simone, 2015. “Commitment, Flexibility, and Optimal Screening of Time Inconsistency.” *Econometrica* 83.4, pp. 1425-1465.
- Hogg, Robert V., Joseph W. McKean, and Allen T. Craig, 2005. *Introduction to Mathematical Statistics*. 6th Edition. Upper Saddle River: Prentice Hall.
- John, Anett, 2018. “When Commitment Fails: Evidence from a Field Experiment.” Working paper.
- 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.
- Moser, Christian and Pedro Olea de Souza e Silva, 2017. “Optimal Paternalistic Savings Policies.” Working paper.
- 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.
- Rice, John A., 1995. *Mathematical Statistics and Data Analysis*. 2nd Edition. Belmont: Duxbury Press.

**Online Appendix C to “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

August 29th 2013

This draft: November 26th 2019

1. INTRODUCTION

In this appendix we provide a complete analysis of the mechanism-design problem described in the main body of the paper.

2. PRELIMINARIES

2.1. Functions of Bounded Variation. We begin by discussing the concept of bounded variation. This concept will be used to formulate our assumptions on the distribution function F in the subsection immediately following this one, namely Section 2.2. More importantly, it plays an essential role in our proof of sufficiency in a much later section, namely Section 16.

The simplest definition of a function of bounded variation is probably that given in the main text: a function $f : (0, \infty) \rightarrow \mathbb{R}$ is of bounded variation iff it is the difference of two bounded and non-decreasing functions $f^+, f^- : (0, \infty) \rightarrow \mathbb{R}$. This definition forms the starting point for the definition that we shall use. However, it needs to be developed into a form that is more convenient for the Lagrangean analysis below.

The first step is to collect the functions of bounded variation into equivalence classes. Intuitively speaking, two functions of bounded variation are equivalent iff they differ only at their points of discontinuity. This step is analogous to the first

ONLINE APPENDIX C TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE
MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

step in defining spaces of Lebesgue integrable functions. (In that case, one collects the Lebesgue integrable functions into equivalence classes. Two Lebesgue integrable functions are equivalent if they differ only on a set of measure 0.)

The second step is to place a norm on the resulting equivalence classes in such a way that the limit of a sequence of equivalence classes is again a suitable equivalence classes. (This step is analogous to the second step in defining spaces of Lebesgue integrable functions.) The main idea here is to note that, since f^+ and f^- are non-decreasing, they are effectively the distribution functions of a pair of non-negative bounded measures df^+ and df^- . Of course, neither df^+ nor df^- is unique. But their difference $df = df^+ - df^-$ is. The main component of the norm is therefore the total variation $\|df\|_{TV}$ of df . The other idea is to note that, while $\|df\|_{TV}$ effectively controls the derivative of f , it does not control the level of f . The remaining component of the norm can therefore be taken to be $|f_R(1)|$, where $f_R(1)$ is the limit of f from the right at 1.

The best way of understanding how these ideas work is to note that we can easily reconstruct f from df and $f_R(1)$. For all $\theta \in (0, 1)$, we have

$$f_R(\theta) = f_R(1) - df((\theta, 1])$$

and

$$f_L(\theta) = f_R(1) - df([\theta, 1]) ,$$

where $f_R(\theta)$ and $f_L(\theta)$ are the limits of f from the right and left at θ . And for all $\theta \in (1, \infty)$, we have

$$f_R(\theta) = f_R(1) + df((1, \theta])$$

and

$$f_L(\theta) = f_R(1) + df((1, \theta)) .$$

We also need to work with the space $\mathcal{BV}(\Theta, \mathbb{R})$ of functions of bounded variation on $\Theta = [\underline{\theta}, \bar{\theta}]$. By analogy with our discussion of the space $\mathcal{BV}((0, \infty), \mathbb{R})$, it should be clear that we can endow $\mathcal{BV}(\Theta, \mathbb{R})$ with the norm

$$\|f\|_{\mathcal{BV}} = |f_R(\theta_0)| + \|df\|_{TV} ,$$

where θ_0 is a fixed element of $(\underline{\theta}, \bar{\theta})$ and df is a bounded measure on Θ . There is, however, one surprise: a function $f \in \mathcal{BV}(\Theta, \mathbb{R})$ has a limit on the left at $\underline{\theta}$ and a limit on the right at $\bar{\theta}$. Indeed, we have

$$f_L(\underline{\theta}) = f_R(\theta_0) - df([\underline{\theta}, \theta_0])$$

and

$$f_R(\bar{\theta}) = f_R(\theta_0) + df((\theta_0, \bar{\theta}]) .$$

To summarize, we denote the space of functions of bounded variation on $(0, \infty)$ by $\mathcal{BV}((0, \infty), \mathbb{R})$, and we denote the space of functions of bounded variation on $\Theta = [\underline{\theta}, \bar{\theta}]$ by $\mathcal{BV}(\Theta, \mathbb{R})$. Unless explicitly stated to the contrary, we shall always use the right-continuous representative of a function of bounded variation. We will usually denote this representative simply by f , but we will occasionally denote it by f_R for emphasis. We will denote the left-continuous representative of f by f_L .

2.2. Assumptions on F . We are now in a position to introduce our assumptions on the distribution function F of the taste shock θ . They are:

A1 Both F and F' are functions of bounded variation on $(0, \infty)$.

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

A3 There exists $\theta_M \in [\underline{\theta}, \bar{\theta}]$ such that: (i) $G' \geq 0$ on $(0, \theta_M)$; and (ii) $G' \leq 0$ on (θ_M, ∞) .

Here G is given by the formula $G(\theta) = (1 - \beta)\theta F'(\theta) + F(\theta)$. If A1 holds then G , like F and F' , is a function of bounded variation on $(0, \infty)$.

2.3. The Support of F' is Connected. Fourth, we note that either $\beta = 1$, in which case the analysis is trivial, or $\beta < 1$, in which case the support of F' is connected.¹ More precisely, we have:

¹Notice that F is a distribution function, not a distribution. Also, F' has a dual interpretation. It can be regarded as: either (i) the non-negative finite measure with distribution function F ; or (ii) the density of that measure with respect to Lebesgue measure. By the same token, the support of F' has a dual interpretation. It can be regarded as: either (i) the support of the non-negative finite measure F' ; or (ii) the support of the non-negative function of bounded variation F' . It makes no difference which of these two interpretations is adopted.

Proposition 1. *Suppose that $\beta < 1$ and that A1-A3 are satisfied. Then there exist $\underline{\kappa}, \bar{\kappa} \in [\underline{\theta}, \bar{\theta}]$ such that: (i) $\underline{\kappa} < \bar{\kappa}$; (ii) $F' > 0$ on $(\underline{\kappa}, \bar{\kappa})$; and (iii) $F' = 0$ on $(0, \infty) \setminus [\underline{\kappa}, \bar{\kappa}]$.*

In what follows we shall therefore take it that $\beta < 1$, and that the support of F' is $[\underline{\theta}, \bar{\theta}]$.

Proof. Note first that there exists $\kappa_1 \in (\underline{\theta}, \bar{\theta})$ such that $F'(\kappa_1) > 0$. Otherwise we would have $F' = 0$ everywhere on $(0, \infty)$, by right-continuity of F' . Next, there exists $\kappa_2 \in (\kappa_1, \bar{\theta})$ such that $F' > 0$ on (κ_1, κ_2) , again by right-continuity of F' . Third, put $\underline{\kappa} = \inf \{\theta \mid F'(\theta) > 0\}$ and $\bar{\kappa} = \sup \{\theta \mid F'(\theta) > 0\}$. Then certainly $\underline{\theta} \leq \underline{\kappa} < \bar{\kappa} \leq \bar{\theta}$. Fourth, put $\alpha = \frac{2-\beta}{1-\beta}$. Then $G' \geq 0$ iff $(\theta^\alpha F')' \geq 0$ and $G' \leq 0$ iff $(\theta^\alpha F')' \leq 0$. There are therefore two possibilities. If $G' \geq 0$ at θ_M (i.e. $\Delta G(\theta_M) \geq 0$), then we must have $\theta^\alpha F' > 0$ on $(\underline{\kappa}, \theta_M]$ (because $(\theta^\alpha F')' \geq 0$ on this interval) and $\theta^\alpha F' > 0$ on $(\theta_M, \bar{\kappa})$ (because $(\theta^\alpha F')' \leq 0$ on this interval); and if $G' \leq 0$ at θ_M (i.e. $\Delta G(\theta_M) \leq 0$), then we must have $\theta^\alpha F' > 0$ on $(\underline{\kappa}, \theta_M)$ and $\theta^\alpha F' > 0$ on $[\theta_M, \bar{\kappa})$.² Either way, we see that: (i) $\theta^\alpha F' > 0$, and hence $F' > 0$, on $(\underline{\kappa}, \bar{\kappa})$; (ii) $\theta_M \leq \bar{\kappa}$, for otherwise we would have $F' > 0$ on the non-empty interval $(\bar{\kappa}, \theta_M)$, and this contradicts the choice of $\bar{\kappa}$; and (iii) $\theta_M \geq \underline{\kappa}$, for otherwise we would have $F' > 0$ on the non-empty interval $(\theta_M, \underline{\kappa})$, and this contradicts the choice of $\underline{\kappa}$. ■

2.4. Constraints on the Budget Set. Fifth, recall that self 0 chooses a subset B of the ambient action set A , and that self 1's choice of a consumption pair from B can therefore be described by a consumption curve $(c_1, c_2) : [\underline{\theta}, \bar{\theta}] \rightarrow B$. We consider three possible constraints on B , namely:

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 .³

²If $\theta_M \leq \underline{\kappa}$ then we take the intervals $(\underline{\kappa}, \theta_M)$ and $(\underline{\kappa}, \theta_M]$ to be empty. Similarly, if $\theta_M \geq \bar{\kappa}$, then we take the intervals $(\theta_M, \bar{\kappa})$ and $[\theta_M, \bar{\kappa})$ to be empty.

³In other words, for any given $(c_1, c_2) \in B$ and any $\Delta c_1 \in [0, \frac{1}{1+\pi} c_2]$, self 1 can increase her own consumption c_1 by Δc_1 at a cost of at most $(1 + \pi) \Delta c_1$ in terms of the consumption c_2 of self 2.

Constraint 3. The penalty for transferring consumption from period 1 to period 2 is no greater than π .⁴

Constraint 1 involves no loss of generality. Indeed, it must be possible for all possible types $\theta \in \Theta$ to find an optimum within B . This being the case, we can always take the closure of B without changing the outcome, since the utility function is continuous. Finally, since A itself is compact, so too is the closure of B . Constraint 2 is an essential part of the formulation of our problem. We wish to avoid extreme outcomes in which self 0 imposes an infinite penalty on self 1 for increasing her own consumption at the expense of self 2. Constraint 3 is simply the mirror image of Constraint 2. It eliminates extreme outcomes in which self 0 imposes an infinite penalty on self 1 for increasing the consumption of self 2 at her own expense.

Remark 2. *If we only impose Constraint 2, then the problem is one sided: Constraint 2 places a limit on the cost, in terms of c_2 , of increasing c_1 ; but there is no corresponding limit on the cost, in terms of c_1 , of increasing c_2 . By imposing Constraint 3, we eliminate this asymmetry.*

Now suppose that B must satisfy all three constraints. Then B must take one of two forms: either

1. it consists of the single point $(0, 0)$; or
2. its frontier consists of a curve that begins at some $(0, c_2)$ such that $c_2 > 0$, slopes downwards with slope between $-(1 + \pi)$ and $-(1 + \pi)^{-1}$, and ends at some $(c_1, 0)$ such that $c_1 > 0$.

Self 0 will never choose the first option, since the optimal pooling point on the frontier of the ambient budget set A is preferable. (By the same token, self 0 will never choose a B , the frontier of which is close to $(0, 0)$.) But, if she chooses the second option, then the resulting consumption curve (c_1, c_2) will be interior. That is, we will have $c_1, c_2 > 0$ on Θ .⁵

⁴In other words, for any given $(c_1, c_2) \in B$ and any $\Delta c_2 \in \left[0, \frac{1}{1+\pi} c_1\right]$, self 1 can increase the consumption c_2 of self 2 by Δc_2 at a cost of at most $(1 + \pi) \Delta c_2$ in terms of her own consumption c_1 .

⁵This follows from our assumption that $U'_t(0+) = \infty$.

The ideal approach to our problem would therefore be to impose all three constraints on B , and to solve the optimization problem of self 0 subject to these constraints. One could then verify ex post that Constraint 3 was not binding.⁶

In practice, we shall take a shortcut. Rather than working explicitly with Constraint 3, we shall instead replace it by the weaker requirement that consumption curves are interior. Our analysis could, of course, be reworked in such a way as to incorporate Constraint 3 explicitly. But, in practice, this would simply involve an additional notational burden.

Remark 3. *The situation would be very different if $\beta > 1$. In that case, it would be Constraint 2 that would not bind. We would therefore replace Constraint 2 by the weaker requirement that consumption curves are interior.*

2.5. Utility Curves. Suppose accordingly that we are given a B satisfying Constraints 1 and 2, and that the associated consumption curve is interior. Define a utility curve

$$(u_1, u_2) : [\underline{\theta}, \bar{\theta}] \rightarrow (U_1(0), U_1(\infty)) \times (U_2(0), U_2(\infty))$$

by the formula $(u_1, u_2)(\theta) = (U_1(c_1(\theta)), U_2(c_2(\theta)))$. Then (u_1, u_2) must satisfy the following conditions:

N1 $C_1(u_1(\theta)) + C_2(u_2(\theta)) \leq y$ for all $\theta \in [\underline{\theta}, \bar{\theta}]$.

N2 u_1 is non-decreasing and u_2 is non-increasing.

N3 $\theta du_1 + \beta du_2 = 0$.

N4 $\beta(1 + \pi) U'_2(C_2(u_2(\theta))) \geq \theta U'_1(C_1(u_1(\theta)))$.

Here: $C_t = U_t^{-1}$; and du_1 is a non-negative finite measure and du_2 is a non-positive finite measure.

Conversely, suppose that a utility curve (u_1, u_2) is interior, in the sense that it satisfies $u_1 > U_1(0)$ and $u_2 > U_2(0)$ on Θ , and that it satisfies Conditions N1-N4.

⁶It turns out that the slope of the optimal budget set is at most -1 . So Constraint 3 certainly is not binding!

Define (c_1, c_2) by the formula $(c_1, c_2)(\theta) = (U_1^{-1}(u_1(\theta)), U_2^{-1}(u_2(\theta)))$. Then there exists a B with slope at least $-(1 + \pi)$ such that (c_1, c_2) is the consumption curve arising from B . Moreover (c_1, c_2) is interior.

2.6. The CRRA Case. Suppose now that $U_1 = U_2 = U$ on $(0, \infty)$, and that U has constant relative risk aversion $\rho > 0$. Indeed, suppose for concreteness that U is given by the formula

$$U(c) = \begin{cases} \frac{c^{1-\rho}-1}{1-\rho} & \text{if } \rho \neq 1 \\ \log(c) & \text{if } \rho = 1 \end{cases}.$$

Then N4 is equivalent to

$$\mathbf{N4'} \quad u_2(\theta) \leq -\frac{1}{\rho} a\left(\frac{\theta}{(1+\pi)^\beta}\right) + b\left(\frac{\theta}{(1+\pi)^\beta}\right) u_1(\theta),$$

where a and b are given by the formulae

$$a(z) = \begin{cases} \frac{z^{1-\frac{1}{\rho}}-1}{1-\frac{1}{\rho}} & \text{if } \rho \neq 1 \\ \log(z) & \text{if } \rho = 1 \end{cases} \quad (1)$$

and

$$b(z) = z^{1-\frac{1}{\rho}}. \quad (2)$$

Remark 4. *It is obvious that N4 becomes weaker as π increases. Since N4' is equivalent to N4, N4' likewise becomes weaker as π increases.*

3. THE MAIN PROBLEM

Our strategy will be to study a relaxed version of the problem of self 0 in which we maximize self 0's expected utility $\int (\theta u(\theta) + w(\theta)) dF(\theta)$ subject to N1, N3 and N4', but not N2. Following Luenberger (1969, Sections 8.3 and 8.4, pp. 216-221), we shall need:

1. A vector space⁷ X , which we take to be $\mathcal{C}(\Theta, \mathbb{R})^2$. Here: $\Theta = [\underline{\theta}, \bar{\theta}] \subset (0, \infty)$ is the space of types; and $\mathcal{C}(\Theta, \mathbb{R})$ is the space of continuous functions from Θ to $\mathbb{R}$.

⁷In our case X is actually a Banach space. For the Lagrangean analysis, we only need the fact that it is a vector space. When we later use calculus to find necessary and sufficient conditions for the maximization of the Lagrangean, we shall need the fact that it is a normed space. Cf. Luenberger (1969, Lemma 1, p. 227).

2. A convex set⁸ $\Omega \subset X$, which we take to consist of all

$$(u, w) \in \left(\mathcal{BV}(\Theta, \text{ran}(U)) \cap \mathcal{C}(\Theta, \text{ran}(U)) \right)^2$$

such that

$$\theta du + \beta dw = 0.$$

Here: $\text{ran}(U)$ is the range of U ;⁹ $\mathcal{C}(\Theta, \text{ran}(U))$ is the space of continuous functions from Θ to $\text{ran}(U)$; $\mathcal{BV}(\Theta, \text{ran}(U))$ is the space of all functions of bounded variation from Θ to $\text{ran}(U)$; and du and dw are in general elements of the space $\mathcal{M}(\Theta, \mathbb{R})$ of finite Borel measures on Θ .

3. A concave function¹⁰ $M : \Omega \rightarrow \mathbb{R}$ (the objective function), which we take to be given by the formula

$$M(u, w) = \int \left(\theta u(\theta) + w(\theta) \right) dF(\theta).$$

4. A normed space¹¹ Z , which we take to be $\mathcal{C}(\Theta, \mathbb{R})$.
5. A closed convex cone P in Z with vertex 0 and non-empty interior, which we take to be $\mathcal{C}(\Theta, [0, \infty))$. With this choice of P , $z_1 \geq z_2$ iff $z_1(\theta) \geq z_2(\theta)$ for all $\theta \in \Theta$ and $z_1 > z_2$ iff $z_1(\theta) > z_2(\theta)$ for all $\theta \in \Theta$. In other words, P is the positive cone of Z .
6. The space Z^* of continuous linear functionals on Z . Since $Z = \mathcal{C}(\Theta, \mathbb{R})$, Z^* can be identified with $\mathcal{M}(\Theta, \mathbb{R})$.
7. The positive cone P^* of Z^* . Since $P = \mathcal{C}(\Theta, [0, \infty))$, P^* can be identified with $\mathcal{M}(\Theta, [0, \infty))$ (i.e. the space of non-negative finite Borel measures on Θ).

⁸In our case Ω is actually a cone, the vertex of which is the constant mapping $\frac{1}{\rho-1}$ when $\rho \neq 1$ and the constant mapping 0 when $\rho = 1$.

⁹I.e. $\text{ran}(U)$ is $(\frac{1}{\rho-1}, \infty)$ when $\rho < 1$, $(-\infty, \infty)$ when $\rho = 1$ and $(-\infty, \frac{1}{\rho-1})$ when $\rho > 1$.

¹⁰In our case M is actually defined on the whole of X (and not just on Ω), and it is linear (and not just concave).

¹¹In our case, Z is actually a Banach space, and not just a normed space.

8. Concave mappings $G_1, G_2 : \Omega \rightarrow Z$ (the constraint mappings),¹² which we take to be given by the formulae

$$(G_1(u, w))(\theta) = y - C(u(\theta)) - K(w(\theta))$$

and

$$(G_2(u, w))(\theta) = b\left(\frac{\theta}{(1+\pi)\beta}\right) u(\theta) - \frac{1}{\rho} a\left(\frac{\theta}{(1+\pi)\beta}\right) - w(\theta),$$

where $C = K = U^{-1} : \text{ran}(U) \rightarrow (0, \infty)$, and a and b are given by the formulae (1) and (2).

Our problem is then to

$$\begin{aligned} & \text{maximize} && M(x) \\ & \text{subject to} && \left\{ \begin{array}{l} x \in \Omega \\ G_1(x) \geq 0 \\ G_2(x) \geq 0 \end{array} \right\}. \end{aligned} \tag{3}$$

4. CHARACTERIZING THE OPTIMUM

In our context, the Lagrangean is the mapping $L : \Omega \times Z^* \times Z^* \rightarrow \mathbb{R}$ given by the formula

$$L(x, \lambda_1, \lambda_2) = M(x) + \langle G_1(x), \lambda_1 \rangle + \langle G_2(x), \lambda_2 \rangle,$$

where $\langle G_i(x), \lambda_i \rangle$ denotes the real number obtained when the linear functional $\lambda_i \in Z^*$ is evaluated at the point $G_i(x) \in Z$.

