

Optimal Illiquidity

JOHN BESHEARS* JAMES CHOI CHRISTOPHER CLAYTON
CHRISTOPHER HARRIS DAVID LAIBSON BRIGITTE MADRIAN

September 26, 2014

ABSTRACT: This paper calculates the socially optimal level of illiquidity in a retirement savings system. We study an environment in which time-inconsistent agents face a tradeoff between commitment and flexibility (cf. Amador, Werning and Angeletos 2006). For this analysis, we assume that the agent has access to two accounts: a perfectly liquid account and an illiquid retirement savings account with an early withdrawal penalty ($0\% \leq \pi \leq 100\%$). When agents have homogeneous present-biased preferences (with short-run discount factor β), we find that the socially optimal retirement savings account should have a penalty that is approximately equal to the level of present bias, $1 - \beta$. In this case, the penalty roughly offsets the present bias of the representative agent. For example, if $\beta = 0.7$, then the socially optimal early withdrawal penalty rate is approximately 30%. However, when agents have heterogeneous preferences, with a range of β values, we find that optimal policy disproportionately addresses the needs of low β agents. In an illustrative calibration with β values distributed uniformly between 0.1 and 1 (with a mean β value of 0.55), we find that the optimal savings system is characterized by a retirement savings account that is essentially perfectly illiquid (i.e., with an early withdrawal penalty rate of $\pi \simeq 100\%$). In other words, our analysis with heterogeneous preferences suggests that savings should be divided between two accounts: one account that is completely liquid and one account that is completely illiquid (like a defined benefit pension plan).

*We gratefully acknowledge helpful advice from Matthew Rabin and John Sabelhaus and seminar participants at the University of Wisconsin, Princeton University, the Norwegian School of Economics, ANPEC, the American Economic Association, and the NBER Summer Institute. This research was supported by the U.S. Social Security Administration through grant #RRC08098400-06-00 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the author(s) and do not represent the views of SSA, any agency of the Federal Government, or the NBER.

1. INTRODUCTION

US defined contribution (DC) savings accounts are more liquid than DC accounts in most (if not all) other developed countries. In the US, certain types of pre-retirement withdrawals are allowed without penalty, and, for IRAs, withdrawals may be made for any reason if a 10% penalty is paid. Liquidity allows significant pre-retirement “leakage”: for every \$1 contributed to the accounts of savers under age 55, \$0.40 simultaneously flows out of the 401(k)/IRA system, not counting loans (Argento, Bryant, and Sabelhaus 2014).

This leakage is sometimes desirable (when it funds legitimate spending needs, like a medical emergency or investment in human capital) and sometimes self-defeating (when it is caused by planning mistakes and/or self-control problems). If an outside observer does not know the details of a household’s financial situation, it is not clear whether leakage is occurring for socially optimal or sub-optimal reasons.

This paper evaluates the optimality of a stylized two-account retirement savings system with one liquid savings account and one partially or fully illiquid savings account. We study preferences that includes both legitimate spending shocks and self-control problems. The self-control problems are modeled as the consequence of present bias (Phelps and Pollak 1968, Laibson 1997): i.e., a discount function with weights $\{1, \beta\delta^2, \dots, \beta\delta^t\}$, where the degree of present bias is $1 - \beta$.

Using a mechanism design framework (cf. Angeletos, Werning, and Amador 2006), we computationally generate the socially optimal level of penalties on the illiquid savings account. We study two special cases. In the first special case, we assume that agents have homogeneous time preferences (i.e., all agents have the same β and δ parameters). In this case, our theoretical model implies that the optimal level of illiquidity is an early withdrawal penalty that is approximately $1 - \beta$ (which is the degree of present bias).

In the second case, we assume that agents have heterogeneous β values. In this heterogeneous-preference case, we find that the socially optimal degree of illiquidity caters to the households with the lowest β values. This asymmetric policy protects the subpopulation with the most extreme self-control problems. Completely illiquid retirement savings generates welfare gains for these low- β agents that swamp (by a ratio of 100 to 1) the welfare losses of the high- β agents (who are made slightly

worse off by the illiquidity). Hence, in a world of heterogeneous agents, socially optimal policy caters disproportionately to the agents with the most severe self-control problems (i.e., those with low β values).

Section 2 describes our model of household behavior. Section 3 describes the planner's problem – i.e., account allocations and an early withdrawal penalty that jointly maximize social welfare subject to information asymmetries between the planner and the households. Section 4 analyzes the solution to the planner's problem in the case of homogeneous preferences, including a description of household behavior in the resulting (second best) equilibrium. Section 5 analyzes the solution to the planner's problem in the case of heterogeneous preferences in the present-bias parameter β , including a description of household behavior in the resulting (second best) equilibrium. In section 6, we conclude and discuss the limitations of our existing analysis and goals for future work.

2. MODEL OF HOUSEHOLD BEHAVIOR

2.1. Introduction. To study the tradeoff between commitment and flexibility Amador, Werning and Angeletos (2006; hereafter AWA) use a model with three conceptual ingredients.¹ We first summarize the key ingredients of this model and then explain how we adapt it to our analysis.

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

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

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

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

¹Also see Ambrus and Egorov (forthcoming) for clarifications of some of the arguments in AWA.

²For some examples of related empirical work see Angeletos et al (2001), Ashraf et al (2006), and Augenblick et al (2014).

Dynamically inconsistent preferences generate a motivation for precommitment.

Second, AWA assume that agents experience 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, AWA assume that self 0 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 over-consuming) and the costs of commitment (preventing later selves from responding flexibly to taste shocks that are not contractable).

We make three key changes to the AWA model.

First, we restrict the savings/commitment technology to a two-account system: one completely liquid account and one illiquid account with an early withdrawal penalty, π . (In ongoing work, which we discuss in the conclusion, we are studying the case of three or more types of accounts, each with its own penalty for early withdrawal.)

Second, we allow for interpersonal transfers. Specifically, we assume that early withdrawal penalties paid by one household go into general government revenue and can be transferred to other households. AWA rule out such transfers and instead require money burning: if a household pays a penalty for an early withdrawal, this penalty is destroyed and can't be transferred to other households through the tax system. Their money burning assumption was made for analytic tractability. It has the undesirable consequences that it reduces the social efficiency of penalty-based retirement accounts. Indeed, our analysis shows that the main theorem of AWA³ does not generalize once one allows penalty payments to be transferred across households through government transfers.

Third, we introduce heterogeneity in the β parameter. In contrast, AWA assume that all agents have the same β value.

