Measuring the Valuation of Liquidity with Penalized Withdrawals*

David Coyne[†], Itzik Fadlon[‡], and Tommaso Porzio[§]

Abstract

We introduce penalized withdrawals from retirement savings accounts as a new revealed-preference tool to measure households' valuation of liquidity. This approach addresses key empirical challenges by providing a proxy for marginal utility without requiring assumptions on preferences or data on access to credit. Using U.S. administrative tax data from 1999-2018, we document three main findings. First, geographic differences, driven by local credit supply, account for over 30% of the nationwide variation in the valuation of liquidity across labor markets. Second, areas hit hardest by the Great Recession saw large increases in the valuation of liquidity, with local credit market spillovers explaining two-thirds of the effect. Third, Black households rely more heavily on penalized withdrawals, even after controlling for income and location, consistent with limited access to formal credit. Together, these findings highlight significant scope for welfare gains through more targeted safety-net policies and demonstrate the practical value of penalized withdrawals as a tool for monitoring liquidity needs.

*We thank for useful comments Adrien Auclert, Corina Boar, Marika Cabral, Raj Chetty, James Choi, Taha Choukhmane, Julie Cullen, Ellora Derenoncourt, Amy Finkelstein, Damon Jones, Andreas Kostøl, David Laibson, Brigitte Madrian, Ulrike Malmendier, Ellen McGrattan, Peter Maxted, Davide Melcangi, Ben Moll, Laura Pilossoph, Andreas Schaab, Gianluca Violante, Gal Wettstein, Johannes Wieland, Steve Zeldes, and seminar participants at Arizona State University, Columbia, Princeton, UT Austin, UCSD, the OTA Research Conference, the NBER Public Economics Meeting (Fall 2020), and the NBER Household Finance and Aging Meetings (Summer Institute 2021) for insightful comments and discussions of the paper. Philip Nye provided excellent research assistance. This research was conducted while David Coyne was an employee at the U.S. Department of the Treasury. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors and do not necessarily reflect the views or the official positions of the U.S. Department of the Treasury. Any taxpayer data used in this research was kept in a secured IRS data repository, and all results have been reviewed to ensure that no confidential information is disclosed.

[†]U.S. Department of the Treasury; email: David.Coyne@Treasury.gov.

[‡]University of California, San Diego and NBER; email:fadlon@ucsd.edu.

[§]Columbia University, CEPR, and NBER; email: tommaso.porzio@columbia.edu.

1 Introduction

In a world without borrowing constraints, households would fully smooth their marginal utility of consumption over time. Any deviations from this ideal would reflect only aggregate shocks or permanent changes in income or consumption. In reality, however, households face limited access to liquidity from formal or informal credit sources (Parker 1999; Johnson et al. 2006). As a result, the extent to which marginal utility today exceeds expected marginal utility tomorrow—what we term the valuation of liquidity—may vary across households. This variation signals potential misallocation of resources and opens the door to welfare gains from reallocating liquidity toward those with higher valuations. From a macroeconomic perspective, these gaps in marginal utility map directly into differences in marginal propensities to consume and thus modulates the impact on aggregate demand of stabilization policies, such as quantitative easing or temporary access to illiquid assets (Cui and Sterk 2021).

Measuring differences in households' valuation of liquidity is empirically challenging for two main reasons. *First*, the utility-based valuation of liquidity is not directly observable, despite recent major advances in consumption measurement.[1] The reason is that, even with perfect consumption data, linking fluctuations in consumption to utility requires the strong assumption that preferences are fixed. Yet, preferences likely evolve with life events—such as health shocks—that change households' circumstances. While this kind of state dependence in preferences has important welfare implications, it remains extremely difficult to estimate.[2]

Second, the valuation of liquidity is fundamentally an equilibrium object. A household's willingness to pay for an extra dollar of liquidity depends not only on its own circumstances but also on the supply of credit they face. Thus, observing consumption or income shocks alone provides an incomplete picture, as these observations miss supply-side changes (e.g., credit crunches). Fully capturing a household's valuation would require granular data on the credit they could access at any given time—a level of detail that, to our knowledge, is unavailable at a national scale.

This paper introduces a revealed-preference approach that sidesteps these challenges. The core idea is simple: when a household pays a known cost for a financial product—in our case, borrows at a known marginal cost—they reveal that their valuation of liquidity is at least as high. We identify one such product: penalized early withdrawals from retirement accounts. These withdrawals are widely accessible and carry an explicit and observable marginal price: an early withdrawal penalty 10%. Based on this insight, we then use US administrative tax

[1] These include methods that infer consumption from budget residuals using rich income and wealth data (e.g., De Giorgi et al. 2019; Kolsrud et al. 2024) and from bank transaction data (e.g., Ganong et al. 2024).

[2] See, e.g., Finkelstein et al. (2009, 2013); Chetty and Finkelstein (2013); Fadlon and Nielsen (2019); Landais and Spinnewijn (2021); Coyne et al. (2024).

records from 1999-2018 to empirically study American households' valuation of liquidity and its variation across time and space.

We start by developing a simple theoretical framework linking penalized withdrawals to the valuation of liquidity. A two-asset heterogeneous agent model provides a structural interpretation of our empirical measures—namely, the frequency and size of penalized withdrawals. We show that observing a penalized withdrawal reflects instances where households' marginal valuation of liquidity exceeds the penalty. We then characterize the amount of liquidity that would provide sufficient insurance to keep a household's marginal valuation of liquidity bounded by the penalty. We refer to this amount as "penalized liquidity," which is a simple money-metric of under-insurance. The model also emphasizes the need for empirical strategies that incorporate both household and market-level drivers, since liquidity valuation is shaped by equilibrium forces.

Motivated by this simple framework, we proceed in three empirical steps.

In the first step, we examine how household-level factors shape the valuation of liquidity. Using within-household variation over time, we identify life events that represent shocks to liquidity needs, such as unemployment or large income declines, and study how these events trigger penalized withdrawals.[3] We find that adverse shocks sharply increase reliance on penalized withdrawals for liquidity. For instance, unemployment leads to a 10.4 percentage point increase in withdrawal likelihood and about $1,600 in additional withdrawn funds. Even wealthy households, consistent with the wealthy hand-to-mouth concept (Kaplan et al. 2014), show meaningful but attenuated responses. Using the recently developed, and extensively validated, race imputation in IRS data (Cronin et al. 2023, Fisher 2023), we also document racial disparities: Black households are more likely than White households to rely on penalized withdrawals to self-insure. Following unemployment, Black households increase their take-up by 35% more than White households.

In the second step, we explore how market-level conditions influence the valuation of liquidity. We exploit spatial variation in the severity of the Great Recession, a major episode of credit supply contraction.[4] We find that households in more affected areas (as measured by unemployment shocks) increased their use of penalized withdrawals significantly more than those in less affected areas. Decomposing the effect, we show that about two-thirds are due to an indirect channel—e.g., tighter local credit markets—rather than direct household-level

[3] While prior work has studied general leakages from retirement accounts after household shocks (e.g., Goodman et al. 2021), our approach focuses on penalized withdrawals as the theoretical object that reveals the valuation of liquidity. Similarly, we complement studies that focus on withdrawal behavior around age 59.5, which is the age when the penalty disappears (Goda et al. 2018; Rong 2023; Stuart and Bryant 2024).

[4] See Argento et al. (2015) for aggregate evidence on retirement account leakages during the Great Recession. We build on Yagan (2019) by leveraging local labor market shocks across Commuting Zones. Again, we focus on penalized withdrawals, as dictated by our theoretical framework.

shocks to income and employment.

In the third and final step, we bring together household- and market-level factors in a unified framework. Using two-way fixed effects models and a mover design, we estimate how much each set of factors contributes to differences in liquidity valuation. We find that local *place effects* explain roughly a third of the spatial variation in penalized withdrawals. We rule out several alternative explanations—such as changes in the household's economic conditions, tax optimization, or learning—by leveraging the timing and nature of moves. We then estimate and interpret both household and location fixed effects. Location effects correlate strongly with proxies for credit supply, such as local credit insecurity and median home values (which can provide collateral). Meanwhile, household effects are more closely tied to race: even after controlling for geography and income, households with a Black primary earner are 30% (2.9 pp) more likely than White households to rely on penalized withdrawals. This underscores the role of structural barriers to credit access beyond geography.

Taken together, our findings show that penalized withdrawals offer a powerful, yet under-utilized lens into the valuation of liquidity of U.S. households. Our approach reveals that local credit supply plays a central role in shaping these valuations, providing strong support for place-based redistributive policies that go above and beyond income-based redistribution to households (Gaubert et al. 2021). We also uncover persistent racial disparities that cannot be explained by income or location alone, pointing to systemic inequities in credit access.

More broadly, by quantifying when and where households are willing to pay to convert illiquid wealth into liquid resources, our results provide micro-founded guidance for macro stabilization policies that seek to improve consumption smoothing during recessions.

Related Literature. Interest in the valuation of liquidity spans several fields, including public finance and macroeconomics. It connects to foundational questions around insurance and capital market imperfections, liquidity constraints, households' ability to smooth marginal utility, and the optimal design of social insurance (see, e.g., Zeldes 1989; Parker 1999; Souleles 1999; Johnson et al. 2006; Card et al. 2007; Chetty and Finkelstein 2013). It also relates to a growing macro literature that emphasizes the liquidity composition of wealth as a key determinant of policy effectiveness, in both models and quasi-experiments (Cui and Sterk 2021; Kreiner et al. 2019). This paper contributes to this broad literature in three main ways.

First, we propose and validate a new tool to assess households' valuation of liquidity—an essential input for both policy design and calibration of structural models. Our approach addresses two longstanding challenges. One is the difficulty of mapping observed behavior to underlying preferences in a way that allows for heterogeneity and state dependence (see, e.g., Landais and Spinnewijn 2021). The other is that the valuation of liquidity is an equilibrium object, shaped not only by households' needs but also by their access to credit—information that is rarely available at the household level on a population-wide scale. We overcome these challenges by focusing on a clear, measurable household choice that can be observed with a population-level coverage: the decision to make a penalized early withdrawal from a retirement account. This decision directly reveals a household's valuation of liquidity, requires no assumptions about stability of preferences, and naturally reflects underlying heterogeneity—whether due to aging, health shocks, or other life-changing events.

Second, we offer a comprehensive empirical characterization of the valuation of liquidity across the U.S., and we identify its key drivers. In doing so, we contribute to several strands of the literature. We add to the growing body of work highlighting the central role of place in shaping outcomes—including education, earnings, health, and mobility (e.g., Chetty and Hendren 2018a,b; Finkelstein et al. 2016, 2021; Card et al. 2023). We show that place effects explain roughly one-third of the spatial variation in households' valuation of liquidity, offering new evidence that access to liquidity—a core input into household well-being—is shaped by geography.[5] We also contribute to the literature on racial disparities in economic outcomes (e.g., Bayer and Charles 2018; Chetty et al. 2020; Derenoncourt and Montialoux 2021; Derenoncourt et al. 2021; Bartscher et al. 2021). We find that, even after accounting for income and location, Black households display systematically higher valuation of liquidity—consistent with limited access to affordable credit and persistent racial inequities in credit markets.[6]

Third, we provide a new set of empirical moments on how economic shocks relate to the valuation of liquidity—valuable for the growing quantitative macro literature with heterogeneous agents (e.g., Krueger et al. 2016; Kaplan et al. 2018; Auclert 2019; Auclert et al. 2020; Laibson et al. 2021). Our findings reinforce the idea that even high-wealth households can be liquidity constrained (Kaplan et al. 2014) and offer externally validated, targeted moments for model calibration. Moments based on the valuation of liquidity—i.e., deviations from the Euler equation—have a key advantage over those based on consumption or income: they are robust to preference heterogeneity and state dependence. This robustness is especially valuable given recent evidence that accounting for preference differences is essential to match the joint distribution of household-level changes in income and consumption (Parker 2017; Aguiar et al. 2020).

[5] Keys et al. (2020) study geographic variation in financial distress (e.g., collections, defaults, bankruptcy), shedding light on channels that may underlie our findings on spatial differences in the valuation of liquidity.

[6] This aligns with Ganong et al. (2020), who document higher consumption elasticities to income shocks for Black and Hispanic households, pointing to racial differences in smoothing capacity.

Structure of the paper. Section 2 provides institutional background on penalized withdrawals, describes our data, and presents preliminary evidence on households' withdrawal behavior. Section 3 lays out the conceptual framework that links penalized withdrawals to the valuation of liquidity and guides our empirical strategy. The core empirical analysis follows in three parts: Section 4 examines how household-level events shape the valuation of liquidity; Section 5 studies the impact of market-level economic shocks; and Section 6 brings these perspectives together to explore the deeper, structural drivers of liquidity valuation. We discuss the policy implications of our findings in Section 7, and conclude in Section 8.

2 Background, Data, and Motivating Facts

We begin by describing the institutional details underlying penalized early withdrawals from retirement accounts and how we measure them in administrative data. We then present key baseline facts about their use, which motivate our approach of using penalized withdrawals to study households' valuation of liquidity.

2.1 Institutional Setting

A variety of savings vehicles in the U.S. impose restrictions on early access to funds, including Health Savings Accounts (HSAs), Certificates of Deposit (CDs), and, most prominently, retirement accounts—either employer-sponsored 401(k) plans or Individual Retirement Accounts (IRAs). While early withdrawals are permitted, they typically incur a 10 percent penalty if taken before age 59.5, in addition to regular income taxes. This penalty forms the basis of our revealed-preference approach: by choosing to incur it, households signal that the liquidity value of funds today exceeds the cost of accessing them early.

Some early withdrawals are exempt from the penalty, such as rollovers between accounts, permanent disability, death of the account holder, higher education expenses, large medical bills, or first-time home purchases.[7] In our analysis, we focus on distributions that are explicitly coded as penalized and not linked to such exceptions.

Penalized withdrawals are one of many tools that households use to meet urgent liquidity needs. Survey evidence from Lusardi et al. (2011) shows that 11 percent of households report they would draw on retirement savings—even with a penalty—if faced with a $2,000 emergency; ranking above other options like pawn loans, borrowing from friends, or taking a second mortgage.[8] This highlights the relevance of penalized withdrawals as a real-world

[7]See IRS website: https://www.irs.gov/retirement-plans/plan-participant-employee/retirement-topics-tax-on-early-distributions.

[8]The different tools and the share of households who expect to use each tool (given in parentheses)

margin of liquidity access.

2.2 Data

We provide a brief description of our data. More details are in Appendix A.

Main data sources and sample construction. We utilize U.S. administrative tax records from 1999 to 2018, using a 10 percent random sample of Social Security Numbers (SSNs). We link each SSN to tax returns (Form 1040) and construct household-level panels, merging records for spouses where applicable.[9] The tax returns are augmented with third-party information returns, such as Form W-2 (wage income), Form 1099-R (retirement distributions), and Form 5498 (retirement account balances and contributions).

Our main analysis focuses on households who have a primary filer aged 45-59 and for whom we have indication for holding a retirement account. We identify households as having a retirement account in a given year if up to that year (within our sample period of 20 years) they report making a contribution to a 401(k) or an IRA account on Form W-2 or Form 5498, or if they have outstanding balances in IRA accounts as reported on Form 5498. This yields a core sample of roughly 10.5 million households.

Variable definitions. Our key outcome variable is penalized withdrawals from retirement savings accounts—specifically, early distributions from 401(k) plans or IRAs that incur the 10 percent penalty for being taken before age 59.5. We observe these distributions through Form 1099-R, using Box 1 to capture the distribution amount and Box 7 to identify the type of distribution based on its code. Several codes indicate a penalized withdrawal, but identifying whether the penalty was ultimately applied requires additional care as follows.

In some cases, distributions coded as penalized may qualify for exceptions. Plan administrators may lack full information on the reason for withdrawal and, in the absence of further documentation from the account holder, may default to assigning a penalty code. However, taxpayers can later report eligible exceptions (e.g., large medical expenses, higher education costs) via Form 5329. We use this form to adjust our classification and ensure we capture only those withdrawals that were truly subject to the 10 percent penalty. Appendix A provides full details on our classification rules and corrections.

To capture household economic conditions, we mainly rely on Form 1040. We define a household's total income as Adjusted Gross Income (AGI), net of any penalized withdrawals.

are: savings (52.4), family (29.6), work more (22.9), credit cards (20.9), sell possession (18.8), liquidate retirement investments even if penalty is required (11.1), pawn assets (7.7), friends (7.4), unsecured loan (7.1), home equity line of credit (HELOC)/second mortgage (4.3), payday/payroll advance loan (3.6), liquidate investments (2.3), sell home (0.4).

[9] Specifically, in cases where spouses indicate that they are married filing separately, we combine their data to build a single household return comparable to those married filing jointly.

AGI includes wages and salaries, capital income, retirement income, and taxable Social Security benefits. We define employment as having positive labor earnings in a given year, and we identify job separations or switches using employer identifiers (EINs) reported on Form W-2. We also use Form 1099-G to flag unemployment events and Schedule D of Form 1040 to measure capital income.

We extract information on retirement account balances from Form 5498, which reports the fair market value of all IRA holdings at year-end (Box 5). This includes all account assets, regardless of whether they are traded on public markets or have easily observable prices. Form 5498 also provides data on IRA contributions and rollovers.

Location is determined using mailing address information reported annually on Form 1040. We construct household panels that track location across years, allowing us to study geographic variation in both economic conditions and withdrawal behavior.

Finally, we incorporate administrative imputations of race and Hispanic origin using the methodology of Fisher (2023). This approach uses a combination of name, location, family structure, and tax variables to estimate racial and ethnic probabilities. Households are assigned to a race/ethnicity category based on the highest predicted probability. This method has been validated and performs well for identifying Black and Hispanic taxpayers.[10]

Additional data sources. To validate our main sample and perform a few robustness exercises, we draw on the Health and Retirement Study (HRS), a nationally representative longitudinal survey of adult households. We use HRS waves 7-14 (years 2004-2018), focusing on households aged 45-59, to compare overall patterns of prevalence of retirement account ownership and withdrawal behavior with our administrative estimates.

2.3 Baseline Facts on Penalized Withdrawals

We document a few empirical facts about how U.S. households use penalized early withdrawals from retirement savings accounts. These patterns provide the empirical foundation for our approach and motivate the conceptual framework that we develop. As we show, many households make penalized withdrawals, but only rarely and typically in response to adverse economic conditions; thus supporting their use as a revealed-preference tool for recovering the valuation of liquidity. We revisit and formalize this revealed-preference interpretation in our model in Section 3, and return to discussing its behavioral underpinnings in Section 7. Appendix B provides the details of the patterns that we summarize here.

[10]See Cronin et al. 2023 and Costello et al. 2024 for details on data validation and several applications.

Prevalence of retirement accounts. First, we establish that the vast majority of households in our sample have retirement accounts. Among households with a primary filer aged 45-59, nearly 90 percent have at least one retirement account within the sample period. This high prevalence reflects the fact that our unit of analysis is the household, rather than individuals. Using data from the Health and Retirement Study (HRS), we corroborate this estimate: 84.1 percent of households in the same age range report having a defined-contribution account. Since our analysis is based on tax filers, it excludes the lowest-income non-filers who are less likely to hold retirement accounts. Indeed, when we incorporate non-filers, overall prevalence rates decline (modestly) to an average of 83.8 percent within ages 45-59 (see Appendix Figure D.3).

Frequency of withdrawals. Second, we find that penalized withdrawals are used by a non-trivial share of households. Roughly 10 percent of households in our sample make a penalized withdrawal in any given year. However, such withdrawals are typically infrequent: most households only withdraw occasionally, rather than drawing on these accounts routinely. This pattern is consistent with the interpretation of penalized withdrawals as a tool to meet acute short-run liquidity needs, rather than a default source of funding.

Size of withdrawals. Third, penalized withdrawals are substantial in size. The average withdrawal is approximately $5,000. Importantly, most withdrawals are not associated with depleting the entire account balance. This suggests that households are operating at interior solutions with respect to the withdrawal margin, consistent with deliberate decision-making rather than account liquidation. In Appendix B, we show that the distribution of withdrawal amounts closely aligns between our administrative tax data and the HRS survey responses.

