Optimal Illiquidity

John Beshears
Harvard University and NBER

James Choi
Yale University and NBER

Christopher Clayton
Harvard University

Christopher Harris
University of Cambridge

David Laibson
Harvard University and NBER

Brigitte C. Madrian
Harvard Kennedy School and NBER

Prepared for the 19th Annual Joint Meeting of the Retirement Research Consortium
August 3-4, 2017
Washington, DC

The research reported herein was pursuant to a grant from the U.S. Social Security Administration (SSA), funded as part of the Retirement Research Consortium. The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the federal government, Harvard University, Yale University, the University of Cambridge, or the NBER Retirement Research Center.

We calculate the socially optimal level of illiquidity in a stylized retirement savings system. We solve the planner's problem in an economy in which time-inconsistent households face a tradeoff between commitment and flexibility (Amador, Werning and Angeletos, 2006). We assume that the planner can set up multiple accounts for households: a perfectly liquid account and partially illiquid retirement savings accounts with early-withdrawal penalties. Revenue from the penalties is collected by the government and redistributed through the tax system. We solve for the socially optimal values of these penalties and the socially optimal allocations to these accounts. When agents have heterogeneous present-biased preferences, social optimality is achieved with three accounts: (1) a liquid account; (2) an account with an early-withdrawal penalty of ≈100 percent; and (3) an account with an early-withdrawal penalty of ≈10 percent. With heterogeneous preferences, the socially optimal retirement savings system in our stylized model looks surprisingly like the existing U.S. system: (1) a liquid account; (2) an illiquid Social Security account (and defined benefit pensions); and (3) a 401(k)/IRA account with a 10-percent penalty. The socially optimal allocations to these accounts and the predicted equilibrium flows of early withdrawals – "leakage" – also match the U.S. system.

Summary of Paper

How much liquidity should be built into a socially optimal savings system? On one hand, flexibility allows households to consume in ways that reflect their idiosyncratic preferences – i.e., households can respond to taste shocks and taste shifters. However, liquidity allows households with self-control problems (and other biases or errors) to over-consume.

If illiquidity is optimal, how should it be implemented? Possible forms of illiquidity include a perfectly illiquid retirement claim (like a typical defined benefit pension or Social Security) or a partially illiquid account (like an IRA or 401(k) plan). In theory, an optimal system might combine different types of illiquid accounts.

In the domain of practical policies, there is a partial consensus on these questions. Almost all developed countries have some form of compulsory saving that is *completely* illiquid (e.g., Social Security in the United States).

But that's where agreement ends. For example, in most developed countries, defined contribution (DC) savings accounts are usually completely illiquid before age 55 (Beshears et al.

2015). By contrast, in the United States, certain types of withdrawals from DC accounts are allowed without penalty, and, for IRAs, withdrawals may be made for any reason if a 10-percent penalty is paid. Liquidity allows significant pre-retirement "leakage:" for every $1 contributed to the accounts of U.S. savers under age 55, $0.40 simultaneously flows out of the 401(k)/IRA system, not counting loans (Argento, Bryant, and Sabelhaus 2014). Until now, no normative model has been used to determine whether such leakage is good or bad from the perspective of overall social welfare. Nevertheless, most policy analysis bemoans leakage (e.g., Hewitt Associates, 2009).

Our paper evaluates the optimality of an *N*-account retirement savings system with a combination of liquid, partially illiquid, and/or fully illiquid accounts. Within this framework, we focus on two special cases: systems with two accounts and systems with three accounts. In all of our analyses, we will assume that the first account is fully liquid, so our two-account system has a fully liquid account and a partially (or fully) illiquid account. Likewise, our three-account system has a fully liquid account and two partially illiquid accounts (one of which might be fully illiquid). We show that the three-account system is a good approximation (with respect to expected welfare) for a completely general system with an arbitrary number of accounts.

We study preferences that include both normatively legitimate taste shifters and normatively undesirable 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 $\{1, \beta\delta^2, \ldots, \beta\delta^t\}$, where the degree of present bias is $1-\beta$. Present bias is the propensity to overweight the present relative to the future. Our model is an aggregate version (with interpersonal transfers) of the flexibility/commitment framework of Amador, Werning, and Angeletos (2006).