In view of our assumptions, the maximum is achieved at $x_0 \in \Omega$ if and only if there exist $\lambda_1, \lambda_2 \in Z^*$ such that:

1. $L(x_0, \lambda_1, \lambda_2) \geq L(x, \lambda_1, \lambda_2)$ for all $x \in \Omega$;
2. $G_1(x) \geq 0$, $\lambda_1 \geq 0$ and

$$\langle G_1(x), \lambda_1 \rangle = 0; \tag{4}$$

¹²In our case G_1 is actually defined on $\Xi = \mathcal{C}(\Theta, \text{ran}(U))^2$ (and not just on Ω). This will be useful when we later want to do calculus. Furthermore G_2 is defined on the whole of X (and not just on Ω), and it is linear (and not just concave).

3. $G_2(x) \geq 0$, $\lambda_2 \geq 0$ and

$$\langle G_2(x), \lambda_2 \rangle = 0. \quad (5)$$

In other words, there exists multipliers λ_1 and λ_2 such that: x_0 maximizes $L(\cdot, \lambda_1, \lambda_2)$ over Ω ; complementary slackness holds for the first constraint; and complementary slackness holds for the second constraint.

Since P^* can be identified with $\mathcal{M}(\Theta, [0, \infty))$, we have the following explicit representations of $M(x)$, $\langle G_1(x), \lambda_1 \rangle$ and $\langle G_2(x), \lambda_2 \rangle$:

$$M(x) = \int \left(\theta u(\theta) + w(\theta) \right) dF(\theta), \quad (6)$$

$$\langle G_1(x), \lambda_1 \rangle = \int \left(y - C(u(\theta)) - K(w(\theta)) \right) d\Lambda_1(\theta) \quad (7)$$

and

$$\langle G_2(x), \lambda_2 \rangle = \int \left(b\left(\frac{\theta}{(1+\pi)\beta}\right) u(\theta) - \frac{1}{\rho} a\left(\frac{\theta}{(1+\pi)\beta}\right) - w(\theta) \right) d\Lambda_2(\theta), \quad (8)$$

where Λ_1 and Λ_2 are the distribution functions of λ_1 and λ_2 respectively.

Remark 5. *In the interests of clarity and consistency, all integrals in this Appendix are Lebesgue-Stieltjes integrals, i.e. integrals with respect to functions of bounded variation.*

5. THE LAGRANGEAN IS FRÉCHET DIFFERENTIABLE

It is immediate from the formulae (6), (7) and (8) that $L(x, \lambda_1, \lambda_2)$ is in fact well defined for all $x \in \Xi = \mathcal{C}(\Theta, \text{ran}(U))^2$. Let us consider accordingly any $x_0 = (u_0, w_0) \in \Xi$ and any $x_1 = (u_1, w_1) \in X$. Because Ξ is open, $x_0 + \varepsilon x_1 \in \Xi$ for all $\varepsilon > 0$ sufficiently small. Furthermore, it can be verified that the directional derivative $\nabla_{x_1} L(x_0, \lambda_1, \lambda_2)$ of L at x_0 in the direction x_1 takes the form

$$\int \left(\theta u_1 + w_1 \right) dF - \int \left(C'(u_0) u_1 + K'(w_0) w_1 \right) d\Lambda_1 + \int \left(b\left(\frac{\theta}{(1+\pi)\beta}\right) u_1 - w_1 \right) d\Lambda_2. \quad (9)$$

This is easily seen to define a continuous linear functional

$$\nabla L(x_0, \lambda_1, \lambda_2) : x_1 \mapsto \nabla_{x_1} L(x_0, \lambda_1, \lambda_2)$$

on X . That is, $L(\cdot, \lambda_1, \lambda_2)$ is Gâteaux differentiable at x_0 with gradient $\nabla L(x_0, \lambda_1, \lambda_2) \in X^*$. Finally, $\nabla L(\cdot, \lambda_1, \lambda_2) : \Xi \rightarrow X^*$ can be shown to be continuous. It follows that $L(\cdot, \lambda_1, \lambda_2)$ is Fréchet differentiable on Ξ .

6. MAXIMIZING THE LAGRANGEAN

Since $L(\cdot, \lambda_1, \lambda_2)$ is convex and Fréchet differentiable on Ξ , and since Ω is convex, the maximum of $L(\cdot, \lambda_1, \lambda_2)$ over Ω is achieved at $x_0 \in \Omega$ iff

$$\nabla_{x-x_0} L(x_0, \lambda_1, \lambda_2) \leq 0$$

for all $x \in \Omega$. In this section we shall identify the restrictions that this places on λ_1 and λ_2 .

To this end, put

$$Y = \left(\mathcal{BV}(\Theta, \mathbb{R}) \cap \mathcal{C}(\Theta, \mathbb{R}) \right) \times \mathbb{R};$$

and consider the affine transformation

$$S : Y \rightarrow \left(\mathcal{BV}(\Theta, \mathbb{R}) \cap \mathcal{C}(\Theta, \mathbb{R}) \right)^2$$

that maps $y = (u, r)$ to $x = (u_0 + u, w_0 + w)$, where w is the unique solution of the equation $\theta du + \beta dw = 0$ with boundary condition $w(\bar{\theta}) = r$.

For any $y \in Y$, we have

$$\begin{aligned} \nabla_{S(y)-x_0} L(x_0, \lambda_1, \lambda_2) &= \int (\theta u + w) dF - \int \left(C'(u_0) u + K'(w_0) w \right) d\Lambda_1 \\ &\quad + \int \left(b \left(\frac{\theta}{(1+\pi)\beta} \right) u - w \right) d\Lambda_2 \\ &= \int (\theta u + w) dF - \int \left(\frac{C'(u_0)}{K'(w_0)} u + w \right) d\tilde{\Lambda}_1 \\ &\quad + \int \left(b \left(\frac{\theta}{(1+\pi)\beta} \right) u - w \right) d\Lambda_2 \end{aligned}$$

(where $d\tilde{\Lambda}_1 = K'(w_0) d\Lambda_1$). Furthermore, integrating by parts, we have

$$\int w dF = [w F]_{\underline{\theta}}^{\bar{\theta}} - \int F dw = w(\bar{\theta}) F(\bar{\theta}) + \int F \frac{\theta}{\beta} du$$

ONLINE APPENDIX C TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

(because $F(\underline{\theta}-) = 0$ and $dw = -\frac{\theta}{\beta} du$)

$$= r F(\bar{\theta}) + \int F \frac{\theta}{\beta} du.$$

Moreover

$$\begin{aligned} \int F \theta du &= \int F (\theta du + u d\theta) - \int F u d\theta \\ &= [\theta u F]_{\underline{\theta}-}^{\bar{\theta}} - \int \theta u dF - \int F u d\theta \\ &= \bar{\theta} u(\bar{\theta}) F(\bar{\theta}) - \int \theta u dF - \int F u d\theta \end{aligned}$$

(where we have again used the fact that $F(\underline{\theta}-) = 0$). Hence

$$\int w dF = \left(\frac{\bar{\theta}}{\beta} u(\bar{\theta}) + r \right) F(\bar{\theta}) - \frac{1}{\beta} \int u (\theta dF + F d\theta).$$

Similarly,

$$\int w d\tilde{\Lambda}_1 = \left(\frac{\bar{\theta}}{\beta} u(\bar{\theta}) + r \right) \tilde{\Lambda}_1(\bar{\theta}) - \frac{1}{\beta} \int u (\theta d\tilde{\Lambda}_1 + \tilde{\Lambda}_1 d\theta)$$

and

$$\int w d\Lambda_2 = \left(\frac{\bar{\theta}}{\beta} u(\bar{\theta}) + r \right) \Lambda_2(\bar{\theta}) - \frac{1}{\beta} \int u (\theta d\Lambda_2 + \Lambda_2 d\theta).$$

Overall, then,

$$\begin{aligned} \nabla_{S(y)-x_0} L(x_0, \lambda_1, \lambda_2) &= \left(\frac{\bar{\theta}}{\beta} u(\bar{\theta}) + r \right) \left(F(\bar{\theta}) - \tilde{\Lambda}_1(\bar{\theta}) - \Lambda_2(\bar{\theta}) \right) \\ &\quad - \frac{1}{\beta} \int u \left((1 - \beta) \theta dF + F d\theta \right) \\ &\quad + \frac{1}{\beta} \int u \left(\left(\theta - \beta \frac{C'(u_0)}{K'(w_0)} \right) d\tilde{\Lambda}_1 + \tilde{\Lambda}_1 d\theta \right) \\ &\quad + \frac{1}{\beta} \int u \left(\left(\theta + \beta b \left(\frac{\theta}{(1+\pi)\beta} \right) \right) d\Lambda_2 + \Lambda_2 d\theta \right). \end{aligned}$$

Next, it is easy to see that the mapping

$$y \mapsto \nabla_{S(y)-x_0} L(x_0, \lambda_1, \lambda_2)$$

defines a continuous linear functional on Y . Since it does not depend on the derivatives of y , this functional extends uniquely to a continuous linear functional

$$y^* : \mathcal{C}(\Theta, \mathbb{R}) \times \mathbb{R} \rightarrow \mathbb{R}.$$

Indeed, we have

$$y^* = (u^*, r^*) \in \mathcal{M}(\Theta, \mathbb{R}) \times \mathbb{R} = (\mathcal{C}(\Theta, \mathbb{R}) \times \mathbb{R})^*,$$

where

$$\begin{aligned} du^* = & -\frac{1}{\beta} \left((1 - \beta) \theta dF + F d\theta \right) + \frac{1}{\beta} \left(\left(\theta - \beta \frac{C'(u_0)}{K'(w_0)} \right) d\tilde{\Lambda}_1 + \tilde{\Lambda}_1 d\theta \right) \\ & + \frac{1}{\beta} \left(\left(\theta + \beta b \left(\frac{\theta}{(1+\pi)\beta} \right) \right) d\Lambda_2 + \Lambda_2 d\theta \right) + \frac{\bar{\theta}}{\beta} \left(F(\bar{\theta}) - \tilde{\Lambda}_1(\bar{\theta}) - \Lambda_2(\bar{\theta}) \right) dI \end{aligned} \quad (10)$$

and

$$r^* = F(\bar{\theta}) - \tilde{\Lambda}_1(\bar{\theta}) - \Lambda_2(\bar{\theta})$$

and I is the distribution function of the unit mass at $\bar{\theta}$.

Remark 6. In reading (10), note that F , $\tilde{\Lambda}_1$, Λ_2 and θ are functions of bounded variation (with θ being simply the identity function). Hence dF , $d\tilde{\Lambda}_1$, $d\Lambda_2$ and $d\theta$ are measures, and the equation as a whole holds in terms of measures.

Finally, it is easy to see that there exists $\varepsilon > 0$ such that $S(y) \in \Omega$ for all $y \in Y \cap B_\varepsilon(0)$, where $B_\varepsilon(0)$ is the open ball in $\mathcal{C}(\Theta, \mathbb{R}) \times \mathbb{R}$ with radius ε and centre 0. It follows that $\langle y, y^* \rangle \leq 0$ for all $y \in Y \cap B_\varepsilon(0)$. But $Y \cap B_\varepsilon(0)$ is dense in $B_\varepsilon(0)$.

Hence $\langle y, y^* \rangle \leq 0$ for all $y \in B_\varepsilon(0)$. Hence $y^* = 0$. In other words, we have

$$\begin{aligned} 0 = & -\frac{1}{\beta} \left((1 - \beta) \theta dF + F d\theta \right) + \frac{1}{\beta} \left(\left(\theta - \beta \frac{C'(u_0)}{K'(w_0)} \right) d\tilde{\Lambda}_1 + \tilde{\Lambda}_1 d\theta \right) \\ & + \frac{1}{\beta} \left(\left(\theta + \beta b \left(\frac{\theta}{(1+\pi)\beta} \right) \right) d\Lambda_2 + \Lambda_2 d\theta \right) + \frac{\bar{\theta}}{\beta} \left(F(\bar{\theta}) - \tilde{\Lambda}_1(\bar{\theta}) - \Lambda_2(\bar{\theta}) \right) dI \end{aligned} \quad (11)$$

and

$$0 = F(\bar{\theta}) - \tilde{\Lambda}_1(\bar{\theta}) - \Lambda_2(\bar{\theta}). \quad (12)$$

Taking advantage of (12), (11) simplifies to

$$\begin{aligned} 0 = & -G d\theta + \left(\left(\theta - \beta \frac{C'(u_0)}{K'(w_0)} \right) d\tilde{\Lambda}_1 + \tilde{\Lambda}_1 d\theta \right) \\ & + \left(\left(\theta + \beta b \left(\frac{\theta}{(1+\pi)\beta} \right) \right) d\Lambda_2 + \Lambda_2 d\theta \right), \end{aligned} \quad (13)$$

where G is given by the formula $G(\theta) = (1 - \beta) \theta F'(\theta) + F(\theta)$.

7. A ONE-PARAMETER FAMILY OF UTILITY CURVES

We shall consider a family of utility curves depending on the single parameter $\theta_1 \in (0, \bar{\theta}]$. For each θ_1 , we begin by finding the point $(c^*(\theta_1), k^*(\theta_1))$ that would be chosen by a self 1 of type θ_1 from the ambient budget set A . The utility curve corresponding to θ_1 is then the utility curve associated with a two-part budget set with slopes of -1 and $-(1 + \pi)$ to the left and right of a kink at $(c^*(\theta_1), k^*(\theta_1))$. Notice that we specifically allow for the possibility that $\theta_1 < \underline{\theta}$.

Put $\theta_2 = (1 + \pi) \theta_1$. Knife-edge cases apart, there are then five possible cases arising from the relative positions of the two non-trivial intervals $[\underline{\theta}, \bar{\theta}]$ and $[\theta_1, \theta_2]$:

Case 1 $[\underline{\theta}, \bar{\theta}]$ contains $[\theta_1, \theta_2]$;

Case 2 $[\theta_1, \theta_2]$ contains $[\underline{\theta}, \bar{\theta}]$;

Case 3 the two intervals overlap, with $[\theta_1, \theta_2]$ lying to the left and $[\underline{\theta}, \bar{\theta}]$ to the right;

Case 4 the two intervals overlap, with $[\underline{\theta}, \bar{\theta}]$ lying to the left and $[\theta_1, \theta_2]$ to the right;

Case 5 $[\theta_1, \theta_2]$ lies entirely to the left of $[\underline{\theta}, \bar{\theta}]$.

(The case in which $[\underline{\theta}, \bar{\theta}]$ lies entirely to the left of $[\theta_1, \theta_2]$ cannot occur, because we are confining θ_1 to the interval $(0, \bar{\theta}]$.)

7.1. Case 1. If the utility curve corresponding to θ_1 is to be an optimum, then the associated multipliers λ_1 and λ_2 must satisfy the three necessary conditions (4), (5) and (13) (i.e. complementary slackness for the first constraint, complementary slackness for the second constraint and the measure equation). In this section, we show that these three necessary conditions tie down λ_1 and λ_2 uniquely. The fourth necessary condition, namely the boundary condition (12), is not needed at this stage. (It will be used below to tie down θ_1 .)

By construction, the maximum-penalty constraint is strictly slack on $[\underline{\theta}, \theta_2)$ and the budget constraint is strictly slack on $(\theta_2, \bar{\theta}]$. Hence $d\Lambda_2 = 0$ on the former interval and $d\tilde{\Lambda}_1 = 0$ on the latter. Furthermore (13) implies that

$$(1 - \beta) \theta_2 \Delta F(\theta_2) = (\theta_2 - \theta_1) \Delta \tilde{\Lambda}_1(\theta_2) + \left(\theta_2 + \beta b \left(\frac{\theta_2}{(1+\pi)\beta} \right) \right) \Delta \Lambda_2(\theta_2),$$

where $\Delta F(\theta_2)$, $\Delta \tilde{\Lambda}_1(\theta_2)$ and $\Delta \Lambda_2(\theta_2)$ are the atoms of dF , $d\tilde{\Lambda}_1$ and $d\Lambda_2$ at θ_2 . But Assumption A1 implies that $\Delta F(\theta_2) = 0$. Since all the terms on the right-hand side of the equation are non-negative, it follows that $\Delta \tilde{\Lambda}_1(\theta_2) = \Delta \Lambda_2(\theta_2) = 0$. Hence $\Delta \Lambda_2(\theta_2) = 0$ (and therefore $d\Lambda_2 = 0$ on $[\underline{\theta}, \theta_2]$) and $\Delta \tilde{\Lambda}_1(\theta_2) = 0$ (and therefore $d\tilde{\Lambda}_1 = 0$ on $[\theta_2, \bar{\theta}]$).

Now let us consider the three intervals $[\underline{\theta}, \theta_1]$, $[\theta_1, \theta_2]$ and $[\theta_2, \bar{\theta}]$ in turn. On $[\underline{\theta}, \theta_1]$, we have $\frac{C'(u_0)}{K'(w_0)} = \frac{\theta}{\beta}$, $\Lambda_2 = 0$ and $d\Lambda_2 = 0$. Hence (13) becomes

$$0 = -G d\theta + \tilde{\Lambda}_1 d\theta.$$

It follows that $\tilde{\Lambda}_1 = G$ almost everywhere with respect to Lebesgue measure $d\theta$. Since both $\tilde{\Lambda}_1$ and G are functions of bounded variation, it then follows (bearing in mind the convention that functions of bounded variation are right continuous) that $\tilde{\Lambda}_1 = G$ everywhere on $[\underline{\theta}, \theta_1)$, and hence that $\tilde{\Lambda}_1(\theta_1-) = G(\theta_1-)$.

ONLINE APPENDIX C TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

On $[\theta_1, \theta_2]$, we have $\frac{C'(u_0)}{K'(w_0)} = \frac{\theta_1}{\beta}$, $\Lambda_2 = 0$ and $d\Lambda_2 = 0$. Hence (13) becomes

$$0 = -G d\theta + (\theta - \theta_1) d\tilde{\Lambda}_1 + \tilde{\Lambda}_1 d\theta.$$

This implies that $\tilde{\Lambda}_1$ takes the form

$$\tilde{\Lambda}_1(\theta) = \frac{1}{\theta - \theta_1} \int_{\theta_1}^{\theta} G(t) dt$$

for all $\theta \in (\theta_1, \theta_2)$, that $\tilde{\Lambda}_1(\theta_1) = G(\theta_1)$ (since $\tilde{\Lambda}_1$ and G are right continuous) and that $\tilde{\Lambda}_1(\theta_2) = \tilde{\Lambda}_1(\theta_2-) = \frac{1}{\theta_2 - \theta_1} \int_{\theta_1}^{\theta_2} G(\theta) d\theta$ (since $\tilde{\Lambda}_1$ cannot have a jump at θ_2).

On $[\theta_2, \bar{\theta}]$, we have $d\tilde{\Lambda}_1 = 0$. Hence (13) becomes

$$0 = -G d\theta + \tilde{\Lambda}_1 d\theta + \left(\theta + \beta b\left(\frac{\theta}{(1+\pi)\beta}\right) \right) d\Lambda_2 + \Lambda_2 d\theta.$$

Furthermore, we have the boundary condition $\Lambda_2(\theta_2) = 0$ (since Λ_2 cannot have a jump at θ_2). Putting $\tilde{\Lambda}_2 = \Lambda_2 + \tilde{\Lambda}_1(\theta_2)$, this equation simplifies slightly to

$$0 = -G d\theta + \left(\theta + \beta b\left(\frac{\theta}{(1+\pi)\beta}\right) \right) d\tilde{\Lambda}_2 + \tilde{\Lambda}_2 d\theta,$$

with boundary condition $\tilde{\Lambda}_2(\theta_2) = \tilde{\Lambda}_1(\theta_2)$.