Link to income shocks. Finally, we document a strong association between penalized withdrawals and income declines. Households that make penalized withdrawals are significantly more likely to have experienced an income loss in the same year. Among withdrawing households, nearly 60 percent experienced a year-over-year decline in income, and they are more than twice as likely as non-withdrawing households to have suffered a loss of 50 percent or more. This correlation holds even after accounting for observable characteristics and supports the interpretation of withdrawals as a response to liquidity needs triggered by adverse shocks.

3 Conceptual Framework

We develop a simple framework with two goals. First, we formalize the idea that penalized withdrawals measure households' valuation of liquidity, explicitly stating the assumptions needed to map withdrawal behavior into liquidity valuation. Second, we motivate our empirical analysis by illustrating how liquidity valuation emerges as an equilibrium object shaped by market conditions and households' liquidity needs.

3.1 Model Setup

We consider household i in region z, who makes life-cycle consumption decisions. Each period, the household earns income $y_{i,t}$, which it allocates between current consumption $c_{i,t}$ and savings in liquid or retirement accounts. An additional share φ of earnings is automatically deposited into retirement savings. Every period, the household faces a liquidity shock $\varepsilon_{i,t}$ drawn from distribution $F(\varepsilon)$, representing unexpected consumption needs (e.g., unemployment or health shocks; Dobkin et al. 2018; Fadlon and Nielsen 2021).

To fund consumption beyond current income, the household may borrow liquid assets at a cost equal to the risk-free rate r plus a premium $\rho_{i,z}(b_{i,t})$. This premium is specific to household i and region z, and rises with the amount borrowed in that period, $b_{i,t}$. The function $\rho_{i,z}(b_{i,t})$ captures the household's perceived shadow cost of borrowing, accounting for household-specific credit constraints and their knowledge of available financial alternatives.

Households may also withdraw from retirement savings, but withdrawals made before the statutory retirement age (t^*) incur a marginal penalty rate τ, leaving only $1-\tau$ dollars available for consumption per dollar withdrawn. Due to the penalty τ the retirement savings account is effectively illiquid.

We denote balances in liquid and illiquid accounts at the start of period t by $a_{i,t}$ and $k_{i,t}$, respectively, with $\Delta a_{i,t}$ and $\Delta k_{i,t}$ representing net flows across periods (e.g. if a household withdraws money from the illiquid account, then $\Delta k_{i,t} < 0$). If the household starts the period with zero or negative liquid assets, any further reduction in liquid wealth must come from borrowing; thus, borrowing is given by $b_{i,t} = \max\{0; a_{i,t-1} - \Delta a_{i,t}; -\Delta a_{i,t}\}$.

Flow utility, $u(c_{i,t}; h_{i,t})$, depends on a household- and time-specific state vector $h_{i,t}$. This vector may include histories of shocks and factors affecting consumption preferences, such as marital status, fertility, and health. Such flexibility accommodates state-dependent preferences, allowing liquidity demand to respond directly to preference shocks rather than solely through income changes. This level of flexibility highlights a strength of our approach: we directly reveal valuation of liquidity from household behavior without having to rely on structural assumptions that would map behaviors (such as consumption choices) to preferences.

We let $V_t(a_{i,t-1}, k_{i,t-1}; h_{i,t})$ be the value of the problem, which is given by

$$V_t(a_{i,t-1}, k_{i,t-1}; h_{i,t}) = \max_{\Delta k_{i,t}, \Delta a_{i,t}} u(c_{i,t}; h_{i,t}) + \beta E_t [V_{t+1}(a_{i,t}, k_{i,t}; h_{i,t+1})]$$

subject to

$$\begin{aligned}
c_{i,t} &= (1-\varphi) y_{i,t} - \varepsilon_{i,t} - \Delta k_{i,t} - \Delta a_{i,t} + \tau \Delta k_{i,t} \mathbb{I}_{(\Delta k_{i,t}<0)} \mathbb{I}_{(t<t^*)} - \rho_{i,z}(b_{i,t}) \mathbb{I}_{(b_{i,t}>0)} \\
a_{i,t} &= (1+r)[a_{i,t-1} + \Delta a_{i,t}] \\
k_{i,t} &= (1+r)[k_{i,t-1} + \Delta k_{i,t} + \varphi y_{i,t}], \\
b_{i,t} &= \max\{0; a_{i,t-1} - \Delta a_{i,t}; -\Delta a_{i,t}\}
\end{aligned}$$

where β is the discount factor. Importantly, we index value functions by both time and the household state vector $h_{i,t}$, as the problem's value varies across households and periods—even conditional on liquid and illiquid asset balances $(a_{i,t-1}, k_{i,t-1})$.

3.2 Valuation of Liquidity and Penalized Withdrawals

Next, we define our primary object of interest: the household's valuation of liquidity. It captures how much more the household values a liquid dollar today relative to a liquid dollar tomorrow.

Definition 1: Equilibrium Valuation of Liquidity. *The equilibrium valuation of liquidity for household i consuming c_t at time t is given by*

$$\theta_{i,t}(c_{i,t}; h_{i,t}) \equiv u'(c_{i,t}; h_{i,t}) \left(E_t \left[\frac{\partial V_{t+1}(a_{i,t}, k_{i,t}; h_{i,t+1})}{\partial a_{i,t+1}} \right] \right)^{-1} (\beta(1+r))^{-1}.$$

It is the ratio between the marginal value of a liquid dollar today (in terms of consumption) at consumption level c_t and the expected value of a marginal liquid dollar tomorrow.

To build intuition, we highlight two observations about $\theta_{i,t}(c_{i,t}; h_{i,t})$. First, under perfect credit markets, the valuation of liquidity equals 1, as the household's Euler equation remains undistorted. Second, even in perfect credit markets, consumption can fluctuate due to changing circumstances, consumption shocks, or changes in preferences. These fluctuations complicate attempts to infer marginal valuations from data on consumption alone—a challenge that our revealed-preference approach directly addresses.

The empirical analysis in this paper builds on the idea that individuals' withdrawal behavior offers a valuable revealed-preference tool for learning about our key object of interest,

$\theta_{i,t}(c_t; h_{i,t})$. This insight is formalized in the Lemma below, which we prove and generalize to other time periods in Appendix C.

Lemma 1: Household Valuation of Liquidity at Withdrawal. *Consider a household i that at time $t = t^* - 1$ makes a penalized withdrawal from the illiquid account without fully depleting it, i.e. $\Delta k_{i,t} < 0$ and $k_{i,t} > 0$.*[11] *The equilibrium valuation of liquidity for this household satisfies*

$$\theta_{i,t}(c_{i,t}; h_{i,t}) \quad = \quad \frac{1}{1-\tau} \geq \frac{1}{1-\rho'_{i,z}(b)}. \tag{1}$$

The intuition is straightforward. In period $t = t^* - 1$, the household knows it will have unrestricted access to the illiquid account in the following period. If it nevertheless chooses to pay the penalty τ to withdraw today, this reveals that a liquid dollar today is worth at least $\frac{1}{1-\tau}$ times a liquid dollar tomorrow. Furthermore, if the household chooses to withdraw from the illiquid account rather than borrowing, it indicates that the marginal cost of borrowing exceeds the effective cost of accessing illiquid funds, so that $\frac{1}{1-\tau} \geq \frac{1}{1-\rho'_{i,z}(b)}$.

This logic extends to earlier periods ($t < t^* - 1$), but with an important caveat. A penalized withdrawal always implies that $\theta_{i,t} > 1$, but it is not necessarily the case that $\theta_{i,t} = \frac{1}{1-\tau}$. This is because the observed withdrawal behavior reflects a trade-off between a liquid dollar today and an illiquid dollar tomorrow: the act of paying the penalty reveals that the household values a liquid dollar today at least $\frac{1}{1-\tau}$ times more than an illiquid dollar tomorrow. Our object of interest, $\theta_{i,t}$, instead, captures the trade-off between liquid dollars at different points in time. The two coincide only if, in the next period, the household values liquid and illiquid dollars equally.[12]

In general, however, illiquid dollars are worth less than liquid ones. This implies that for

[11]The assumption that $k_{i,t} > 0$ guarantees that the Euler equation is satisfied with equality. If the household fully depletes the retirement savings account, which we show in Appendix B to be infrequent in the data, then we would get $\theta_{i,t}(c_{i,t}; h_{i,t}) > \frac{1}{1-\tau}$. Throughout the characterization, we focus on the empirically typical case in which $k_{i,t} > 0$.

[12]This condition holds when the household does not expect to make additional penalized withdrawals before retirement, at which point all illiquid funds become accessible. Reassuringly, this pattern is supported by the data: most households make only a single penalized withdrawal (Appendix B). Consistent with this, our results are nearly identical when focusing on households aged 55-59, who are close to the statutory "retirement" age in the context of withdrawal penalties (i.e., 59.5) and thus proxy for households who make withdrawal decisions at the vicinity of period $t = t^* - 1$ (see Appendix Figures D.15, D.16, and D.17).

households making early withdrawals, the following inequality must hold:[13]

$$1 \leq \theta_{i,t}(c_{i,t}; h_{i,t}) \leq \frac{1}{1-\tau}. \tag{2}$$

This bound captures a central insight of our setting: as long as a household has retirement savings available, its valuation of liquidity is capped by the cost of accessing those funds prematurely. Penalized withdrawals thus serve as a form of self-insurance against liquidity shocks. Households that do not make such withdrawals, despite having access, must place a relatively low value on liquid dollars today.

In practice, penalized withdrawals are infrequent. At any point in time, we can compute the share of households with certain characteristics who make a penalized withdrawal. This share is also the probability that an individual household with those characteristics makes such a withdrawal—providing a natural measure of their average valuation of liquidity.

Definition 2: Probability of Making a Penalized Withdrawal. *Consider a set of N households denoted by $\mathscr{I}$. The average probability that a household in this group makes a penalized withdrawal at time t is given by:*

$$\mathcal{P}_t(\mathscr{I}) = 1 - \frac{1}{N} \sum_{i \in \mathscr{I}} F\left(\bar{\varepsilon}_{i,t}\right),$$

where $\bar{\varepsilon}_{i,t}(h_{i,t})$ is the threshold shock such that household i makes a penalized withdrawal at time t if and only if $\varepsilon_{i,t} \geq \bar{\varepsilon}_{i,t}$.

To further capture the extent of self-insurance behavior, we define an empirical measure—"penalized liquidity"—which quantifies how much liquidity households extract via penalized withdrawals.

Definition 3: Penalized Liquidity. *Consider a set of N households denoted by $\mathscr{I}$. Their average penalized liquidity from time t to t', denoted $\Lambda_{t,t'}(\mathscr{I})$, is defined as:*

$$\Lambda_{t,t'}(\mathscr{I}) \equiv \frac{1}{N} \sum_{k=t}^{t'} \sum_{i \in \mathscr{I}} \Delta k_{i,k},$$

where $\Delta k_{i,k}$ is the amount withdrawn from the illiquid account with a penalty.

These two measures guide our empirical analysis. Frequent and sizable penalized withdrawals indicate a high valuation of liquidity. Conversely, households with available balances

[13] This inequality holds as long as households retain positive retirement balances, a condition broadly satisfied in our sample.

who do not withdraw are likely not facing severe liquidity constraints. Together, $\mathcal{P}_t(\mathscr{I})$ and $\Lambda_{t,t'}(\mathscr{I})$ provide a direct way to assess how much liquidity is needed to keep households well-insured—i.e., to ensure that $\theta_{i,t} \leq \frac{1}{1-\tau}$. Moreover, comparing $\mathcal{P}_t(\mathscr{I})$ and $\Lambda_{t,t'}(\mathscr{I})$ across demographic groups (e.g., by race or income) allows us to infer disparities in access to formal credit markets or alternative insurance mechanisms.

3.3 Liquidity Demand and Supply

The valuation of liquidity is an equilibrium object that depends on both household-level demand for funds and market-level credit supply. Accordingly, as we discuss next, shocks to either side can raise households' valuation of liquidity and trigger a penalized withdrawal.

We represent demand and supply curves as inverse relationships between the quantity of funds demanded or supplied and the corresponding "price," which in our setting is the effective marginal interest rate.[14]

A household's (inverse) demand for liquidity specifies the marginal interest rate, $\mathbb{D}_i(\bar{b})$, at which the household is willing to borrow (or save) an amount $\bar{b} \equiv [\Delta k_{i,t} + \Delta a_{i,t} + \tau \Delta k_{i,t} \mathbb{I}_{(\Delta k_{i,t}<0)} \mathbb{I}_{(t<t^*)} + \rho_{i,z}(b)\mathbb{I}_{(b>0)}] - (1+r)a_{i,t-1}$ to finance consumption.[15]

Definition 4: Demand for Liquidity. *The demand for liquidity is a function $\mathbb{D}_i(\bar{b})$ that solves*

$$\theta_{i,t}(x_{i,t} + \bar{b}; h_{i,t}) \equiv \frac{\mathbb{D}_i(\bar{b})}{1+r},$$

where $x_{i,t} \equiv (1-\varphi)y_{i,t} + (1+r)a_{i,t-1} - \varepsilon_{i,t}$ is household i's cash-on-hand at time t, net of the liquidity shock.

By definition, $\bar{b}$ is the total amount of funds borrowed to satisfy consumption, since $c_{i,t} = x_{i,t} + \bar{b}$. The demand for liquidity can thus be interpreted as the level of consumption a household would choose if it could borrow at the marginal interest rate given by $\mathbb{D}_i(\bar{b})$. This function is decreasing, as the valuation of liquidity declines with higher consumption for any utility function with diminishing marginal utility.

[14] For this graphical illustration, we focus on the final period before t^*, when the marginal cost of a penalized withdrawal simplifies to $\frac{1}{1-\tau}$. We also focus on households with available illiquid funds (i.e., those for whom $k_{i,t} > 0$), which are the overwhelming majority in our data. The insights generalize straightforwardly.

[15] $\bar{b}$ captures total funds used for consumption from all sources, liquid and illiquid. We distinguish it from b, which refers only to funds borrowed from the liquid account.

Definition 5: Supply of Liquidity. *The supply of liquidity is a function $\mathbb{S}_{i,z}(\bar{b})$ defined as:*

$$\mathbb{S}_{i,z}(\bar{b}) \equiv \begin{cases} 1+r & \text{if } \bar{b} \le 0, \\ \frac{1+r}{1-\rho'_{i,z}(\bar{b})} & \text{if } \bar{b} > 0 \text{ and } \rho'_{i,z}(\bar{b}) < \tau, \\ \frac{1+r}{1-\tau} & \text{if } \Delta k_{i,t} < 0 \text{ so that } \rho'_{i,z}(\bar{b}) = \tau. \end{cases}$$

The supply of liquidity gives the marginal interest rate a household must pay to access funds of the amount $\bar{b}$, whether from liquid or illiquid sources.[16]

In equilibrium, the household's borrowing decision satisfies $\mathbb{S}_{i,z}(\bar{b}) = \mathbb{D}_i(\bar{b})$. At this point, the valuation of liquidity also satisfies $\theta_{i,t}(c_{i,t}; h_{i,t}) = \frac{\mathbb{D}_i(\bar{b})}{1+r}$, which implies that the valuation equals 1 if and only if the shadow cost of capital equals the risk-free rate.

Shocks to Demand and Supply of Liquidity and Penalized Withdrawals To fix ideas and illustrate how shocks to demand or supply of liquidity affect the equilibrium valuation of liquidity and trigger penalized withdrawals, Figure 1 presents several scenarios.

In the upper panels (a)-(c), we plot demand and supply curves and highlight the resulting equilibrium valuation of liquidity in three cases. Panel (a) depicts perfect credit markets: the supply curve is flat at the risk-free interest rate, so demand shifts are absorbed entirely through changes in borrowing. Liquidity needs are fully insured, and the Euler equation remains undistorted. Panel (b) shows imperfect credit markets for a household without access to a retirement savings account ($k_{i,t} = 0$). The supply curve is upward-sloping due to the convex borrowing cost $\rho_{i,z}(b)$. Borrowing becomes expensive, reducing credit demand and increasing the equilibrium valuation of liquidity, denoted by θ_1. Panel (c) introduces the option of a penalized withdrawal from a retirement account. As long as households hold illiquid assets, they can access those funds at a cost τ, facing a marginal cost of liquidity that equals $\frac{1+r}{1-\tau}$. This option lowers the valuation of liquidity to $\theta_2 < \theta_1$, and reduces liquid borrowing ($b_{3,a} < b_2$). The gap $b_{3,b} - b_{3,a}$ represents the "penalized liquidity" withdrawn at cost τ.

The lower panels (d)-(f) illustrate how shifts in demand or supply can trigger penalized withdrawals. Panel (d) shows our starting point: a household who is borrowing only from liquid funds. In panel (e), a credit crunch—modeled as an upward shift in the borrowing cost function $\rho_{i,z}(b)$—raises the marginal cost of credit.[17] Although credit market conditions tighten, access to illiquid savings remains unaffected. The household thus reduces liquid borrowing (from b_1 to $b_{2,a}$) and partially offsets the decline with a penalized withdrawal

[16]The definition of $\mathbb{S}_{i,z}(\bar{b})$ spans the full support of $\bar{b}$ under the assumption that $k_{i,t} > 0$.

[17]For simplicity, we assume the household has sufficient illiquid savings ($k_{i,t}$) to cover the entire support of b shown in the figure.

$(b_{2,b} - b_{2,a})$. Liquidity valuation rises, but the availability of illiquid funds dampens this increase via self-insurance. In panel (f), we consider a positive demand shock, such as a drop in income $y_{i,t}$. Total borrowing rises (from b_1 to $b_{3,b}$), again triggering a penalized withdrawal $(b_{3,b} - b_{3,a})$. While in this scenario borrowing from liquid savings rises, in both scenarios of panels (e) and (f) the ability to make a penalized withdrawal reduces the sensitivity of the equilibrium valuation of liquidity to shocks.

These scenarios highlight how, in response to either type of shock, the amount withdrawn—i.e., the "penalized liquidity"—serves as a useful proxy for the shock's magnitude and the household's residual need to self-insure. Larger shocks correspond to larger withdrawals.

From Model to Data. We conclude this section by summarizing how the model informs our empirical analysis. First, it shows that penalized withdrawals reveal a high valuation of liquidity. Second, it demonstrates that such high valuations can be driven by shocks to either household-level demand or market-level credit supply. Accordingly, the first two parts of our empirical analysis exploit variation in *shocks* on both sides of the market. Section 4 studies household-level events. Focusing on transitory variation within households over time, we identify life-cycle events that affect liquidity demand. This maps to $\mathbb{D}_i$ at time t, primarily driven by fluctuations in $y_{i,t}$. Section 5 examines shocks to local credit markets, leveraging spatial variation in the impact of the Great Recession. This maps to $\mathbb{S}_{i,z}$ and reflects credit conditions summarized by $\rho_{i,z}(b)$. Finally, Section 6 provides a comprehensive analysis of the *permanent* components influencing the valuation of liquidity. We use a standard movers design to unpack these components into household-specific and location-specific factors.

4 Valuation of Liquidity after Household-Level Events

In this section, we study how adverse household-level economic events affect valuation of liquidity. This analysis provides empirical support for our model—confirming that shocks trigger penalized withdrawals—and quantifies how strongly these shocks impact the valuation of liquidity. It also establishes household-level benchmarks for comparison with the market-level shocks that we study in the next section.

Estimating Equation. The event study estimating equation takes the form:

$$y_{i,t} = \alpha_i + x_{i,t}\lambda + \sum_{r \neq -2, r=-5}^{r=10} \beta_r \times I_r + \varepsilon_{i,t}, \tag{3}$$

where $y_{i,t}$ is withdrawal behavior of household i at time t, $x_{i,t}$ is a full set of age fixed effects for the primary-filer and (cyclical) calendar year fixed effects, and α_i are household fixed effects.[18] We let $r(i,t)$ denote time relative to the year of the event, so that $I_r = \mathbb{I}_{r(i,t)=r}$ represent a set of relative time indicators. We take the baseline year to be -2 to capture changes in trends that could happen toward the realization of the event,[19] and we run the analysis over the 16-year horizon from year -5 to year +10. We plot β_r around different events to trace the evolution of households' withdrawal behavior.