2.2. Timing and Preferences. We assume a continuum of households indexed on the unit interval, $i \in [0, 1]$.

³AWA's key theorem can be summarized as follows: In the socially optimal system, there are no penalties paid in equilibrium, so that the illiquid account is essentially perfectly illiquid.

To simplify notation, and without loss of generality, we assume that the interest rate is deterministic and set the gross real interest to $R = 1$ (so the net real interest rate is 0).

The simplest model that elicits a tradeoff between commitment and flexibility has three periods: an initial period in which the planner/government establishes savings accounts; a following period in which a consumption/savings choice is made by the household with immediate utility consequences; and a final (retirement) period in which residual wealth is consumed.

Period 0. On behalf of each household, the planner (i.e., the government) puts x dollars of savings into a liquid account and z dollars of savings into an illiquid account with early withdrawal penalty π .⁴ We explain how the government sets x , z , and π in the next section of the paper. The early withdrawal penalty is only paid if the household withdraws money from the illiquid account *before* period 2 (i.e., withdrawals in period 1 are penalized). Any withdrawals made in period 2 are not penalized.

Period 1. A taste shock $\theta_{1,i} \in \Theta = [\underline{\theta}, \bar{\theta}]$ is realized. Self 1 observes $\theta_{1,i}$ and makes a consumption/savings decision, $c_{1,i}$. If $c_{1,i} \leq x$, no early withdrawal is made from the illiquid account. If $c_{1,i} > x$, then the household partially funds this consumption with a withdrawal, w_i , from the illiquid account such that

$$w_i(1 - \pi) + x = c_{1,i}.$$

The government revenue obtained from withdrawals in period 1, is πw_i , where $w_i \in [0, z]$ is the withdrawal made by household i in period 1.

Period 2. A taste shock $\theta_{2,i} \in \Theta = [\underline{\theta}, \bar{\theta}]$ is realized. Self 2 observes $\theta_{2,i}$ and consumes all remaining wealth, such that

$$c_{2,i} = \begin{cases} x - c_{1,i} + z & \text{if } c_{1,i} \leq x \\ z - w_i & \text{if } c_{1,i} > x \end{cases}$$

⁴We discuss a generalization of the number of accounts in the conclusion.

Household preferences defined at dates 1 and 2 follow, where β is the parameter that reflects present-bias and δ is the standard discount factor:

$$\begin{aligned} \text{utility of self 1} &= \theta_{1,i} u(c_{1,i}) + \theta_{2,i} \beta \delta u(c_{2,i}) \\ \text{utility of self 2} &= \theta_{2,i} u(c_2) \end{aligned} \quad (\text{Households' Objectives})$$

Here u is the instantaneous utility function. We will assume that u is in the class of constant relative risk aversion so that

$$u(c) = \frac{c^{1-\gamma} - 1}{1 - \gamma}.$$

Parameter γ is the coefficient of relative risk aversion.

3. PLANNER'S PROBLEM

So far, we've described the decision problem faced by households. Now we introduce a planner/government that is not prone to present bias. The assumption on the planner's preferences is not an empirical/*positive* assumption, but instead made to answer the *normative* question of what policy should be. (Actual governments may be prone to present bias, reducing their ability to create socially optimal savings mechanisms.)

In our analysis, the planner's objective is to optimize average lifetime utility, which includes only δ -discounting and excludes discounting driven by the present bias parameter, β .

$$E \int_i [\theta_{1,i} u(c_{1,i}) + \theta_{2,i} \delta u(c_{2,i})] di. \quad (\text{Planner's Objective})$$

Here the expectation is taken over all of the stochastic variables, which the planner does not directly control: $\{\theta_{1,i}, c_{1,i}, \theta_{2,i}, c_{2,i}\}_{i \in [0,1]}$.

The planner picks x (liquid savings), z (illiquid savings), and π (the early withdrawal penalty) to maximize the Planner's Objective subject to three sets of constraints:

- (1) Self 1 picks $c_{1,i}$ to maximize its own objective function, $\theta_{1,i} u(c_{1,i}) + \theta_{2,i} \beta \delta u(c_{2,i})$, which reflects present bias;

- (2) Self 2 consumes all residual wealth;
- (3) The planner cannot violate its own budget constraint:

$$x + z = y + \pi E \int_i w_i di$$

where y is per-capita resources. To simplify notation, we use the following shorthand: $Ew = E \int_i w_i di$.

Note that the third constraint embodies an externality that arises from interpersonal transfers. When agents pay penalties in equilibrium, those penalties relax the government's budget constraint and enable the government to transfer resources among agents in the economy. Specifically, if $\pi > 0$, then total account allocations, $x + z$, can exceed society's total per-capita endowment, y . Without loss of generality we normalize $y = 1$.

4. SOLUTION TO THE PLANNER'S PROBLEM IN THE CASE OF HOMOGENEOUS TIME PREFERENCES (AND HETEROGENEOUS TASTE SHOCKS)

We begin our analysis by studying an economy comprised of a homogeneous population of agents, in the sense that all agents have a common value of δ and a common value of β . However, our agents remain heterogeneous in the sense that they receive idiosyncratic taste shocks (which we model with a truncated normal distribution). In Exhibit 1, we plot the socially optimal level of the penalty on the illiquid account as a function of the homogeneous (population-wide) value of β . For this analysis (and all analysis that follows) we hold δ fixed at unity. Exhibit 1 reports β values from 0.6 to 0.8, which represents the range that is most frequently estimated in empirical analyses that assume a homogeneous population. We find that the optimal level of penalty is approximately equal to $1 - \beta$, so the penalty approximately offsets the degree of present bias. For example, if all agents have $\beta = 0.7$, then the optimal early withdrawal penalty is 0.32 when the coefficient of relative risk aversion is one. This numerical result is not very sensitive to the coefficient of relative risk aversion.

Exhibit 1 also reports the same optimal penalty when the coefficient of relative risk aversion is 2. The optimal penalty lines are nearly identical in the two cases.