7.2. Cases 2 to 5. In order to handle the remaining cases, we need a unified construction. (This construction includes Case 1 too.) It is more convenient to work in terms of the distribution function $\Psi = \Psi(\cdot; \theta_1)$ of the total multiplier $d\tilde{\Lambda}_1 + d\Lambda_2$, and to view this as a function on $[0, \bar{\theta}]$. The construction is then very simple. For all $\theta_1 \in (0, \bar{\theta}]$:

1. put $\Psi = G$ on $[0, \theta_1]$;
2. if $\theta_1 < \bar{\theta}$ (so that Ψ is not yet defined on the whole of $[0, \bar{\theta}]$), then let Ψ be the unique bounded solution of the o.d.e.

$$0 = -G + (\theta - \theta_1) \Psi' + \Psi$$

on $(\theta_1, \theta_2 \wedge \bar{\theta}]$, i.e. put

$$\Psi(\theta) = \frac{1}{\theta - \theta_1} \int_{\theta_1}^{\theta} G(t) dt;$$

3. if $\theta_2 < \bar{\theta}$ (so that Ψ is still not defined on the whole of $[0, \bar{\theta}]$), then let Ψ be the unique solution of the o.d.e.

$$0 = -G + \left(\theta + \beta b\left(\frac{\theta}{(1+\pi)\beta}\right) \right) \Psi' + \Psi$$

on $(\theta_2, \bar{\theta}]$ with boundary condition

$$\Psi(\theta_2) = \frac{1}{\theta_2 - \theta_1} \int_{\theta_1}^{\theta_2} G(t) dt.$$

Then, using arguments similar to those of the preceding section, it is easy to verify that the three necessary conditions (4), (5) and (13) together imply that, for all $\theta_1 \in (0, \bar{\theta}]$, Ψ must take the given form.

Remark 7. We can easily extend the construction of Ψ to include the case $\theta_1 = 0$. Indeed, in line with the construction above, we can let $\Psi(\cdot; 0)$ be the unique bounded solution of the o.d.e.

$$0 = -G + \left(\theta + \beta b\left(\frac{\theta}{(1+\pi)\beta}\right) \right) \Psi' + \Psi$$

on $(0, \bar{\theta}]$.

Remark 8. With this definition of $\Psi(\cdot; 0)$, $\Psi(\cdot; \theta_1)$ is independent of θ_1 for $\theta_1 \in [0, \frac{1}{1+\pi} \bar{\theta}]$.

8. EXISTENCE OF AN OPTIMUM

For all $\theta_1 \in (0, \bar{\theta}]$, we have shown that there exists a unique $\Psi = \Psi(\cdot; \theta_1)$ satisfying the two complementary slackness conditions (4) and (5) and the measure equation (13). This $\Psi(\cdot; \theta_1)$ does not in general satisfy the boundary condition (12). The

purpose of the current section is to establish that there is at least one choice of θ_1 for which (12) is satisfied.

For all $\theta_1 \in (0, \bar{\theta}]$, put $\psi(\theta_1) = \Psi(\bar{\theta}; \theta_1)$. Now consider G . We certainly have $G = 0$ on $(0, \underline{\theta})$ and $G = F(\bar{\theta})$ on $[\bar{\theta}, \infty)$. Furthermore Assumption A3 tells us that G is first increasing (on $(0, \theta_M)$) and then decreasing (on (θ_M, ∞)). It is therefore obvious that there exists $\underline{\theta}_F \in [\underline{\theta}, \bar{\theta}]$ such that $G < F(\bar{\theta})$ on $(0, \underline{\theta}_F)$ and $G \geq F(\bar{\theta})$ on $[\underline{\theta}_F, \bar{\theta}]$. Our first lemma sharpens this observation.

Lemma 9. *There exists $\bar{\theta}_F \in [\underline{\theta}, \bar{\theta})$ such that $G \leq F(\bar{\theta})$ on $(0, \bar{\theta}_F)$ and $G > F(\bar{\theta})$ on $(\bar{\theta}_F, \bar{\theta})$.*

Proof. Suppose first that $\theta_M < \bar{\theta}$. Suppose further that there exists $\xi_0 \in (\theta_M, \bar{\theta})$ such that $G(\xi_0) = G(\bar{\theta})$. Since $G' \leq 0$ on $(\theta_M, \infty) \supset (\xi_0, \bar{\theta}]$, it follows that $G' = 0$ on $(\xi_0, \bar{\theta}]$. We also know that $G = F(\bar{\theta})$ on $[\bar{\theta}, \infty)$, and therefore that $G' = 0$ on $(\bar{\theta}, \infty)$. Overall, then, $G' = 0$ on (ξ_0, ∞) . Hence $\theta^\alpha F'(\theta)$ is constant on (ξ_0, ∞) , where $\alpha = \frac{2-\beta}{1-\beta}$. (Cf. the proof of Proposition 1.) Hence $F' = 0$ on (ξ_0, ∞) . But this contradicts the fact that $\bar{\theta}$ is the maximum of the support of F . We may therefore conclude that $G > G(\bar{\theta}) = F(\bar{\theta})$ on $(\theta_M, \bar{\theta})$.

Suppose second that $\theta_M = \bar{\theta}$. Then $G' \geq 0$ on $(0, \bar{\theta})$. Hence $\theta^\alpha F'(\theta)$ is non-decreasing on $(0, \bar{\theta})$. In particular, if we fix $\xi_1 \in (\underline{\theta}, \bar{\theta})$, then we will have $\theta^\alpha F'(\theta) \geq \xi_1^\alpha F'(\xi_1)$ for all $\theta \in (\xi_1, \bar{\theta})$. Letting $\theta \uparrow \bar{\theta}$, it follows that $\bar{\theta}^\alpha F'(\bar{\theta}-) \geq \xi_1^\alpha F'(\xi_1)$. But $F'(\xi_1) > 0$, since ξ_1 lies in the interior of the support of F . Hence $F'(\bar{\theta}-) > 0$. Hence $G(\bar{\theta}-) = (1 - \beta) \bar{\theta} F'(\bar{\theta}-) + F(\bar{\theta}) > F(\bar{\theta})$. Hence there exists $\varepsilon > 0$ such that $G > F(\bar{\theta})$ on $(\bar{\theta} - \varepsilon, \bar{\theta})$.

Overall, then, we have the following picture: $G = 0$ on $(0, \underline{\theta})$; there exists $\xi_2 \in [\underline{\theta}, \bar{\theta})$ such that $G > F(\bar{\theta})$ on $(\xi_2, \bar{\theta})$; and $G = F(\bar{\theta})$ on $[\bar{\theta}, \infty)$. Furthermore Assumption A3 ensures that $\{\theta \mid G(\theta) > F(\bar{\theta})\}$ is an interval. We may therefore put $\bar{\theta}_F = \inf \{\theta \mid \theta \in [\underline{\theta}, \bar{\theta}), G(\theta) > F(\bar{\theta})\}$. ■

Now, it follows from the construction of Ψ given in Section 7.2 that $\Psi(\bar{\theta}; \theta_1)$ is a convex combination of the values of G on the interval $(\theta_1, \bar{\theta})$. Combining this observation with Lemma 9, we obtain:

Lemma 10. *$\psi > F(\bar{\theta})$ on $[\bar{\theta}_F, \bar{\theta})$.* ■

We also have:

Lemma 11. $\psi < F(\bar{\theta})$ on $(0, \frac{1}{1+\pi} \underline{\theta}]$.

Proof. We begin by noting that $\Psi = \Psi(\cdot; \theta_1)$ is independent of θ_1 for $\theta_1 \in [0, \frac{1}{1+\pi} \underline{\theta}]$. It is therefore the unique bounded solution of the o.d.e.

$$0 = -G + \left(\theta + \beta b\left(\frac{\theta}{(1+\pi)\beta}\right) \right) \Psi' + \Psi \quad (14)$$

on $(0, \bar{\theta}]$. We compare Ψ with the function Φ which is the unique bounded solution of the o.d.e.

$$0 = -G + \theta \Phi' + \Phi \quad (15)$$

on $(0, \bar{\theta}]$. Now

$$\Phi(\theta) = \frac{1}{\theta} \int_0^\theta G(t) dt = (1 - \beta) F(\theta) + \beta \frac{1}{\theta} \int_0^\theta F(t) dt.$$

Hence: $\Phi = 0$ on $(0, \underline{\theta}]$; and $0 < \Phi < F$ on $(\underline{\theta}, \bar{\theta}]$. Hence

$$\Phi' = \frac{G - \Phi}{\theta} \geq \frac{F - \Phi}{\theta} \geq 0$$

on $(0, \bar{\theta}]$, with strict inequality on $(\underline{\theta}, \bar{\theta}]$. Furthermore, Φ is a supersolution of the equation for Ψ . Indeed, we have

$$\begin{aligned} & -G + \left(\beta b\left(\frac{\theta}{(1+\pi)\beta}\right) + \theta \right) \Phi' + \Phi \\ = & -G + \beta b\left(\frac{\theta}{(1+\pi)\beta}\right) \Phi' + \theta \Phi' + \Phi \end{aligned}$$

(on rearranging)

$$= \beta b\left(\frac{\theta}{(1+\pi)\beta}\right) \Phi'$$

(since Φ satisfies (15))

$$\geq 0$$

on $(0, \bar{\theta}]$, with strict inequality on $(\underline{\theta}, \bar{\theta}]$. That is, Φ is a supersolution of the equation for Ψ , and a strict supersolution on $(\underline{\theta}, \bar{\theta}]$. Hence $\Phi > \Psi$ on $(\underline{\theta}, \bar{\theta}]$. In

particular, $\psi(0) = \Psi(\bar{\theta}) < \Phi(\bar{\theta}) < F(\bar{\theta})$. The general case now follows on noting that $\Psi(\cdot; \theta_1) = \Psi(\cdot; 0)$ for all $\theta_1 \in (0, \frac{1}{1+\pi}\underline{\theta}]$. Cf. the remark at the end of Section 7.2. ■

Since ψ is continuous, we can combine Lemmas 10 and 11 to obtain:

Proposition 12. *There exists $\theta_1 \in (\frac{1}{1+\pi}\underline{\theta}, \bar{\theta}_F)$ such that $\psi(\theta_1) = F(\bar{\theta})$. ■*

That is, there exists $\theta_1 \in (\frac{1}{1+\pi}\underline{\theta}, \bar{\theta}_F)$ such that equation (12) is satisfied. However, we still need to verify that the multipliers associated with any such θ_1 are non-negative.

9. NON-NEGATIVITY OF THE MULTIPLIER

In this section we show that, if $\theta_1 \in (0, \bar{\theta}_F)$ is such that $\psi(\theta_1) \leq F(\bar{\theta})$, then $\Psi = \Psi(\cdot; \theta_1)$ is non-decreasing on $[0, \bar{\theta}]$. We treat the intervals $[0, \theta_1]$, $(\theta_1, \bar{\theta}_F)$ and $[\bar{\theta}_F, \bar{\theta}]$ in turn. We begin with a simple lemma.

Lemma 13. $\bar{\theta}_F \leq \theta_M$.

Proof. In the light of Lemma 9, $\sup G > F(\bar{\theta})$. Moreover it follows from the definition of θ_M that $\sup G = \max\{G_L(\theta_M), G(\theta_M)\}$. There are therefore two possibilities. Either $G_L(\theta_M) > F(\bar{\theta})$, in which case there is a left neighbourhood of θ_M on which $G > F(\bar{\theta})$, and therefore $\bar{\theta}_F < \theta_M$. Or $G(\theta_M) > F(\bar{\theta})$, in which case necessarily $\bar{\theta}_F \leq \theta_M$. ■

Lemma 14. $G' \geq 0$ on $[0, \bar{\theta}_F]$.

Proof. From Lemma 13 we know that $[0, \bar{\theta}_F) \subset [0, \theta_M)$. Hence $G' \geq 0$ on $[0, \bar{\theta}_F)$. However, $G \leq F(\bar{\theta})$ on $[0, \bar{\theta}_F)$ and $G > F(\bar{\theta})$ on $(\bar{\theta}_F, \bar{\theta})$. Hence $\Delta G(\bar{\theta}_F) \geq 0$. Hence $G' \geq 0$ on $[0, \bar{\theta}_F]$. ■

Proposition 15. *Suppose that $\psi(\theta_1) \leq F(\bar{\theta})$. Then $\Psi' \geq 0$ on $[0, \theta_1]$.*

Proof. Since $\psi(\theta_1) \leq F(\bar{\theta})$, Lemma 10 implies that $\theta_1 < \bar{\theta}_F$. Hence $[0, \theta_1] \subset [0, \bar{\theta}_F]$. But $\Psi = G$ on $[0, \theta_1]$ by construction of Ψ , and Lemma 14 tells us that $G' \geq 0$ on $[0, \bar{\theta}_F]$. It follows that $\Psi' \geq 0$ on $[0, \theta_1]$. ■

ONLINE APPENDIX C TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

In order to show that $\Psi' \geq 0$ on $(\theta_1, \bar{\theta}]$, we use the fact that Ψ solves

$$0 = -G + (\theta - \theta_1) \Psi' + \Psi \quad (16)$$

on $(\theta_1, \theta_2 \wedge \bar{\theta}]$ and

$$0 = -G + \left(\theta + \beta b \left(\frac{\theta}{(1+\pi)\beta} \right) \right) \Psi' + \Psi \quad (17)$$

on $(\theta_2 \wedge \bar{\theta}, \bar{\theta}]$. We also make use of the corresponding o.d.e. for θ , namely

$$\dot{\theta} = -(\theta - \theta_1) \quad (18)$$

on $(\theta_1, \theta_2 \wedge \bar{\theta}]$ and

$$\dot{\theta} = - \left(\theta + \beta b \left(\frac{\theta}{(1+\pi)\beta} \right) \right) \quad (19)$$

on $(\theta_2 \wedge \bar{\theta}, \bar{\theta}]$. Specifically, for all $g, h \in (\theta_1, \bar{\theta}]$ such that $h < g$, let $T(h, g)$ denote the time at which the solution of the o.d.e. (18-19) for θ starting from g hits h , and put $S(h, g) = \exp(-T(h, g))$. Notice that $S(\cdot, g)$ increases from 0 at θ_1 to 1 at g .

Lemma 16. *Suppose that $\psi(\theta_1) \leq F(\bar{\theta})$. Then $\Psi \leq G$ on (θ_1, θ_M) .*

Proof. Since $\psi(\theta_1) \leq F(\bar{\theta})$, Lemma 10 implies that $\theta_1 < \bar{\theta}_F$. Furthermore Lemma 13 tells us that $\bar{\theta}_F \leq \theta_M$. For all $g \in (\theta_1, \theta_M)$, we therefore have

$$\Psi(g) = \int_{\theta_1}^g \frac{\partial S}{\partial h}(h, g) G(h) dh \leq \int_{\theta_1}^g \frac{\partial S}{\partial h}(h, g) G(g-) dh$$

(with equality iff $G(g-) = G(\theta_1)$)

$$= G(g-).$$

Taking limits from the right (and using the continuity of Ψ and the right continuity of G) then yields $\Psi \leq G$. ■

Lemma 17. *The sign of Ψ' coincides with that of $G - \Psi$ on $(\theta_1, \bar{\theta}]$.*

Proof. We have

$$\Psi' = \frac{G - \Psi}{\theta - \theta_1}$$

on $(\theta_1, \theta_2 \wedge \bar{\theta}]$ (from equation (16)) and

$$\Psi' = \frac{G - \Psi}{\theta + \beta b\left(\frac{\theta}{(1+\pi)\beta}\right)}$$

on $(\theta_2 \wedge \bar{\theta}, \bar{\theta}]$ (from equation (17)). Bearing in mind that we have $\theta - \theta_1 > 0$ on $(\theta_1, \theta_2 \wedge \bar{\theta}]$, it follows that the sign of Ψ' coincides with that of $G - \Psi$ on $(\theta_1, \theta_2 \wedge \bar{\theta}] \cup (\theta_2 \wedge \bar{\theta}, \bar{\theta}] = (\theta_1, \bar{\theta}]$. ■

Proposition 18. *Suppose that $\psi(\theta_1) \leq F(\bar{\theta})$. Then $\Psi' \geq 0$ on $(\theta_1, \bar{\theta}_F)$.*

Proof. From Lemma 16 we know that $\Psi \leq G$ on (θ_1, θ_M) and therefore on $(\theta_1, \bar{\theta}_F) \subset (\theta_1, \theta_M)$. Lemma 17 then implies that $\Psi' \geq 0$ there. ■

Proposition 19. *Suppose that $\psi(\theta_1) \leq F(\bar{\theta})$. Then $\Psi' > 0$ on $[\bar{\theta}_F, \bar{\theta})$.*

Proof. For all $\theta \in [\bar{\theta}_F, \bar{\theta})$, we have

$$\Psi(\bar{\theta}) = S(\theta, \bar{\theta}) \Psi(\theta) + \int_{\theta}^{\bar{\theta}} \frac{\partial S}{\partial h}(h, \bar{\theta}) G(h) dh. \quad (20)$$

Since

$$\int_{\theta}^{\bar{\theta}} \frac{\partial S}{\partial h}(h, \bar{\theta}) dh = 1 - S(\theta, \bar{\theta}),$$

this means that $\Psi(\bar{\theta})$ is a convex combination of $\Psi(\theta)$ and the values of G on $(\theta, \bar{\theta}]$. But $\Psi(\bar{\theta}) = \psi(\theta_1) \leq F(\bar{\theta})$ and $G > F(\bar{\theta})$ on $(\theta, \bar{\theta})$. So we must have $\Psi(\theta) < F(\bar{\theta})$. On the other hand, $G(\theta) \geq F(\bar{\theta})$ since $\theta \geq \bar{\theta}_F$. Lemma 17 therefore implies that $\Psi'(\theta) > 0$. ■

10. UNIQUENESS OF THE OPTIMUM

At this point we have established that the utility curve associated with θ_1 solves the main maximization problem (3) iff $\psi(\theta_1) = F(\bar{\theta})$. Furthermore $\psi < F(\bar{\theta})$ on $[0, \frac{1}{1+\pi}\underline{\theta}]$ and $\psi > F(\bar{\theta})$ on $[\bar{\theta}_F, \bar{\theta})$. Hence there exists $\theta_1 \in (\frac{1}{1+\pi}\underline{\theta}, \bar{\theta}_F)$ such that $\psi(\theta_1) = F(\bar{\theta})$. In the current section, we show that the set of θ_1 for which $\psi(\theta_1) = F(\bar{\theta})$ is an interval. Furthermore, if we strengthen Assumption A3 by

requiring that G is strictly increasing to the left of its peak, then $\psi' > 0$ over the relevant range. It then follows that there is a unique θ_1 for which $\psi(\theta_1) = F(\bar{\theta})$. This result is limited: it shows that – under the strengthened version of A3 – there is exactly one solution to problem (3) within our one-parameter family of utility curves; but it does not show that there is exactly one solution to problem (3) in Ω . It is, however, very suggestive.

The main idea of the proof is to find an explicit formula for ψ' , and then use this formula to determine the sign of ψ' . Of course, the formula depends on whether $\theta_1 < \frac{1}{1+\pi}\bar{\theta}$ or $\theta_1 > \frac{1}{1+\pi}\bar{\theta}$. In the former case: $\theta_2 = (1+\pi)\theta_1 < \bar{\theta}$; the maximum-penalty constraint is strictly binding; and the types in the range $[\theta_2, \bar{\theta}]$ will choose to incur the consumption penalty. In the latter case: $\theta_2 = (1+\pi)\theta_1 > \bar{\theta}$; the maximum-penalty constraint is strictly slack; and no type will choose to incur the consumption penalty.

In order to state the formula for ψ' , it will be helpful to introduce the functions ϕ , χ , ζ and η given by the formulae

$$\phi(\theta_1) = \frac{1}{\theta_2 - \theta_1} \int_{\theta_1}^{\theta_2} G(\theta) d\theta \quad (21)$$

for all $\theta_1 \in (0, \infty)$ (where we have suppressed the dependence of θ_2 on θ_1),

$$\chi(\theta_1) = \frac{1}{\bar{\theta} - \theta_1} \int_{\theta_1}^{\bar{\theta}} G(\theta) d\theta \quad (22)$$

for all $\theta_1 \in (0, \bar{\theta})$,

$$\zeta(\theta_1) = \frac{\frac{\beta}{1+\pi} b\left(\frac{\theta_1}{\beta}\right)}{\theta_1 \left(\theta_1 + \frac{\beta}{1+\pi} b\left(\frac{\theta_1}{\beta}\right)\right)} (G(\theta_2) - \phi(\theta_1)) + \frac{1}{\pi \theta_1} (G(\theta_2) - G(\theta_1)) \quad (23)$$

for all $\theta_1 \in (0, \infty)$ (where we have suppressed the dependence of θ_2 on θ_1) and

$$\eta(\theta_1) = \frac{\chi(\theta_1) - G(\theta_1)}{\bar{\theta} - \theta_1} \quad (24)$$

for all $\theta_1 \in (0, \bar{\theta})$.