4.1 Unemployment Event

We define an unemployment event as the first period we observe at least one of the household members receiving unemployment benefits. In Figure 2, we plot the event study coefficients β_r, estimated when the outcome is either an indicator for making a penalized withdrawal (in panel (a)) or the amount withdrawn (in panel (b)). As the event approaches, we see an increase in penalized withdrawals that is then followed by a large spike at the year of the event. Through the lens of our model, these patterns imply an increase in the valuation of liquidity, which in turn maps to under-insurance of unemployment shocks.[20]

In terms of magnitudes, panel (a) shows that the share of households with sufficiently high valuation of liquidity to trigger a penalized withdrawal doubles at the onset of the event: at baseline (in $t = -2$) households make penalized withdrawals at a rate of 9.9 percentage points (pp), which increases at the unemployment event (in $t = 0$) by 10.4 pp. Panel (b) shows that households make additional penalized withdrawals of an average of approximately \$1,600 in that same year. This amount exactly maps to the concept of "penalized liquidity" that we defined in the theoretical framework; that is, it is the average amount of liquidity

[18]In the samples on which we run these regressions, we include all households to help with identification of non-event coefficients and we accordingly add to $x_{i,t}$ a dummy for households who do not experience an event. In addition, in Appendix Figure D.5 we repeat the analysis among households who stay in the same commuting zone around the events that we study with almost identical results.

[19]We note that the year -1 coefficient can incorporate anticipation but also potential effects of the onset of an event. This is due to the annual frequency of the data at the end of a calendar year and the defined timing of the event. For example, households who experience an event of a large decline in income (which we take to be at least 30 percent) between the end of period -1 and period 0 would be assigned a "large income decline" event at time 0, but the process of a decline in income could have already (and likely) began throughout year -1.

[20]The findings go hand-in-hand with the important literature on the effects of unemployment on earnings and consumption, which has shown large declines in consumption in the short run with lingering effects on earnings in the long run (See, e.g., Sullivan and Von Wachter 2009; Kolsrud et al. 2018; Schmieder et al. 2018; Ganong and Noel 2019; Gerard and Naritomi 2021.). In comparison to these assessments of income or consumption, our investigation of the valuation of liquidity is robust to the possibility that preferences are themselves affected by employment status, e.g., from complementarities between consumption and leisure. This could be the case, for example, if employment leads to different consumption needs, such as the classic substitution to cooking meals at home while unemployed and the corresponding reduction in time and monetary costs involved in commuting. Indeed, a key advantage of our framework is that it freely allows for state dependence in preferences for any shock we would consider.

injection needed to keep the marginal valuation of liquidity at or below the withdrawal penalty.

Two observations are useful for interpreting the magnitude of the penalized liquidity. Comparing the withdrawn amounts to the decline in household income around the event, we find that penalized withdrawals compensate on average for less than 8 percent of the average income decline (which is approximately $20,900 at the onset of the event, see panel (a) of Appendix Figure D.4). This suggests that, on average when including zeros of non-withdrawing households, the households in our sample are relatively well-insured. However, the relatively small average masks substantial heterogeneity when we consider comparing households that are induced to withdraw to those who do not. Scaling the effect on amounts by the effect on take-up, we find that the withdrawn amount averages to about $19,000 per household who makes a penalized withdrawal at the year of the event.

Heterogeneous Effects. We next examine how reliance on retirement withdrawals for self-insurance against unemployment varies across household types. This sheds light on the determinants of liquidity valuation and helps validate our approach by linking it to observable characteristics such as race, wealth, and location. While our heterogeneity analysis is descriptive, it is informative for identifying households with higher liquidity needs. We focus on withdrawal frequencies, which have a clearer interpretation than amounts, as the latter may reflect differences in income.

We begin by examining heterogeneity by race. Panel (c) of Figure 2 shows event study estimates separately for households with Black and White primary earners. Despite experiencing smaller income declines (Appendix Figure D.4, panel (b)), Black households are significantly more likely to withdraw upon unemployment. This suggests more limited access to alternative, lower-cost liquidity sources.

Panel (d) of Figure 2 turns to differences by capital income, a proxy for non-housing wealth. We group households into six bins: separating households with negative wealth, zero wealth, and then creating four equally-sized bins among those with positive capital income. Consistent with our revealed-preference interpretation, households with greater financial means show smaller withdrawal responses, reflecting lower marginal valuation of liquidity. Yet even households in the top capital income quartile—with an average of nearly $40,000—display sizable increases in withdrawals, underscoring that even wealthy households may be liquidity constrained (e.g., Kaplan et al. 2014).

4.2 Income Changes

We next study how changes in household income affect the valuation of liquidity. We define a large income loss event as the first year we observe a household's overall income falling by more than 30 percent relative to the previous year. Panels (a) and (b) of Figure 3 display the event study coefficients for withdrawal frequency and amounts, respectively. We observe significant spikes at the event year, indicating that households, on average, require approximately $2,000 in penalized liquidity to keep their marginal valuation of liquidity within bounds. Similar to unemployment events, Black households exhibit a notably larger increase in withdrawals compared to White households upon experiencing large income losses (panel (e), Figure 3).

We further refine the analysis by studying households' withdrawals as a function of the deviation of their income flow from their average income across our data period. We split households by whether a member of the household switched jobs in that year. We do this because job changes themselves, as displayed in Appendix Figure D.8, lead to increased take-up.[21]

Panels (c) and (d) of Figure 3 reveal several clear patterns. First, we observe a strong gradient in withdrawal frequency as income losses become larger, aligning with our model's prediction that penalized withdrawals serve as a short-term self-insurance mechanism. Second, we find stark asymmetry around zero: withdrawal behaviors flatten completely when incomes increase. This asymmetry supports our self-insurance interpretation and rules out strategic tax-driven withdrawals, which would imply withdrawal sensitivity across both income losses and gains. Third, even households experiencing income gains make non-negligible penalized withdrawals, averaging around 1 percent of annual income. This suggests that equilibrium liquidity valuation reflects not only income shocks but also expenditure-driven consumption needs (e.g., health or child-related expenses). This result shows that even full insurance against negative income shocks could not be sufficient to fully smooth marginal utility over time.

[21] Job changes could increase withdrawals due to higher liquidity valuation during transition periods or other factors such as salience or cashing out small balances. We exclude account rollovers from employer-sponsored accounts, which are just mechanical transfers of funds and could be common upon job separation. That said, upon job separation, low balances below a certain threshold can be automatically paid out in cash to the departing employee, with thresholds of $5,000 prior to 2005 and $1,000 thereafter. To account for negligible balances and these automatic passive penalized distributions, Appendix Figure D.7 replicates our event study analyses but where the outcome variables are indicators for taking penalized withdrawals that are higher than given thresholds.

5 Valuation of Liquidity during the Great Recession

We now examine how broad economic shocks influence households' valuation of liquidity, focusing on the case of the Great Recession.

Estimating Equation. We estimate specifications of the following event study type:

$$(4)\qquad y_{i,z,t} = \mathit{\Gamma}_z + \alpha_i + x_{i,t}\lambda + \sum_{r\neq 2006, r=2000}^{r=2017} \beta_r \times I_r + \sum_{r\neq 2006, r=2000}^{r=2017} \theta_r \times I_r \times Treat_z + \varepsilon_{i,t}.$$

In this specification, $y_{i,z,t}$ is the outcome for household i in Commuting Zone (CZ) z in year t; I_r are calendar year indicators running from years 2000 to 2017, where year 2006 is taken to be the baseline year; $x_{i,t}$ is a full set of primary-filer age fixed effects, (cyclical) calendar year fixed effects, and potential household-specific economic controls; $\mathit{\Gamma}_z$ are CZ fixed effects; and α_i are household fixed effects. $Treat_z$ is the treatment intensity of location z in terms of unemployment shock. Following Yagan (2019), we measure treatment intensity as the change in a CZ's unemployment rate from 2007 to 2009. Our parameters of interest are θ_r, which capture the relative change in behavior in a locality that is exposed to a 1 pp larger local unemployment shock.

Results. The solid lines in panels (a) and (b) of Figure 4 plot the θ_r estimates from equation (4) for either the frequency or the amounts of withdrawals. We first observe that locations, which were about to be hit differentially by the Great Recession, were on similar trajectories prior to the event (in support of the research design). Then, we find that Commuting Zones that were more severely affected by the Great Recession, as measured by unemployment rates, have seen a larger increase in penalized withdrawals. The response peaks around the height of the Great Recession with an increase in frequency of local penalized withdrawals of 0.403 pp per 1 pp in local unemployment shock in year 2009, which results in a total increase of \$63.4 in penalized liquidity per household (including zeros). Calculating the cumulative effects from 2007-2012, we find an overall increase of 1.47 pp in the probability of taking a penalized withdrawal and a corresponding overall increase of \$251 in penalized liquidity.

To make the magnitudes comparable with the household-level unemployment event, we need to normalize the coefficients for the market-level event so that they reflect the incidence of unemployment. This implies looking at a counterfactual increase of 100 pp in unemployment to be comparable to the event of an individual becoming unemployed (which is equivalent to a transition from 0 to 100 pp). We find that the flow effect on making a penalized withdrawal of a locality-level unemployment shock is about 4 times as large as the direct effect of a household-level unemployment shock that we have estimated (40.3 pp vs. 10.4 pp). This suggests that the effect of a 1 pp of local unemployment on the valuation of liquidity is about one-quarter due to increases in household demand (0.258=10.4/40.3) and about three-quarters due to economic conditions that led to a decrease in local supply. We reach a very similar conclusion using the results for the amount withdrawn: multiplying the point estimate by 100, we get $6,340 (from the Great Recession results in Figure 4) to be compared with $1,590 (from the unemployment event study results in Figure 2).

In light of these patterns, we break down the cumulative impact of the Great Recession into: (1) a direct effect arising from individual households' economic circumstances, such as changes in employment status and income; and (2) an indirect effect stemming from broader market-level spillovers. To achieve this, we augment the estimating equation (4) by flexibly controlling for household-level economic indicators—including unemployment, wage earnings, and gross income—using their lagged, current, and lead values, along with interactions with calendar year dummies. The estimates are reported in the dashed lines in panels (a) and (b) of Figure 4. Consistent with the previous argument, we find that the indirect impact accounts for approximately three-quarters of the overall effect of the Great Recession. Our findings are therefore closely consistent with a tightening of local credit conditions for all workers in distressed locations. Through the lens of our model, the results imply that a contraction in the local credit supply primarily drives the observed increase in the equilibrium valuation of liquidity.

In sum, the Great Recession highlights the significant role of market-wide shocks in shaping households' valuation of liquidity—both through direct income disruptions and indirect credit market spillovers.

6 Unpacking the Drivers of Valuation of Liquidity

We now turn to the permanent drivers of liquidity valuation. Are differences in penalized withdrawals across markets primarily driven by household characteristics, or does the location itself matter? To answer this, we first leverage a movers design to quantify the share of the large spatial differences in penalized withdrawals across local markets attributable to locations versus households. Then, we explore how the estimated location effects and average household effects within a location correlate with a set of locality-level observables. This analysis sheds light on the local factors that shape households' liquidity valuation.

As a preliminary step, we document that there are large differences in penalized withdrawals across locations. Panel (a) of Figure 5 plots a map of the average annual share of households that make a penalized withdrawal by commuting zones, aggregated over the full sample period (1999-2018). We find large differences across regions, with a mean of 9.8 pp and a standard deviation of 1.7 pp.[22]

6.1 Movers' Design

We use the following two-way fixed effects specification for household outcomes (Abowd et al. 1999):

$$y_{i,t} = \alpha_i + \mathit{\Gamma}_{z(i,t)} + x_{i,t}\lambda + \varepsilon_{i,t}, \tag{5}$$

where $y_{i,t}$ is an indicator for whether household i makes a penalized withdrawal at time t; α_i is a household fixed effect; $\mathit{\Gamma}_{z(i,t)}$ are location fixed effects in determining the household's outcome, where $z(i,t)$ indexes the location of household i in year t; $x_{i,t}$ is a vector of potential time-varying controls, including indicators for age of primary filer, (cyclical) calendar year fixed effects, and household-level economic conditions. The household fixed effects capture non-transitory household characteristics that affect a household's valuation of liquidity, e.g., limited credit-product choice set due to systematically low average credit scores. The location fixed effects can reflect access to traditional financial institutions, such as banks, as well as local social networks and informal support, such as religious organizations.

To identify the causal contribution of location to differences in the valuation of liquidity, we exploit variation among households who move across commuting zones. Specifically, we implement a movers design following Finkelstein et al. (2016), which allows us to decompose spatial differences in penalized withdrawals into components attributable to households versus their local environment.

The core idea is simple: if a household moves from a low- to high-withdrawal area and subsequently increases its use of penalized withdrawals, it suggests that location plays a causal role (under the identifying assumption we discuss below). To formalize this, we define Δ_i as the difference in average withdrawal rates between the destination CZ (z_D) and origin CZ (z_O):

$$\Delta_i = E[y_{i,t} \mid z(i,t) = z_D] - E[y_{i,t} \mid z(i,t) = z_O].$$

These expectations are computed using data on *non-movers* to avoid mechanical correlation (what is known as leave-out means).

We then define $r(i,t)$ as the time relative to the move (e.g., $r = 0$ is the year of the

[22] To address known biases in plug-in estimates of second moments due to sampling errors (see, e.g., Andrews et al. 2008), we estimate the standard deviations of location-level statistics based on a split-sample approach (as in, e.g., Finkelstein et al. 2021 and Card et al. 2023). Specifically, with a random sample split, we conduct the estimation on each subsample separately and assess the variation based on the covariance of estimates across subsamples.

move), and let $I_{r(i,t)>0}$ indicate post-move periods. To quantify the share of the difference in withdrawal rates that is due to location, as opposed to household traits, we define:

$$\theta = \frac{\Gamma_{z_D} - \Gamma_{z_O}}{y_{z_D} - y_{z_O}},$$

where $y_{z_j} \equiv E[y_{i,t} \mid z(i,t) = z_j]$. This θ parameter captures the average *passthrough* of location-level differences into household behavior after the move. That is, how much of the CZ-level difference in penalized withdrawals is "transmitted" to movers.

This leads us to the following empirical specification for movers:

$$y_{i,t} = \mu_i + \theta \times I_{r(i,t)>0} \times \Delta_i + x_{i,t}\lambda + \varepsilon_{i,t} \tag{6}$$

where $\mu_i = \alpha_i + \Gamma_{z_O}$, and $x_{i,t}$ includes a set of time-varying controls.

Dynamic Specification. As a baseline specification we take an extended version of equation (6) to the data to allow for flexible dynamics by estimating the following event-study specification:

$$y_{i,t} = \mu_i + \sum_{r \neq -2} \beta_r \times I_r + \sum_{r \neq -2} \theta_r \times I_r \times \Delta_i + x_{i,t}\lambda + \varepsilon_{i,t}, \tag{7}$$

where $I_r = \mathbb{I}_{r(i,t)=r}$ are indicators for time relative to the move. To be consistent with the event-study analyses from Section 4, our baseline period is taken to be two years prior to the move ($r = -2$). The event-study specification in equation (7) allows us to test for parallel trends in the pre-move period (based on θ_r for $r < -1$) and to investigate dynamics in location effects in the post-move period (based on θ_r for $r > 0$). In these estimations robust standard errors are clustered at the origin CZ level.

Results. We estimate equation (7) on a balanced sample of households observed for at least 9 periods: from -3 to $+5$. We start with a specification in which the vector $x_{i,t}$ includes primary-filer age fixed effects and (cyclical) calendar year fixed effects. Panel (a) of Figure 7 plots the θ_r coefficients when withdrawal take-up is the outcome variable. First, the figure shows that there are no differential pre-trends across households who move to differential intensity locations, in support of the design as we discuss below. Second, it shows a sharp increase at the time of the move that persists over time, indicating a lasting effect of location: permanent location characteristics pass through to household withdrawals with an average rate of 0.34 in the post-move years (periods 1 to 5). This result implies that place effects account for a third of the overall spatial differentials that we have found in penalized withdrawals. This stands as one of our main findings and highlights that the local environment is a crucial determinant of the valuation of liquidity.

Panel (b) of Figure 7 provides estimates when we run the same specification but with amounts withdrawn as the outcome, which allows us to measure the additional penalized liquidity needed by households when moving to locations with more frequent penalized withdrawals. The estimates average in the post-move years to \$5,750, implying that if a household moves across CZs from the 5th percentile (6.8 pp) to the 95th percentile (12.7 pp) of withdrawal frequency, they would need approximately \$340 ($= \$5,750 \times (0.127 - 0.068)$) a year in additional penalized liquidity to keep their valuation of liquidity within bounds. In turn, this implies that "compliers" who increase their take-up upon such a move, make additional withdrawals that amount to an average of \$16,900 (=\$5,750/0.34).

Interpretation and Robustness. Our interpretation of these results—guided by the conceptual framework from Section 3—is that households who move to areas with weaker credit supply rely more on penalized withdrawals to meet their liquidity needs.

In what follows, we conduct a series of empirical checks to support this interpretation. Specifically, we test the validity of the movers design, examine key threats to identification, and assess alternative explanations for the observed patterns.

First, we test our identifying assumption of parallel trends: our design requires that households' underlying trends in withdrawals do not systematically differ by Δ_i. The standard testable implication of this assumption is whether there are differential trends in the pre-move period across households with differential Δ_i. Reassuringly, even when we estimate equation (7) for an extended window that runs from year -5 to $+10$ (on an unbalanced sample of households) we find that there are virtually no pre-trends—see panel (c) of Figure 7 for take-up and Appendix Figure D.11 for amounts.

Another aspect to consider is that mover designs cannot account for shocks that simultaneously differ across households with varying treatment intensity Δ_i *and* align exactly with the timing of moves. We therefore ask: can the results be explained by differential changes to household economic conditions interacted with timing of move to differentially intense locations? Two pieces of evidence are inconsistent with this idea. In panel (c) of Figure 7 we see a high degree of persistence for up to 10 years in the estimates for the role of location, θ_r. This is in contrast to the effects of household-level shocks, which have been shown above (in Section 4) to be transitory with clear dissipating dynamics. Similarly, Appendix Figure D.9 shows a comparable pattern of transitory dynamics for the move event itself, as captured by the "event study" coefficients β_r from the movers equation (7). These combined findings are hard to reconcile with patterns in passthroughs being driven by shocks aligned with the time of move.[23] We then directly account for household-specific economic conditions that may change around the time of the move by including a flexible set of (endogenous) controls: unemployment, wage earnings, and gross income. For each, we include lagged, current, and lead values, as well as interactions with time relative to the move. The results show that the estimates hardly change in terms of either dynamics or magnitudes—see panel (d) of Figure 7.

Second, moving to alternative explanations, households might learn about withdrawals from peers upon moving to higher-intensity locations. However, this explanation is inconsistent with the immediate jump in withdrawals at the time of the move and limited dynamics afterward (similar to the logic in Finkelstein et al. 2022). To directly test the learning hypothesis, we analyze households who had previously made penalized withdrawals before moving. Panel (e) of Figure 7 shows very similar results, though noisier due to fewer households in the subsample, again suggesting learning does not drive our findings.

Third, another explanation could be tax optimization.[24] To test this in the context of moves, we control flexibly for location-specific top marginal tax rates interacted with time relative to the move. The slight attenuation in estimates in panel (f) of Figure 7 indicates that tax motives likely play only a modest role.

6.2 Drivers of Location and Household Effects

Finally, we turn to investigate the potential drivers underlying the permanent place effects (Γ_z) and the permanent household effects (α_i) in determining the valuation of liquidity. We now directly estimate equation (5) to provide us with these estimates. In these estimations we include as household-level economic controls in $x_{i,t}$ the following variables: unemployment, wage earnings, and gross income, with lagged, current, and lead values.

Panels (b) and (c) of Figure 5 first display maps of the estimated location effects (Γ_z) and the estimated household fixed effects (α_i) averaged within a CZ. There is meaningful geographic variation in both elements, with a 1 pp standard deviation for location effects and a 1.8 pp standard deviation for CZ-level average household fixed effects. We then investigate their correlations with a range of CZ-level characteristics (taken from Chetty et al. 2016 unless noted otherwise). Figure 6 reports all the normalized regression coefficients from a

[23] We note that the moderate decline in the estimates in the extended window of post-move years is attributable to attrition and return moves (see panel (a) of Appendix Figure D.10). They attenuate the persistence in the effects since we assign a household the same destination location for the entire post-move period, whether they subsequently moved or not, because these behaviors could be endogenous to the initial move. Panel (b) of Appendix Figure D.10 illustrates this point: when we scale the estimates by the share of movers that are still at the assigned destination, the dynamics flattens out.