Exhibit 2 explores the $\beta = 0.7$ case more thoroughly (i.e., the case in which all consumers have a β value of 0.7). Now we illustrate how variation in the penalty affects welfare (assuming that the amount of money in the liquid and illiquid accounts is optimized for the given penalty level). Welfare is optimized at a penalty of 0.32 (as noted before) and welfare does not vary symmetrically around optimum. Instead, there is a pronounced asymmetry, with the welfare function concave everywhere to the left of the optimum and convex sufficiently far to the right of the optimum. This convex region arises because sufficiently high penalties eliminate almost all withdrawals from the illiquid account and further increases in the penalty make little (and ultimately no) difference. This asymmetry implies that penalties far above the optimum do far less damage to welfare than penalties far below the optimum, an observation that we will come back to later. Exhibit 2 can also be used to measure the welfare consequences of setting the penalty rate suboptimally. When the penalty rate is set suboptimally, the agent loses less than 0.01 utils, which, from a money metric perspective represents only 1% of the agent's lifetime resources. In other words, suboptimal penalties have only modest welfare consequences in an economy with $\beta = 0.7$ agents.

Exhibit 3 is closely related to Exhibit 2. In Exhibit 3, as in Exhibit 2, we plot the penalty, π , on the horizontal axis. But now we plot the planner's optimal allocation to the liquid (x) and illiquid accounts (z) on the vertical axis. We also plot the equilibrium value of expected (paid) penalties: $\pi \times Ew$. Note that the government's budget constraint implies that $x + z = \pi Ew + 1$. Exhibit 3 shows that higher values of π lead the planner to put fewer resources in the illiquid account. Exhibit 3 also shows that the amount of expected penalties is non-monotonic. This hump-shaped pattern for expected penalties arises because there are two offsetting effects.

$$\frac{d[\pi Ew]}{d\pi} = Ew + \pi \frac{d[Ew]}{d\pi}.$$

The first effect is positive, since $Ew \geq 0$ by definition. The second effect is negative, since $\frac{d[Ew]}{d\pi} \leq 0$. The first effect dominates when π is in a neighborhood of 0. The

second effect dominates when Ew is small (and π is large). This generates the hump-shaped pattern for expected penalty payments. Note too that penalty payments are always a very small fraction of economic activity in this calibrated example, even for calibrations that maximize the amount of paid penalties (i.e., $\pi = 0.2$).

Exhibit 4 explores an extreme case $\beta = 0.1$ (for illustrative purposes). Here too we assume homogeneous preferences, so we assume that all agents have this extreme parameter value. Again we illustrate how variation in the penalty affects welfare. Welfare is optimized at a penalty of essentially 100%. Exhibit 4 can also be used to measure the welfare consequences of setting the penalty rate suboptimally. When the penalty rate is set suboptimally, the agent loses 0.8 utils, which, from a money metric perspective represents 80% of the agent's lifetime resources. Hence, for agents with extreme impatience (e.g., $\beta = 0.1$), the wrong early withdrawal penalty has dire consequences.

In Exhibit 5 we continue to study the extreme case $\beta = 0.1$. As in Exhibit 3, we plot the planner's optimal allocation to the liquid (x) and illiquid accounts (z) on the vertical axis (the jagged parts of the curve are generated by computational distortions resulting from inadequately fine partition sizes). We also plot the equilibrium value of expected (paid) penalties: $\pi \times Ew$. Exhibit 5 shows that higher values of π lead the planner to put fewer resources in the illiquid account. Exhibit 5 also shows that the amount of expected penalties is non-monotonic (for the same reasons discussed with respect to Exhibit 3). Finally, note that penalty payments are now a more substantial fraction of economic activity. For calibrations that maximize the amount of paid penalties (i.e., $\pi = 0.6$), penalty payments represent approximately 20% of total lifetime economic resources.

In Exhibit 6, we study a wide range of homogeneous preference cases. Each line represents a separate case, with β varying from 0.1 to 1, in steps of 0.1. The bottom line reproduces the $\beta = 0.1$ case that we have already discussed. The fourth line from the top represents the $\beta = 0.7$ case, which we have also already discussed. The top line represents the case of dynamically consistent preferences ($\beta = 1$). This figure illustrates that the welfare gains for low β agents swamp the welfare gains for high β agents. From a welfare perspective, it barely matters where the penalty is set for agents with $\beta \geq 0.7$. But it is enormously costly to set the penalty incorrectly

for agents with $\beta \leq 0.4$. We'll return to these issues when we consider heterogeneous economies next.

5. SOLUTION TO THE PLANNER'S PROBLEM IN THE CASE OF HETEROGENEOUS PREFERENCES IN THE PRESENT-BIAS PARAMETER β (AND HETEROGENEOUS TASTE SHOCKS).

We now study the case of a heterogeneous economy. We use a crude benchmark, which gives uniform weight to each of the 10 types that we studied in the previous section: $\beta \in \{0.1, 0.2, \dots, 1.0\}$. We assume that the government doesn't know who is who, or, even if it can discriminate in principle, won't discriminate in practice. Such uniformity is the norm in modern savings systems: in other words, people who are deemed to have the worst self-control problems are not singled out for different treatment.

For the aggregate (heterogeneous) population, the government will pick a single triple $\{x, z, \pi\}$, where, as before x is the amount allocated to the liquid account, z is the amount allocated to the illiquid account, and π is the penalty rate for early withdrawals from the illiquid account.

Exhibit 7 plots the welfare level for each agent as the penalty rate, π , is varied from 0 to 1. For each level of π , the socially optimal levels of $x(\pi)$ and $z(\pi)$ are chosen. These are plotted in Exhibit 8, which is analogous to Exhibits 3 and 5. Exhibit 8 also plots the equilibrium expected penalty payment (for each level of π). Exhibit 7, which plots the heterogeneous consumer case, is subtly different from Exhibit 6, which plots the homogeneous consumer case. For example, in Exhibit 7, high- β households are made significantly better off as the penalty rate, π , is raised from 0 to 0.5. This effect is due to the fact they are being cross-subsidized by the low- β households, who are paying substantial penalties, as shown in Exhibit 9. Those penalty payments relax the government's budget constraint, enabling the government to give more resources to all households. The substantial cross-subsidies also explain why the welfare gains for low- β households are so muted for low levels of π in Exhibit 7. As π rises, the low- β households have access to a better commitment technology, but they are funneling money (on net) to the high- β members of society.

These cross-subsidies begin to wain as the penalty rate gets high enough to even

discourage the low- β households from using the illiquid account in period 1. Exhibit 9 illustrates the low- β households essentially stop making early withdrawals when the penalty rate, π , rises about 0.9.