Lemma 20. $\psi'(\theta_1) = S(\theta_2, \bar{\theta}) \zeta(\theta_1)$ for $\theta_1 \in (0, \frac{1}{1+\pi} \bar{\theta})$.

Proof. Equation (20) gives

$$\psi(\theta_1) = \int_{\theta_2}^{\bar{\theta}} \frac{\partial S}{\partial h}(h, \bar{\theta}) G(h) dh + S(\theta_2, \bar{\theta}) \phi(\theta_1),$$

where we have used the fact that $\Psi(\bar{\theta}; \theta_1) = \psi(\theta_1)$ and $\Psi(\theta_2; \theta_1) = \phi(\theta_1)$. Hence

$$\begin{aligned} \psi' &= -\frac{\partial S}{\partial h}(\theta_2, \bar{\theta}) G(\theta_2) \theta_2' + \frac{\partial S}{\partial h}(\theta_2, \bar{\theta}) \theta_2' \phi + S(\theta_2, \bar{\theta}) \phi' \\ &= \exp(-T(\theta_2, \bar{\theta})) \left(\frac{\partial T}{\partial h}(\theta_2, \bar{\theta}) \theta_2' (G(\theta_2) - \phi) + \phi' \right), \end{aligned}$$

where we have suppressed the dependence of ϕ and ψ on θ_1 , and where θ_2' and ϕ' denote the derivatives of θ_2 and ϕ with respect to θ_1 . Furthermore

$$\phi = \frac{1}{\theta_2 - \theta_1} \int_{\theta_1}^{\theta_2} G(t) dt$$

and

$$T(\theta_2, \bar{\theta}) = \int_{\theta_2}^{\bar{\theta}} \frac{dt}{t + \beta b\left(\frac{t}{(1+\pi)\beta}\right)}.$$

Hence

$$\begin{aligned} \phi' &= -\frac{\theta_2' - 1}{(\theta_2 - \theta_1)^2} \int_{\theta_1}^{\theta_2} G(t) dt + \frac{1}{\theta_2 - \theta_1} (G(\theta_2) \theta_2' - G(\theta_1)) \\ &= \frac{1}{\theta_2 - \theta_1} (-\pi \phi + ((1 + \pi) G(\theta_2) - G(\theta_1))) \\ &= \frac{1}{\theta_1} (G(\theta_2) - \phi) + \frac{1}{\pi \theta_1} (G(\theta_2) - G(\theta_1)) \end{aligned}$$

and

$$\frac{\partial T}{\partial h}(\theta_2, \bar{\theta}) = -\frac{1}{\theta_2 + \beta b\left(\frac{\theta_2}{(1+\pi)\beta}\right)}. \quad (25)$$

Overall, then,

$$\begin{aligned}
 \exp(T(\theta_2, \bar{\theta})) \psi' &= \frac{\partial T}{\partial h}(\theta_2, \bar{\theta}) \theta_2' (G(\theta_2) - \phi) + \phi' \\
 &= -\frac{1 + \pi}{\theta_2 + \beta b\left(\frac{\theta_2}{(1+\pi)\beta}\right)} (G(\theta_2) - \phi) + \frac{1}{\theta_1} (G(\theta_2) - \phi) + \frac{1}{\pi \theta_1} (G(\theta_2) - G(\theta_1)) \\
 &= \left(\frac{1}{\theta_1} - \frac{1}{\theta_1 + \frac{\beta}{1+\pi} b\left(\frac{\theta_1}{\beta}\right)} \right) (G(\theta_2) - \phi) + \frac{1}{\pi \theta_1} (G(\theta_2) - G(\theta_1))
 \end{aligned}$$

(collecting terms in $(G(\theta_2) - \phi)$ and $(G(\theta_2) - G(\theta_1))$, and using the fact that $\theta_2 = (1 + \pi) \theta_1$)

$$= \frac{\frac{\beta}{1+\pi} b\left(\frac{\theta_1}{\beta}\right)}{\theta_1 \left(\theta_1 + \frac{\beta}{1+\pi} b\left(\frac{\theta_1}{\beta}\right) \right)} (G(\theta_2) - \phi) + \frac{1}{\pi \theta_1} (G(\theta_2) - G(\theta_1)).$$

Making ψ' the subject of this equation, we obtain the required result. ■

The second of the two formulae for ψ' is given by the following lemma.

Lemma 21. $\psi' = \eta$ on $\left[\frac{1}{1+\pi} \bar{\theta}, \bar{\theta}\right)$.

Proof. We have $\psi = \chi$ on $\left[\frac{1}{1+\pi} \bar{\theta}, \bar{\theta}\right)$. Moreover it is easy to check that

$$\chi'(\theta_1) = \frac{\chi(\theta_1) - G(\theta_1)}{\bar{\theta} - \theta_1}$$

on $(0, \bar{\theta})$. ■

There are now two main cases to consider. The more general of the two main cases occurs when $\frac{1}{1+\pi} \bar{\theta} < \bar{\theta}_F$. In this case, there are three main subcases to consider:

Subcase 1 $\theta_1 \in \left(\frac{1}{1+\pi} \bar{\theta}, \frac{1}{1+\pi} \bar{\theta}_F\right]$. In this subcase, the maximum-penalty constraint is strictly binding in the sense that $\theta_2 < \bar{\theta}$. I.e. all the types in the non-trivial range of $[\theta_2, \bar{\theta}]$ choose to make an early withdrawal from the penalty account.

Subcase 2 $\theta_1 \in \left[\frac{1}{1+\pi} \bar{\theta}_F, \frac{1}{1+\pi} \bar{\theta}\right]$. In this subcase, the maximum-penalty constraint is weakly binding in the sense that $\theta_2 \leq \bar{\theta}$.

Subcase 3 $\theta_1 \in [\frac{1}{1+\pi}\bar{\theta}, \bar{\theta}_F)$. In this subcase, the maximum-penalty constraint is weakly slack in the sense that $\theta_2 \geq \bar{\theta}$.

The less general of the two main cases occurs when $\frac{1}{1+\pi}\bar{\theta} \geq \bar{\theta}_F$. In this case, the third subcase does not arise.

We deal with both of the two main cases simultaneously. The first subcase is settled by the following lemma.

Lemma 22. *Suppose that $\theta_1 \in (0, \frac{1}{1+\pi}\bar{\theta}_F]$. Then $\psi'(\theta_1) \geq 0$.*

Proof. The proof relies on the formula $\psi'(\theta_1) = S(\theta_2, \bar{\theta})\zeta(\theta_1)$ given in Lemma 20. This formula is valid for $\theta_1 \in (0, \frac{1}{1+\pi}\bar{\theta}) \supset (0, \frac{1}{1+\pi}\bar{\theta}_F]$.

We have $[\theta_1, \theta_2] \subset (0, \bar{\theta}_F]$ and therefore $G' \geq 0$ on $[\theta_1, \theta_2]$. Hence $G(\theta_2) \geq \phi(\theta_1)$ and $G(\theta_2) \geq G(\theta_1)$. It then follows from formula (23) that $\zeta(\theta_1) \geq 0$ (with equality iff $G(\theta_2) = G(\theta_1)$), and thence that $\psi'(\theta_1) \geq 0$ (with equality iff $G(\theta_2) = G(\theta_1)$). ■

We now turn to the second subcase.

Lemma 23. *Suppose that $\theta_1 \in [\frac{1}{1+\pi}\bar{\theta}_F, \frac{1}{1+\pi}\bar{\theta}]$ and that $\psi(\theta_1) \leq F(\bar{\theta})$. Then $G(\theta_2) - \phi(\theta_1) \geq 0$.*

Proof. We have $\theta_2 \in [\bar{\theta}_F, \bar{\theta}]$ and therefore

$$G(\theta_2) \geq F(\bar{\theta})$$

(with strict inequality if $\theta_1 \in (\frac{1}{1+\pi}\bar{\theta}_F, \frac{1}{1+\pi}\bar{\theta})$, because then $\theta_2 \in (\bar{\theta}_F, \bar{\theta})$)

$$\geq \psi(\theta_1) = \Psi(\bar{\theta}; \theta_1)$$

(by assumption and by definition of ψ respectively)

$$\geq \Psi(\theta_2; \theta_1)$$

(with strict inequality if $\theta_1 \in [\frac{1}{1+\pi}\bar{\theta}_F, \frac{1}{1+\pi}\bar{\theta})$, because then $\theta_2 < \bar{\theta}$)

$$= \phi(\theta_1)$$

(by definition of ϕ). ■

Lemma 24. *Suppose that $\theta_1 \in [\frac{1}{1+\pi} \bar{\theta}_F, \frac{1}{1+\pi} \bar{\theta}]$ and that $\psi(\theta_1) \leq F(\bar{\theta})$. Then $G(\theta_2) - G(\theta_1) > 0$.*

Proof. We have

$$G(\theta_2) \geq F(\bar{\theta})$$

(with strict inequality if $\theta_1 \in (\frac{1}{1+\pi} \bar{\theta}_F, \frac{1}{1+\pi} \bar{\theta})$)

$$\geq \psi(\theta_1) = \Psi(\bar{\theta}; \theta_1) > \Psi(\theta_1; \theta_1) = G(\theta_1)$$

(by assumption, by definition of ψ , because $\theta_1 < \bar{\theta}$ and by construction of Ψ respectively). ■

Combining Lemmas 23 and 24, we obtain the following result about the right-hand derivative of ψ .

Lemma 25. *Suppose that $\theta_1 \in [\frac{1}{1+\pi} \bar{\theta}_F, \frac{1}{1+\pi} \bar{\theta})$ and that $\psi(\theta_1) \leq F(\bar{\theta})$. Then $\psi'(\theta_1) > 0$.*

Proof. The proof again relies on the formula $\psi'(\theta_1) = S(\theta_2, \bar{\theta}) \zeta(\theta_1)$ given in Lemma 20. In view of this formula, $\psi'(\theta_1) > 0$ if $G(\theta_2) \geq \phi(\theta_1)$ and $G(\theta_2) > G(\theta_1)$ with at least one strict inequality. But Lemmas 23 and 24 show that $G(\theta_2) \geq \phi(\theta_1)$ and $G(\theta_2) > G(\theta_1)$ respectively. ■

We also need the corresponding result about the left-hand derivative of ψ .

Lemma 26. *Suppose that $\theta_1 \in (\frac{1}{1+\pi} \bar{\theta}_F, \frac{1}{1+\pi} \bar{\theta}]$ and that $\psi(\theta_1) \leq F(\bar{\theta})$. Then $\psi'_L(\theta_1) > 0$.*

Proof. The proof parallels that of Lemma 25, with minor changes. First of all, bearing in mind that ϕ is continuous, we have

$$\zeta_L(\theta_1) = \frac{\frac{\beta}{1+\pi} b\left(\frac{\theta_1}{\beta}\right)}{\theta_1 \left(\theta_1 + \frac{\beta}{1+\pi} b\left(\frac{\theta_1}{\beta}\right)\right)} (G_L(\theta_2) - \phi(\theta_1)) + \frac{1}{\pi \theta_1} (G_L(\theta_2) - G_L(\theta_1)) \quad (26)$$

for all $\theta_1 \in (0, \infty)$. As in Lemma 20, we then have $\psi'_L(\theta_1) = S(\theta_2, \bar{\theta}) \zeta_L(\theta_1)$ for $\theta_1 \in (0, \frac{1}{1+\pi} \bar{\theta}]$. Next, just as the single peakedness of G implies that $G > F(\bar{\theta})$

on $(\bar{\theta}_F, \bar{\theta})$, so it also implies that $G_L > F(\bar{\theta})$ on $(\bar{\theta}_F, \bar{\theta})$. Arguing as in Lemma 23, we can therefore show that $G_L(\theta_2) - \phi(\theta_1) \geq 0$ for $\theta_1 \in (\frac{1}{1+\pi} \bar{\theta}_F, \frac{1}{1+\pi} \bar{\theta}]$, with strict inequality if $\theta_1 < \frac{1}{1+\pi} \bar{\theta}$. (We cannot however extend this to $\theta_1 \in [\frac{1}{1+\pi} \bar{\theta}_F, \frac{1}{1+\pi} \bar{\theta}]$, since we cannot deduce from the fact that $G_L > F(\bar{\theta})$ on $(\bar{\theta}_F, \bar{\theta})$ that $G_L \geq F(\bar{\theta})$ at $\bar{\theta}_F$.) Next, as in Lemma 24, we have $G_L(\theta_2) - G_L(\theta_1) > 0$ on $(\frac{1}{1+\pi} \bar{\theta}_F, \frac{1}{1+\pi} \bar{\theta}]$. Indeed, as in the proof of that lemma, we can show that

$$G_L(\theta_2) \geq F(\bar{\theta})$$

(with strict inequality if $\theta_1 \in (\frac{1}{1+\pi} \bar{\theta}_F, \frac{1}{1+\pi} \bar{\theta})$)

$$\geq \psi(\theta_1) = \Psi(\bar{\theta}; \theta_1) > \Psi(\theta_1; \theta_1) = G(\theta_1).$$

In particular, $G(\theta_1) < F(\bar{\theta})$; and therefore $\theta_1 < \bar{\theta}_F \leq \theta_M$; and therefore $G_L(\theta_1) \leq G(\theta_1)$. Finally, applying (26) yields the required result. ■

Next, we prove a lemma that will be needed for the third subcase.

Lemma 27. $\chi > G$ on $[0, \bar{\theta}_F)$.

Proof. For all $\theta_1 \in [0, \bar{\theta}_F)$, we have

$$\chi(\theta_1) = \frac{1}{\bar{\theta} - \theta_1} \int_{\theta_1}^{\bar{\theta}} G(\theta) d\theta = \frac{1}{\bar{\theta} - \theta_1} \int_{\theta_1}^{\bar{\theta}_F} G(\theta) d\theta + \frac{1}{\bar{\theta} - \theta_1} \int_{\bar{\theta}_F}^{\bar{\theta}} G(\theta) d\theta.$$

Moreover:

$$\int_{\theta_1}^{\bar{\theta}_F} G(\theta) d\theta \geq (\bar{\theta}_F - \theta_1) G(\theta_1),$$

since $G' \geq 0$ on $[0, \bar{\theta}_F]$ by Lemma 14; and

$$\int_{\bar{\theta}_F}^{\bar{\theta}} G(\theta) d\theta > (\bar{\theta} - \bar{\theta}_F) F(\bar{\theta}) \geq (\bar{\theta} - \bar{\theta}_F) G(\theta_1),$$

since $G > F(\bar{\theta})$ on $(\bar{\theta}_F, \bar{\theta})$ and (by Lemma 9) $G \leq F(\bar{\theta})$ on $(0, \bar{\theta}_F)$. Hence

$$\chi(\theta_1) > \frac{\bar{\theta}_F - \theta_1}{\bar{\theta} - \theta_1} G(\theta_1) + \frac{\bar{\theta} - \bar{\theta}_F}{\bar{\theta} - \theta_1} G(\theta_1) = G(\theta_1),$$

as required. ■

We can now deal with the third subcase, which arises only in the first scenario.

Lemma 28. *Suppose that $\frac{1}{1+\pi}\bar{\theta} < \bar{\theta}_F$ – i.e. that we are in the first scenario – and that $\theta_1 \in [\frac{1}{1+\pi}\bar{\theta}, \bar{\theta}_F)$. Then $\psi'(\theta_1) > 0$.*

Proof. Since $\theta_1 \geq \frac{1}{1+\pi}\bar{\theta}$, we may apply Lemma 21 to obtain

$$\psi'(\theta_1) = \frac{\chi(\theta_1) - G(\theta_1)}{\bar{\theta} - \theta_1}.$$

Since $\theta_1 < \bar{\theta}_F$, we may apply Lemma 27 to obtain $\chi(\theta_1) - G(\theta_1) > 0$. The result follows. ■

Combining Lemmas 22, 25, 26 and 28, we obtain:

Proposition 29. *The set of $\theta_1 \in (0, \bar{\theta}_F)$ such that $\psi(\theta_1) = F(\bar{\theta})$ is a closed interval. ■*

The idea behind the proof of the proposition is straightforward. We know from Proposition 12 that all solutions to the equation $\psi = F(\bar{\theta})$ lie in $(\frac{1}{1+\pi}\underline{\theta}, \bar{\theta}_F)$. Hence, to prove the proposition, we need only show that $\psi' \geq 0$ on this interval. Furthermore this is what Lemma 22 (for the interval $(\frac{1}{1+\pi}\underline{\theta}, \frac{1}{1+\pi}\bar{\theta}_F]$), Lemmas 25 and 26 (for the interval $(\frac{1}{1+\pi}\bar{\theta}_F, \frac{1}{1+\pi}\bar{\theta})$) and Lemma 28 (for the interval $[\frac{1}{1+\pi}\bar{\theta}, \bar{\theta}_F)$) seem to tell us. The only complication is that Lemmas 25 and 26 both require the side condition $\psi \leq F(\bar{\theta})$. However, they make up for this by providing strict rather than weak inequalities. The proof of the Proposition does therefore go through.

Indeed, we actually have $\psi' > 0$ on the interval $(\frac{1}{1+\pi}\bar{\theta}_F, \bar{\theta}_F)$. Hence, the only way in which non-uniqueness can occur at all is if there is a non-trivial interval, contained in $(\frac{1}{1+\pi}\underline{\theta}, \frac{1}{1+\pi}\bar{\theta}_F]$, on which $\psi = F(\bar{\theta})$. Unfortunately, it is possible to construct an example in which precisely this form of non-uniqueness occurs. The spirit of the example is that there exist $\theta_3, \theta_4 \in [\underline{\theta}, \bar{\theta}_F)$ such that: (i) $\theta_4 > (1 + \pi)\theta_3$ (i.e. it is possible that the entire interval of types associated with the kink lies within $[\theta_3, \theta_4)$); and (ii) $G_L(\theta_4) = G(\theta_3)$ (i.e. G is constant on $[\theta_3, \theta_4)$). It then follows that, if there exists $\theta_1 \in [\theta_3, \frac{1}{1+\pi}\theta_4]$ such that $\psi(\theta_1) = F(\bar{\theta})$, then $\psi(\theta_1) = F(\bar{\theta})$ for all $\theta_1 \in [\theta_3, \frac{1}{1+\pi}\theta_4]$.

There are two ways to eliminate this possibility. The first way is to ensure that G cannot have a “flat” of the type envisaged. The following assumption is more than sufficient to ensure this:

A4 G is strictly increasing on $[\underline{\theta}, \theta_M)$.¹³

We then have:

Proposition 30. *Suppose that Assumptions A1-A4 hold. Then there is a unique $\theta_1 \in (0, \bar{\theta}_F)$ such that $\psi(\theta_1) = F(\bar{\theta})$. ■*

Working with Assumption A4 certainly simplifies our comparative statics. See Sections 11-14 below. However, we can still obtain satisfactory comparative-statics results without it. See Section 15 below.

The second way to eliminate the possibility of non-uniqueness is ensure that G cannot have a long enough flat:

Proposition 31. *Suppose that Assumptions A1-A3 hold, and that $\pi > \frac{\bar{\theta}-\underline{\theta}}{\underline{\theta}}$. Then there is a unique $\theta_1 \in (0, \bar{\theta}_F)$ such that $\psi(\theta_1) = F(\bar{\theta})$. ■*

In particular, if $\pi = \infty$, then we there is a unique optimum within our one-parameter family of candidate optima.

11. COMPARATIVE STATICS WITH A4

The analysis of Sections 8-10 shows that, for all $\pi \in [0, \infty)$, the set of solutions of the equation

$$\psi(\theta_1, \pi) = F(\bar{\theta}) \tag{27}$$

is a non-empty interval. We denote this interval by $\tau(\pi) = [\underline{\tau}(\pi), \bar{\tau}(\pi)]$. The purpose of the current section is to investigate the dependence of τ on π .

In order to simplify the exposition, it will be helpful to assume for the time being that A4 holds. This ensures that the interval $\tau(\pi)$ collapses to a single point, which we shall denote by $\tau_1(\pi)$. It also ensures that $\frac{\partial \psi}{\partial \theta_1}(\tau_1(\pi), \pi) > 0$.

¹³Notice that G is identically 0 on $(0, \underline{\theta})$. It does not therefore make sense to require that G is strictly increasing on $(0, \theta_M)$.