[24] We have already seen evidence inconsistent with this conjecture in the analysis of household events, where we find no gradient in the region of positive income changes.

series of univariate OLS regressions, while Appendix Figures D.12 and D.13 include all the corresponding scatter plots. We focus here on discussing the characteristics which display particularly strong correlations with the estimates from the statistical model.

Location Effects. We consider the Credit Insecurity Index, which is a measure developed by the Federal Reserve Bank of New York for assessing American communities' credit health and well-being (Hamdani et al. 2019). We find that areas with a higher Credit Insecurity Index exhibit higher propensities of using penalized withdrawals, consistent with lower availability of alternative sources of credit.

We also consider the correlation with a location's median house value. We find that locations with higher home values display less reliance on withdrawals, consistent with the notion that high home values can provide collateral that reduces risk in the credit market. Both of these findings provide consistent support for our theory of how withdrawals reveal the valuation of liquidity.

Household Effects. Household fixed effects most notably correlate with measures of racial composition: households who live in communities with a high share of Black residents are significantly more likely to make penalized withdrawals irrespective of their current location. How should we interpret this result? Withdrawals are a financial instrument that, conditional on having a retirement account, does not discern (or discriminate) across households of different social groups. Thus, differences in withdrawal behavior among households who hold retirement accounts can reveal differential access to alternative credit. Account-holding households in predominantly Black communities reveal a higher valuation of liquidity, suggesting they have more limited access to alternative channels of credit.[25]

The observed disparities could reflect limited credit access for Black households or for non-Black households in predominantly Black communities. Three pieces of evidence support the notion that the lack of access to credit among Black families is the more likely explanation.

First, location effects themselves are uncorrelated with the share of Black families. A randomly selected household moving into an area with many Black families would not experience increased withdrawals, indicating that limited credit access follows Black families wherever they move.

Second, repeating the analysis at the (finer) ZIP Code level and correlating household and ZIP Code fixed effects with the share of Black households (controlling for CZ fixed

[25] Similar patterns arise when we explore a location's percent of children living with single mothers. Single mothers typically face economic disadvantages, and our results suggest they too may have limited credit access.

effects), we find nearly identical correlations (Figure 6).[26]

Third, Table 1 directly examines correlations between household fixed effects and primary earners' imputed race, with and without economic controls. Black households systematically rely more on penalized withdrawals than White households, with mean fixed effects over 30% higher (12.44 pp vs. 9.5 pp, see column 2). The table further shows this gap persists after controlling for economic circumstances (column 1 vs. column 2).

Overall, the findings indicate that Black households face significant barriers to accessing credit, which go above and beyond their financial circumstances or where they live. This result provides important information for policies that aim to increase economic equity as it underscores meaningful disparities in access to credit

7 Discussion

Lastly, this section reflects on the soundness of our empirical approach and the broader implications of our findings. We discuss the conceptual foundations of our approach, specifically the use of penalized withdrawals as a revealed-preference tool, we then consider the scope and relevance of this tool going forward, and we conclude with key policy implications.

Using Withdrawals as a Revealed-Preferences Tool. Penalized withdrawals from retirement accounts offer a tractable and robust revealed-preference approach to studying households' valuation of liquidity. When a household chooses to pay a 10 percent penalty to access funds early, it reveals that the marginal value of liquidity exceeds the marginal cost imposed by the penalty. As such, this behavior provides a lower bound on the household's willingness to pay for liquidity.

This approach is grounded in a long tradition of using behavioral responses to price changes to identify economic preferences, particularly in public finance and welfare analysis (e.g., Chetty 2008; Chetty and Finkelstein 2013; Landais and Spinnewijn 2021; Kolsrud et al. 2018; Fadlon and Nielsen 2019; Finkelstein et al. 2019; Finkelstein and Hendren 2020; Hendren and Sprung-Keyser 2020; Coyne et al. 2024; Kolsrud et al. 2024). As with any revealed-preference strategy, the key identifying assumption is that households are making active, optimizing choices on the relevant margin.

While this assumption has been questioned in the context of retirement savings—that is, during the accumulation phase, where behavior is often shaped by inertia and defaults (e.g., Skinner 2007; Poterba 2014; Beshears et al. 2018; Fadlon and Laibson 2022)—our focus

[26]Spatial segregation creates substantial variation in the share of Black households both across and within CZs (Appendix Figure D.14).

is on the decumulation phase, where the evidence for optimization is stronger. We showed in Section 2.3 that penalized withdrawals are relatively infrequent, rarely involve depleting accounts, and are strongly associated with adverse income shocks. These patterns are consistent with forward-looking behavior in response to liquidity needs, which we discuss in further detail in Appendix B. Appendix B also discusses possible behavioral interpretations, describing how the evidence is inconsistent with some leading behavioral biases studied in the context of retirement savings.

Beyond the patterns that we show, the structure of the penalty itself—a cost that is incurred contemporaneously with the benefits from more liquidity—makes the marginal tradeoff salient and may limit the influence of present-biased preferences. Present bias becomes problematic in the accumulation phase with the timing decoupling of the liquidity benefits enjoyed today and the cost of reduced consumption incurred only later in the future after retirement (Laibson 1997). Theoretical work on optimal illiquidity and flexible commitment devices provides further support for this view (Amador et al. 2006; Beshears et al. 2020a).

Finally, recent empirical work also supports the interpretation that households are optimizing on the decumulation margin. While Chetty et al. (2014) find limited responsiveness to tax incentives during accumulation, studies of behavior around policy thresholds during decumulation (e.g., age 59.5 penalty removal, required minimum distributions) show large responses to price changes (Goda et al. 2018; Rong 2023; Stuart and Bryant 2024; Leganza 2024).

Overall, these patterns reinforce our view that penalized withdrawals represent a meaningful behavioral margin for revealing households' valuation of liquidity. Even with the possibility that behavioral biases may influence decision-making, their presence would primarily introduce *additional* terms into a full welfare analysis (such as "internalities" in Mullainathan et al. 2012 and Spinnewijn 2015). This crucially implies that the marginal valuations that we estimate continue to capture the household's key tradeoffs and remain economically informative. A full incorporation of behavioral factors is beyond the scope of this paper, but we note that even in models with present bias, such as Maxted (2020), withdrawal behavior still reflects the relative value of liquidity. Our framework also accommodates limited awareness or imperfect knowledge of the household's alternative credit options, so long as the household actively chooses to pay the penalty—revealing that the household's valuation of the product exceeds its "price."

Relevance and Reach of the Tool. Defined-contribution retirement accounts have become the dominant retirement savings vehicle in the United States. Access to these accounts is widespread and growing, due to automatic enrollment policies, the availability of IRAs, and new state- and federal-level initiatives (e.g., auto-IRA programs in California, Oregon, and Illinois; SECURE Act 2.0; see, e.g., Bloomfield et al. 2023). Households across the entire income distribution maintain balances in these accounts, and participation is increasingly the norm rather than the exception (Siliciano and Wettstein 2021).

This means that penalized withdrawals are not only conceptually useful, but also empirically scalable. The margin we study applies to a large share of U.S. households, and the administrative data allow us to observe behavior precisely and comprehensively. Moreover, as retirement balances continue to grow and as the coverage of defined-contribution plans expands, this margin will likely become even more relevant for understanding liquidity behavior.

Policy Implications. Our findings offer several lessons for policy design, especially for efforts to enhance liquidity access and improve social insurance targeting.

First, our results suggest scope for refining the tax treatment of early withdrawals. The current system already waives penalties for certain household-level events, and Congress has occasionally offered broader relief during systemic crises (e.g., the COVID-19 pandemic). A more structured approach could adjust the penalty based on observable household characteristics, local credit conditions, or macroeconomic indicators. Doing so could improve insurance against shocks while preserving the commitment value of retirement accounts.

Second, place-based policies could address the geographic disparities in liquidity access that we document. Just as Empowerment Zones target business activity in distressed areas (Gaubert et al. 2021), similar place-based incentives could be directed toward financial services provision in underserved communities.

Third, the substantial spatial variation in valuation of liquidity suggests that emerging financial technologies could play a role in bridging credit access gaps. With appropriate regulation, digital lending tools and FinTech platforms could offer households in "credit deserts" alternatives to high-penalty borrowing from retirement accounts.

Overall, penalized withdrawals provide a unique lens on liquidity constraints in the U.S. household sector. By quantifying how the value of liquidity varies across people and places, this tool can help guide more targeted and effective policy design.

8 Conclusion

This paper introduces conceptually and validates empirically penalized withdrawals from retirement savings accounts as a robust tool that carries information on households' valuation of liquidity. Using population tax records, this tool allows us to characterize the anatomy of the equilibrium valuation of liquidity among American families, providing several sets of findings. First, we find that location effects can explain over 30 percent of the nationwide differences in the valuation of liquidity across labor markets. Second, analyzing the Great Recession, we find that market-level shocks lead to large increases the valuation of liquidity, where spillovers in local credit tightening accounted for three-thirds of the overall effect. Third, while we show that the use of penalized withdrawals for liquidity needs in pervasive, we find that Black households rely on self-insurance from penalized withdrawals to a larger extent than White households with similar economic conditions and regardless of where they live. These results provide novel evidence suggesting that Black American families are systematically underserved by formal credit markets and have limited access to cheaper means of securing liquidity throughout the country.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High Wage Workers and High Wage Firms. *Econometrica 67*(2), 251–333.

Aguiar, M. A., M. Bils, and C. Boar (2020). Who Are the Hand-to-Mouth? Technical report, National Bureau of Economic Research.

Amador, M., I. Werning, and G.-M. Angeletos (2006). Commitment vs. flexibility. *Econometrica 74*(2), 365–396.

Andrews, M. J., L. Gill, T. Schank, and R. Upward (2008). High wage workers and low wage firms: negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society: Series A (Statistics in Society) 171*(3), 673–697.

Argento, R., V. L. Bryant, and J. Sabelhaus (2015). Early Withdrawals from Retirement Accounts During the Great Recession. *Contemporary Economic Policy 33*(1), 1–16.

Auclert, A. (2019). Monetary Policy and the Redistribution Channel. *American Economic Review 109*(6), 2333–67.

Auclert, A., M. Rognlie, and L. Straub (2020). Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model. Technical report, National Bureau of Economic Research.

Bartscher, A. K., M. Kuhn, M. Schularick, and P. Wachtel (2021). Monetary Policy and Racial Inequality.

Bayer, P. and K. K. Charles (2018). Divergent paths: A new perspective on earnings differences between black and white men since 1940. *The Quarterly Journal of Economics 133*(3), 1459–1501.

Beshears, J., J. J. Choi, C. Clayton, C. Harris, D. Laibson, and B. C. Madrian (2020a, July). Optimal illiquidity. Working Paper 27459, National Bureau of Economic Research.

Beshears, J., J. J. Choi, C. Clayton, C. Harris, D. Laibson, and B. C. Madrian (2020b). Optimal Illiquidity. Technical report, National Bureau of Economic Research.

Beshears, J., J. J. Choi, D. Laibson, and B. C. Madrian (2018). Behavioral household finance. In *Handbook of Behavioral Economics: Applications and Foundations 1*, Volume 1, pp. 177–276. Elsevier.

Bloomfield, A., K. M. Lee, J. Philbrick, and S. Slavov (2023). How do firms respond to state retirement plan mandates? Technical report, National Bureau of Economic Research.

Card, D., R. Chetty, and A. Weber (2007). Cash-on-Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market. *Quarterly Journal of Economics 122*(4), 1511–1560.

Card, D., J. Rothstein, and M. Yi (2023, August). Location, location, location. Working Paper

31587, National Bureau of Economic Research.

Chetty, R. (2008). Moral Hazard vs. Liquidity and Optimal Unemployment Insurance. *Journal of Political Economy 116*(2), 173–234.

Chetty, R. and A. Finkelstein (2013). Social insurance: Connecting theory to data. In *Handbook of public economics*, Volume 5, pp. 111–193. Elsevier.

Chetty, R., J. N. Friedman, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from denmark. *The Quarterly Journal of Economics 129*(3), 1141–1219.

Chetty, R. and N. Hendren (2018a). The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects. *Quarterly Journal of Economics 133*(3), 1107–1162.

Chetty, R. and N. Hendren (2018b). The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates. *Quarterly Journal of Economics 133*(3), 1163–1228.

Chetty, R., N. Hendren, M. R. Jones, and S. R. Porter (2020). Race and economic opportunity in the united states: An intergenerational perspective. *The Quarterly Journal of Economics 135*(2), 711–783.

Chetty, R., M. Stepner, S. Abraham, S. Lin, B. Scuderi, N. Turner, A. Bergeron, and D. Cutler (2016). The Association Between Income and Life Expectancy in the United States, 2001-2014. *Journal of the American Medical Association 315*(16), 1750–1766.

Costello, R., P. DeFilippes, R. Fisher, B. Klemens, and E. Y. Lin (2024). *Marriage Penalties and Bonuses by Race and Ethnicity: An Application of Race and Ethnicity Imputation.* US Department of the Treasury, Office of Tax Analysis.

Coyne, D., I. Fadlon, S. P. Ramnath, and P. K. Tong (2024, May). Household labor supply and the value of social security survivors benefits. *American Economic Review 114*(5), 1248â80.

Cronin, J.-A., P. DeFilippes, and R. Fisher (2023). Tax expenditures by race and hispanic ethnicity: An application of the u.s. treasury department's race and hispanic ethnicity imputation. Technical report, Office of Tax Analysis Working Paper 122.

Cui, W. and V. Sterk (2021). Quantitative easing with heterogeneous agents. *Journal of Monetary Economics 123*, 68–90.

De Giorgi, G., A. Frederiksen, and L. Pistaferri (2019, 05). Consumption network effects. *The Review of Economic Studies 87*(1), 130–163.

Derenoncourt, E., C. H. Kim, M. Kuhn, and M. Schularick (2021). The Racial Wealth Gap, 1860-2020. *Manuscript, Princeton University and University of Bonn.*

Derenoncourt, E. and C. Montialoux (2021). Minimum wages and racial inequality. *The Quarterly Journal of Economics 136*(1), 169–228.

Diamond, P. A. (1977). A framework for social security analysis. *Journal of Public Economics 8*(3), 275–298.

Dobkin, C., A. Finkelstein, R. Kluender, and M. J. Notowidigdo (2018). The economic consequences of hospital admissions. *American Economic Review 108*(2), 308–352.

Fadlon, I. and D. Laibson (2021). Paternalism and Pseudo-Rationality: An Illustration Based on Retirement Savings. Technical report.

Fadlon, I. and D. Laibson (2022). Paternalism and pseudo-rationality: An illustration based on retirement savings. *Journal of Public Economics 216*, 104763.

Fadlon, I. and T. H. Nielsen (2019). Household labor supply and the gains from social insurance. *Journal of Public Economics 171*, 18–28.

Fadlon, I. and T. H. Nielsen (2021). Family labor supply responses to severe health shocks: Evidence from danish administrative records. *American Economic Journal: Applied Economics 13*(3), 1–30.

Feldstein, M. (1985). The Optimal Level of Social Security Benefits. *Quarterly Journal of Economics 100*(2), 303–320.

Finkelstein, A., M. Gentzkow, D. Li, and H. L. Williams (2022). What drives risky prescription opioid use? evidence from migration. Technical report, National Bureau of Economic Research.

Finkelstein, A., M. Gentzkow, and H. Williams (2016). Sources of Geographic Variation in Health Care: Evidence From Patient Migration. *Quarterly Journal of Economics 131*(4), 1681–1726.

Finkelstein, A., M. Gentzkow, and H. Williams (2021). Place-Based Drivers of Mortality: Evidence from Migration. *American Economic Review 111*(8), 2697–2735.

Finkelstein, A. and N. Hendren (2020). Welfare analysis meets causal inference. *Journal of Economic Perspectives 34*(4), 146–167.

Finkelstein, A., N. Hendren, and M. Shepard (2019, April). Subsidizing health insurance for low-income adults: Evidence from massachusetts. *American Economic Review 109*(4), 1530â67.

Finkelstein, A., E. F. Luttmer, and M. J. Notowidigdo (2013). What good is wealth without health? the effect of health on the marginal utility of consumption. *Journal of the European Economic Association 11*(suppl_1), 221–258.

Finkelstein, A., E. F. P. Luttmer, and M. J. Notowidigdo (2009). Approaches to estimating the health state dependence of the utility function. *American Economic Review 99*(2), 116–121.

Fisher, R. (2023). Estimation of race and ethnicity by re-weighting tax data. *The Department of Treasury Office of Tax Analysis. Technical Working Paper 11.*

Ganong, P., F. Greig, P. Noel, D. M. Sullivan, and J. Vavra (2024, September). Spending and job-finding impacts of expanded unemployment benefits: Evidence from administrative micro data. *American Economic Review 114*(9), 2898â2939.

Ganong, P., D. Jones, P. Noel, F. Greig, D. Farrell, and C. Wheat (2020). Wealth, Race, and

Consumption Smoothing of Typical Income Shocks. *National Bureau of Economic Research Working Paper* (w27552).

Ganong, P. and P. Noel (2019). Consumer Spending during Unemployment: Positive and Normative Implications. *American Economic Review 109*(7), 2383–2424.

Gaubert, C., P. M. Kline, D. Vergara, and D. Yagan (2021). Place-based redistribution. Technical report, National Bureau of Economic Research.

Gerard, F. and J. Naritomi (2021, March). Job Displacement Insurance and (the Lack of) Consumption-Smoothing. *American Economic Review 111*(3), 899–942.

Goda, G. S., D. Jones, and S. Ramnath (2018). How do distributions from retirement accounts respond to early withdrawal penalties? evidence from administrative tax returns. Technical report, Working paper.

Goodman, L., J. Mortenson, K. Mackie, and H. R. Schramm (2021). Leakage From Retirement Savings Accounts In The United States. *National Tax Journal 74*(3), 689–719.

Hamdani, K., C. K. Mills, E. Reyes, and J. Battisto (2019). Unequal Access to Credit: The Hidden Impact of Credit Constraints. *Federal Reserve Bank of New York*.

Hendren, N. and B. Sprung-Keyser (2020). A unified welfare analysis of government policies. *The Quarterly Journal of Economics 135*(3), 1209–1318.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). Household Expenditure and the Income Tax Rebates of 2001. *American Economic Review 96*(5), 1589–1610.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary Policy According to HANK. *American Economic Review 108*(3), 697–743.

Kaplan, G., G. L. Violante, et al. (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica 82*(4), 1199–1239.

Keys, B. J., N. Mahoney, and H. Yang (2020). What Determines Consumer Financial Distress? Place- and Person-Based Factors. Technical report, National Bureau of Economic Research.

Kolsrud, J., C. Landais, P. Nilsson, and J. Spinnewijn (2018, April). The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden. *American Economic Review 108*(4-5), 985–1033.

Kolsrud, J., C. Landais, D. Reck, and J. Spinnewijn (2024, January). Retirement consumption and pension design. *American Economic Review 114*(1), 89â133.

Kreiner, C. T., D. D. Lassen, and S. Leth-Petersen (2019). Liquidity constraint tightness and consumer responses to fiscal stimulus policy. *American Economic Journal: Economic Policy 11*(1), 351–379.

Krueger, D., K. Mitman, and F. Perri (2016). Macroeconomics and Household Heterogeneity. In *Handbook of Macroeconomics*, Volume 2, pp. 843–921. Elsevier.

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics 112*(2), 443–478.

Laibson, D., P. Maxted, and B. Moll (2021). Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy. Technical report, National Bureau of Economic Research.

Landais, C. and J. Spinnewijn (2021). The Value of Unemployment Insurance. *Review of Economic Studies 88*(6), 3041–3085.

Leganza, J. M. (2024). The effect of required minimum distributions on intergenerational transfers. *Journal of Public Economics 232*, 105091.

Lusardi, A., D. J. Schneider, and P. Tufano (2011). Financially Fragile Households: Evidence and Implications. Working Paper 17072, National Bureau of Economic Research.

Maxted, P. (2020). Present Bias in Consumption-Saving Models: A Tractable Continuous-Time Approach. Technical report, Mimeo.

Mullainathan, S., J. Schwartzstein, and W. J. Congdon (2012). A reduced-form approach to behavioral public finance. *Annu. Rev. Econ. 4*(1), 511–540.