Exhibit 10 puts everything together by calculating total welfare (with population weights) as a function of the penalty rate π . Social welfare is maximized at a very high penalty rate: essentially $\pi = 1$. In other words, social welfare is maximized when the system is tuned to serve the interests of the households with the most severe present bias. Of course, high penalty rates reduce the welfare of the high- β households, but this effect is small. For example, moving a $\beta = 1.0$ household from a $\pi = 0$ economy to a $\pi = 1.0$ economy, reduces the welfare of the $\beta = 1.0$ agent by approximately 1% (on a money metric basis). But the welfare of the $\beta = 0.1$ rises by over 80% (on a money metric basis). Hence, the welfare gains of the low- β types swamp the welfare losses of the high- β types.

6. CONCLUSIONS

This paper studies the socially optimal level of illiquid financial accounts in a retirement savings system. We study an environment in which time-inconsistent agents face a tradeoff between commitment and flexibility (cf. Amador, Werning and Angeletos 2006). For this analysis, we assume that the agent has access to two accounts: a perfectly liquid account and an illiquid retirement savings account with an early withdrawal penalty (π).

When agents have homogeneous present-biased preferences, we find that the socially optimal retirement savings account should be relatively liquid, with a penalty that is approximately equal to the level of present bias, $1 - \beta$. In this case, the penalty is set to offset the present bias of the representative agent. For example, if $\beta = 0.7$, then the socially optimal early withdrawal penalty rate is approximately 30%. Likewise, if $\beta = 0.55$, then the socially optimal early withdrawal penalty rate is approximately 45%.

However, when agents have heterogeneous preferences, with a range of β values, we find that optimal policy disproportionately addresses the needs of low β agents. In an illustrative simulation with ln utility and β values distributed uniformly between 0.1 and 1 (with a mean β value of 0.55), we find that the optimal system is charac-

terized by a retirement savings account that is essentially perfectly illiquid (i.e., with a penalty rate of $\pi \simeq 100\%$). In other words, our analysis with heterogeneous preferences suggests that savings should be divided between two accounts: one account that is completely liquid and one account that is completely illiquid (like a defined benefit pension plan).

If our theoretical results prove to be robust, it might be beneficial to create a new type of completely illiquid (defined contribution) savings account that is used in parallel with the existing low-or-no-penalty retirement savings account. On the other hand, Social Security might already provide a socially optimal level of such completely illiquid savings. More work is needed to quantitatively evaluate the adequacy of highly illiquid savings in the current U.S. retirement savings system.

For many reasons, it is premature to apply our findings to the design of a practical retirement savings system. Various aspects of our theoretical analysis may not generalize to practical retirement savings decisions. Future work should (i) extend our 3-period model to a general N -period model, (ii) add additional accounts (so that an agent could have a liquid account, a partially illiquid account, and a fully illiquid account, much like the current U.S. retirement savings system)⁵; (iii) consider the possibility that present-bias, $1 - \beta$, is correlated with observable characteristics, like income; and (iv) add other kinds of intertemporal taste shifters (e.g., rather than assuming that period-by-period utility is given by $\theta_t u(c_t)$, we could instead assume that the taste shock is ‘inside’ the utility function, so that $u(c_t - \theta_t)$). Most importantly, future work should also consider other kinds of self-control problems and/or cognitive errors.⁶

We hope that this paper encourages further research on the question of how the retirement savings system should be designed to maximize social welfare.

⁵In preliminary analysis on this issue, we have found that adding a partially illiquid account generates no additional welfare benefit (in the heterogeneous β case). In other words, it is not socially optimal to introduce the partially illiquid account.

⁶E.g., Bernheim and Rangel (2004), Fudenberg and Levine (2006), and Gul and Pesendorfer (2001).

7. REFERENCES

- Amador, Manuel, Ivan Werning, and George-Marios Angeletos, 2006. "Commitment vs. Flexibility." *Econometrica* 74, pp. 365-396.
- Ambrus, Attila and Georgy Egorov, forthcoming. "Commitment-Flexibility Trade-off and Withdrawal Penalties." *Econometrica*.
- Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, Stephen Weinberg, 2001. "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives* 15(3), pp. 47-68.
- Argento, Robert, Victoria L. Bryant and John Sabelhaus, forthcoming, "Early Withdrawals from Retirement Accounts During the Great Recession," *Contemporary Economic Policy*.
- Ariely, Dan and Klaus Wertenbroch, 2002. "Procrastination, Deadlines, and Performance: Self-Control by Precommitment." *Psychological Science* 13, pp. 219-224.
- Ashraf, Nava, Dean Karlan, and Wesley Yin, 2006. "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Phillipines." *Quarterly Journal of Economics* 121, pp. 635-672.
- Augenblick, Ned, Muriel Niederle, and Charles Sprenger, 2012. "Working Over Time: Dynamic Inconsistency in Real Effort Tasks." Working paper.
- Bernheim, B. Douglas and Antonio Rangel, 2004. "Addiction and Cue-Triggered Decision Processes." *American Economic Review* 94, pp. 1558-1590.
- Colin Camerer, Samuel Issacharoff, George Loewenstein, Ted O'Donoghue & Matthew Rabin. 2003. "Regulation for Conservatives: Behavioral Economics and the Case for 'Asymmetric Paternalism'." 151 *University of Pennsylvania Law Review* 101: 1211-1254.
- Chow, Vinci Y. C. "Demand for a Commitment Device in Online Gaming." Working paper.

- DellaVigna, Stefano and Ulrike Malmendier, 2006. "Paying Not to Go to the Gym." *American Economic Review* 96, pp. 694-719.
- DellaVigna, Stefano and Daniele Paserman, 2005. "Job Search and Impatience." *Journal of Labor Economics* 23, pp. 527-588.
- Dupas, Pascaline and Jonathan Robinson, forthcoming. "Why Don't the Poor Save More? Evidence from Health Savings Experiments." *American Economic Review*.
- Fudenberg, Drew and David K. Levine, 2006. "A Dual-Self Model of Impulse Control." *American Economic Review* 96, pp. 1449-1476.
- Gine, Xavier, Dean Karlan, and Jonathan Zinman, 2010. "Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation." *American Economic Journal: Applied Economics* 2, pp. 213-235.
- Gul, Faruk and Wolfgang Pesendorfer, 2001. "Temptation and Self-Control." *Econometrica* 69, pp. 1403-1435.
- 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, 2010. "Temptation and Commitment in the Laboratory." Working paper.
- Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan, 2010. "Self-Control at Work: Evidence from a Field Experiment." Working paper.
- Laibson, David, 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112, pp. 443-477.
- Laibson, David, Andrea Repetto, and Jeremy Tobacman, 2007. "Estimating Discount Functions with Consumption Choices over the Lifecycle." Working paper.