ONLINE APPENDIX C TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE
MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

If we assume further that all the functions involved are sufficiently smooth, then we can apply the Implicit-Function Theorem to the equation

$$\psi(\tau_1(\pi), \pi) = F(\bar{\theta})$$

to conclude that

$$\tau'_1 = - \frac{\frac{\partial \psi}{\partial \pi}}{\frac{\partial \psi}{\partial \theta_1}}. \quad (28)$$

In particular: τ_1 will be increasing (decreasing) in π iff $\frac{\partial \psi}{\partial \pi} < 0$ ($\frac{\partial \psi}{\partial \pi} > 0$); and the allocation to the illiquid account will be increasing (decreasing) in π iff $\frac{\partial \psi}{\partial \pi} > 0$ ($\frac{\partial \psi}{\partial \pi} < 0$).

Motivated by these observations, we look first at the case in which the maximum-penalty constraint is strictly binding. More precisely, we put $\tau_2(\pi) = (1 + \pi) \tau_1(\pi)$, and we consider the case in which $\tau_2(\pi) < \bar{\theta}$. In other words, there is a non-trivial interval of types $(\tau_2(\pi), \bar{\theta})$ who choose to consume out of the illiquid account and therefore pay the penalty for doing so. In this case we begin by finding explicit formulae for $\frac{\partial \psi}{\partial \pi}$ and $\frac{\partial \psi}{\partial \theta_1}$. We then go on to find conditions under which $\frac{\partial \psi}{\partial \pi} > 0$ and $\frac{\partial \psi}{\partial \theta_1} > 0$, thereby ensuring that $\tau'_1 < 0$ (and hence that the allocation to the illiquid account will be strictly increasing in π).

We look second at the case in which the maximum-penalty constraint is strictly slack. More precisely, we consider the case in which $\tau_2(\pi) > \bar{\theta}$. In other words, even the highest type is not tempted to consume out of the illiquid account. In this case we again begin by finding explicit formulae for $\frac{\partial \psi}{\partial \pi}$ and $\frac{\partial \psi}{\partial \theta_1}$. It turns out that $\frac{\partial \psi}{\partial \pi} = 0$. We therefore concentrate on finding conditions under which $\frac{\partial \psi}{\partial \theta_1} > 0$, thereby ensuring that $\tau'_1 = 0$ (and hence that the allocation to the illiquid account does not change with π).

We look third at the intermediate case in which $\tau_2(\pi) = \bar{\theta}$. This case is important because it is $\tau_2(\pi)$ that determines whether we are in the strictly binding case $\tau_2(\pi) < \bar{\theta}$ or the strictly slack case $\tau_2(\pi) > \bar{\theta}$. Our analysis of the comparative statics of our problem is not therefore complete until we have understood how the transition between these two cases occurs.

12. THE STRICTLY BINDING CASE

In this section we focus on the set V of (θ_1, π) such that

1. $\theta_1 \in (0, \bar{\theta})$,
2. $\pi \in (0, \infty)$ and
3. $\theta_2 = (1 + \pi) \theta_1 < \bar{\theta}$.

In other words, we do not impose the requirement that $\theta_1 = \tau_1(\pi)$ (i.e. that θ_1 be optimal for the given π), but we do require that the maximum-penalty constraint is binding (in the sense that types in the non-empty interval $(\theta_2, \bar{\theta})$ are choosing to pay the penalty).

12.1. The formula for $\frac{\partial \psi}{\partial \pi}$. Consider the o.d.e.

$$\dot{\theta} = - \left(\theta + \beta b \left(\frac{\theta}{(1+\pi)\beta} \right) \right) \quad (29)$$

on $[\theta_2, \bar{\theta}]$, with initial condition $\theta(0) = \bar{\theta}$. Let $T(h; \pi)$ denote the first hitting time of $h \in [\theta_2, \bar{\theta}]$, and put $S(h; \pi) = \exp(-T(h; \pi))$. Then the formula for $\frac{\partial \psi}{\partial \pi}$ is given by the following proposition.

Proposition 32. *Suppose that $\theta_2 < \bar{\theta}$. Then*

$$\begin{aligned} \frac{\partial \psi}{\partial \pi}(\theta_1, \pi) &= \left(\frac{1}{\pi} S(\theta_2, \pi) - \theta_1 \frac{\partial S}{\partial h}(\theta_2, \pi) - \frac{\partial S}{\partial \pi}(\theta_2, \pi) \right) (G(\theta_2) - \phi(\theta_1, \pi)) \\ &\quad - \int_{(\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) dG(h). \end{aligned}$$

Proof. Equation (20) can be written

$$\psi(\theta_1, \pi) = \int_{[\theta_2, \bar{\theta}]} \frac{\partial S}{\partial h}(h, \pi) G(h) dh + S(\theta_2, \pi) \phi(\theta_1, \pi).$$

Hence

$$\begin{aligned} \frac{\partial \psi}{\partial \pi} &= \int_{[\theta_2, \bar{\theta}]} \frac{\partial^2 S}{\partial h \partial \pi}(h, \pi) G(h) dh - \frac{\partial S}{\partial h}(\theta_2, \pi) G(\theta_2) \frac{\partial \theta_2}{\partial \pi} \\ &\quad + \left(\frac{\partial S}{\partial h}(\theta_2, \pi) \frac{\partial \theta_2}{\partial \pi} + \frac{\partial S}{\partial \pi}(\theta_2, \pi) \right) \phi + S(\theta_2, \pi) \frac{\partial \phi}{\partial \pi}, \end{aligned} \quad (30)$$

where we have suppressed the dependence of ψ and ϕ on θ_1 and π . Now:

$$\begin{aligned} \int_{[\theta_2, \bar{\theta}]} \frac{\partial^2 S}{\partial h \partial \pi}(h, \pi) G(h) dh &= \int_{[\theta_2, \bar{\theta}]} \frac{\partial^2 S}{\partial \pi \partial h}(h, \pi) G(h) dh \\ &= \left[\frac{\partial S}{\partial \pi}(h, \pi) G(h) \right]_{\theta_2-}^{\bar{\theta}} - \int_{[\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) dG(h) \\ &= -\frac{\partial S}{\partial \pi}(\theta_2, \pi) G(\theta_2-) - \int_{[\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) dG(h) \\ &= -\frac{\partial S}{\partial \pi}(\theta_2, \pi) G(\theta_2) - \int_{(\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) dG(h), \end{aligned}$$

where we have used the fact that $\frac{\partial S}{\partial \pi}(\bar{\theta}) = 0$;

$$\frac{\partial \phi}{\partial \pi} = \frac{G(\theta_2) - \phi}{\theta_2 - \theta_1} \frac{\partial \theta_2}{\partial \pi} = \frac{G(\theta_2) - \phi}{\pi};$$

and

$$\frac{\partial \theta_2}{\partial \pi} = \theta_1.$$

Substituting into (30), we therefore obtain

$$\begin{aligned} \frac{\partial \psi}{\partial \pi} &= -\frac{\partial S}{\partial \pi}(\theta_2, \pi) G(\theta_2) - \int_{(\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) dG(h) - \frac{\partial S}{\partial h}(\theta_2, \pi) G(\theta_2) \theta_1 \\ &\quad + \left(\frac{\partial S}{\partial h}(\theta_2, \pi) \theta_1 + \frac{\partial S}{\partial \pi}(\theta_2, \pi) \right) \phi + S(\theta_2, \pi) \frac{G(\theta_2) - \phi}{\pi} \\ &= -\frac{\partial S}{\partial \pi}(\theta_2, \pi) (G(\theta_2) - \phi(\theta_1, \pi)) - \int_{(\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) dG(h) \\ &\quad + \left(\frac{1}{\pi} S(\theta_2, \pi) - \theta_1 \frac{\partial S}{\partial h}(\theta_2, \pi) \right) (G(\theta_2) - \phi(\theta_1, \pi)). \end{aligned}$$

The required formula now follows on rearranging. ■

In view of Proposition 32, it is clear that there are three main contributions to $\frac{\partial \psi}{\partial \pi}$, namely:

1. $\frac{1}{\pi} S(\theta_2, \pi) - \theta_1 \frac{\partial S}{\partial h}(\theta_2, \pi) - \frac{\partial S}{\partial \pi}(\theta_2, \pi);$
2. $G(\theta_2) - \phi(\theta_1, \pi);$
3. $-\int_{(\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) dG(h).$

We discuss these contributions in turn.

The first contribution can be signed quite generally:

Proposition 33. *Suppose that $\theta_2 \leq \bar{\theta}$. Then*

$$\frac{1}{\pi} S(\theta_2, \pi) - \theta_1 \frac{\partial S}{\partial h}(\theta_2, \pi) - \frac{\partial S}{\partial \pi}(\theta_2, \pi) > 0.$$

In other words, Contribution 1 is strictly positive.

Proof. Explicit calculation shows that

$$\frac{1}{\pi} S(\theta_2, \pi) - \theta_1 \frac{\partial S}{\partial h}(\theta_2, \pi) - \frac{\partial S}{\partial \pi}(\theta_2, \pi) = \frac{N}{D},$$

where

$$\begin{aligned} N = 1 + (1 + \pi) & \left(\frac{\theta_2}{\beta(1 + \pi)} \right)^{1/\rho} + \left(\frac{\bar{\theta}}{\beta(1 + \pi)} \right)^{1/\rho} \\ & + (1 + \pi) \left(\frac{\theta_2}{\beta(1 + \pi)} \right)^{1/\rho} \left(\frac{\bar{\theta}}{\beta(1 + \pi)} \right)^{1/\rho} \\ & + \rho \pi \left(\left(\frac{\bar{\theta}}{\beta(1 + \pi)} \right)^{1/\rho} - \left(\frac{\theta_2}{\beta(1 + \pi)} \right)^{1/\rho} \right). \end{aligned}$$

and

$$D = \pi \left(1 + (1 + \pi) \left(\frac{\theta_2}{\beta(1 + \pi)} \right)^{1/\rho} \right)^{1-\rho} \left(1 + (1 + \pi) \left(\frac{\bar{\theta}}{\beta(1 + \pi)} \right)^{1/\rho} \right)^{1+\rho}.$$

Now the last term in the formula for N is non-negative, since $\theta_2 \leq \bar{\theta}$. Hence $N > 0$. Finally, it is clear that $D > 0$. ■

The second contribution can only be signed when $\theta_1 = \tau_1(\pi)$ (or, more generally, when $\theta_2 \in (\underline{\theta}, \bar{\theta})$ and $\psi(\theta_1) \leq F(\bar{\theta})$). This, however, is enough for the purpose of our comparative statics.

Proposition 34. *Suppose that:*

1. $\theta_2 \in (\underline{\theta}, \bar{\theta})$;
2. $\psi(\theta_1) \leq F(\bar{\theta})$;
3. *Assumption A4 holds.*

Then $G(\theta_2) - \phi(\theta_1, \pi) > 0$. In other words, Contribution 2 is strictly positive.

Proof. We break the proof down into the cases $\theta_2 \in (\underline{\theta}, \bar{\theta}_F)$ and $\theta_2 \in [\bar{\theta}_F, \bar{\theta})$. In the first case, the proof parallels that of Lemma 22. We have $[\theta_1, \theta_2] \subset (\frac{1}{1+\pi}\underline{\theta}, \bar{\theta}_F) \subset (0, \theta_M)$ and $\theta_2 > \underline{\theta}$. Assumption A4 therefore implies that $G(\theta_2) > \phi(\theta_1, \pi)$. In the second case, it follows from the proof of Lemma 23 that $G(\theta_2) - \phi(\theta_1, \pi) > 0$. ■

The third contribution cannot be signed under our primary assumptions. It is, however, worth drawing attention to three special cases in which it can be signed. In all three cases, the comparative statics end up going the same way: $\frac{\partial \psi}{\partial \pi} > 0$, and therefore the allocation to the illiquid account will be increasing in π . We state these three cases as separate propositions, corresponding to the cases $\rho < 1$, $\rho = 1$ and $\rho > 1$.

Proposition 35. *Suppose that:*

1. $\rho < 1$;
2. $G' \leq 0$ on $(\underline{\theta}, \infty)$ (i.e. $\theta_M = \underline{\theta}$);
3. $\theta_2 \in (\underline{\theta}, \bar{\theta})$.

Then $-\int_{(\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) dG(h) \geq 0$. In other words, Contribution 3 is non-negative.

Proof. It is easy to show that we have

$$\frac{\partial S}{\partial \pi}(h, \pi) \begin{cases} > 0 & \text{if } \rho < 1 \\ = 0 & \text{if } \rho = 1 \\ < 0 & \text{if } \rho > 1 \end{cases}$$

for all $h \in [\theta_2, \bar{\theta}]$. Furthermore, we have

$$\frac{\partial S}{\partial \pi}(\bar{\theta}, \pi) = 0 \quad \text{for all } \rho,$$

because $S(\bar{\theta}, \pi) = 1$. We can therefore proceed as follows.

First, we know that $\theta_2 \in (\underline{\theta}, \bar{\theta})$. Hence $G' \leq 0$ on $(\theta_2, \bar{\theta}] \subset (\underline{\theta}, \infty)$. Second, $\rho < 1$. Hence $\frac{\partial S}{\partial \pi}(\cdot, \pi) \geq 0$ on $(\theta_2, \bar{\theta}] \subset [\theta_2, \bar{\theta}]$. Putting these two observations together gives us the required conclusion. ■

Remark 36. If $G' \leq 0$ on $(\underline{\theta}, \infty)$ then necessarily $\Delta G(\underline{\theta}) > 0$. Hence it is essential for the proof of Proposition 35 that we restrict attention to $\theta_2 > \underline{\theta}$.

Proposition 37. Suppose that:

1. $\rho = 1$;
2. $\theta_2 \in (0, \bar{\theta})$.

Then $-\int_{(\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) dG(h) = 0$. In other words, Contribution 3 is zero.

Proof. This follows at once from the fact that $\frac{\partial S}{\partial \pi}(\cdot, \pi) = 0$ on $[\theta_2, \bar{\theta}]$. ■

Proposition 38. Suppose that:

1. $\rho > 1$;
2. $G' \geq 0$ on $(0, \bar{\theta})$ (i.e. $\theta_M = \bar{\theta}$);
3. $\theta_2 \in (0, \bar{\theta})$.

Then $-\int_{(\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) dG(h) \geq 0$. In other words, Contribution 3 is non-negative.

Proof. Note first that $G' \geq 0$ on $(\theta_2, \bar{\theta}) \subset (0, \bar{\theta})$. Second, $\rho > 1$. Hence $\frac{\partial S}{\partial \pi}(\cdot, \pi) < 0$ on $(\theta_2, \bar{\theta}) \subset [\theta_2, \bar{\theta}]$. (Cf. the proof of Proposition 35.) Third, $\frac{\partial S}{\partial \pi}(\bar{\theta}, \pi) = 0$. Putting these three observations together, we obtain

$$\begin{aligned} \int_{(\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) dG(h) &= \int_{(\theta_2, \bar{\theta})} \frac{\partial S}{\partial \pi}(h, \pi) dG(h) + \int_{[\bar{\theta}, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) dG(h) \\ &= \int_{(\theta_2, \bar{\theta})} \frac{\partial S}{\partial \pi}(h, \pi) dG(h) \leq 0, \end{aligned}$$

as required. ■

Remark 39. If $G' \geq 0$ on $(0, \bar{\theta})$ then necessarily $\Delta G(\bar{\theta}) < 0$. The fact that $\frac{\partial S}{\partial \pi}(\bar{\theta}, \pi) = 0$ therefore plays an essential role in the proof of Proposition 38.

12.2. The formula for $\frac{\partial \psi}{\partial \theta_1}$. Let $T(h; \pi)$ and $S(h; \pi) = \exp(-T(h; \pi))$ be defined as in the preceding section. Then the formula for $\frac{\partial \psi}{\partial \theta_1}$ is given by the following proposition.

Proposition 40. Suppose that $\theta_2 < \bar{\theta}$. Then

$$\begin{aligned} \frac{\partial \psi}{\partial \theta_1}(\theta_1, \pi) &= \left(\frac{\beta b\left(\frac{\theta_1}{\beta}\right)}{\theta_1 \left(\theta_2 + \beta b\left(\frac{\theta_1}{\beta}\right)\right)} (G(\theta_2) - \phi(\theta_1, \pi)) \right. \\ &\quad \left. + \frac{1}{\pi \theta_1} (G(\theta_2) - G(\theta_1)) \right) S(\theta_2, \bar{\theta}). \end{aligned}$$

Proof. This is simply a restatement of Lemma 20. ■

In view of Proposition 40, there are two main contributions to $\frac{\partial \psi}{\partial \theta_1}$, namely

1. $G(\theta_2) - \phi(\theta_1, \pi)$;
2. $G(\theta_2) - G(\theta_1)$.

We have already given conditions under which the first is strictly positive (in Proposition 34). The second is strictly positive under the same conditions:

Proposition 41. *Suppose that:*

1. $\theta_2 \in (\underline{\theta}, \bar{\theta})$;
2. $\psi(\theta_1) \leq F(\bar{\theta})$;
3. *Assumption A4 holds.*

Then $G(\theta_2) - G(\theta_1) > 0$.

Proof. We break the proof down into the cases $\theta_1 \in (\frac{1}{1+\pi}\underline{\theta}, \frac{1}{1+\pi}\bar{\theta}_F)$ and $\theta_1 \in [\frac{1}{1+\pi}\bar{\theta}_F, \frac{1}{1+\pi}\bar{\theta})$. In the first case, the proof parallels that of Lemma 22. We have $[\theta_1, \theta_2] \subset (\frac{1}{1+\pi}\bar{\theta}_F, \bar{\theta}_F) \subset (0, \theta_M)$, and moreover $\theta_2 > \underline{\theta}$. Assumption A4 therefore implies that $G(\theta_2) > G(\theta_1)$. In the second case, Lemma 24 implies directly that that $G(\theta_2) > G(\theta_1)$. ■

13. THE STRICTLY SLACK CASE

In this section we focus on the set W of (θ_1, π) such that

1. $\theta_1 \in (0, \bar{\theta})$,
2. $\pi \in (0, \infty)$ and
3. $\theta_2 = (1 + \pi)\theta_1 > \bar{\theta}$.

In other words, we do not impose the requirement that $\theta_1 = \tau_1(\pi)$ (i.e. that θ_1 be optimal for the given π), but we do require that the maximum-penalty constraint is slack in the sense that no types are choosing to pay the penalty.

13.1. The formula for $\frac{\partial \psi}{\partial \pi}$. The formula for $\frac{\partial \psi}{\partial \pi}$ is given by the following proposition.

Proposition 42. *Suppose that $\theta_2 > \bar{\theta}$. Then*

$$\frac{\partial \psi}{\partial \pi}(\theta_1, \pi) = 0.$$

Proof. As in the proof of Lemma 21, we have $\psi(\theta_1, \pi) = \chi(\theta_1)$ for $\theta_1 \in (\frac{1}{1+\pi}\bar{\theta}, \bar{\theta})$, where

$$\chi(\theta_1) = \frac{1}{\bar{\theta} - \theta_1} \int_{\theta_1}^{\bar{\theta}} G(\theta) d\theta.$$

Hence ψ is independent of π for such θ_1 . ■

13.2. The formula for $\frac{\partial \psi}{\partial \theta_1}$. The formula for $\frac{\partial \psi}{\partial \theta_1}$ is given by the following proposition.

Proposition 43. *Suppose that $\theta_2 > \bar{\theta}$. Then*

$$\frac{\partial \psi}{\partial \theta_1}(\theta_1, \pi) = \frac{\chi(\theta_1) - G(\theta_1)}{\bar{\theta} - \theta_1}.$$

Proof. As already noted, we have $\psi(\theta_1, \pi) = \chi(\theta_1)$ for $\theta_1 \in (\frac{1}{1+\pi}\bar{\theta}, \bar{\theta})$. Moreover

$$\chi'(\theta_1) = \frac{\chi(\theta_1) - G(\theta_1)}{\bar{\theta} - \theta_1},$$

as in the proof of Lemma 21. ■

In view of Proposition 43, there is really only one contribution to $\frac{\partial \psi}{\partial \theta_1}$, namely $\chi(\theta_1) - G(\theta_1)$. It is not possible to sign $\chi(\theta_1) - G(\theta_1)$ for all θ_1 , but it is possible to sign it when $\theta_1 = \tau_1(\pi)$, and indeed much more generally when $\theta_1 \in (0, \bar{\theta}_F)$. As before, this is enough for the purpose of our comparative statics.