O'Donoghue, T. and M. Rabin (1999). Doing It Now or Later. *American Economic Review 89*(1), 103.

Parker, J. A. (1999). The Reaction of Household Consumption to Predictable Changes in Social Security Taxes. *American Economic Review 89*(4), 959–973.

Parker, J. A. (2017). Why Don't Households Smooth Consumption? Evidence from a $25 Million Experiment. *American Economic Journal: Macroeconomics 9*(4), 153–83.

Poterba, J. M. (2014). Retirement security in an aging population. *American Economic Review 104*(5), 1–30.

Read, D., G. Loewenstein, and M. Rabin (1999). Choice Bracketing. *Journal of Risk and Uncertainty 19*(1-3), 171–97.

Rong, M. (2023). Savings liquidity and consumption insurance: Evidence from ira early withdrawal penalties. Technical report.

Schmieder, J., T. von Wachter, and J. Heining (2018). The Costs of Job Displacement Over the Business Cycle and its Sources: Evidence from Germany. Technical report, Technical Report, UCLA, Mimeo.

Siliciano, R. L. and G. Wettstein (2021). Can the drawdown patterns of earlier cohorts help predict boomers behavior?

Skinner, J. (2007). Are you sure you're saving enough for retirement? *Journal of Economic Perspectives 21*(3), 59–80.

Souleles, N. S. (1999). The Response of Household Consumption to Income Tax Refunds. *American Economic Review 89*(4), 947–958.

Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association 13*(1), 130–167.

Stuart, E. and V. L. Bryant (2024). The impact of withdrawal penalties on retirement savings. *Journal of Public Economics 232*, 105083.

Sullivan, D. and T. Von Wachter (2009). Job Displacement and Mortality: An Analysis Using Administrative Data. *Quarterly Journal of Economics 124*(3), 1265–1306.

Thaler, R. H. (1999). Mental Accounting Matters. *Journal of Behavioral Decision Making 12*(3), 183–206.

Yagan, D. (2019). Employment Hysteresis from the Great Recession. *Journal of Political Economy 127*(5), 2505–2558.

Zeldes, S. P. (1989). Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy 97*(2), 305–346.

Figure 1: Conceptual Framework: Illustration

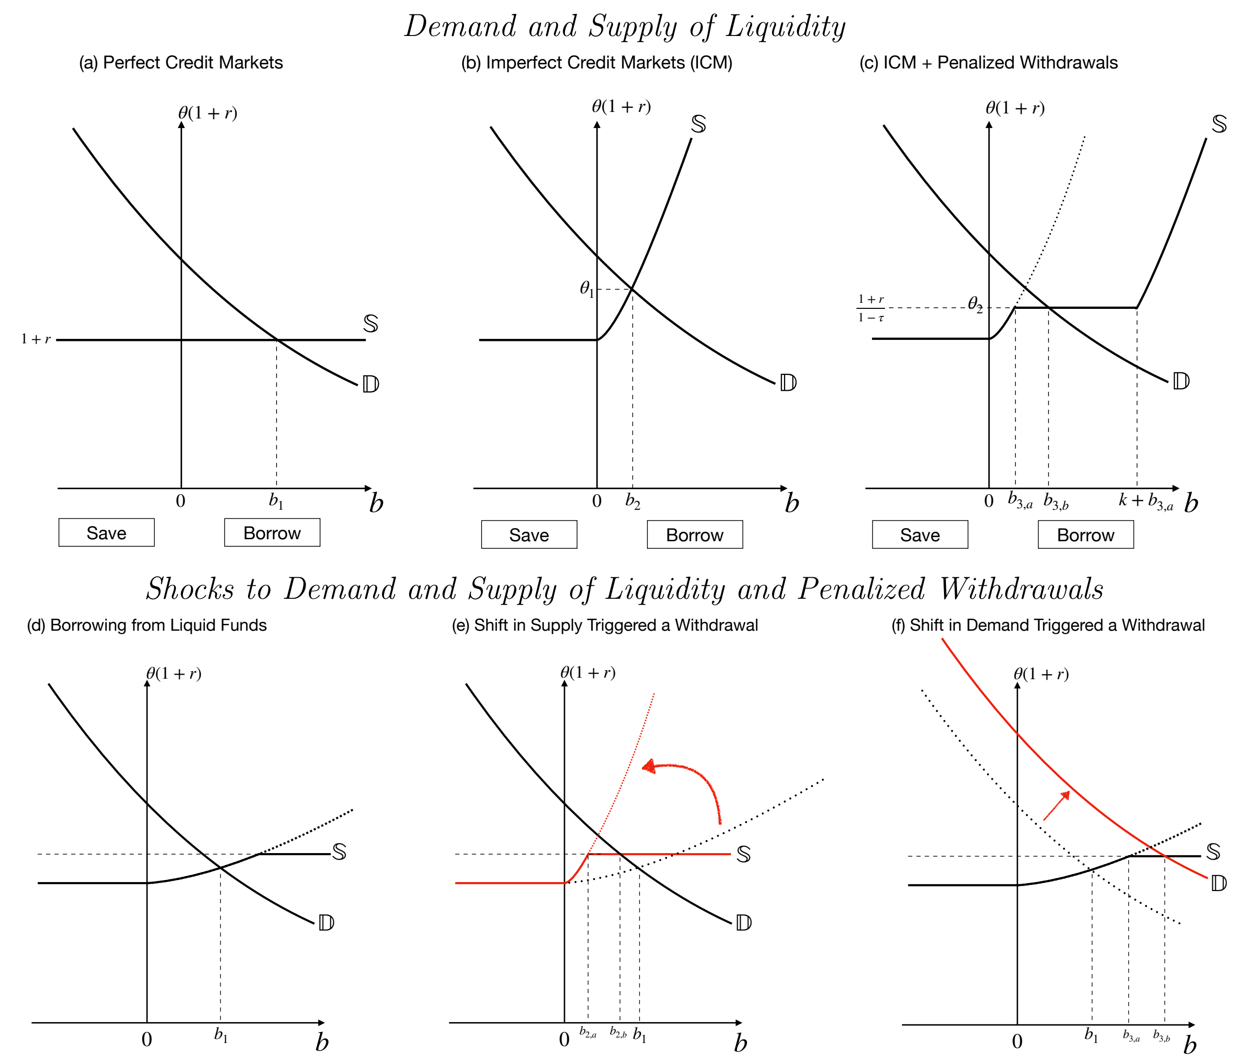

Notes: These figures illustrate how the supply and demand for liquidity, as defined in the main text, determine the equilibrium valuation of liquidity. The top three panels consider three different financial markets: perfect market (left), imperfect (middle), imperfect with the possibility to make a penalized withdrawal of at most an amount k (right). The bottom three panels study the impact of a shock to either the supply or demand of liquidity. The left panel shows the starting equilibrium, in which a household has access to a retirement savings account, but chooses to borrow from liquid funds. The middle panel shows a shock to the supply of liquidity (tightening credit conditions) which triggers a penalized withdrawal. The right panel shows a shock to the demand for liquidity (e.g., a negative income shock) which also triggers a penalized withdrawal.

Figure 2: Unemployment and Penalized Withdrawals

(a) Frequency of Withdrawals

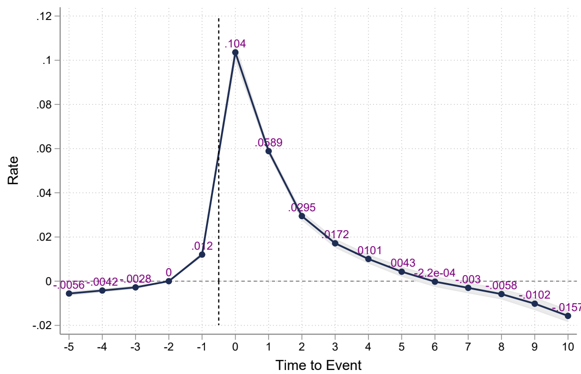

(b) Withdrawn Amounts

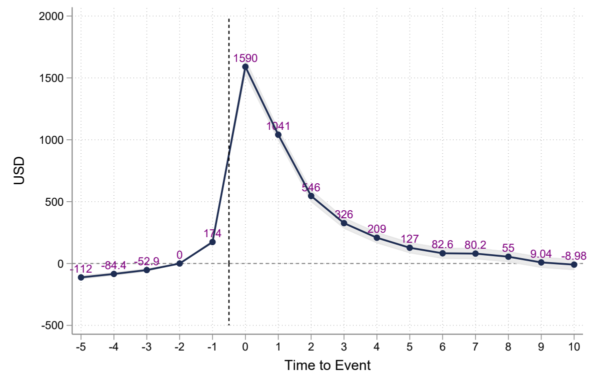

(c) Event Response by Race

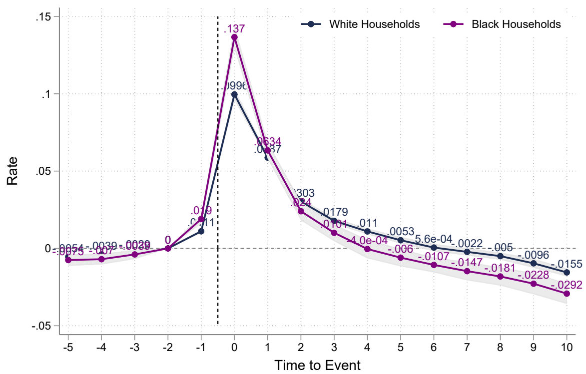

(d) Event Response by Capital Income

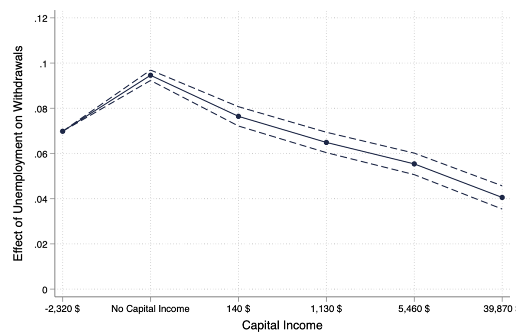

Notes: This figure studies penalized withdrawals around the event of household unemployment, defined as the first period we observe at least one of the household members receiving unemployment benefits. Panels A and B plot the event study coefficients from specification (3) when the outcome variables are take-up and amounts of penalized withdrawals, respectively. Panels C and D study heterogeneity by household characteristics. Panel C plots the unemployment event study of the take-up of penalized withdrawals, split by whether the household's primary filer is Black or White. Panel D plots how the point estimates at time 0 (i.e., at the onset of the unemployment event) vary as a function of household capital income. We split households into those with negative, zero, and positive capital income, and we bin the ones with positive capital income into four equal-sized groups. The labels on the x-axis indicate the average capital income within the corresponding bin.

Figure 3: Income Changes and Penalized Withdrawals

(a) Withdrawal Frequency after Large Income Loss

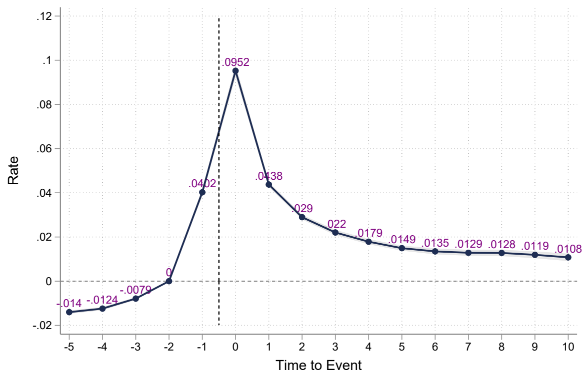

(b) Withdrawn Amounts after Large Income Loss

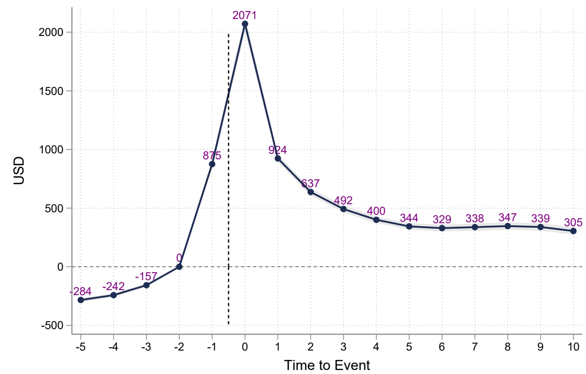

(c) Take-Up as a Function of Income Changes

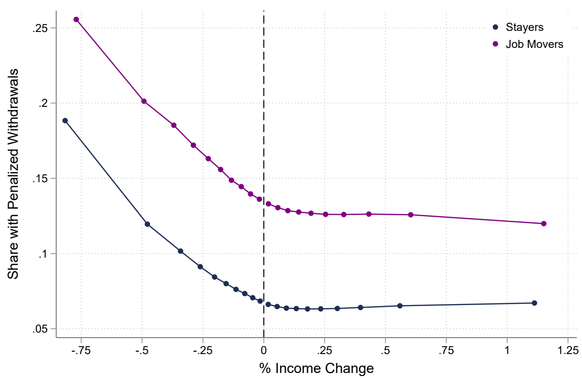

(d) Amounts as a Function of Income Changes

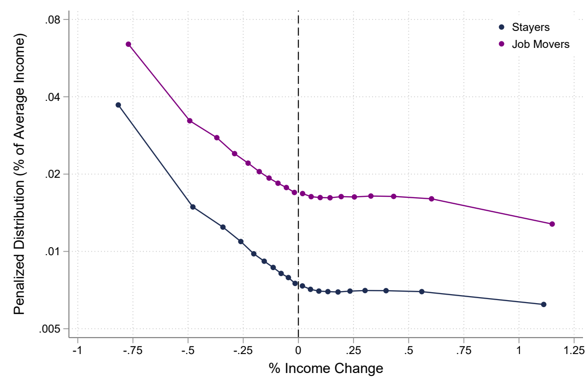

(e) Responses to Large Income Loss by Race

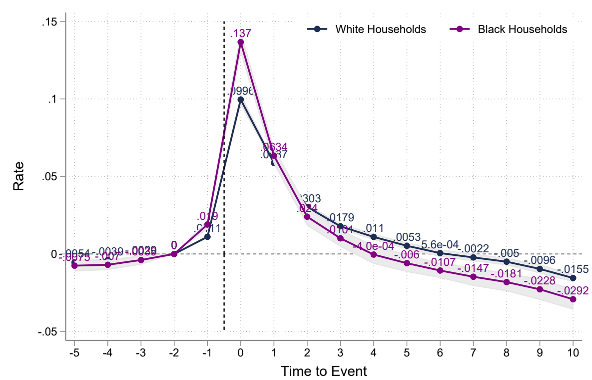

Notes: This figure studies penalized withdrawals around changes in household income. Panels A and B plot the event study coefficients from specification (3) when the outcome variables are take-up and amounts of penalized withdrawals, respectively, around the event of a large income loss. Large income loss is defined as the first period we observe a household experiencing a decline in overall income of more than 30 percent (relative to a previous year). Panels C-D study households' take-up and amounts of withdrawals as a function of the deviation of their income flow from their average income across our data period. We split households by whether a member of the household switched jobs that year because job changes themselves, as displayed in Appendix Figure D.8, lead to increased take-up. Panel E plots the event study of take-up of penalized withdrawals, split by whether the household's primary filer is Black or White.

Figure 4: Penalized Withdrawals and Local Unemployment during the Great Recession

(a) Frequency of Withdrawals

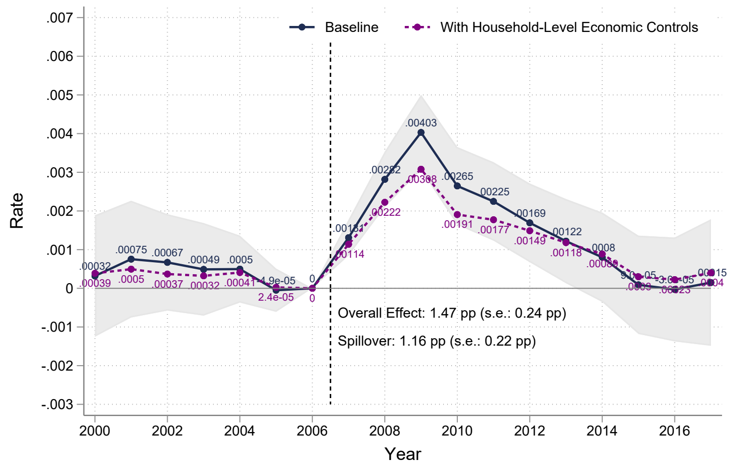

(b) Withdrawn Amounts

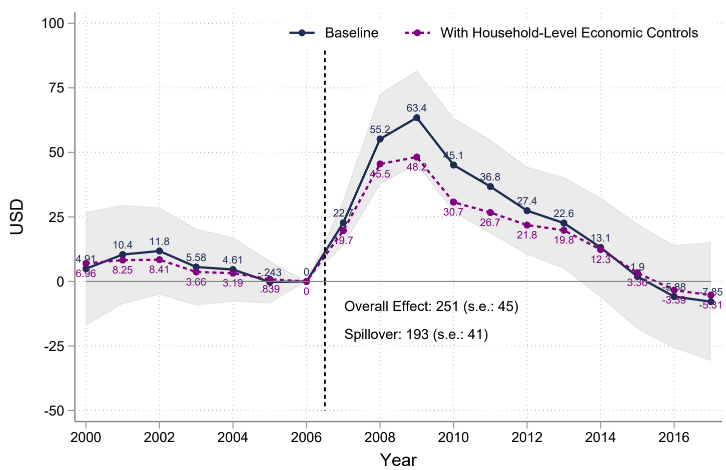

Notes: This figure displays estimates of the effect of the Great Recession on penalized withdrawals using equation (4). It provides estimates for the relative change in behavior in a locality that was exposed to a 1 percentage point larger local unemployment shock. Panel A analyzes the frequency of withdrawals, and panel B analyzes withdrawal amounts.

Figure 5: Geography of Withdrawals and Valuation of Liquidity

(a) Overall Variation

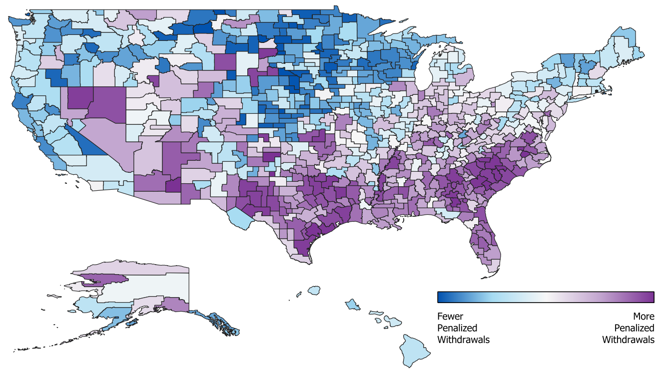

(b) Location Fixed Effects

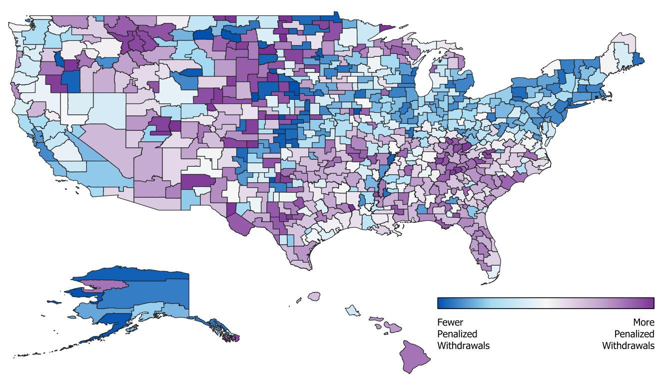

(c) Household Fixed Effects

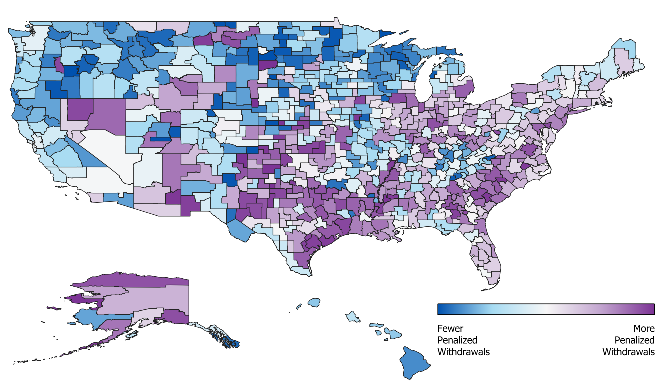

Notes: This figure studies the geography of penalized withdrawals and the valuation of liquidity. Panel A plots a map of the average annual share of households that have made a penalized withdrawal by Commuting Zones (CZs). Then, based on estimation of equation (5) with household-level economic controls, panel B plots a map of the location fixed effects, Γ_z, and panel C plots a map of the household fixed effects, α_i, collapsed at the CZ level. The economic controls include unemployment, wage earnings, and gross income, with lagged, current, and lead values.