- Loewenstein, George, 1996. "Out of Control: Visceral Influences on Behavior." *Organizational Behavior and Human Decision Processes* 65, pp. 272-292.
- Loewenstein, George and Drazen Prelec, 1992. "Anomalies in Intertemporal Choice: Evidence and an Interpretation." *Quarterly Journal of Economics* 107, pp. 573-597.
- McClure, Samuel M., David Laibson, George Loewenstein, and Jonathan D. Cohen, 2004. "Separate Neural Systems Value Immediate and Delayed Monetary Rewards." *Science* 306, pp. 503-507.
- McClure, Samuel M., Keith M. Ericson, David Laibson, George Loewenstein, and Jonathan D. Cohen, 2007. "Time Discounting for Primary Rewards." *Journal of Neuroscience* 27, pp. 5796-5804.
- Meier, Stephan and Charles Sprenger, 2010. "Present-Biased Preferences and Credit Card Borrowing." *American Economic Journal: Applied Economics* 2, pp. 193-210.
- Milkman, Katherine L., Julia A. Minson, and Kevin G. M. Volpp, 2012. "Holding the Hunger Games Hostage at the Gym: An Evaluation of Temptation Bundling." Working paper.
- Milkman, Katherine L., Todd Rogers, and Max H. Bazerman, 2008. "Harnessing Our Inner Angels and Demons: What We Have Learned About Want/Should Conflicts and How That Knowledge Can Help Us Reduce Short-Sighted Decision Making." *Perspectives on Psychological Science* 3, pp. 324-338.
- Milkman, Katherine L., Todd Rogers, and Max H. Bazerman, 2009. "Highbrow Films Gather Dust: Time-Inconsistent Preferences and Online DVD Rentals." *Management Science* 55, pp. 1047-1059.
- Milkman, Katherine L., Todd Rogers, and Max H. Bazerman, 2010. "I'll Have the Ice Cream Soon and the Vegetables Later: A Study of Online Grocery Purchases and Order Lead Time." *Marketing Letters* 21, pp. 17-36.

- O'Connor, Kathleen M., Carsten K. W. De Dreu, Holly Schroth, Bruce Barry, Terri R. Lituchy, and Max H. Bazerman, 2002. "What We Want To Do Versus What We Think We Should Do: An Empirical Investigation of Intrapersonal Conflict." *Journal of Behavioral Decision Making* 15, pp. 403-418.
- O'Donoghue, Ted and Matthew Rabin, 1999. "Doing It Now or Later." *American Economic Review* 89, pp. 103-124.
- Oster, Sharon M. and Fiona M. Scott Morton, 2005. "Behavioral Biases Meet the Market: The Case of Magazine Subscription Prices." *Berkeley Electronic Press: Advances in Economic Analysis & Policy* 5(1).
- Phelps, E. S. and R. A. Pollak, 1968. "On Second-Best National Saving and Game-Equilibrium Growth." *Review of Economic Studies* 35, pp. 185-199.
- Read, Daniel and Barbara van Leeuwen, 1998. "Predicting Hunger: The Effects of Appetite and Delay on Choice." *Organizational Behavior and Human Decision Processes* 76, pp. 189-205.
- Read, Daniel, George Loewenstein, and Shobana Kalyanaraman, 1999. "Mixing Virtue and Vice: Combining the Immediacy Effect and the Diversification Heuristic." *Journal of Behavioral Decision Making* 12, pp. 257-273.
- Redelmeier, Donald A. and Eldar Shafir, 1995. "Medical Decision Making in Situations That Offer Multiple Alternatives." *Journal of the American Medical Association* 273, pp. 302-305.
- Reuben, Ernesto G., Paola Sapienza, and Luigi Zingales, 2009. "Procrastination and Impatience." Working paper.
- Royer, Heather, Mark Stehr, and Justin Sydnor, 2012. "Incentives, Commitments and Habit Formation in Exercise: Evidence from a Field Experiment with Workers at a Fortune-500 Company." Working paper.
- Shafir, Eldar, Itamar Simonson, and Amos Tversky. "Reason-based choice." *Cognition* 49.1 (1993): 11-36.

Shapiro, Jesse, 2005. "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle." *Journal of Public Economics* 89, pp. 303-325.

Simonson, Itamar, 1989. "Choice Based on Reasons: The Case of Attraction and Compromise Effects." *Journal of Consumer Research* 16, 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=2011040401> accessed April 18, 2011.

Strotz, R. H., 1955. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies* 23, pp. 165-180.

Thaler, Richard H. and H. M. Shefrin, 1981. "An Economic Theory of Self-Control." *Journal of Political Economy* 89, pp. 392-406.

Wertenbroch, Klaus, 1998. "Consumption Self-Control by Rationing Purchase Quantities of Virtue and Vice." *Marketing Science* 17, pp. 317-337.

Exhibit 1: Socially optimal penalty (π) on early withdrawals from the illiquid account as a function of the present bias parameter (β) for a homogeneous population of agents.

Simulation with homogeneous population of agents with β parameter given by horizontal axis and taste shocks distributed with Gaussian density truncated at zero.