We divide our analysis into the cases of homogeneous present bias and heterogeneous present bias. In the homogeneous case, we assume that all agents have the same degree of present bias – in other words, the same value of β. Under a homogeneous β, our model implies that partially illiquid accounts with early-withdrawal penalties $\pi \approx 1-\beta$ play an economically significant role in improving social welfare.

We then relax the homogeneity assumption, and assume that agents have *heterogeneous* present-bias. In this heterogeneous-preference case, we find that fully illiquid accounts play an important role in improving welfare, whereas partially illiquid accounts matter relatively little.

We show that the socially optimal degree of illiquidity mostly caters to households with the lowest β values (i.e., the households with the largest amount of present bias, $1 - \beta$). Completely illiquid retirement savings generates large welfare gains for these low-β agents, and these welfare gains swamp the welfare losses of the high-β agents (who are made slightly worse off by shifting some of their wealth from perfectly liquid accounts to perfectly illiquid accounts).

To the extent that there is *also* a role for partially illiquid accounts in the heterogeneous-β economy, we find that they should have low early-withdrawal penalties – approximately 10 percent. This implies that the partially illiquid accounts look much like a typical 401(k) account. Moreover, these partially illiquid accounts display a high level of leakage in equilibrium. In other words, early withdrawals (i.e., pre-retirement withdrawals) are common from this partially illiquid account. This leakage is a two-edged sword: it results in part from legitimate taste shocks and in part from self-control problems (i.e., low β). The costs of the partially illiquid account to low β types (who end up paying most of the early withdrawal penalties) and benefits to high β types (who are net recipients of these penalties) are nearly off-setting, although the net effect for all of society is slightly positive. This analysis suggests that the U.S. system, which exhibits high leakage in practice, is not necessarily suboptimal (though it is "second-best" because of information asymmetries).

In summary, three findings emerge from the analysis of our stylized two-period model (which allows for mechanisms that admit interpersonal transfers and incorporates heterogeneity in present bias):

1. The constrained-efficient social optimum is well-approximated by a two-account system, with one account that is completely liquid and a second account that is completely illiquid. Little welfare gain is obtained by moving beyond this simple two-account system.
2. If a third account is added, its optimized early-withdrawal penalty is between 6 percent and 13 percent.
3. In equilibrium, the leakage rate from this (partially illiquid) third account ranges from 65 percent to 90 percent.

These properties have analogs in the U.S. retirement savings systems. The United States has fully liquid accounts (i.e., a standard checking account), fully illiquid accounts (i.e., Social

Security), and a partially illiquid defined contribution system with a 10-percent penalty for early withdrawals. This partially illiquid DC system has a leakage rate of 40 percent.

Despite these similarities, it is inappropriate to conclude that our findings demonstrate the social optimality of the U.S. system. Our simulation framework is highly stylized. For example, we assume only two periods (e.g., working life and retirement). We assume a particular form of multiplicative taste shifters.[1] We assume that households are naive with respect to their present bias. We study a limited set of distributions of the present bias parameter, β.[2] We only study N-account retirement savings systems (instead of studying a fully general mechanism design framework).[3]

Much more work is needed to interrogate our three main findings. It is not yet clear whether our results – which, to our surprise, seem to rationalize the fundamental structure of the U.S. retirement savings system – will continue to hold as we enrich and expand our analysis.

References

Amador, Manuel, Ivan Werning, and George-Marios Angeletos. 2006. "Commitment vs. Flexibility." *Econometrica* 74: 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): 1-16.

Beshears, John, James J. Choi, Joshua Hurwitz, David Laibson, Brigitte C. Madrian. 2015 "Liquidity in Retirement Savings Systems: An International Comparison." *The American Economic Review* 105(5): 420-425.

Hewitt Associates. 2009. "The Erosion of Retirement Security from Cash-outs: Analysis and Recommendations." Lincolnshire, IL.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112: 443-477.

Phelps, E. S. and R. A. Pollak. 1968. "On Second-Best National Saving and Game-Equilibrium Growth." *Review of Economic Studies* 35: 185-199.

[1] We assume $\theta u(c)$, but we could have instead assumed $u(c - \theta)$. In ongoing work, we are studying this case.
[2] Little is known about the empirical distribution of present bias.
[3] Future drafts of this paper will contain such analysis.