Figure 6: Withdrawals and Locality Characteristics

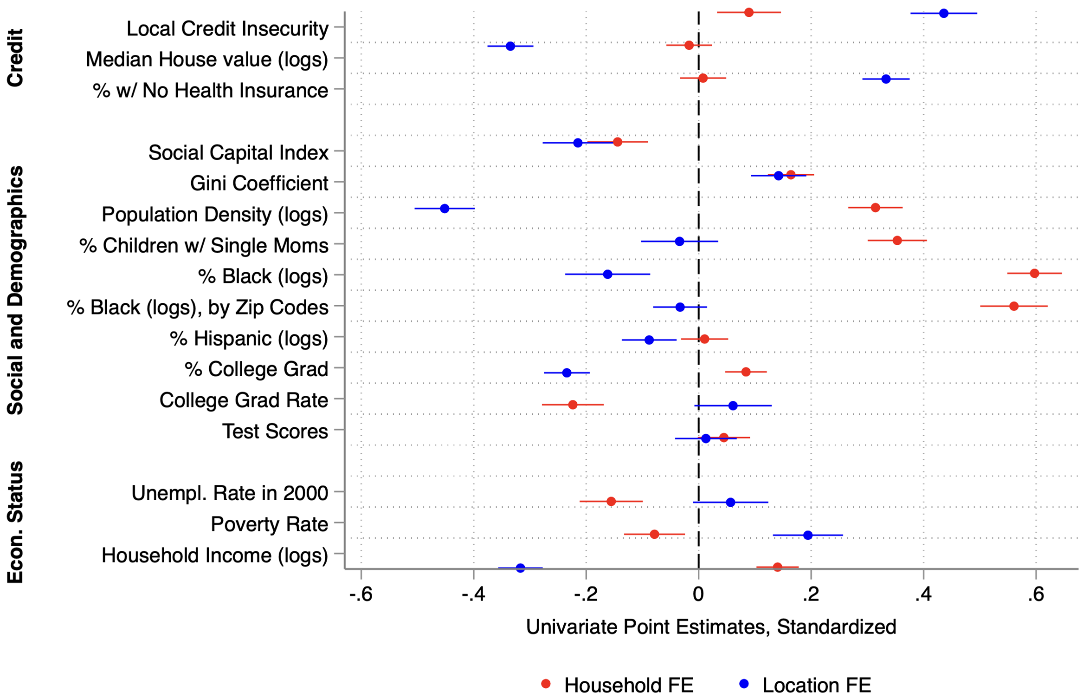

Notes: This figure display correlations of the regional differences across CZs with CZ-level social and economic characteristics. We display correlations of these characteristics separately for the location fixed effects, Γ_z, and for the household fixed effects, α_i, collapsed at the CZ level.

Figure 7: Movers Analysis

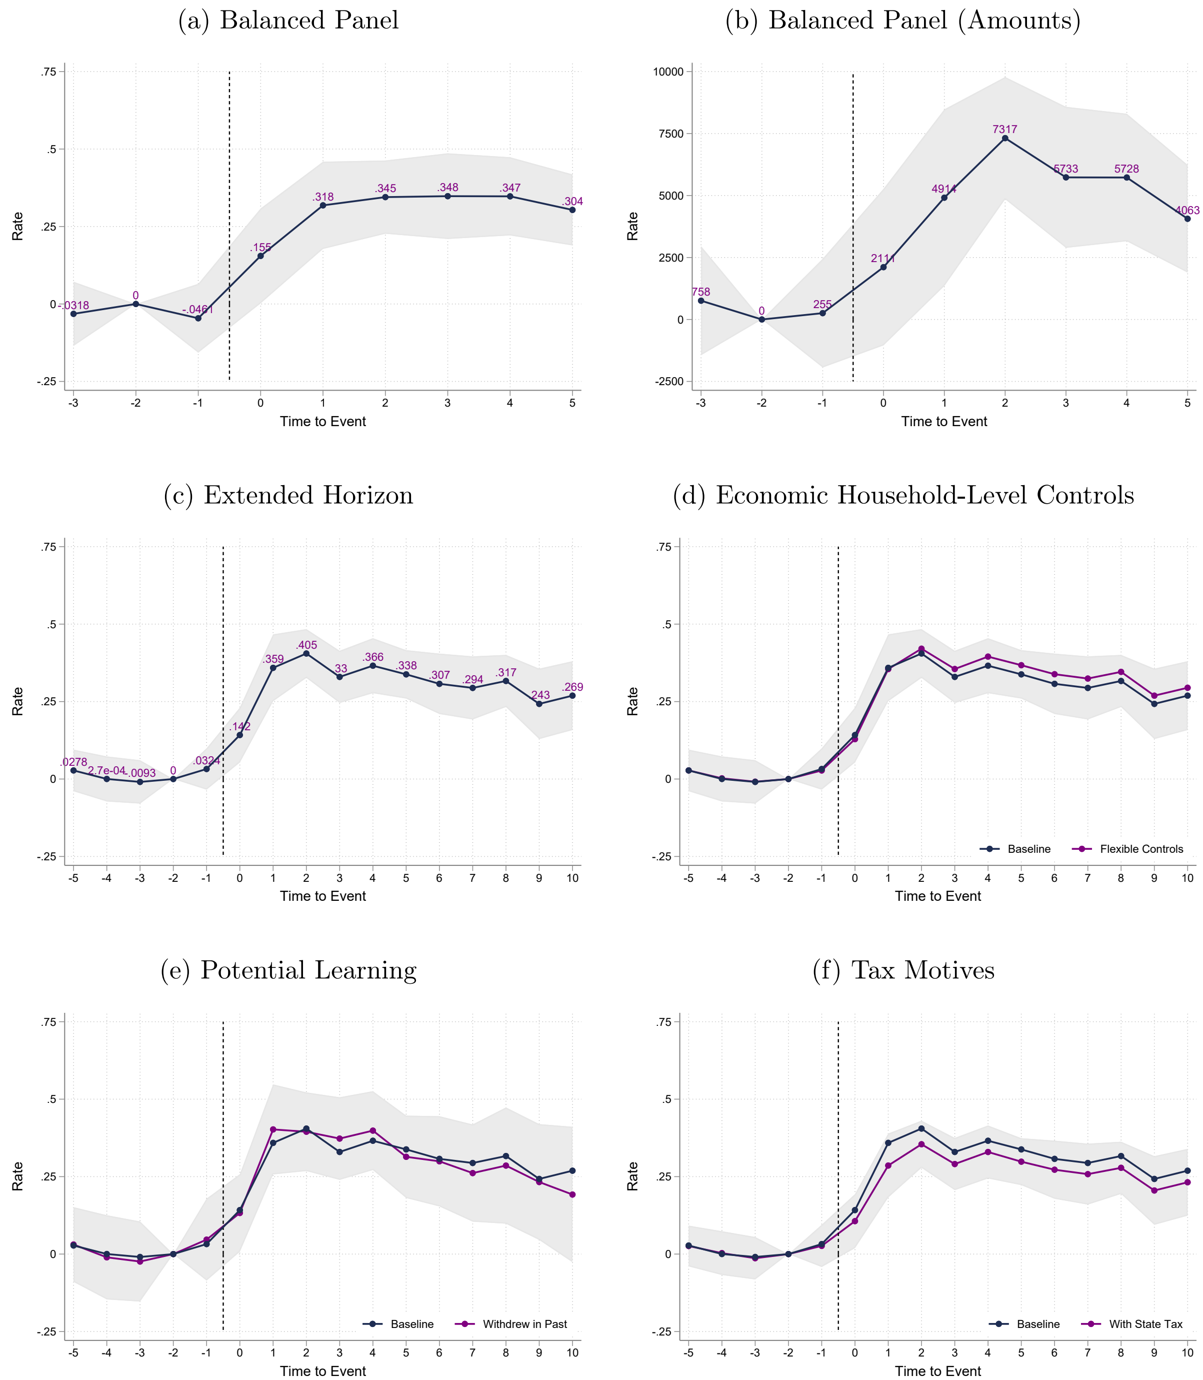

Notes: These figures display estimates for the share of spatial differentials in withdrawals that can be attributed to location, using the movers design specification of equation (7). Panels A and B show the estimates from a balanced panel of households we observe in the window [-3, +5] years around the move for the outcomes take-up and amounts of penalized withdrawals, respectively. Panel C shows the estimates from an unbalanced panel of households on an extended time window that spans the years [-5, +10] around the move for withdrawal take-up. The corresponding plot for withdrawal amounts for an extended time horizon around the move is reported in Appendix Figure D.11. Panels D-F provide a series of robustness investigations. Panel D runs a specification that includes flexible (endogenous) economic controls: unemployment, wage earnings, and gross income, with lagged, current, and lead values, including interactions of all these variables with time with respect to the move. Panel E studies learning as a potential channel by focusing on the sample of households who had already made a penalized withdrawal in the pre-move periods. Panel F tests the explanation of tax optimization by including controls for a location's top marginal tax rate (that varies over state and time) flexibly interacted with time relative to the move. In all estimations, we include as controls household fixed effects, a full set of primary-filer age fixed effects, and (cyclical) calendar year fixed effects. Robust standard errors are clustered at the origin CZ level.

Table 1: Penalized Withdrawals and Race

	Making a Penalized Withdrawal	
	(1)	(2)
Black	0.0314	0.0294
	(0.0019)	(0.0018)
Constant	0.1018	0.0950
	(0.0009)	(0.0008)
Economic Controls		X
Number of Households	7,317,958	7,317,958

Notes: This table studies the association between the frequency of making a penalized withdrawal and race. The sample includes all households in which the primary filer is either Black or White. We study the correlation between the primary earner's race and the estimated household-level fixed effect from specifications of equation (5) without economic controls in column 1 and with economic controls in column 2. The economic controls include unemployment, wage earnings, and gross income, with lagged, current, and lead values. Robust standard errors are clustered at the CZ level.

Online Appendix

A Data Description

Sample construction. We start by randomly selecting 10 percent of US. individuals. This selection is based on the last 4 digits of an individual's social security number. While social security numbers historically have not been completely random, the last 4 digits have always been randomly assigned. For these individuals, we collect date of birth and (possible) date of death from Social Security Administration administrative data.

We next build a sample of US taxpayers by selecting a taxpayer if either the primary filer or the spouse is included in the 10 percent sample of individuals. We identify spouses for those who file using the status "Married filing separately" or "Married filing jointly," and we aggregate the data to create a consistent panel at the household level throughout our data's time range. Our data span the tax years 1999 through 2018, and we focus on "prime age" households with primary filers aged 45-59.

Data related to retirement accounts and withdrawals. Form 1099-R provides information on gross distributions in Box 1, and taxable amounts in Box 2. Importantly, Box 7 provides a code that describes the character of the distribution. This code helps to indicate whether a distribution would be subject to the additional tax penalty. It also provides a check box next to Box 7 that indicates whether the distribution was from an IRA/SEP IRA/SIMPLE IRA.

While Form 1099-R provides important information on the character of a retirement distribution, it does not provide information on contributions nor account balances. We can add this information for IRAs using Form 5498. Box 1 provides traditional IRA contributions, while Box 10 contains Roth IRA contributions. Boxes 8 and 9 show the amount of SEP and SIMPLE IRA contributions, respectively. As an additional check to the type of account, Box 7 includes check boxes that denote the character of the account. Box 5 provides the fair market value of the account, which we can use as a measure of the retirement resources available at a point in time.

Furthermore, since early withdrawals of Roth IRA contributions are not subject to the penalty, we also collect information provided on Form 8606. We collect the taxable Roth IRA distribution amount reported in Line 25c, which reports only distributions of earnings. This portion of the distribution is the only part that would be subject to the 10 percent penalty if not corrected on Form 5329 (see below).

The information provided on Form 1099-R is also subject to the information available to the fund manager at the time of the withdrawal. The fund manager is unlikely to know if a withdrawal made with no known exception is later rolled into another qualified account manually within 60 days. In instances such as these, taxpayers are instructed to fill out Form 5329, which allows taxpayers to essentially provide information on what portion of their early distributions are not subject to the additional tax penalty. For example, a taxpayer may fill out Form 5329 and claim an exception from the early distribution penalty by indicating the distribution was made for qualified expenses, such as medical expenses, health insurance premiums, qualified higher education expenses, first-time home purchase, qualified reservist distribution, or qualified birth or adoption distributions. We use this information, reported in Part I of Form 5329 Lines 1-4, to better identify which early distributions are subject to the additional tax penalty.

We measure a penalized distribution as one that is reported on Form 1099-R with distribution codes 1, J, or S that has not been otherwise corrected by the taxpayer as a non-penalized distribution on Form 5329 or Form 8606. We do this first by reducing the amount of seemingly penalized distributions with code J to the updated amount from Form 8606 when a Form 8606 is present. Then we reduce the total amount of penalized distribution to the amount reported on Form 5329 if Form 5329 is present. If no Form 5329 is present, then we assume the taxpayer pays a penalty on the full amount of distributions labeled with distribution codes 1, J, and S. Together, these changes capture the actions available to taxpayers to rectify Forms 1099-R that may incorrectly categorize distributions as being subject to a penalty.

Data related to access to penalized early distributions. In our main analysis, we condition on ever having access to a retirement account. We measure having access to a retirement account if the taxpayer (primary filer or spouse) reports: (a) contributions to a retirement account, (b) a positive balance for an IRA, or (c) a retirement account distribution. Contributions can be reported on either Form W-2 for employer-sponsored plans (Box 12 includes a check box for employers to indicate whether the employee is an active participant in a retirement plan), or on Form 5498 for IRAs. Fair market value of IRA accounts is also reported on Form 5498.

While contributions or positive IRA balances reported on Form 5498 clearly indicate access to a retirement fund, the information on Form W-2 is more ambiguous. The check box in Box 12 includes both defined benefit plans and defined contribution plans. For the purposes of our analysis, we want to condition on those who have access to retirement funds and could withdraw those funds, which most generally only includes those participants in a defined contribution plan. We use information from Form 5500 compiled by the Center for Retirement Research at Boston College to identify which of the employers in our sample offered a defined contribution plan. We can match about 20 percent of our sample's employers and find that over 90 percent offered a defined contribution plan. Thus, while we do not directly observe whether an individual taxpayer subscribes to a defined contribution plan with the employer, we at least know that most of the taxpayers identified as having access to retirement funds by our instrument participated in a retirement plan with an employer that offered a defined contribution plan.

Finally, if we see in our sample period a taxpayer taking a distribution from a retirement account but fail to see any retirement fund contributions or balances as noted above, then we assume those contributions were made prior to the beginning of our sample and thus assume the taxpayer has had access for our entire sample.

Data related to demographics and economic circumstances. Form 1099-G reports unemployment insurance (UI) payments made to individuals. We define an unemployment event based on receiving UI payments..[27] We define a large negative income shock as a deviation of 30 percent or more from a rolling average adjusted gross income less penalized distributions over the sample period. Comparing to a rolling average helps to prevent coding the year after a positive shock as a negative shock in income.

We say that a taxpayer moves if the address reported on their tax return places them in a different Commuting Zone (CZ) than in the year prior. Note that this omits local moves within the same CZ. We infer a change of primary job by seeing if the highest paying W-2 switches from one payer to a different payer between two years.

Finally, we impute race using the methodology described in Fisher (2023). This method uses information on a taxpayer's name, location at a given time, family characteristics, and income characteristics to predict race and ethnicity. Dummy variables for race are then created based on which estimated probability is highest for each taxpayer.[28]

[27] Note that tax data are reported annually, so there are potential timing issues where UI payments can span across years. In our data this would appear as 2 straight years with unemployment spells, but we cannot distinguish between a single spell that spans December to the following January and two separate unemployment spells.

[28] Note that Fisher (2023) includes Hispanic origin as a mutually exclusive category from other races. This differs from other data sources (e.g., the Census Bureau) which include separate indicators for race and Hispanic origin.

B Preliminary Facts and Their Implications

We document four sets of key facts about US households' use of penalized withdrawals. These facts offer support for the hypothesis that penalized withdrawals are used as self-insurance for short-run liquidity needs, and they accordingly motivate the focus of the model in Section 3 and our core empirical analysis thereafter.

Fact 1: Most households have retirement accounts. Appendix Figure D.1 shows the prevalence of retirement savings accounts across US households, by age and income, focusing on all households whose primary filer is between ages 25 and 70. Panel (a) shows that, for our selected age group (ages 45-59), almost 90 percent of households have at least one account. Panel (b) shows that, among households with income above the median (marked by the vertical line), almost every household has an account. Accounts are instead less prevalent, as expected, for lower-income households. Nonetheless, even among the households with low levels of annual income, e.g., between \$10,000 and \$20,000, approximately half have an account. We note that the high prevalence is reflective of our analysis unit of interest, that is, a household, rather than individuals. We further corroborate the prevalence of defined-contribution retirement accounts that we impute from our data using the Health and Retirement Study (HRS). The HRS is a longitudinal panel study that surveys a representative sample of approximately 20,000 people in America and is widely-used in retirement related research in the US. We use data from waves 7-14, which cover the years 2004-2018, and focus on households with primary respondents between the ages of 45-59 for whom we can identify an account type (DC or DB) or whether the household reported not having an account. Among these households, we calculate that 14,392 have at least one defined contribution account, which amounts to a share of 84.14%. We note that the HRS is a representative sample of overall households in the U.S., whereas we focus on tax filers and thereby exclude non-filers who have less resources and could be expected to have accounts at lower rates. Indeed, in Appendix Figure D.3 we find that overall prevalence rates shift downwards moderately when non-filers are included, with an average account prevalence rate of 83.8% over ages 45-59.

Fact 2: Penalized withdrawals are widely used but infrequently. Next, panels (c) and (d) of Appendix Figure D.1 show that penalized withdrawals are widely used by households throughout the age and income distributions. Almost 10 percent of households within our age group make a penalized withdrawal in any given year. Penalized withdrawals are prevalent across the age distribution, but they fall, as expected, after age 55, when separation from employers becomes an excepted event. They are also prevalent across the income distribution, along with a declining frequency as household income increases. This is consistent with the idea that higher-income households have alternative cheaper sources of short-run liquidity to insure against economic shocks. Importantly, penalized withdrawals are not concentrated among a few households, but are a prevalent liquidity tool across the whole population. Panel (e) of Appendix Figure D.1 shows that almost half of all households observed for 15 consecutive years in our sample take a penalized withdrawal in at least one year. Moreover, the typical household withdraws infrequently, consistent with the hypothesis that households use penalized withdrawals as a tool to access liquidity when the need arises. Finally, panel (f) of Appendix Figure D.1 shows among households who made a withdrawal in some period, the distribution of subsequent years within our data frame that the household made additional withdrawals.[29] The figure displays a large mass at zero, consistent with penalized withdrawals reflecting temporary financial constraints that require short-run liquidity.

Fact 3: Withdrawn amounts are sizable, yet accounts are not fully depleted. Panel (a) of Appendix Figure D.2 shows the CDF of the dollar amounts of penalized distributions. The typical withdrawal is approximately $5,000. Importantly, penalized withdrawals are usually not associated with an account closure and they deplete only a relatively small fraction of the available funds. Here, we leverage the fact that the data include outstanding balances for IRA accounts. We look at households who have an IRA account at time $t-1$ and who make a penalized withdrawal from an IRA account between periods $t-1$ and t. Panel (b) of Appendix Figure D.2 shows that the share of households who deplete funds is consistently below half throughout the account balance distribution and that it is much lower, as expected, among households with non-trivial amounts in their accounts. Second, in panel (c) of Appendix Figure D.2 we plot the CDF of the ratio of penalized IRA distributions out of balances for households that do not fully deplete their accounts: the median withdrawal depletes approximately 25 percent of outstanding IRA balances. Overall, in the context of IRAs where we have information on balances, the evidence shows that most households are within an interior solution with respect to their withdrawal decision margin.[30] This evidence is consistent with the interpretation that penalized distributions are a result of households withdrawing the necessary amount of money to self-insure a shock rather than closing old or secondary accounts, which could have been, in principle, a concern in using penalized withdrawals as a revealed preference tool.