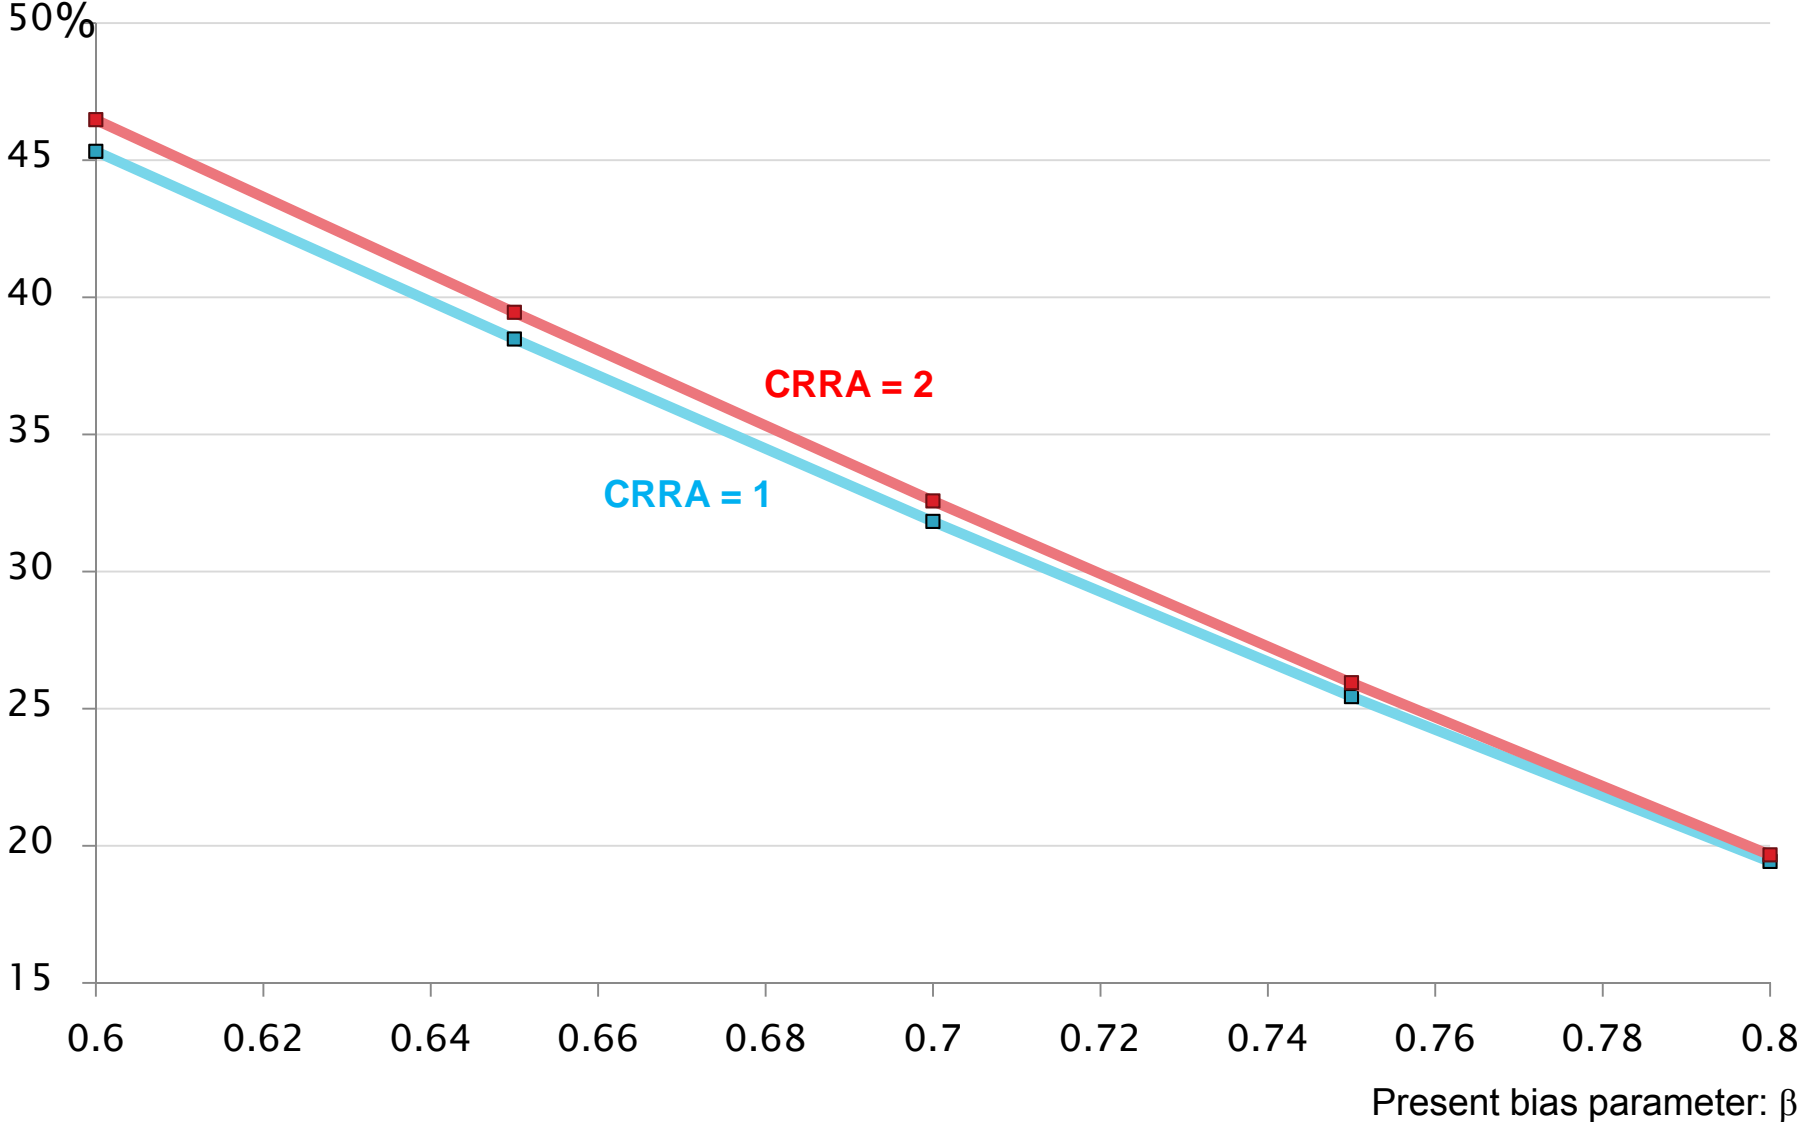


Exhibit 2: Expected utility as a function of the early withdrawal penalty (π) on the illiquid account.

Simulation with In utility and $\beta=0.7$




Exhibit 3: Allocations to the liquid account (blue) and the illiquid account (green) as a function of the early withdrawal penalty (π) on the illiquid account. Expected penalties paid are reported in the bottom line (red).

(Simulation with ln utility and $\beta=0.7$)




Exhibit 4: Expected utility as a function of the early withdrawal penalty (π) on the illiquid account.

(Simulation with ln utility and $\beta=0.1$)




Exhibit 5: Allocations to the liquid account (blue) and the illiquid account (green) as a function of the early withdrawal penalty (π) on the illiquid account. Expected penalties paid are reported in the bottom line (red).