Proposition 44. *Suppose that $\theta_1 \in (0, \bar{\theta}_F)$. Then $\chi(\theta_1) - G(\theta_1) > 0$.*

Proof. This is simply a special case of Lemma 27. ■

14. THE INTERMEDIATE CASE

Up to now we have focussed on the comparative statics of τ_1 . For example, we have shown that if A4 is satisfied and $\rho = 1$ then: (i) $\tau_1' < 0$ when $\tau_2(\pi) < \bar{\theta}$; and (ii) $\tau_1' = 0$ when $\tau_2(\pi) > \bar{\theta}$. However, this leaves open the question of what happens at the transition between the two cases. For example, does τ_1 jump up when $\tau_2(\pi) = \bar{\theta}$? Does it jump down? Or is there more than one value of π for which $\tau_2(\pi) = \bar{\theta}$?

In order to address these questions, we need to understand the comparative statics of τ_2 . These comparative statics are quite complex in the binding case. However, they simplify as the borderline between the two cases is approached. Moreover they are simpler still in the slack case.

14.1. Comparative Statics of τ_2 in the Weakly Binding Case. We begin this section by looking at the comparative statics of τ_2 when the maximum-penalty constraint is strictly binding (in the sense that $\tau_2(\pi) < \bar{\theta}$). More precisely, we show that $\tau_2'(\pi)$ satisfies a simple linear equation. We then go on to check whether this equation remains valid when the maximum-penalty constraint is only weakly binding (in the sense that $\tau_2(\pi) \uparrow \bar{\theta}$).

Proposition 45. *Suppose that $\tau_2(\pi) < \bar{\theta}$. Then*

$$D(\tau_1(\pi), \pi) \tau_2'(\pi) = N(\tau_1(\pi), \pi), \quad (31)$$

where

$$D(\theta_1, \pi) = (1 + \pi) \frac{\theta_1 + \beta b \left(\frac{\theta_1}{\beta} \right)}{\theta_2 + \beta b \left(\frac{\theta_1}{\beta} \right)} (G(\theta_2) - \phi(\theta_1, \pi)) + (\phi(\theta_1, \pi) - G(\theta_1)),$$

$$N(\theta_1, \pi) = \left((\phi(\theta_1, \pi) - G(\theta_1)) + \frac{\pi(1 + \pi)}{S(\theta_2, \pi)} \int_{[\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h; \pi) d\tilde{G}(h) \right) \theta_1$$

and $\tilde{G} \in \mathcal{BV}([\theta_2, \bar{\theta}], \mathbb{R})$ is given by the formulae $\tilde{G}_L(\theta_2) = \phi(\theta_1, \pi)$ and $\tilde{G} = G$ on $(\theta_2, \bar{\theta}]$.

Proof. We have

$$\tau_2(\pi) = (1 + \pi) \tau_1(\pi)$$

and therefore

$$\tau_2'(\pi) = \tau_1(\pi) + (1 + \pi) \tau_1'(\pi). \quad (32)$$

Now,

$$\frac{\partial \psi}{\partial \theta_1}(\tau_1(\pi), \pi) \tau_1'(\pi) + \frac{\partial \psi}{\partial \pi}(\tau_1(\pi), \pi) = 0.$$

Hence, multiplying (32) through by $\frac{\partial \psi}{\partial \theta_1}(\tau_1(\pi), \pi)$, we obtain

$$\begin{aligned} \frac{\partial \psi}{\partial \theta_1} \tau_2' &= \frac{\partial \psi}{\partial \theta_1} \tau_1 + (1 + \pi) \frac{\partial \psi}{\partial \theta_1} \tau_1' \\ &= \frac{\partial \psi}{\partial \theta_1} \tau_1 - (1 + \pi) \frac{\partial \psi}{\partial \pi}, \end{aligned}$$

where we have suppressed the dependence of $\frac{\partial \psi}{\partial \theta_1}$ and $\frac{\partial \psi}{\partial \pi}$ on $\tau_1(\pi)$ and π , and the dependence of τ_1 and τ_2 on π . We may therefore put

$$D(\theta_1, \pi) = \frac{\pi \theta_1}{S(\theta_2, \bar{\theta})} \frac{\partial \psi}{\partial \theta_1}(\theta_1, \pi)$$

and

$$N(\theta_1, \pi) = \frac{\pi \theta_1}{S(\theta_2, \bar{\theta})} \left(\frac{\partial \psi}{\partial \theta_1}(\theta_1, \pi) \theta_1 - (1 + \pi) \frac{\partial \psi}{\partial \pi}(\theta_1, \pi) \right).$$

Equation (31) now follows on applying the formulae for $\frac{\partial \psi}{\partial \pi}(\theta_1, \pi)$ and $\frac{\partial \psi}{\partial \theta_1}(\theta_1, \pi)$ given in Propositions 32 and 40. ■

Equation (31) can be solved for $\tau_2'(\pi)$ under the conditions of Proposition 34, namely that: (i) $\tau_2(\pi) \in (\underline{\theta}, \bar{\theta})$; (ii) $\psi(\tau_1(\pi)) \leq F(\bar{\theta})$; and (iii) Assumption A4 holds. This is not, however, enough for our current purposes: we need to make sure that it can still be solved for $\tau_2'(\pi)$ in the limiting case $\tau_2(\pi) \uparrow \bar{\theta}$. To this end, recall that

$$V = \{(\theta_1, \pi) \mid \theta_1 \in (0, \bar{\theta}), \pi \in (0, \infty), \theta_2 < \bar{\theta}\},$$

and put

$$\partial V = \{(\theta_1, \pi) \mid \theta_1 \in (0, \bar{\theta}), \pi \in (0, \infty), \theta_2 = \bar{\theta}\}.$$

Furthermore, for all $(\theta_1, \pi) \in V \cup \partial V$, put

$$\begin{aligned} \underline{D}(\theta_1, \pi) &= (1 + \pi) \frac{\theta_1 + \beta b\left(\frac{\theta_1}{\beta}\right)}{\theta_2 + \beta b\left(\frac{\theta_1}{\beta}\right)} (G_L(\theta_2) - \phi(\theta_1, \pi)) \\ &\quad + (\phi(\theta_1, \pi) - \max\{G(\theta_1), G_L(\theta_1)\}) \end{aligned}$$

and

$$\begin{aligned} \overline{D}(\theta_1, \pi) &= (1 + \pi) \frac{\theta_1 + \beta b\left(\frac{\theta_1}{\beta}\right)}{\theta_2 + \beta b\left(\frac{\theta_1}{\beta}\right)} (G_L(\theta_2) - \phi(\theta_1, \pi)) \\ &\quad + (\phi(\theta_1, \pi) - \min\{G(\theta_1), G_L(\theta_1)\}). \end{aligned}$$

Then we have:

Lemma 46. *Suppose that $(\tilde{\theta}_1, \tilde{\pi}) \in V \rightarrow (\theta_1, \pi) \in \partial V$. Then*

$$\underline{D}(\theta_1, \pi) \leq \liminf D(\tilde{\theta}_1, \tilde{\pi}) \leq \limsup D(\tilde{\theta}_1, \tilde{\pi}) \leq \overline{D}(\theta_1, \pi).$$

Proof. Note first that b and ϕ are both continuous. Hence $b\left(\frac{\tilde{\theta}_1}{\beta}\right) \rightarrow b\left(\frac{\theta_1}{\beta}\right)$ and $\phi(\tilde{\theta}_1, \tilde{\pi}) \rightarrow \phi(\theta_1, \pi)$. Next, put $\tilde{\theta}_2 = (1 + \pi)\tilde{\theta}_1$ and $\theta_2 = (1 + \pi)\theta_1$. Then $\tilde{\theta}_2 \uparrow \theta_2$, and therefore $G(\tilde{\theta}_2) \rightarrow G_L(\theta_2)$. Finally,

$$\begin{aligned} \min\{G(\theta_1), G_L(\theta_1)\} &\leq \liminf G(\tilde{\theta}_1) \\ &\leq \limsup G(\tilde{\theta}_1) \\ &\leq \max\{G(\theta_1), G_L(\theta_1)\}. \end{aligned}$$

The result follows. ■

The next step is to sign $\underline{D}$. This cannot be done everywhere on $V \cup \partial V$. But it can be done when $\theta_2 = \bar{\theta}$ and $\theta_1 = \tau_1(\pi)$. Indeed, it is enough to require that $\theta_2 \in (\bar{\theta}_F, \bar{\theta}]$ (i.e. we do not actually have to be on the boundary) and that $\psi(\theta_1, \pi) \leq F(\bar{\theta})$ (i.e.

we do not actually have to be at an optimum). We begin with a lemma.

Lemma 47. *Suppose that:*

1. $\theta_2 \in (\bar{\theta}_F, \bar{\theta}]$;
2. $\psi(\theta_1) \leq F(\bar{\theta})$.

Then $G_L(\theta_2) > G(\theta_1) \geq G_L(\theta_1)$.

Proof. The proof is similar to that of Lemma 24. Note first that

$$G_L(\theta_2) \geq F(\bar{\theta})$$

(with strict inequality if $\theta_2 < \bar{\theta}$)

$$\geq \psi(\theta_1) = \Psi(\bar{\theta}; \theta_1) > \Psi(\theta_1; \theta_1) = G(\theta_1)$$

(by assumption, by definition of ψ , because $\theta_1 < \theta_2 \leq \bar{\theta}$ and by construction of Ψ respectively). Second, Lemma 10 tells us that $\psi > F(\bar{\theta})$ on $[\bar{\theta}_F, \bar{\theta})$. But we have $\psi(\theta_1) \leq F(\bar{\theta})$. Hence $\theta_1 < \bar{\theta}_F$ and therefore $G' \geq 0$ at $\theta_1 \in (0, \bar{\theta}_F) \subset (0, \theta_M)$. That is, $G(\theta_1) - G_L(\theta_1) = \Delta G(\theta_1) \geq 0$. ■

We can now sign $\underline{D}$.

Proposition 48. *Suppose that:*

1. $\theta_2 \in (\bar{\theta}_F, \bar{\theta}]$;
2. $\psi(\theta_1) \leq F(\bar{\theta})$.

Then $\underline{D}(\theta_1, \pi) > 0$.

Proof. Two things follow from Lemma 47. First, $G(\theta_1) \geq G_L(\theta_1)$. Hence the formula for $\underline{D}(\theta_1, \pi)$ simplifies to

$$\underline{D}(\theta_1, \pi) = (1 + \pi) \frac{\theta_1 + \beta b \left(\frac{\theta_1}{\beta} \right)}{\bar{\theta} + \beta b \left(\frac{\theta_1}{\beta} \right)} (G_L(\theta_2) - \phi(\theta_1, \pi)) + (\phi(\theta_1, \pi) - G(\theta_1)).$$

In particular, $\underline{D}(\theta_1, \pi)$ is a strictly positive linear combination of the two terms $G_L(\theta_2) - \phi(\theta_1, \pi)$ and $\phi(\theta_1, \pi) - G(\theta_1)$. Second, $G_L(\theta_2) - G(\theta_1) > 0$. Hence the sum of the two terms $G_L(\theta_2) - \phi(\theta_1, \pi)$ and $\phi(\theta_1, \pi) - G(\theta_1)$ is strictly positive. It therefore suffices to show that each of these two terms is non-negative. We have

$$G_L(\theta_2) \geq F(\bar{\theta}) \geq \psi(\theta_1) = \Psi(\bar{\theta}; \theta_1)$$

(as in the proof of Lemma 47)

$$\geq \Psi(\theta_2; \theta_1) > \Psi(\theta_1; \theta_1)$$

(since $\Psi' \geq 0$ on $(\theta_1, \bar{\theta}_F)$ (by Proposition 18) and $\Psi' > 0$ on $[\bar{\theta}_F, \bar{\theta})$ (by Proposition 19))

$$= G(\theta_1)$$

(again as in the proof of Lemma 47). In particular, since $\Psi(\theta_2; \theta_1) = \phi(\theta_1, \pi)$, we have $G_L(\theta_2) \geq \phi(\theta_1, \pi)$ and $\phi(\theta_1, \pi) > G(\theta_1)$. ■

Since $\underline{D} > 0$, finding the sign of N and finding the sign of $\tau'_2(\pi)$ amount to the same thing. Note first that

$$\tau_2(\pi) = (1 + \pi) \tau_1(\pi)$$

and hence

$$\tau'_2(\pi) = (1 + \pi) \tau'_1(\pi) + \tau_1(\pi).$$

We therefore face a tension. On the one hand, we are mainly interested in the case in which $\tau'_1(\pi) < 0$. For our purposes, then, the first contribution to $\tau'_2(\pi)$ (namely $(1 + \pi) \tau'_1(\pi)$) is negative. However, the second contribution (namely $\tau_1(\pi)$) is necessarily positive. The net effect is therefore ambiguous. Worse still, what we really need to show for the purposes of comparative statics is that $\tau'_2(\pi) > 0$ (so that the curve $(\tau_1(\pi), \pi)$ crosses the boundary $\theta_2 = \bar{\theta}$ in a simple way). This is directly at odds with our interest in the case in which $\tau'_1(\pi) < 0$.

Fortunately, the problem of signing $\tau'_2(\pi)$ at the boundary is much simpler than the problem of signing $\tau'_2(\pi)$ in V . With this end in mind, for all $(\theta_1, \pi) \in \partial V$, put

$$\underline{N}(\theta_1, \pi) = (\phi(\theta_1, \pi) - \max \{G(\theta_1), G_L(\theta_1)\}) \theta_1$$

and

$$\overline{N}(\theta_1, \pi) = (\phi(\theta_1, \pi) - \min \{G(\theta_1), G_L(\theta_1)\}) \theta_1.$$

Then we have the following lemma.

Lemma 49. *Suppose that $(\tilde{\theta}_1, \tilde{\pi}) \in V \rightarrow (\theta_1, \pi) \in \partial V$. Then*

$$\underline{N}(\theta_1, \pi) \leq \liminf N(\tilde{\theta}_1, \tilde{\pi}) \leq \limsup N(\tilde{\theta}_1, \tilde{\pi}) \leq \overline{N}(\theta_1, \pi).$$

Proof. The proof is similar to that of Lemma 46. Put $\tilde{\theta}_2 = (1 + \pi) \tilde{\theta}_1$ and $\theta_2 = (1 + \pi) \theta_1$. Then $\tilde{\theta}_2 \uparrow \theta_2 = \bar{\theta}$, and therefore $\frac{\partial S}{\partial \pi}(\tilde{\theta}_2, \pi) \rightarrow 0$ and $\int_{(\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h; \pi) dG(h) \rightarrow 0$. Furthermore $\phi(\tilde{\theta}_1, \tilde{\pi}) \rightarrow \phi(\theta_1, \pi)$ and

$$\begin{aligned} \min \{G(\theta_1), G_L(\theta_1)\} &\leq \liminf G(\tilde{\theta}_1) \\ &\leq \limsup G(\tilde{\theta}_1) \\ &\leq \max \{G(\theta_1), G_L(\theta_1)\}. \end{aligned}$$

Passing to the limit in the formula given for N in the statement of Proposition 45, we therefore obtain the required result. ■

Combining Lemma 49 with the earlier Lemma 47, we obtain:

Proposition 50. *Suppose that:*

1. $\theta_2 = \bar{\theta}$;
2. $\psi(\theta_1) \leq F(\bar{\theta})$.

Then $\underline{N}(\theta_1, \pi) > 0$.

Proof. The proof is similar to that of Proposition 48. First, because $\psi(\theta_1) \leq F(\bar{\theta})$ and therefore $\theta_1 < \bar{\theta}_F$, the formula for $\underline{N}(\theta_1, \pi)$ simplifies to

$$\underline{N}(\theta_1, \pi) = (\phi(\theta_1, \pi) - G(\theta_1)) \theta_1.$$

Second, we have

$$\Psi(\bar{\theta}; \theta_1) > \Psi(\theta_1; \theta_1) = G(\theta_1).$$

It remains only to note that, because $\theta_2 = \bar{\theta}$, we have $\phi(\theta_1, \pi) = \Psi(\bar{\theta}; \theta_1)$. ■

Combining Propositions 48 and 50, we see that $\tau'_2(\pi) > 0$ on ∂V . In other words, whatever the behaviour of the curve $(\tau_1(\pi), \pi)$ in V , it points out of V at ∂V . I.e. it can exit, but not enter, V at ∂V . In particular, there exists $\pi_1 \in (0, \infty)$ such that $\tau_2(\pi) < \bar{\theta}$ iff $\pi \in [0, \pi_1)$.

14.2. Comparative Statics of τ_2 in the Weakly Slack Case. We begin this section by looking at the comparative statics of τ_2 when the maximum-penalty constraint is strictly slack (in the sense that $\tau_2(\pi) > \bar{\theta}$). More precisely, we show that $\tau'_2(\pi)$ satisfies a simple linear equation. We then go on to check whether this equation remains valid when the maximum-penalty constraint is only weakly slack (in the sense that $\tau_2(\pi) \downarrow \bar{\theta}$). Our first proposition is analogous to Proposition 45.

Proposition 51. *Suppose that $\tau_2(\pi) > \bar{\theta}$. Then*

$$D(\tau_1(\pi), \pi) \tau'_2(\pi) = N(\tau_1(\pi), \pi), \tag{33}$$

where

$$D(\theta_1, \pi) = \frac{\chi(\theta_1) - G(\theta_1)}{\bar{\theta} - \theta_1}$$

and

$$N(\theta_1, \pi) = \frac{\chi(\theta_1) - G(\theta_1)}{\bar{\theta} - \theta_1} \theta_1.$$

Notice that, if $\chi(\theta_1) - G(\theta_1) > 0$, then we can divide through by $D(\tau_1(\pi), \pi)$ to conclude that $\tau'_2(\pi) = \theta_1$. Furthermore $\chi(\theta_1) - G(\theta_1) > 0$ if $\theta_1 = \tau_1(\pi)$, and indeed much more generally if $\theta_1 \in (0, \bar{\theta}_F)$. Cf. Proposition 44. But it does not hold for all $(\theta_1, \pi) \in W$.

Proof. As in the proof of Proposition 45, we have

$$\frac{\partial \psi}{\partial \theta_1} \tau'_2 = \frac{\partial \psi}{\partial \theta_1} \tau_1 - (1 + \pi) \frac{\partial \psi}{\partial \pi}.$$

We may therefore put

$$D(\theta_1, \pi) = \frac{\partial \psi}{\partial \theta_1}(\theta_1, \pi)$$

and

$$N(\theta_1, \pi) = \frac{\partial \psi}{\partial \theta_1}(\theta_1, \pi) \theta_1 - (1 + \pi) \frac{\partial \psi}{\partial \pi}(\theta_1, \pi).$$

Equation (33) now follows on applying the formulae for $\frac{\partial \psi}{\partial \pi}(\theta_1, \pi)$ and $\frac{\partial \psi}{\partial \theta_1}(\theta_1, \pi)$ given in Propositions 42 and 43. ■

The next step is to ensure that equation (33) can still be solved for $\tau'_2(\pi)$ in the limiting case $\tau_2(\pi) \downarrow \bar{\theta}$. To this end, recall that

$$W = \{(\theta_1, \pi) \mid \theta_1 \in (0, \bar{\theta}), \pi \in (0, \infty), \theta_2 > \bar{\theta}\},$$

and put

$$\partial W = \{(\theta_1, \pi) \mid \theta_1 \in (0, \bar{\theta}), \pi \in (0, \infty), \theta_2 = \bar{\theta}\}.$$

Furthermore, for all $(\theta_1, \pi) \in W \cup \partial W$, put

$$\begin{aligned} \underline{D}(\theta_1, \pi) &= \frac{\chi(\theta_1) - \max\{G(\theta_1), G_L(\theta_1)\}}{\bar{\theta} - \theta_1}, \\ \overline{D}(\theta_1, \pi) &= \frac{\chi(\theta_1) - \min\{G(\theta_1), G_L(\theta_1)\}}{\bar{\theta} - \theta_1} \end{aligned}$$

and

$$\begin{aligned} \underline{N}(\theta_1, \pi) &= \frac{\chi(\theta_1) - \max\{G(\theta_1), G_L(\theta_1)\}}{\bar{\theta} - \theta_1} \theta_1, \\ \overline{N}(\theta_1, \pi) &= \frac{\chi(\theta_1) - \min\{G(\theta_1), G_L(\theta_1)\}}{\bar{\theta} - \theta_1} \theta_1. \end{aligned}$$

Then we have:

Lemma 52. *Suppose that $(\tilde{\theta}_1, \tilde{\pi}) \in W \rightarrow (\theta_1, \pi) \in \partial W$. Then*

$$\underline{D}(\theta_1, \pi) \leq \liminf D(\tilde{\theta}_1, \tilde{\pi}) \leq \limsup D(\tilde{\theta}_1, \tilde{\pi}) \leq \overline{D}(\theta_1, \pi)$$

and

$$\underline{N}(\theta_1, \pi) \leq \liminf N(\tilde{\theta}_1, \tilde{\pi}) \leq \limsup N(\tilde{\theta}_1, \tilde{\pi}) \leq \overline{N}(\theta_1, \pi).$$

Proof. Note first that χ is continuous. Hence $\chi(\tilde{\theta}_1) \rightarrow \chi(\theta_1)$. On the other hand, as in the proof of Lemma 46,

$$\begin{aligned} \min \{G(\theta_1), G_L(\theta_1)\} &\leq \liminf G(\tilde{\theta}_1) \\ &\leq \limsup G(\tilde{\theta}_1) \\ &\leq \max \{G(\theta_1), G_L(\theta_1)\}. \end{aligned}$$

The result follows. ■

The next step is to sign $\underline{D}$. This cannot be done everywhere on $W \cup \partial W$. But it can be done when $\theta_2 = \bar{\theta}$ and $\theta_1 = \tau_1(\pi)$. Indeed, it is enough to require that $\theta_2 \in [\bar{\theta}, \infty)$ (i.e. we do not actually have to be on the boundary) and that $\psi(\theta_1, \pi) \leq F(\bar{\theta})$ (i.e. we do not actually have to be at an optimum). We begin with a lemma.