Fact 4: Penalized withdrawals are strongly associated with income losses. Lastly, panel (d) of Appendix Figure D.2 shows that households who make a penalized withdrawal are more likely to have suffered an income loss. We plot the CDF of annual income changes, separating households according to whether they are making a penalized withdrawal in a given year. Among households who make a penalized withdrawal, almost 60 percent have experienced an income loss.

[29] We provide two versions of this distribution for different definitions of the withdrawal periods, one that uses a one-year period and another that uses a three-year period (to allow for a longer period of "consecutive" liquidity needs).

[30] This goes in tandem with the patterns in panel (a) of Appendix Figure D.2 where penalized distributions are lower compared to any distribution, consistent with the idea that households limit the amount withdrawn due to the presence of the marginal penalty.

Moreover, they are more likely to have experienced large income losses. For example, they are twice as likely to have suffered an income loss larger than 50 percent relative to households who have not made a penalized withdrawal.

Taken together, these four facts provide evidence that households use penalized withdrawals as a mean to mitigate short-run needs for liquidity. This evidence thus motivates us to use penalized withdrawals as a revealed-preference tool to characterize the needs and valuation of liquidity across American households. Yet, we address in what follows two potential concerns with our approach. The first is that in our main dataset we cannot observe how households use their funds, and hence we cannot directly show that these funds are used for self-insurance. The second is that any revealed-preference approach relies on the assumption that agents are maximizing choices on the studied margin.

Evidence from Health and Retirement Study (HRS). We complement our data with information on premature withdrawals among American families from the Health and Retirement Study (HRS). Despite small samples, the key benefit from doing so is that households are asked to provide the reasons they withdrew funds prematurely. To get closest to our population, we use survey waves 7-14 which cover the years 2004-2018, and we further restrict the sample to respondents who have defined contribution pension plans and are of ages 45-59. The survey does not separate penalized from non-penalized withdrawals, so we provide statistics that pertain to any withdrawal that occurs prior to age 59.5 upon which the penalty is waived.

We rely on two main questions in the HRS that relate to a household's experience between consecutive waves which are typically two years apart. The first question pertains to withdrawals and asks: "Not including any money you rolled into an IRA, not including any money you used to purchase an annuity. How much money in total did you 'withdraw'/'receive in payments' 'since you left that business or employer'/'since we last talked to you in [Previous Wave Interview Month] [Previous Wave Interview Year]'?" The second question pertains to the usage of withdrawn funds and asks: "What did you do with the money?" where respondents can choose among the options: bought durables (house, car, etc.), spent it, saved/invested, paid off debt, rolled into IRA, gave it away, other, as well as don't know and refuse to answer. The information on the usage of withdrawals that we use is based on the first usage indicated by the household.

Appendix Table B.1 summarizes these statistics. Panel (a) first provides the distribution of amounts of balances in defined contribution accounts and withdrawals from them, with numbers that are broadly in line with total withdrawals in Appendix Figure D.2 from the tax data. Second, the taxonomy of uses of funds from early withdrawals in panel (b) aligns well with the notion that these funds are used to finance concurrent expenditure needs or repay outstanding debt. These results corroborate the indirect evidence provided from the tax data that early withdrawals are a signal of liquidity needs.

Table B.1: Health and Retirement Study (HRS): Defined Contribution Accounts

(a) Distribution of Balances and Withdrawal Amounts

	Mean	10th	25th	50th	75th	90th
Balances	147,456	3,000	12,000	50,000	154,900	370,000
Withdrawals	20,489	1,220	3,000	8,000	20,000	43,600

(b) Use of Withdrawals

	Number of Observations	Percent
Bought durables	578	14
Spent it	1,306	31
Saved/invested	661	16
Paid debt	985	24
Rolled into IRA	141	3
Gave it away	104	3
Other	249	6
Don't know	56	1
Refused	69	2
Total	4,149	100

Notes: These tables display summary statistics on defined contribution accounts from the Health and Retirement Study (HRS). We use HRS data from waves 7-14, which cover the years 2004-2018. The sample is restricted to respondents who are between the ages of 45-59.5. We focus on the 14,392 households who have defined contribution pension plans, who represent a population share of 84.14%. For these households, the average and median ages are 59.3 and 58.2 years old, respectively. The first line in panel (a) displays the distribution of balances in their retirement accounts. We then use two main questions in the HRS, which relate to a household's experience between consecutive waves that are typically two years apart. The first question pertains to withdrawals and asks: "Not including any money you rolled into an IRA, not including any money you used to purchase an annuity. How much money in total did you 'withdraw'/'receive' in payments 'since you left that business or employer'/'since we last talked to you in [Previous Wave Interview Month] [Previous Wave Interview Year]'?" The second question pertains to the usage of funds and asks: " What did you do with the money?" where respondents can choose among the options: bought durables (house, car, etc.), spent it, saved/invested, paid off debt, rolled into IRA, gave it away, other, as well as don't know and refuse to answer. Combining the responses to the two questions, we identified 3,279 unique households with withdrawal episodes. Among them, only 222 observations have non-missing positive values. The second row in panel (a) displays the distribution of these withdrawn amounts. The table in panel (b) displays the usages of withdrawals among the 3,279 households identified a with withdrawal episode. We count multiple usages if funds were used for more than one reason within a withdrawal episode.

Possible Behavioral Interpretations. Revealed-preference approaches rely on households' ability to optimize on the margin investigated. The regularities we have seen above are closely consistent with various predictions of a model in which households optimize on the margin of taking penalized withdrawals. Still, it is important to assess the degree to which alternative explanations could drive the observed behavior. Indeed, economists justify the existence of illiquid accounts, that are either fully illiquid such as Social Security or partially illiquid such as 401(k)s/IRAs, with a trade-off between taste shocks (e.g., a realization of a real consumption need) and present biases that may lead them to over-consume (Amador et al. 2006; Beshears et al. 2020b; Fadlon and Laibson 2021).[31] In our context, the main concern is that the observed behavior could be generated by behavioral biases, such as *narrow bracketing* (e.g., Thaler 1999), *mental accounting*

[31] In fact, one traditional rationale for government intervention in retirement savings (particularly in the form of Social Security) has been that some individuals lack the foresight to save for their retirement years (Diamond 1977; Feldstein 1985).

(e.g., Read et al. 1999), or *myopia/present bias* (e.g., Laibson 1997; O'Donoghue and Rabin 1999), and may not convey information on the underlying valuation of liquidity. Reassuringly, as we next discuss, the evidence presented in the beginning of this appendix is not consistent with these interpretations. Of course, while the evidence is inconsistent with these behavioral explanations governing the results, they could still play a role.

We first consider *narrow bracketing*, whereby households do not integrate their entire portfolios into their decision making. The facts that most households withdraw sizable amounts and that the penalized withdrawals are only infrequently linked to the closure of a specific account (when we observe balances) mitigate this concern. With narrow bracketing we would have expected withdrawals to be the result of households disregarding some small amounts left in isolated accounts, which they then might close down without a direct link to their actual liquidity needs.

Second, under *mental accounting*, households' behavior would involve some assignment of activities to specific accounts, thereby potentially avoiding the liquidation of funds that are mentally designated for consumption later in the future. In contrast, we have seen that withdrawals are prevalent across the whole population and that they are increasingly used exactly when large income losses occur.

Third, if penalized withdrawals were driven by *myopic* behavior among a particular share of the population with present bias, we would expect to observe that most of the withdrawals are due to repeated take-up by the same set of households. Instead, panel (e) of Appendix Figure D.1 shows that withdrawals are rare for any given household and widespread across the population. While the data are inconsistent with the particular margin of penalized withdrawals being driven by myopia, some households are naturally present-biased and the infrequency of penalized withdrawals certainly does not preclude their presence. However, in such a case, observing a penalized withdrawal would still inform us about the relative valuation of liquidity in a given period among optimizing ("non-naive") present-biased households as implied by the properties of their value functions developed in Maxted (2020).

C Proof of Lemma 1 and Generalization

We first characterize the solution of the model described in Section 3 to prove Lemma 1. We then offer a generalization of the Lemma to the case in which households are further from the statutory retirement age.

Proof of Lemma 1. We start from the recursive formulation of the problem

$$V_t\left(a_{i,t-1},k_{i,t-1};h_{i,t}\right) = \max_{\Delta k_{i,t},\Delta a_{i,t}} u\left(c_{i,t};h_{i,t}\right)+\beta E_t\left[V_{t+1}\left(a_{i,t},k_{i,t};h_{i,t+1}\right)\right]$$

subject to

$$\begin{aligned}
c_{i,t} &= (1-\varphi)\,y_{i,t}-\varepsilon_{i,t}-\Delta k_{i,t}-\Delta a_{i,t}+\tau\Delta k_{i,t}\mathbb{I}_{(\Delta k_{i,t}<0)}\mathbb{I}_{(t<t^*)}-\rho_{i,z}\left(b_{i,t}\right)\mathbb{I}_{(b_{i,t}>0)}\\
a_{i,t} &= (1+r)\left[a_{i,t-1}+\Delta a_{i,t}\right]\\
k_{i,t} &= (1+r)\left[k_{i,t-1}+\Delta k_{i,t}+\varphi y_{i,t}\right],\\
b_{i,t} &= \max\{0;a_{i,t-1}-\Delta a_{i,t};-\Delta a_{i,t}\}
\end{aligned}$$

First, notice that the household would never deposit into the illiquid account, i.e., $\Delta k_{i,t}\leq 0$, since the illiquid account pays the same interest rate as the liquid account but it leads to a penalty in the case of a withdrawal, hence it is strictly dominated as a saving instrument. For this same reason, we know that $\frac{\partial V_{t+1}(a_{i,t+1},k_{i,t+1};h_{i,t+1})}{\partial k_{i,t+1}}\leq\frac{\partial V_{t+1}(a_{i,t+1},k_{i,t+1};h_{i,t+1})}{\partial a_{i,t+1}}$ with an equality sign if and only if the household knows with certainty that she is not going to make a penalized withdrawal nor borrow from the liquid account (which would entail them paying the marginal cost $\rho_{iz}\left(b_{i,t}\right)$) before date t^*. In this latter case, all dollars deposited in the illiquid account will become liquid with certainty and would not be used before t^* since the household is not expecting to need liquidity from any source before t^*. As a result, $\frac{\partial V_{t+1}(a_{i,t+1},k_{i,t+1};h_{i,t+1})}{\partial k_{i,t+1}}=\frac{\partial V_{t+1}(a_{i,t+1},k_{i,t+1};h_{i,t+1})}{\partial a_{i,t+1}}$.

Next, take the first order conditions with respect to $\Delta k_{i,t}$ and $\Delta a_{i,t}$, taking into account that the derivative is different depending on whether the values of this choice variables are positive or negative (and excluding the non-relevant case $\Delta k_{i,t}>0$), we get:

$$\{b_{i,t}=0\}:\qquad u'\left(c_{i,t};h_{i,t}\right)\left(\beta\left(1+r\right)\right)^{-1}=E_t\left[\frac{\partial V_{t+1}\left(a_{i,t+1},k_{i,t+1};h_{i,t+1}\right)}{\partial a_{i,t+1}}\right] \tag{8}$$

$$\{\Delta a_{i,t}<0,b_{i,t}>0\}:\qquad u'\left(c_{i,t};h_{i,t}\right)\left(\beta\left(1+r\right)\right)^{-1}=\frac{E_t\left[\frac{\partial V_{t+1}\left(a_{i,t+1},k_{i,t+1};h_{i,t+1}\right)}{\partial a_{i,t+1}}\right]}{1-\rho'_{i,z}\left(b\right)} \tag{9}$$

$$(10)\qquad \{\Delta k_{i,t}<0\}: \qquad u'\left(c_{i,t};h_{i,t}\right)\left(\beta\left(1+r\right)\right)^{-1} \geq \frac{E_t\left[\frac{\partial V_{t+1}\left(a_{i,t+1},k_{i,t+1};h_{i,t+1}\right)}{\partial k_{i,t+1}}\right]}{1-\tau}.$$

$$(11)\qquad \{\Delta k_{i,t}=0\}: \qquad u'\left(c_{i,t};h_{i,t}\right)\left(\beta\left(1+r\right)\right)^{-1} \leq \frac{E_t\left[\frac{\partial V_{t+1}\left(a_{i,t+1},k_{i,t+1};h_{i,t+1}\right)}{\partial k_{i,t+1}}\right]}{1-\tau}$$

Lemma 1 is then derived directly from the first order conditions, the definition of $\theta_{i,t}\left(c_{i,t};h_{i,t}\right)$, and the argument we just made on the relationship between $\frac{\partial V_{t+1}}{\partial k_{i,t+1}}$ and $\frac{\partial V_{t+1}}{\partial a_{i,t+1}}$, which are identical under the assumption of Lemma 1.

Valuation of Liquidity and Access to Penalized Withdrawals. Next, we describe how we derive the general bound for the valuation of liquidity, which we use to obtain equation 2 and to prove the generalized version of Lemma 1 below.

To see why the valuation of liquidity is bounded from above for households with positive balances in their retirement accounts, consider the following argument. As long as the household has funds in the illiquid account (i.e. $k_{i,t+1}>0$ as assumed by Lemma 3) one of two things must be true: either (i) $\theta_{i,t}\left(c_{i,t};h_{i,t}\right)<\frac{1}{1-\tau}$ and the household does not make a penalized withdrawal; or (ii) the household makes a penalized withdrawal but remains at an interior solution, thereby satisfying the first-order condition (10) with equality. In the latter case, $\theta_{i,t}\left(c_{i,t};h_{i,t}\right)=\frac{1}{1-\tau}\pi\left(h_{i,t}\right)$, where

$$\pi\left(h_{i,t}\right) \equiv \frac{E_t\left[\frac{\partial V_{t+1}\left(a_{i,t+1},k_{i,t+1};h_{i,t+1}\right)}{\partial k_{i,t+1}}\right]}{E_t\left[\frac{\partial V_{t+1}\left(a_{i,t+1},k_{i,t+1};h_{i,t+1}\right)}{\partial a_{i,t+1}}\right]}$$

is the relative value of an illiquid dollar. As discussed above, this ratio must be weakly less than one.

Generalization of Lemma 1. Lemma 1 considered the case in which households find themselves just before the statutory retirement age. This assumption is convenient because it implies that the marginal value of liquid and illiquid funds tomorrow are identical. In general, however, illiquid funds are less valuable, and thus we are only able to bound the valuation of liquidity at a time of a withdrawal.

Lemma 1b: Withdrawals and Equilibrium Valuation of Liquidity. *If a household withdraws from the illiquid account at time $t<t^*$, then:*

$$\frac{1}{1-\tau} \geq \theta_{i,t}\left(c_{i,t};h_{i,t}\right) \;\; \geq \;\; \left(\frac{1}{1-\tau}\right)\pi\left(h_{i,t}\right),$$

where $\pi\left(h_{i,t}\right) \in [1-\tau, 1]$. *Furthermore, for all* $h_{i,t}$, $\pi\left(h_{i,t}\right) = 1$ *if either* $t = t^{*}$ *or if the perceived probability that a household makes a penalized withdrawal or borrows from the liquid account is zero in the window* $\left(t, t^{*}\right]$.

We have already discussed above that $\pi\left(h_{i,t}\right) \leq 1$ since illiquid dollars are less valuable than liquid ones (unless no withdrawal is made before t^{*}). Next, consider the lower bound of $\pi\left(h_{i,t}\right)$. To prove the lower bound we proceed by contradiction. Assume that the household is maximizing and that $\pi\left(h_{i,t}\right) < (1-\tau)$. Then, build an alternative strategy by withdrawing one dollar from the illiquid account, paying the penalty τ, and transferring $(1-\tau)$ dollars into the liquid account. This deviation generates a total change in the household's value of the problem that is given by:

$$-E_t\left[\frac{\partial V_{t+1}\left(a_{i,t+1}, k_{i,t+1}; h_{i,t+1}\right)}{\partial k_{i,t+1}}\right] + (1-\tau)\, E_t\left[\frac{\partial V_{t+1}\left(a_{i,t+1}, k_{i,t+1}; h_{i,t+1}\right)}{\partial a_{i,t+1}}\right],$$

which would be bigger than 0 if $\pi\left(h_{i,t}\right) < 0$. We have thus found a welfare enhancing deviation and reached a contradiction.

D Appendix Figures

Figure D.1: Prevalence of Retirement Savings Accounts and Penalized Withdrawals

Prevalence of Accounts

(a) By Age of Primary Filer

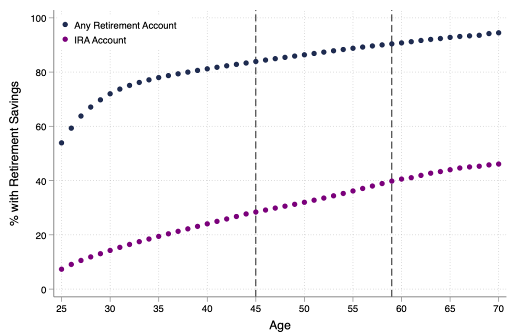

(b) By Household Overall Income

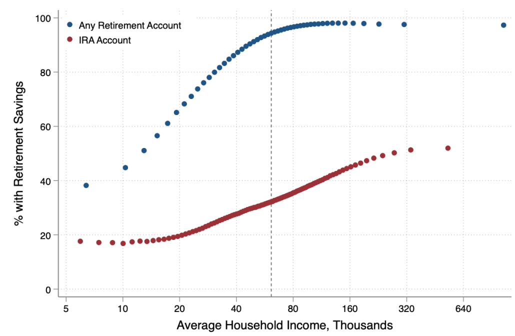

Prevalence of Withdrawals

(c) By Age of Primary Filer

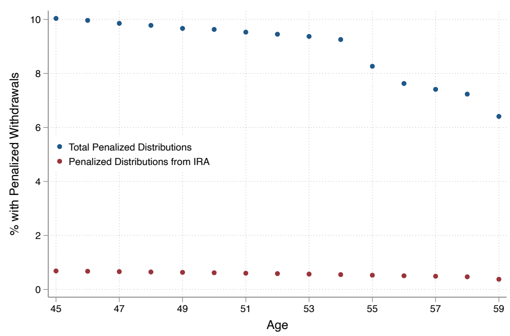

(d) By Household Average Income

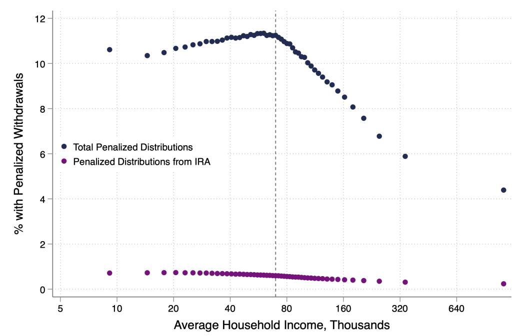

(e) Number of Penalized Withdrawals

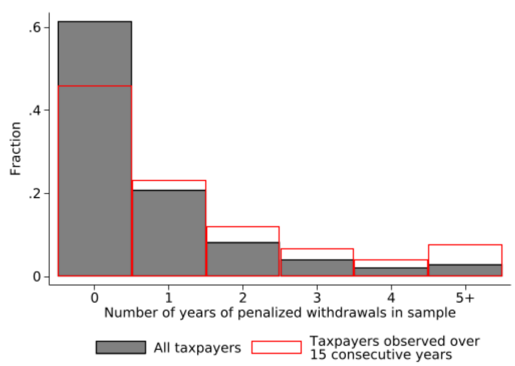

(f) Repeated Penalized Withdrawals

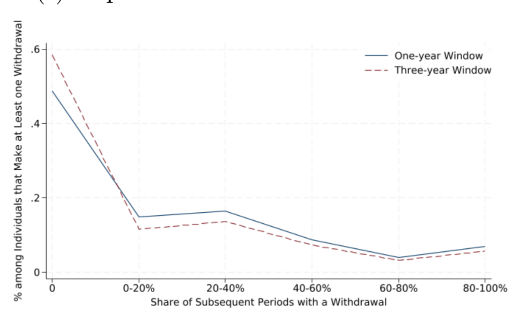

Notes: These figures illustrate the prevalence of retirement accounts and penalized withdrawals. We identify households as having accounts in a given year if up to that year within our sample period of 20 years they have made a contribution to 401(k)/IRA accounts or have balances in IRA accounts. The prevalence of penalized withdrawals is calculated as the share of households that make a penalized withdrawal within the year averaged across all years in our data. We include in the figures information on both any type of account (401(k)/IRA) and IRA accounts only. Panels A-B analyze prevalence of accounts. Panel A plots the share of households with retirement accounts by age. Panel B plots the share of households with retirement accounts by average household income (where the vertical line marks the median value in our sample). Panels C-F analyze the prevalence of penalized withdrawals. Panel C plots the share of households with a penalized withdrawal by age. Panel D shows the distribution of annual withdrawals by household income (where the vertical line marks the median value in our sample). Panel E shows the distribution of the number of years a household has taken a penalized withdrawal. Panel F shows, among households who make a withdrawal in some period, the distribution of subsequent years within our data frame the household made additional withdrawals. We provide two definitions of a withdrawal period as being either one or three years (to allow for a longer period of "consecutive" liquidity needs).