(Simulation with ln utility and $\beta=0.1$)

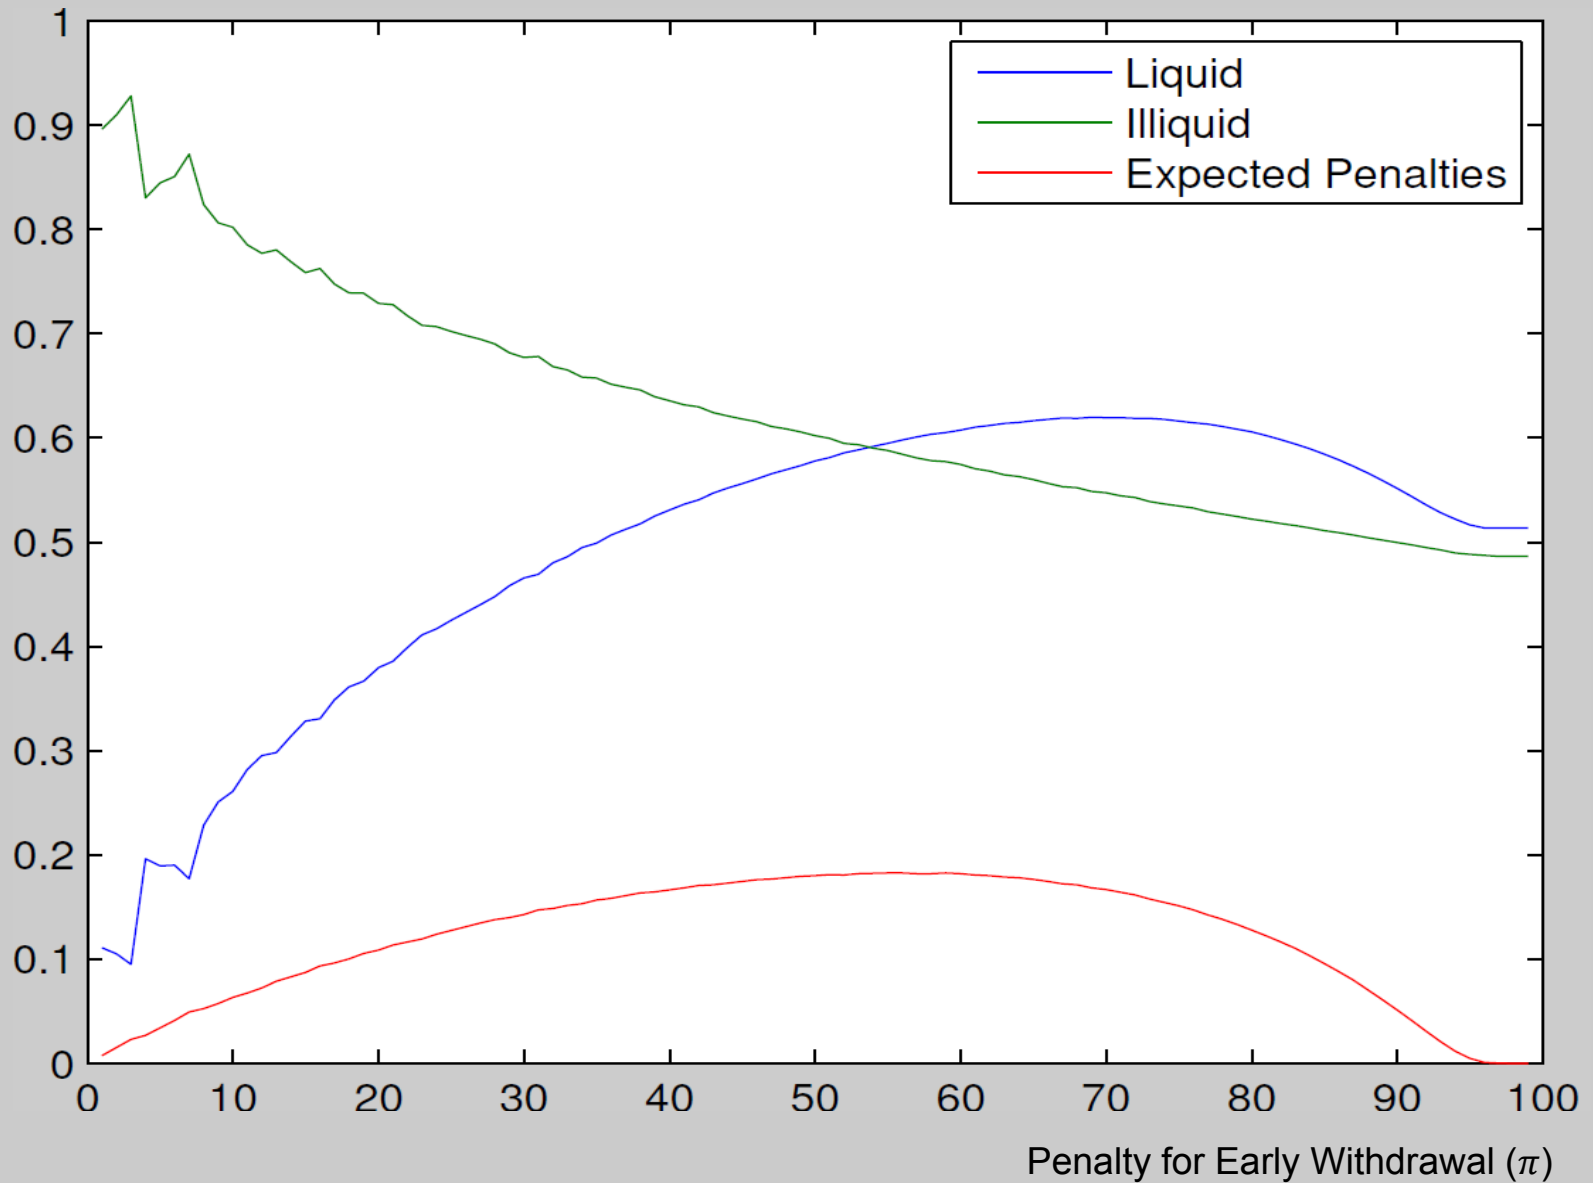


Exhibit 6: Expected utility as a function of the early withdrawal penalty (π) on the illiquid account.

(Simulation with \ln utility and a homogeneous population of agents: β values ordered on right column)




Exhibit 7: Expected utility as a function of the early withdrawal penalty (π) on the illiquid account.

(Simulation with In utility and a heterogeneous population of agents: β values ordered on right column)




Exhibit 8: Allocations to the liquid account (blue) and the illiquid account (green) as a function of the early withdrawal penalty (π) on the illiquid account. Expected penalties paid are reported in the bottom line (red).

(Simulation with ln utility and heterogeneous population with respect to present bias parameter)




Exhibit 9: Expected penalties paid by each β type (ordered in label).