Lemma 53. *Suppose that:*

1. $\theta_2 \in [\bar{\theta}, \infty)$;
2. $\psi(\theta_1) \leq F(\bar{\theta})$.

Then $G(\theta_1) \geq G_L(\theta_1)$.

Proof. The proof is identical to the relevant part of that of Lemma 47. Since $\psi(\theta_1) \leq F(\bar{\theta})$, we must have $\theta_1 < \bar{\theta}_F$. Hence $G' \geq 0$ at $\theta_1 \in (0, \bar{\theta}_F) \subset (0, \theta_M)$. ■

Proposition 54. *Suppose that:*

1. $\theta_2 \in [\bar{\theta}, \infty)$;

2. $\psi(\theta_1) \leq F(\bar{\theta})$.

Then $\underline{D}(\theta_1, \pi), \underline{N}(\theta_1, \pi) > 0$.

Proof. Note first that, in view of Lemma 53, we have

$$\underline{D}(\theta_1, \pi) = \frac{\chi(\theta_1) - G(\theta_1)}{\bar{\theta} - \theta_1}$$

and

$$\underline{N}(\theta_1, \pi) = \frac{\chi(\theta_1) - G(\theta_1)}{\bar{\theta} - \theta_1} \theta_1.$$

Second, since $\psi(\theta_1) \leq F(\bar{\theta})$, we have $\theta_1 < \bar{\theta}_F$. Finally, Proposition 44 tells us that $\chi(\theta_1) - G(\theta_1) > 0$ for $\theta_1 \in (0, \bar{\theta}_F)$. ■

It follows from Proposition 54 that $\tau'_2(\pi) > 0$ on ∂W . In other words, whatever the behaviour of the curve $(\tau_1(\pi), \pi)$ in W , it points into W at ∂W . I.e. it can enter, but not exit, W at ∂W . In particular, there exists $\pi_2 \in (0, \infty)$ such that $\tau_2(\pi) > \bar{\theta}$ iff $\pi \in (\pi_2, \infty)$.

14.3. Comparative Statics of τ_2 in the Remaining Case. At this point we have established that there exist $0 < \pi_1 \leq \pi_2 < \infty$ such that $\tau_2(\pi) < \bar{\theta}$ iff $\pi \in [0, \pi_1)$ and $\tau_2(\pi) > \bar{\theta}$ iff $\pi \in (\pi_2, \infty)$.¹⁴ The remaining question is therefore whether it is possible that $\pi_1 < \pi_2$, in other words that there is a non-trivial interval (π_1, π_2) over which $\tau_2(\pi) = \bar{\theta}$.

Suppose for a contradiction that there is such an interval. Then, over this interval, we must have both

$$\psi(\tau_1(\pi), \pi) = F(\bar{\theta}) \tag{34}$$

(because $\tau_1(\pi)$ is the optimal θ_1) and

$$\tau_2(\pi) = \bar{\theta}. \tag{35}$$

Hence

$$F(\bar{\theta}) = \psi(\tau_1(\pi), \pi) = \Psi(\bar{\theta}; \tau_1(\pi), \pi) = \Psi(\tau_2(\pi); \tau_1(\pi), \pi)$$

¹⁴That $\pi_1 \leq \pi_2$ follows at once from the fact that we cannot have $\tau_2(\pi) < \bar{\theta}$ and $\tau_2(\pi) > \bar{\theta}$ simultaneously.

(by equation (34), by definition of ψ and by equation (35))

$$= \phi(\tau_1(\pi); \pi) = \frac{1}{\tau_2(\pi) - \tau_1(\pi)} \int_{\tau_1(\pi)}^{\tau_2(\pi)} G(\theta) d\theta = \frac{1}{\bar{\theta} - \tau_1(\pi)} \int_{\tau_1(\pi)}^{\bar{\theta}} G(\theta) d\theta$$

(by construction of Ψ , by definition of ϕ , by equation (35) again). Multiplying through by $\bar{\theta} - \tau_1(\pi)$, we therefore obtain

$$\int_{\tau_1(\pi)}^{\bar{\theta}} G(\theta) d\theta = (\bar{\theta} - \tau_1(\pi)) F(\bar{\theta}).$$

Differentiating with respect to π , we then obtain

$$-G(\tau_1(\pi)) \tau_1'(\pi) = -\tau_1'(\pi) F(\bar{\theta})$$

or

$$(G(\tau_1(\pi)) - F(\bar{\theta})) \tau_1'(\pi) = 0.$$

But equation (35) implies that $(1 + \pi) \tau_1(\pi) = \bar{\theta}$ and therefore

$$\tau_1'(\pi) = -\frac{\theta_1}{1 + \pi} \neq 0.$$

We conclude that $G(\tau_1(\pi)) - F(\bar{\theta}) = 0$. This, however, is impossible. For we have

$$G(\tau_1(\pi)) = \Psi(\tau_1(\pi); \tau_1(\pi), \pi) \leq \Psi(\bar{\theta}_F; \tau_1(\pi), \pi) < \Psi(\bar{\theta}; \tau_1(\pi), \pi)$$

(by construction of Ψ , by Proposition 18 and by Proposition 19)

$$= \psi(\tau_1(\pi), \pi) = F(\bar{\theta})$$

(as above). The only possible conclusion is therefore that $\pi_1 = \pi_2$.

15. COMPARATIVE STATICS WITHOUT A4

We divide our discussion into the same three cases that we considered in Section 12.1, namely:

ONLINE APPENDIX C TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

1. $\rho < 1$ and $G' \leq 0$ on $(\underline{\theta}, \infty)$;
2. $\rho = 1$;
3. $\rho > 1$ and $G' \geq 0$ on $(0, \bar{\theta})$.

Of these, the first is by far the simplest.

Proposition 55. *Suppose that $\rho < 1$ and $G' \leq 0$ on $(\underline{\theta}, \infty)$. Then $\underline{\tau} = \bar{\tau}$ for all $\pi \in (0, \infty)$. Furthermore there exists $\pi_1 \in (0, \infty)$ such that: the maximum-penalty constraint is strictly binding for all $\pi \in (0, \pi_1)$; and the maximum-penalty constraint is strictly slack for all $\pi \in (\pi_1, \infty)$. Finally:*

1. $\underline{\tau} = \bar{\tau}$ is strictly decreasing on $(0, \pi_1)$; and
2. $\underline{\tau} = \bar{\tau}$ is constant on (π_1, ∞) .

In other words, for all values of the maximum penalty $\pi \in [0, \infty)$, there is a unique optimum within our one-parameter family. Furthermore there exists a critical level π_1 of π . Below π_1 , the maximum-penalty constraint is strictly binding and the optimal savings target is strictly increasing in π . Above π_1 , the maximum-penalty constraint is strictly slack and the optimal savings target is independent of π .

Proof. Since $G' \geq 0$ on $(0, \underline{\theta})$ and $G' \leq 0$ on $(\underline{\theta}, \infty)$, we can put $\theta_M = \underline{\theta}$. For this choice of θ_M , A4 holds. Indeed: the interval $[\underline{\theta}, \theta_M)$ is empty, and therefore G is certainly strictly increasing on $[\underline{\theta}, \theta_M)$! We may therefore apply the analysis of Sections 11-14 to conclude that there is a unique $\theta_1 = \tau_1(\pi)$ such that $\psi(\theta_1, \pi) = F(\bar{\theta})$, and that $\tau'_1(\pi) < 0$. ■

Remark 56. *There is also a direct proof of Theorem 55. A sketch of this proof runs as follows. Since $G' \leq 0$ on $(\underline{\theta}, \infty)$, we must have $\bar{\theta}_F = \underline{\theta}$. Furthermore we always have $\theta_1 < \bar{\theta}_F$ and $\theta_2 > \underline{\theta}$; and in the strictly binding case we also have $\theta_2 < \bar{\theta}$. Hence, in the strictly binding case, we have*

$$G(\theta_2) > F(\bar{\theta}) = \psi(\theta_1, \pi) = \Psi(\bar{\theta}; \theta_1, \pi) > \Psi(\theta_2; \theta_1, \pi) = \phi(\theta_1, \pi) > G(\theta_1).$$

In particular,

$$G(\theta_2) - \phi(\theta_1, \pi) > 0.$$

It then follows from the formulae for $\frac{\partial \psi}{\partial \pi}$ and $\frac{\partial \psi}{\partial \theta_1}$ given in Propositions 32 and 40 – both of which feature the term $G(\theta_2) - \phi(\theta_1, \pi)$ – that

$$\frac{\partial \psi}{\partial \pi}, \frac{\partial \psi}{\partial \theta_1} > 0.$$

That is, there is a unique $\theta_1 = \tau_1(\pi)$ such that $\psi(\theta_1, \pi) = F(\bar{\theta})$, and $\tau'_1(\pi) < 0$. (The important point here is the fact that our assumption on G allows us to sign the core term $G(\theta_2) - \phi(\theta_1, \pi)$, and thereby the derivatives $\frac{\partial \psi}{\partial \pi}$ and $\frac{\partial \psi}{\partial \theta_1}$, directly.)

Proposition 57. Suppose that $\rho = 1$. Then there exists $\pi_0 \in [0, \infty)$ such that: $\underline{\tau} < \bar{\tau}$ for all $\pi \in (0, \pi_0)$; and $\underline{\tau} = \bar{\tau}$ for all $\pi \in (\pi_0, \infty)$. Furthermore there exists $\pi_1 \in (\pi_0, \infty)$ such that: the maximum-penalty constraint is strictly binding for all $\pi \in (0, \pi_1)$; and the maximum-penalty constraint is strictly slack for all $\pi \in (\pi_1, \infty)$. Finally:

1. $\underline{\tau}$ is constant on $(0, \pi_0)$ and $\bar{\tau}$ is strictly decreasing on $(0, \pi_0)$;
2. $\underline{\tau} = \bar{\tau}$ is strictly decreasing on (π_0, π_1) ; and
3. $\underline{\tau} = \bar{\tau}$ is constant on (π_1, ∞) .

In other words, there are two critical levels of π , namely π_0 and π_1 . Below π_0 , there is a continuum of optima from within our one-parameter family; and, above π_0 , there is a unique optimum from within our one-parameter family. Below π_1 , the maximum-penalty constraint is strictly binding; and, above π_1 , the maximum-penalty constraint is strictly slack. Furthermore, below π_0 : the smallest of the possible optimal savings targets is strictly increasing in π ; and the largest of the possible optimal savings targets is independent of π . Between π_0 and π_1 : there is only one optimal savings target, and this is strictly increasing in π . And, above π_1 : there is again only one optimal savings target, and this is independent of π .

Proof. This Proposition can be proved in three main steps. Fix $\pi > 0$ and suppose that, for this π , there exist $\theta_3, \theta_4 \in (\frac{1}{1+\pi}\underline{\theta}, \bar{\theta}_F)$ such that $\theta_3 < \theta_4$ and $\{\theta_1 \mid \psi(\theta_1, \pi) = F(\bar{\theta})\} = [\theta_3, \theta_4]$. Then the first step is to show that:

ONLINE APPENDIX C TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

1. $G < G(\theta_3)$ on $(0, \theta_3)$;
2. $G = G(\theta_3)$ on $[\theta_3, (1 + \pi)\theta_4]$;
3. $G > G(\theta_3)$ on $((1 + \pi)\theta_4, \bar{\theta})$.

Furthermore $G(\theta_3) < F(\bar{\theta})$. In other words, if there is a multiplicity of optimal savings targets, then G must have a flat. Moreover the domain of this flat consists precisely of the half-open interval $[\theta_3, (1 + \pi)\theta_4]$, where θ_3 is the smallest possible choice of θ_1 and θ_4 is the highest possible choice of θ_1 .

Now put $\theta_5 = (1 + \pi)\theta_4$ and $\pi_0 = \frac{\theta_5 - \theta_3}{\theta_3}$. Then the second step is to show that, for all $\hat{\pi} \in (0, \pi_0)$,

$$\{\theta_1 \mid \psi(\theta_1, \hat{\pi}) = F(\bar{\theta})\} = [\theta_3, \frac{1}{1+\hat{\pi}}\theta_5].$$

In other words, if there is some $\pi > 0$ for which there is a multiplicity of optimal savings targets, then there is a whole range of π for which there is a multiplicity of optimal savings targets. Furthermore both the multiplicity of optimal savings targets and range of π for which there is a multiplicity of optimal savings targets are associated with the same flat of G .

Finally, fix $\hat{\pi}, \pi > 0$ and suppose that:

1. $\hat{\pi} < \pi$;
2. there exist $\hat{\theta}_3, \hat{\theta}_4 \in (\frac{1}{1+\pi}\underline{\theta}, \bar{\theta}_F)$ such that $\hat{\theta}_3 < \hat{\theta}_4$ and $\{\theta_1 \mid \psi(\theta_1, \hat{\pi}) = F(\bar{\theta})\} = [\hat{\theta}_3, \hat{\theta}_4]$;
3. there exist $\theta_3, \theta_4 \in (\frac{1}{1+\pi}\underline{\theta}, \bar{\theta}_F)$ such that $\theta_3 < \theta_4$ and $\{\theta_1 \mid \psi(\theta_1, \pi) = F(\bar{\theta})\} = [\theta_3, \theta_4]$.

Then the third step is to show that $\hat{\theta}_3 = \theta_3$, $\hat{\theta}_4 = \frac{1+\pi}{1+\hat{\pi}}\theta_4$ and $\pi < \hat{\pi}_0 = \frac{(1+\hat{\pi})\hat{\theta}_4 - \hat{\theta}_3}{\hat{\theta}_3}$. In other words, if there is a multiplicity of optimal savings targets associated with both $\hat{\pi}$ and π , then both multiplicities derive from the same flat of G . ■

Proposition 58. *Suppose that $\rho > 1$ and $G' \geq 0$ on $(0, \bar{\theta})$. Then there exists $\pi_1 \in (0, \infty)$ such that: the maximum-penalty constraint is strictly binding for all $\pi \in (0, \pi_1)$; and the maximum-penalty constraint is strictly slack for all $\pi \in (\pi_1, \infty)$. Furthermore:*

1. $\underline{\tau}$ and $\bar{\tau}$ are both strictly decreasing on $(0, \pi_1)$;
2. $\underline{\tau} = \bar{\tau}$ is constant on (π_1, ∞) .

In other words, there exists a critical level π_1 of π . Below π_1 : the maximum-penalty constraint is strictly binding; and the set of optimal savings targets is strictly increasing in π . Above π_1 : the maximum-penalty constraint is strictly slack; and the optimal savings target is independent of π .

Proof. Let L be the locus of all those $(\theta_1, \pi) \in V$ such that $\psi(\theta_1, \pi) = F(\bar{\theta})$, let L_θ be the projection of L onto the first coordinate, and let L_π be the projection of L onto the second coordinate. Then, in order to prove the result, it suffices to show that there is a non-increasing function $\varpi : L_\theta \rightarrow L_\pi$, the graph of which is L . For then the inverse $\tau : L_\pi \rightarrow L_\theta$ of ϖ is a strictly increasing correspondence.

Note first that, for all $(\theta_1, \pi) \in V$, we have

$$\begin{aligned} \frac{\partial \psi}{\partial \theta_1}(\theta_1, \pi) = & \left(\frac{1 + \pi}{\pi \theta_1} \frac{\theta_1 + \beta b\left(\frac{\theta_1}{\beta}\right)}{\theta_2 + \beta b\left(\frac{\theta_1}{\beta}\right)} (G(\theta_2) - \phi(\theta_1, \pi)) \right. \\ & \left. + \frac{1}{\pi \theta_1} (\phi(\theta_1, \pi) - G(\theta_1)) \right) S(\theta_2, \bar{\theta}) \end{aligned}$$

and

$$\begin{aligned} \frac{\partial \psi}{\partial \pi}(\theta_1, \pi) = & \left(\frac{1}{\pi} S(\theta_2, \pi) - \theta_1 \frac{\partial S}{\partial h}(\theta_2, \pi) \right) (G(\theta_2) - \phi(\theta_1, \pi)) \\ & - \int_{[\theta_2, \bar{\theta}]} \frac{\partial S}{\partial \pi}(h, \pi) d\tilde{G}(h), \end{aligned}$$

where, as above, $\tilde{G} \in \mathcal{BV}([\theta_2, \bar{\theta}], \mathbb{R})$ is given by the formulae $\tilde{G}_L(\theta_2) = \phi(\theta_1, \pi)$ and $\tilde{G} = G$ on $(\theta_2, \bar{\theta}]$.¹⁵

Next, since $G' \geq 0$ on $(0, \bar{\theta})$, we must have $\phi(\theta_1, \pi) - G(\theta_1) \geq 0$, $G(\theta_2) - \phi(\theta_1, \pi) \geq 0$ and $\tilde{G}' \geq 0$ on $[\theta_2, \bar{\theta})$. Hence $\frac{\partial \psi}{\partial \theta_1} \geq 0$.

Third, if in addition if $\psi(\theta_1, \pi) = F(\bar{\theta})$, then we must have $\frac{\partial \psi}{\partial \pi} > 0$. Indeed, it is always the case that $\bar{\theta}_F < \bar{\theta}$ and $G > F(\bar{\theta})$ on $(\bar{\theta}_F, \bar{\theta})$. (See Lemma 9.)

¹⁵For the definition of V , see the beginning of Section 12.

Moreover, if $\psi(\theta_1, \pi) = F(\bar{\theta})$, then we also have $G(\theta_1) < F(\bar{\theta})$. Overall, then, if $\psi(\theta_1, \pi) = F(\bar{\theta})$ then G is non-trivial on $(\theta_1, \bar{\theta})$. Now suppose for a contradiction that $\frac{\partial \psi}{\partial \pi} = 0$. Then we must have $G(\theta_2) - \phi(\theta_1, \pi) = 0$ (which is the same thing as saying that $\tilde{G}' = 0$ on $\{\theta_2\}$) and $\tilde{G}' = 0$ on $(\theta_2, \bar{\theta})$. Moreover the former implies that $G(\theta_2) = G_L(\theta_2) = \phi(\theta_1, \pi) = G(\theta_1)$, and the latter implies that $G_L(\bar{\theta}) = G(\theta_2)$. So G is trivial on $(\theta_1, \bar{\theta})$, which is the required contradiction.