(Simulation with \ln utility and heterogeneous population with respect to present bias parameter)

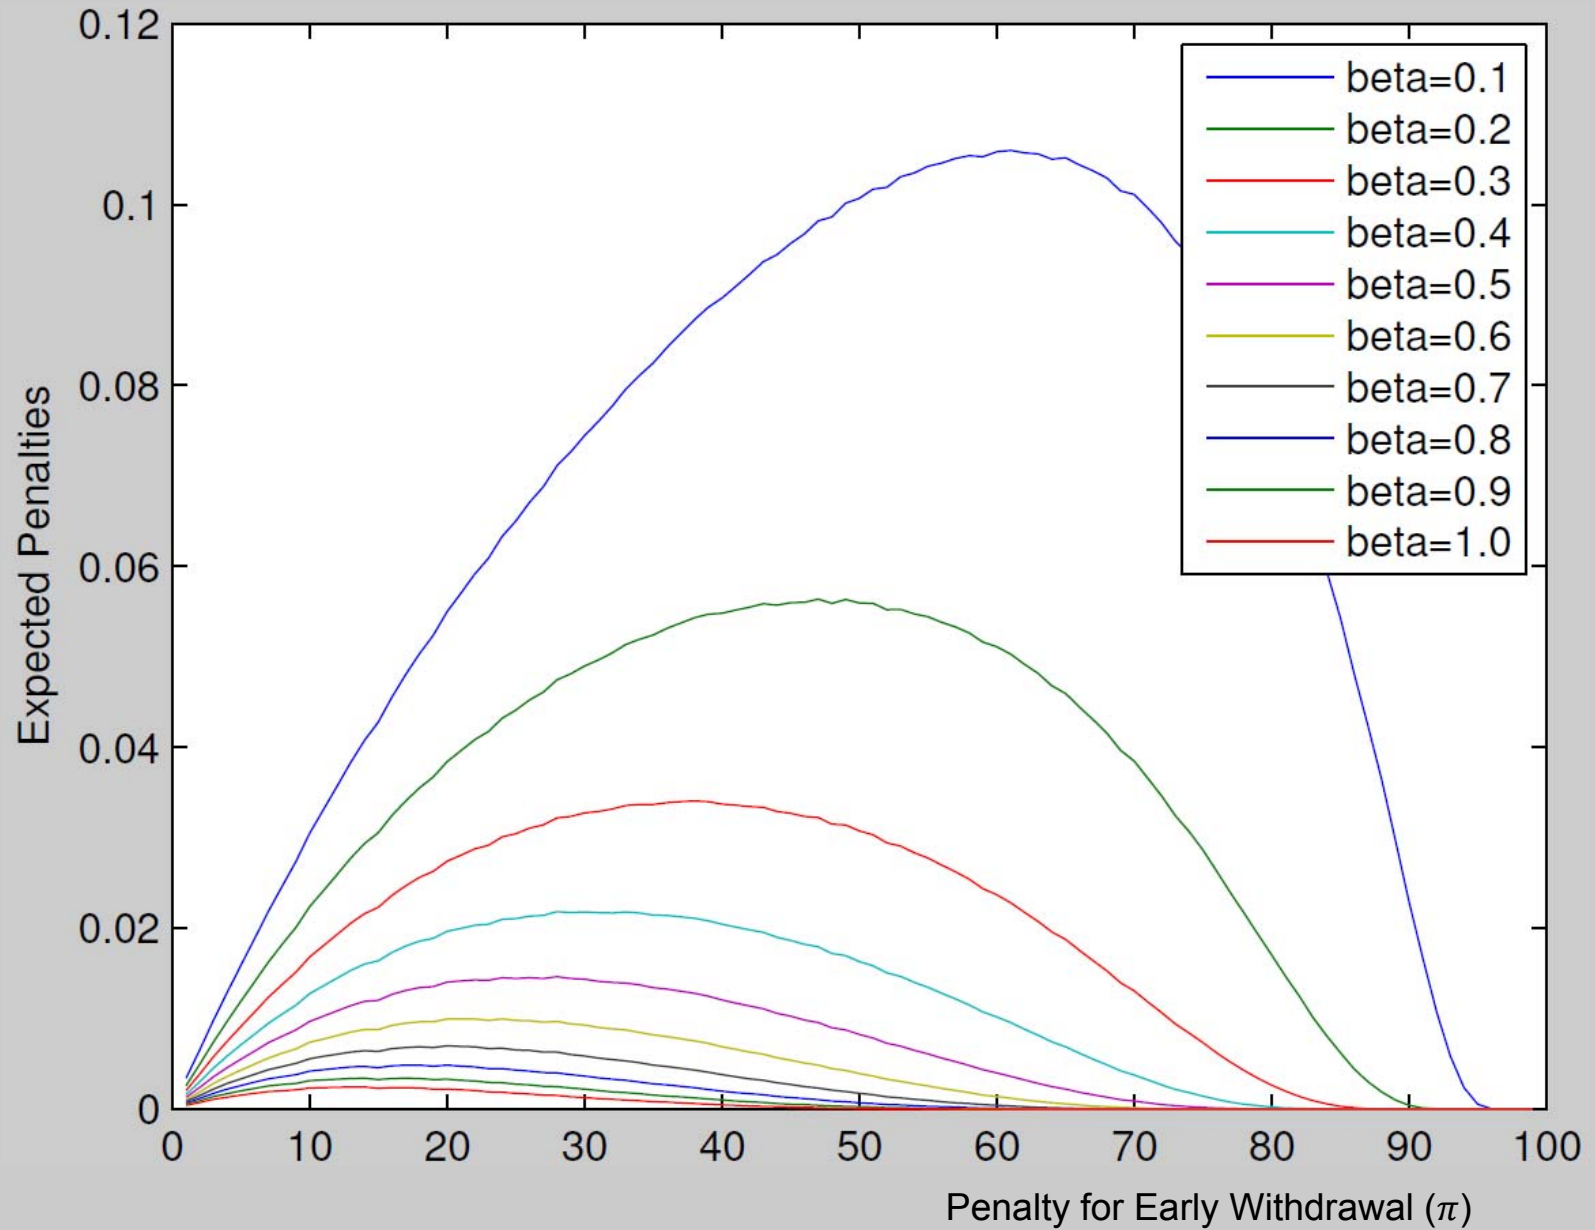


Exhibit 10: Expected Utility For Total Population

(Simulation with ln utility and heterogeneous population with respect to present bias parameter)