Finally, since $\frac{\partial \psi}{\partial \pi} > 0$, there is a unique $\pi = \varpi(\theta_1)$ such that $\psi(\theta_1, \pi) = F(\bar{\theta})$ and moreover

$$\varpi'(\theta_1) = -\frac{\frac{\partial \psi}{\partial \theta_1}}{\frac{\partial \psi}{\partial \pi}} \geq 0.$$

This completes the proof. ■

Remark 59. *It is interesting to compare the levels of uniqueness obtained in Propositions 55, 57 and 58. When $\rho < 1$, we have uniqueness for all $\pi \in (0, \infty)$. When $\rho = 1$, a limited form of non-uniqueness can develop: there exists $\pi_0 \in [0, \pi_1)$ such that there is non-uniqueness on $(0, \pi_0)$ and uniqueness on (π_0, ∞) . And, when $\rho > 1$, non-uniqueness takes the form that one might expect in a convex optimization problem. However, we do at least get strict monotonicity on the whole of $(0, \pi_1)$.*

16. EXISTENCE OF A FULL OPTIMUM

Suppose that self 0 is required to pick a B satisfying Constraints 1 and 2. Then the utility curve (u, w) that results will satisfy the following three conditions:

I (u, w) is interior, in the sense that $u, w > U(0)$ on Θ .

M (u, w) is monotonic, in the sense that u is non-decreasing and w is non-increasing.

DE (u, w) satisfies the differential equation $\theta du + \beta dw = 0$.

If B is also convex, then (u, w) will also satisfy:

C (u, w) is continuous.

Now, the set Ω with which we have worked so far consists of utility curves (u, w) that satisfy I, BV, DE and C, where BV is the condition:

$\mathbf{BV}$ (u, w) is of bounded variation.

Since $\mathbf{BV}$ is weaker than $\mathbf{M}$, this means that Ω contains all the utility curves that can result from convex B , and more besides. We have therefore solved a relaxed version of the convex- B problem. Since the solution of this relaxed problem is feasible in the convex- B problem, we have therefore also solved the convex- B problem. The purpose of the present section is to solve the general problem in which B is not required to be convex.

Suppose accordingly that Ω consists of all $(u, w) \in \mathcal{BV}(\Theta, \text{ran}(U))^2$ such that $\theta du + \beta dw = 0$. In other words, let Ω consist of utility curves (u, w) that satisfy $\mathbf{I}$, $\mathbf{BV}$, $\mathbf{DE}$ but not $\mathbf{C}$. Put $X = \mathcal{BV}(\Theta, \mathbb{R})^2$, $\Xi = \mathcal{BV}(\Theta, \text{ran}(U))^2$ and $Z = \mathcal{BV}(\Theta, \mathbb{R})$. Then the objective function M and the constraint mappings G_1 and G_2 continue to make sense. The analysis of Luenberger (1969) therefore shows that $x_0 \in \Omega$ solves the problem

$$\begin{aligned} & \text{maximize} && M(x) \\ & \text{subject to} && \left\{ \begin{array}{l} x \in \Omega \\ G_1(x) \geq 0 \\ G_2(x) \geq 0 \end{array} \right\} \end{aligned}$$

iff there exist $\lambda_1, \lambda_2 \in Z^*$ such that:

1. $L(x_0, \lambda_1, \lambda_2) \geq L(x, \lambda_1, \lambda_2)$ for all $x \in \Omega$, where

$$L(x, \lambda_1, \lambda_2) = M(x) + \langle G_1(x), \lambda_1 \rangle + \langle G_2(x), \lambda_2 \rangle;$$

2. $G_1(x) \geq 0$, $\lambda_1 \geq 0$ and $\langle G_1(x), \lambda_1 \rangle = 0$;
3. $G_2(x) \geq 0$, $\lambda_2 \geq 0$ and $\langle G_2(x), \lambda_2 \rangle = 0$.

In other words, there exists multipliers λ_1 and λ_2 such that: (1) x_0 maximizes $L(\cdot, \lambda_1, \lambda_2)$ over Ω ; (2) complementary slackness holds for the first constraint; and (3) complementary slackness holds for the second constraint.

At this point, however, we encounter an obstacle. While the dual space $\mathcal{C}(\Theta, \mathbb{R})^*$ of $\mathcal{C}(\Theta, \mathbb{R})$ has a convenient representation as the space $\mathcal{M}(\Theta, \mathbb{R})$, the dual space

$\mathcal{BV}(\Theta, \mathbb{R})^*$ of $\mathcal{BV}(\Theta, \mathbb{R})$ does not have a similarly convenient representation. This makes it difficult to use the necessity part of the Lagrangean characterization of the optimum. We can, however, still hope to use the sufficiency part.

The idea here is to note that the elements of $\mathcal{M}(\Theta, \mathbb{R})$ can be used to induce continuous linear functionals on $\mathcal{BV}(\Theta, \mathbb{R})$. For example, $\mu \in \mathcal{M}(\Theta, \mathbb{R})$ induces $\mu_R \in \mathcal{BV}(\Theta, \mathbb{R})^*$ via the formula

$$\langle z, \mu_R \rangle = \int z_R d\mu,$$

where z_R denotes the right-continuous version of z . However, in pursuing this idea, it is important to note that μ also induces $\mu_L \in \mathcal{BV}(\Theta, \mathbb{R})^*$ via the formula

$$\langle z, \mu_L \rangle = \int z_L d\mu,$$

where z_L denotes the left-continuous version of z . In other words, there is no canonical association between elements of $\mathcal{M}(\Theta, \mathbb{R})$ and continuous linear functionals on $\mathcal{BV}(\Theta, \mathbb{R})$.

Our plan is therefore to start from a θ_1 such that $\Psi(\bar{\theta}; \theta_1) = F(\bar{\theta})$, in the hope that $\Psi(\cdot; \theta_1)$ can be used to generate multipliers that can be used in the sufficiency part of the Lagrangean characterization of an optimum. Indeed, suppose that we are given such a θ_1 . Then, bearing in mind that $\Delta\Psi(\theta_2; \theta_1) = 0$, we may put $d\tilde{\Lambda}_1 = d\Psi(\cdot; \theta_1)$ on $[\underline{\theta}, \theta_2]$, $d\Lambda_2 = d\Psi(\cdot; \theta_1)$ on $[\theta_2, \bar{\theta}]$ and $d\Lambda_1 = \frac{1}{K'(w_0)} d\tilde{\Lambda}_1$. Furthermore, if we let λ_1 and λ_2 be the continuous linear functionals induced on $\mathcal{BV}(\Theta, \mathbb{R})$ by $d\Lambda_1$ and $d\Lambda_2$ using integration with respect to the right-continuous versions of functions, then we have

$$\begin{aligned} L(x, \lambda_1, \lambda_2) &= \int \left(\theta u(\theta) + w(\theta) \right) dF(\theta) \\ &\quad + \int \left(y - C(u(\theta)) - K(w(\theta)) \right) d\Lambda_1(\theta) \\ &\quad + \int \left(b\left(\frac{\theta}{(1+\pi)^\beta}\right) u(\theta) - \frac{1}{\rho} a\left(\frac{\theta}{(1+\pi)^\beta}\right) - w(\theta) \right) d\Lambda_2(\theta) \end{aligned}$$

for all $x \in X$. Our objective is then to show that the utility curve $x_0 = (u_0, w_0)$ associated with θ_1 maximizes $L(\cdot, \lambda_1, \lambda_2)$.

It suffices to show that, for all $x_1 \in \Omega$, the directional derivative $\nabla_x L(x_0, \lambda_1, \lambda_2)$ of L at x_0 in the direction $x = x_1 - x_0$ is non-positive. As in Section 6, we have

$$\begin{aligned} \nabla_x L(x_0, \lambda_1, \lambda_2) &= \int (\theta u + w) dF - \int (C'(u_0) u + K'(w_0) w) d\Lambda_1 \\ &\quad + \int \left(b\left(\frac{\theta}{(1+\pi)\beta}\right) u - w \right) d\Lambda_2 \\ &= \int (\theta u + w) dF - \int \left(\frac{C'(u_0)}{K'(w_0)} u + w \right) d\tilde{\Lambda}_1 \\ &\quad + \int \left(b\left(\frac{\theta}{(1+\pi)\beta}\right) u - w \right) d\Lambda_2. \end{aligned}$$

Furthermore, notwithstanding the fact that we are now working in a more general context, we can eliminate the terms $\int w dF$, $\int w d\tilde{\Lambda}_1$ and $\int w d\Lambda_2$ using integration by parts.

Indeed, the general formula for integration by parts tells us that

$$\int_{[\underline{\theta}, \bar{\theta}]} w(\theta) dF(\theta) = [w F]_{\underline{\theta}-}^{\bar{\theta}} - \int_{[\underline{\theta}, \bar{\theta}]} F(\theta) dw(\theta) + \sum_{\theta \in [\underline{\theta}, \bar{\theta}]} \Delta w(\theta) \Delta F(\theta),$$

where

$$[w F]_{\underline{\theta}-}^{\bar{\theta}} = w(\bar{\theta}) F(\bar{\theta}) - w(\underline{\theta}-) F(\underline{\theta}-).$$

We therefore have

$$\int w dF = [w F]_{\underline{\theta}-}^{\bar{\theta}} - \int F dw + \sum \Delta w \Delta F$$

(where we have suppressed the dependence on θ and where the domains of all integrals and sums are understood to be the whole of $[\underline{\theta}, \bar{\theta}]$)

$$= w(\bar{\theta}) F(\bar{\theta}) + \int F \frac{\theta}{\beta} du - \sum \frac{\theta}{\beta} \Delta u \Delta F$$

(because $F(\underline{\theta}-) = 0$ and $dw = -\frac{\theta}{\beta} du$)

$$= w(\bar{\theta}) F(\bar{\theta}) + \frac{1}{\beta} \int F \theta du - \frac{1}{\beta} \sum \theta \Delta u \Delta F.$$

ONLINE APPENDIX C TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

Moreover

$$\int F \theta du = [(F \theta) u]_{\underline{\theta}-}^{\bar{\theta}} - \int u d(F \theta) + \sum \Delta(F \theta) \Delta u$$

(applying the general formula for integration by parts to $\int F \theta du$)

$$= \bar{\theta} u(\bar{\theta}) F(\bar{\theta}) - \int u (\theta dF + F d\theta) + \sum \theta \Delta F \Delta u$$

(since $F(\underline{\theta}-) = 0$, $d(F \theta) = \theta dF + F d\theta$ and $\Delta(F \theta) = \theta \Delta F$). Overall, then,

$$\int w dF = \left(\frac{\bar{\theta}}{\beta} u(\bar{\theta}) + w(\bar{\theta}) \right) F(\bar{\theta}) - \frac{1}{\beta} \int u (\theta dF + F d\theta).$$

By the same token, and bearing in mind that we did not use the fact that $\Delta F = 0$ in the derivation of the previous paragraph, we have

$$\int w d\tilde{\Lambda}_1 = \left(\frac{\bar{\theta}}{\beta} u(\bar{\theta}) + w(\bar{\theta}) \right) \tilde{\Lambda}_1(\bar{\theta}) - \frac{1}{\beta} \int u (\theta d\tilde{\Lambda}_1 + \tilde{\Lambda}_1 d\theta)$$

and

$$\int w d\Lambda_2 = \left(\frac{\bar{\theta}}{\beta} u(\bar{\theta}) + w(\bar{\theta}) \right) \Lambda_2(\bar{\theta}) - \frac{1}{\beta} \int u (\theta d\Lambda_2 + \Lambda_2 d\theta).$$

We therefore have

$$\nabla_x L(x_0, \lambda_1, \lambda_2) = \int u du^* + w(\bar{\theta}) r^*$$

where, as in section 6 above,

$$\begin{aligned} du^* = & -\frac{1}{\beta} \left((1 - \beta) \theta dF + F d\theta \right) + \frac{1}{\beta} \left(\left(\theta - \beta \frac{C'(u_0)}{K'(w_0)} \right) d\tilde{\Lambda}_1 + \tilde{\Lambda}_1 d\theta \right) \\ & + \frac{1}{\beta} \left(\left(\theta + \beta b\left(\frac{\theta}{(1+\pi)\beta}\right) \right) d\Lambda_2 + \Lambda_2 d\theta \right) + \frac{\bar{\theta}}{\beta} \left(F(\bar{\theta}) - \tilde{\Lambda}_1(\bar{\theta}) - \Lambda_2(\bar{\theta}) \right) dI, \end{aligned}$$

$$r^* = F(\bar{\theta}) - \tilde{\Lambda}_1(\bar{\theta}) - \Lambda_2(\bar{\theta})$$

and I is the distribution function of the unit mass at $\bar{\theta}$. Finally, by construction of $\tilde{\Lambda}_1$ and Λ_2 , we have $u^* = 0$ and $r^* = 0$. So in fact $\nabla_x L(x_0, \lambda_1, \lambda_2) = 0$. In particular, the utility curve $x_0 = (u_0, w_0)$ associated with θ_1 does indeed maximize $L(\cdot, \lambda_1, \lambda_2)$.

Remark 60. Great care is needed in choosing the space Z . One possible choice is $\mathcal{C}(\Theta, \mathbb{R})$, the space of all continuous functions on Θ endowed with the sup norm. This choice has the advantage that there is a convenient representation for Z^* . However, it also requires that $\Omega \subset \mathcal{C}(\Theta, \mathbb{R})$, and this is not an economically reasonable restriction. Another possible choice is $\mathcal{B}(\Theta, \mathbb{R})$, the space of all bounded functions on Θ endowed with the sup norm. This choice has the advantage that it includes all economically relevant utility curves. Unfortunately, it leads to a different problem: the measures $d\Lambda_1$ and $d\Lambda_2$ associated with $\Psi(\cdot; \theta_1)$ do not induce continuous linear functionals on $\mathcal{B}(\Theta, \mathbb{R})$, since functions in $\mathcal{B}(\Theta, \mathbb{R})$ are not in general measurable. The results of Luenberger (1969) do not therefore apply. Our solution to this double problem is to use $\mathcal{BV}(\Theta, \mathbb{R})$. This space is big enough to include all economically relevant utility curves, but small enough that $d\Lambda_1$ and $d\Lambda_2$ can be used to induce continuous linear functionals on it (albeit not in a canonical way).

17. DISTRIBUTIONS

17.1. Beta Distribution. 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 $0 < a < b$ and $\zeta, \eta > 0$. It is unbounded at a if $\zeta < 1$, in which case we require that $\underline{\theta} \in (a, b)$ in order to ensure that A1 is satisfied, and unbounded at b if $\eta < 1$, in which case we require that $\bar{\theta} \in (a, b)$ in order to ensure that A1 is satisfied.

There are then four main cases. Three of the cases are easy to describe:

Case 1 if $\zeta > 1$ and $\eta \geq 1$ then A3 is satisfied for all choices of $\underline{\theta}, \bar{\theta} \in (a, b)$;

Case 2 if $\zeta > 1$ and $\eta < 1$ then A3 is again satisfied for all choices of $\underline{\theta}, \bar{\theta} \in (a, b)$, albeit for somewhat different reasons;

Case 4 if $\zeta < 1$ and $\eta < 1$, then A3 is violated for some choices of $\underline{\theta}, \bar{\theta} \in (a, b)$.

Case 3 is more involved. If $\zeta < 1$ and $\eta \geq 1$, then A3 is satisfied for all choices of $\underline{\theta}, \bar{\theta} \in (a, b)$ iff

$$\frac{(\sqrt{1-\zeta} + \sqrt{\eta-1}\sqrt{\frac{a}{b}})^2}{1 - \frac{a}{b}} \geq 1 + \frac{1}{1-\beta}. \quad (36)$$

As this inequality makes clear, A3 is more likely to be satisfied if: either (i) ζ is close to 0 (i.e. the spike at a is very pronounced); or (ii) η is large (i.e. the density decays very quickly towards b); or (iii) $\frac{a}{b}$ is close to 1 (i.e. the density is concentrated in a narrow band).¹⁶ It is also worth noting that, as $\frac{a}{b} \rightarrow 0$, the left-hand side of (36) converges to $1 - \zeta < 1$. Hence A3 is violated for some choices of $\underline{\theta}, \bar{\theta} \in (a, b)$ when $\zeta < 1$ and $\frac{a}{b}$ is small. This is in striking contrast with the standard case studied in both Rice and Hogg et al. In that case A3 is satisfied for all $\underline{\theta}, \bar{\theta} \in (a, b)$ when $\zeta < 1$ and $\frac{a}{b} = 0$.

Note finally that the right-hand side of (36) is strictly increasing in β . Hence, if we fix a distribution for which $\zeta < 1$ and $\eta \geq 1$, then the conclusion is that A3 will be satisfied provided that β is *far enough below* 1. I.e. A3 is more likely to be satisfied when the decision maker is *more* time-inconsistent.

17.2. Cauchy Distribution. 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. This distribution satisfies A3 for all $\underline{\theta}, \bar{\theta} \in (0, \infty)$ iff

$$\frac{\mu}{\sigma} \leq \sqrt{\frac{1 - (1 - \beta)^2}{(1 - \beta)^2}}. \quad (37)$$

In other words, taking β as given, A3 is satisfied iff the distribution is not located too far to the right. If (37) does not hold then, for some choices of $\underline{\theta}, \bar{\theta} \in (0, \infty)$, G is first increasing (at $\underline{\theta}$), then decreasing, then increasing again, then finally decreasing again (at $\bar{\theta}$).

We can also make $1 - \beta$ the subject of the inequality (37). Doing so, we find that

¹⁶For the purposes of the present discussion, $\zeta \in (0, 1)$, $\eta \in [1, \infty)$ and $\frac{a}{b} \in (0, 1)$.

A3 is satisfied for all $\underline{\theta}, \bar{\theta} \in (0, \infty)$ iff: either $\mu \leq 0$; or $\mu > 0$ and

$$1 - \beta \leq \left(1 + \frac{\mu^2}{\sigma^2}\right)^{-\frac{1}{2}}.$$

In other words, taking the parameters μ and σ of the Cauchy distribution as given, A3 is satisfied iff: either $\mu \leq 0$; or $\mu > 0$ and β is sufficiently close to 1. I.e. A3 is more likely to be satisfied when the decision maker is less time-inconsistent.

17.3. Log-Gamma Distribution. The density of the Log-Gamma distribution is proportional to

$$x^{-\frac{\eta+1}{\eta}} (\log(x))^{\zeta-1}$$

on $(1, \infty)$, where $\zeta, \eta > 0$. It is unbounded at 1 if $\zeta < 1$, in which case we require that $\underline{\theta} > 1$ in order to ensure that A1 is satisfied. It violates A3 for some choices of $\underline{\theta}, \bar{\theta} \in (1, \infty)$ iff $\zeta < 1$ and $\eta > 1 - \beta$. In other words, taking β as given, A3 is violated iff there is a singularity at 1 and the rate of decay at ∞ is sufficiently slow.

Note finally that, if we fix a distribution for which $\zeta < 1$, then the conclusion is that A3 will be satisfied provided that β is *far enough below* 1. I.e. A3 is more likely to be satisfied when the decision maker is *more* time-inconsistent.

17.4. Pareto Distribution. 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. It violates A3 for some choices of $\underline{\theta}, \bar{\theta} \in (\mu, \infty)$ iff

$$\zeta < \frac{1}{1 - \beta} \tag{38}$$

and

$$\frac{\mu}{\sigma} > \frac{1}{\zeta + 1} \left(1 + \frac{1}{1 - \beta}\right). \tag{39}$$

In other words, it violates A3 iff its right-hand tail is sufficiently fat and, taking the fatness of the tail as given, it is located sufficiently far to the right. In particular, if

ONLINE APPENDIX C TO “WHICH EARLY WITHDRAWAL PENALTY ATTRACTS THE
MOST DEPOSITS TO A COMMITMENT SAVINGS ACCOUNT?”

$\frac{\mu}{\sigma} \leq 1$, then the Pareto type II distribution satisfies A3 for all $\underline{\theta}, \bar{\theta} \in (\mu, \infty)$. For in that case: either (i) $\zeta \geq \frac{1}{1-\beta}$ and therefore (38) is violated; or (ii) $\zeta < \frac{1}{1-\beta}$, in which case $\frac{1}{\zeta+1} \left(1 + \frac{1}{1-\beta}\right) > \frac{1}{\zeta+1} (1 + \zeta) = 1 \geq \frac{\mu}{\sigma}$ and therefore (39) is violated.

We can also make $1 - \beta$ the subject of these inequalities. Doing so, we find that A3 is violated for some $\underline{\theta}, \bar{\theta} \in (\mu, \infty)$ iff

$$\left(\frac{\mu}{\sigma} (\zeta + 1) - 1\right)^{-1} < 1 - \beta < \zeta^{-1}.$$

In particular, if $\frac{\mu}{\sigma} > 1$ and $\zeta > 1$ (so that $\left(\frac{\mu}{\sigma} (\zeta + 1) - 1\right)^{-1} < \zeta^{-1} < 1$), then A3 is satisfied iff β is either close enough to 1 or *far enough below* 1.