Figure D.2: Statistics on Penalized Withdrawals

(a) CDF of Withdrawals

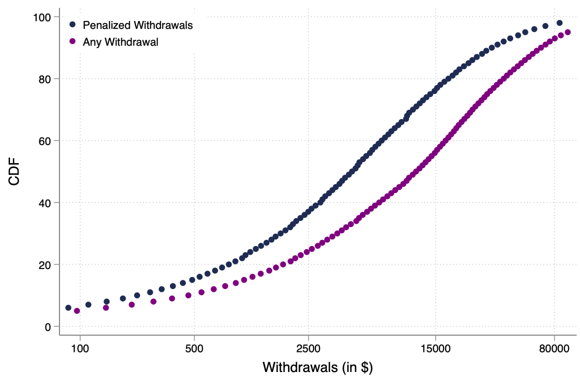

(b) Share of IRA Accounts Fully Depleted

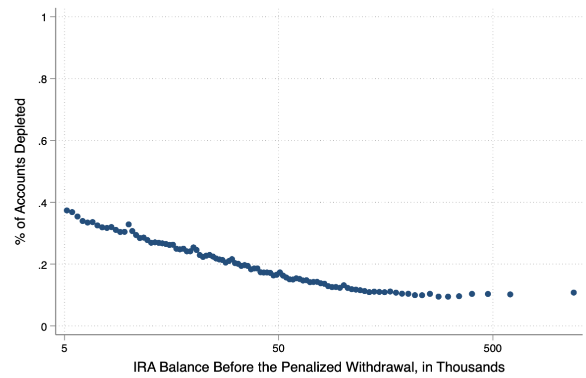

(c) Share of IRA Balances Withdrawn

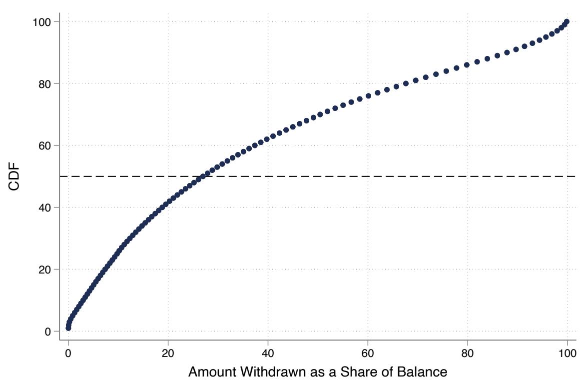

(d) Penalized Withdrawals and Income Changes

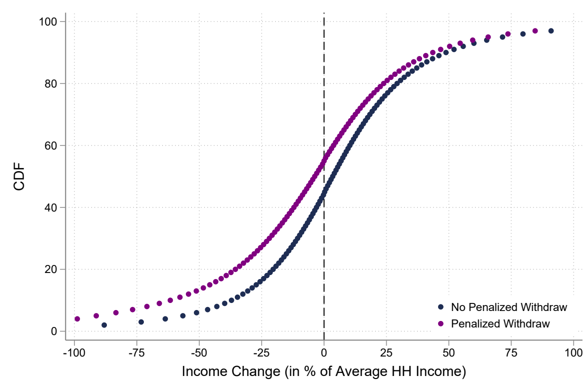

Notes: These figures provide different statistics regarding the behavior of penalized withdrawals. Panel A shows the overall CDF of amounts of penalized withdrawals and compares it with the overall CDF of amounts of withdrawals of any kind. Panels B and C focus on households who have an IRA account and make a penalized withdrawal from such an account. Panel B first computes the share of households who have fully depleted their IRA account after the withdrawal. Panel C then shows, only for households who do not fully deplete their IRA accounts, the CDF of the ratio of the amounts of penalized withdrawals to the previous IRA balances. Panel D plots the CDF of annual income changes, separating households according to whether they made a penalized withdrawal in a given year.

Figure D.3: Prevalence of Accounts: Inclusion of Non-Filers

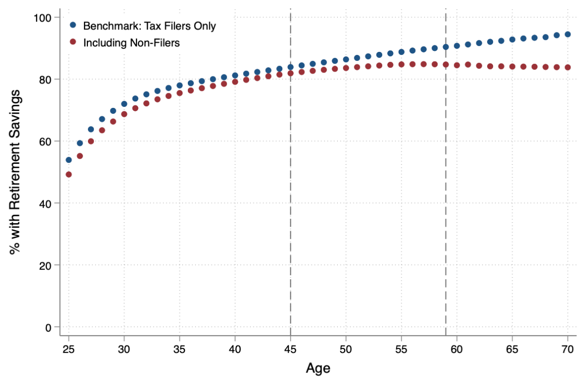

Notes: This figure illustrates the prevalence of retirement accounts by age. We compare our benchmark sample of households that file tax returns to a more inclusive sample, which combines in households that do not file tax returns. The figure shows that households that do not file tax returns are less likely to have a retirement savings account.

Figure D.4: Unemployment Event: Adjusted Gross Income

(a) All Households

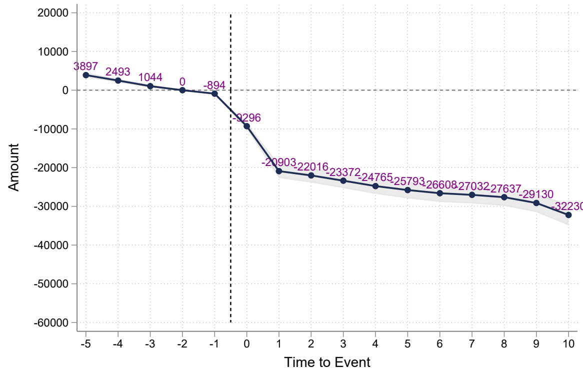

(b) By Race

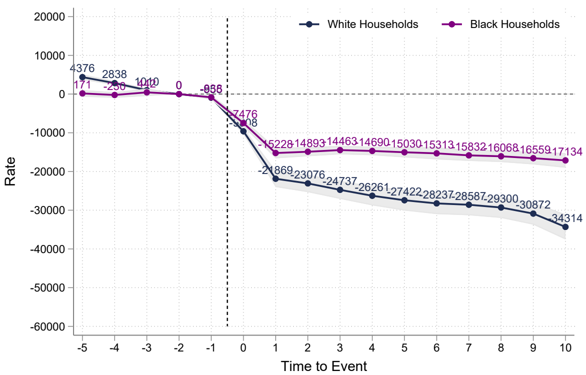

(c) By Race in Percent

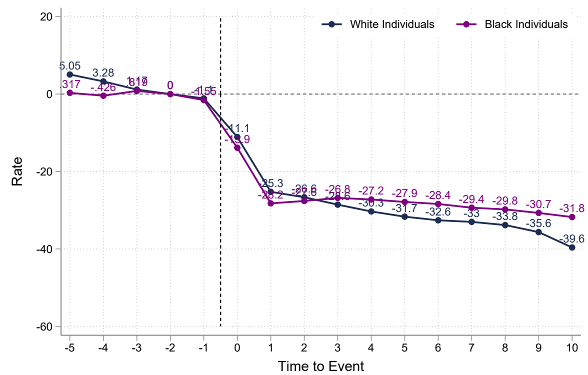

Notes: This figure studies households' adjusted gross income (AGI) around the event of unemployment, defined as the first period we observe at least one of the household members receiving unemployment benefits. Panel A plots the event study coefficients from specification (3) for the entire sample. Panel B plots the event study coefficients from separate specifications of equation (3) for households whose primary filer is Black and for households whose primary filer is White. Panel C plots the coefficients from panel B scaled by the race-specific baseline in period -2.

Figure D.5: Households who Stay in the Same CZ

(a) Unemployment Event

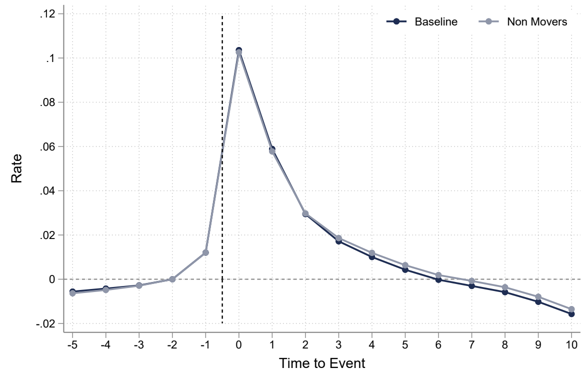

(b) Large Income Loss

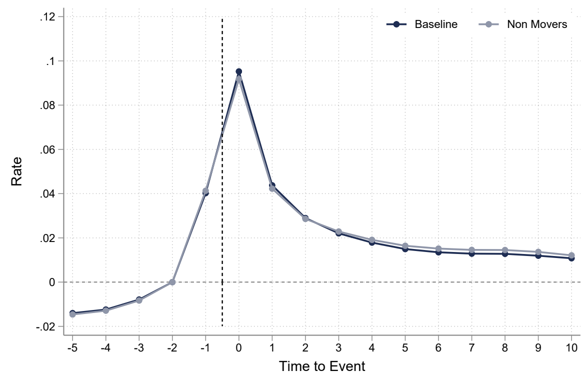

Notes: This figure plots the event study coefficients from specification (3) for the events of unemployment in panel A and large income losses in panel B. We compare the overall sample to a restricted sample in which we include only households that do not change their Commuting Zone around the event. Specifically, we only consider households that are in the same Commuting Zone in periods -1 and 1.

Figure D.6: Unemployment Event: Primary Filers Younger than 55

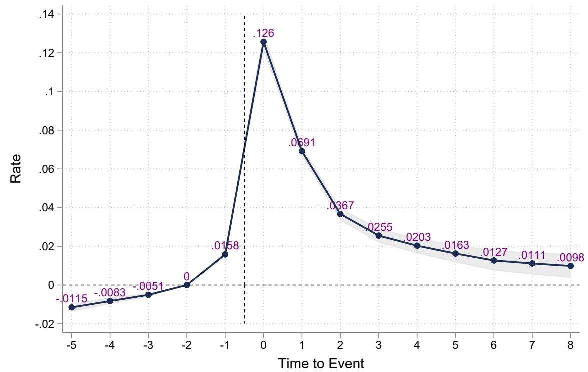

Notes: This figure plots the event study coefficients from specification (3) for the event of unemployment, as defined by the first period we observe at least one of the household members receiving unemployment benefits. We include observations of primary filers younger than 55.

Figure D.7: Event Studies by Amount Withdrawn

Unemployment

(a) Percentage Points

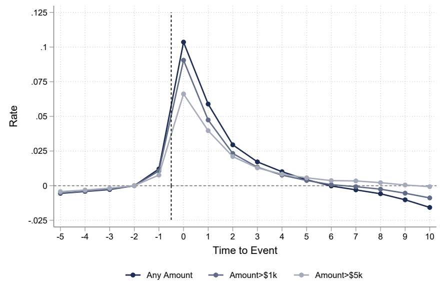

(b) Percent Change

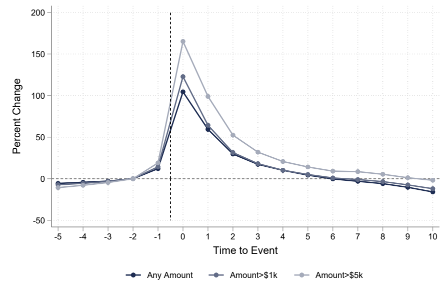

Large Income Losses

(c) Percentage Points

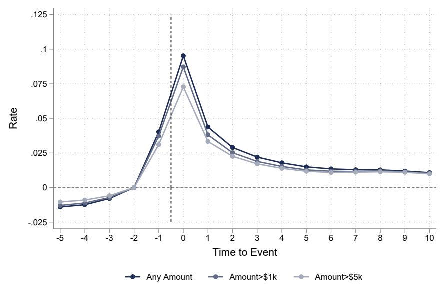

(d) Percent Change

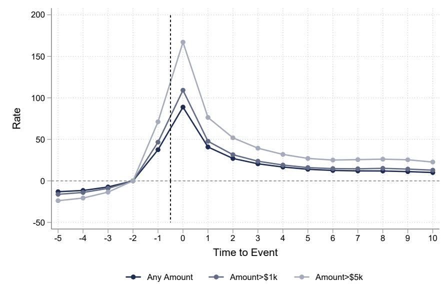

Notes: This figure plots the event study coefficients from specification (3) for the events of unemployment and large income loss. We study indicators for making penalized withdrawals of different amount thresholds: any amount, more than $1,000, and more than $5,000. For each event, the left panel reports estimates in percentage points, and the right panel reports these estimates in percent changes relative to the respective baseline levels in period -2.

Figure D.8: Event Study of Job Switch

(a) Frequency of Withdrawals (b) Withdrawn Amounts

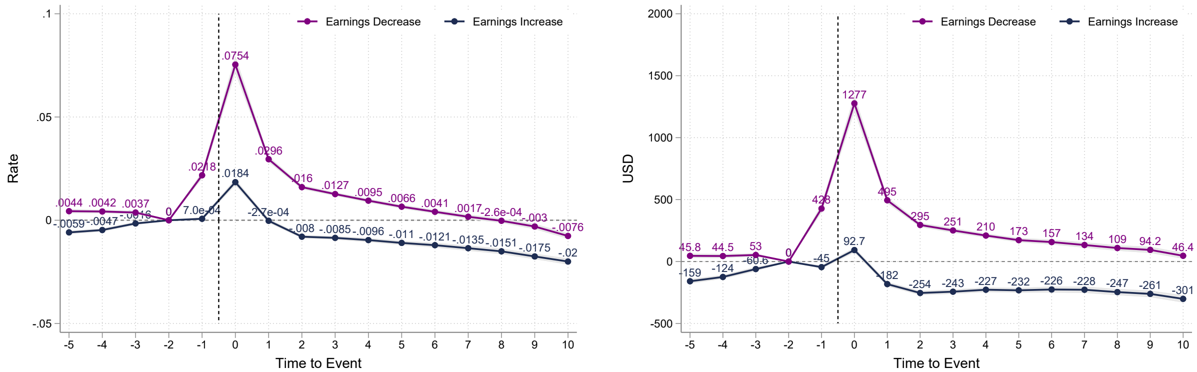

Notes: This figure studies penalized withdrawals around the event of a job switch using specification (3). It focuses on the sample of households for whom we see a change in employer from period $t-1$ to period t without an episode of being on unemployment benefits. We then split households by whether the employee experienced an earnings increase or an earnings decrease upon the switch. Panel A studies frequency of withdrawals, and panel B studies withdrawal amounts.

Figure D.9: Event Study Estimates around the Move Event

(a) Balanced Panel (b) Extended Horizon

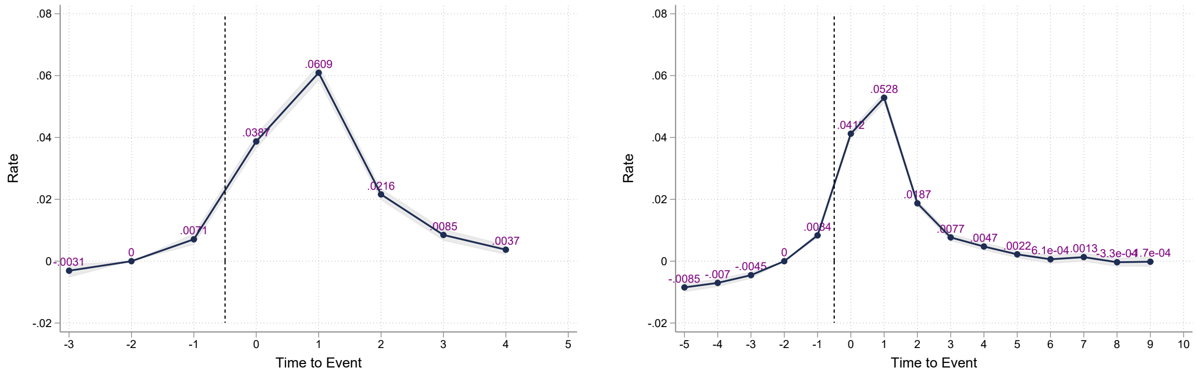

Notes: These figures display estimates for the event study coefficients of a move (β_r) from the estimation of equation (7).

Figure D.10: Movers Design—Attrition and Return Moves

(a) Dynamics

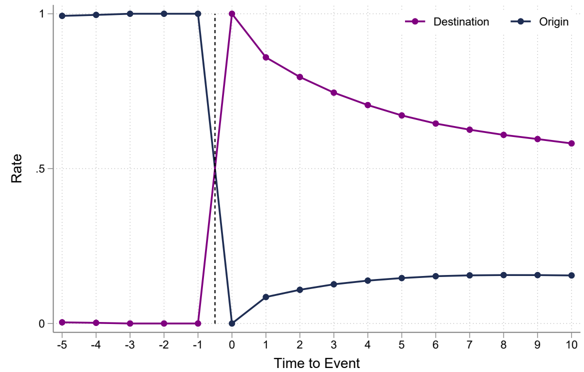

(b) Passthrough Scaled by Movers Still in Destination

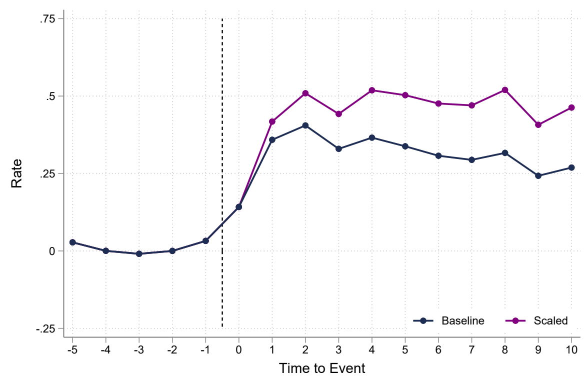

Notes: This figure provides additional analyses for the movers design. Panel A displays indicators for a household's geographic location around the move. In the movers design, we assign a household the same destination location for the entire post-move period. In this figure, we display indicator variables for whether, in a given period, the household remains in the assigned destination unit and whether the household returns to the assigned origin unit. Panel B scales the estimates for the movers analysis from panel C of Figure 7 by the share of movers still at the assigned destination.

Figure D.11: Movers Analysis–Extended Horizon

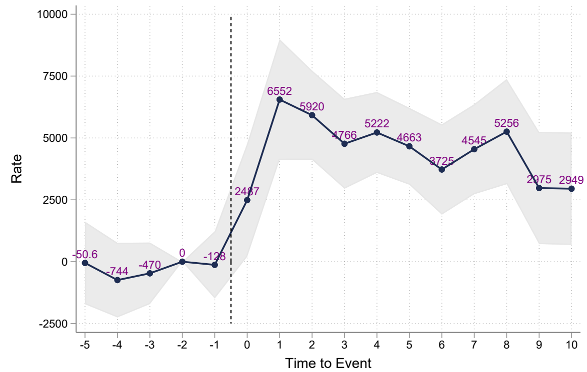

Notes: This figure displays estimates for the share of spatial differentials in withdrawals that can be attributed to location, using the movers design specification of equation (7). We show the estimates from an unbalanced panel of households on an extended time window that spans the years [-5, +10] around the move, where the outcome analyzed is amounts of penalized withdrawals.

Figure D.12: Correlations with Location Fixed Effects

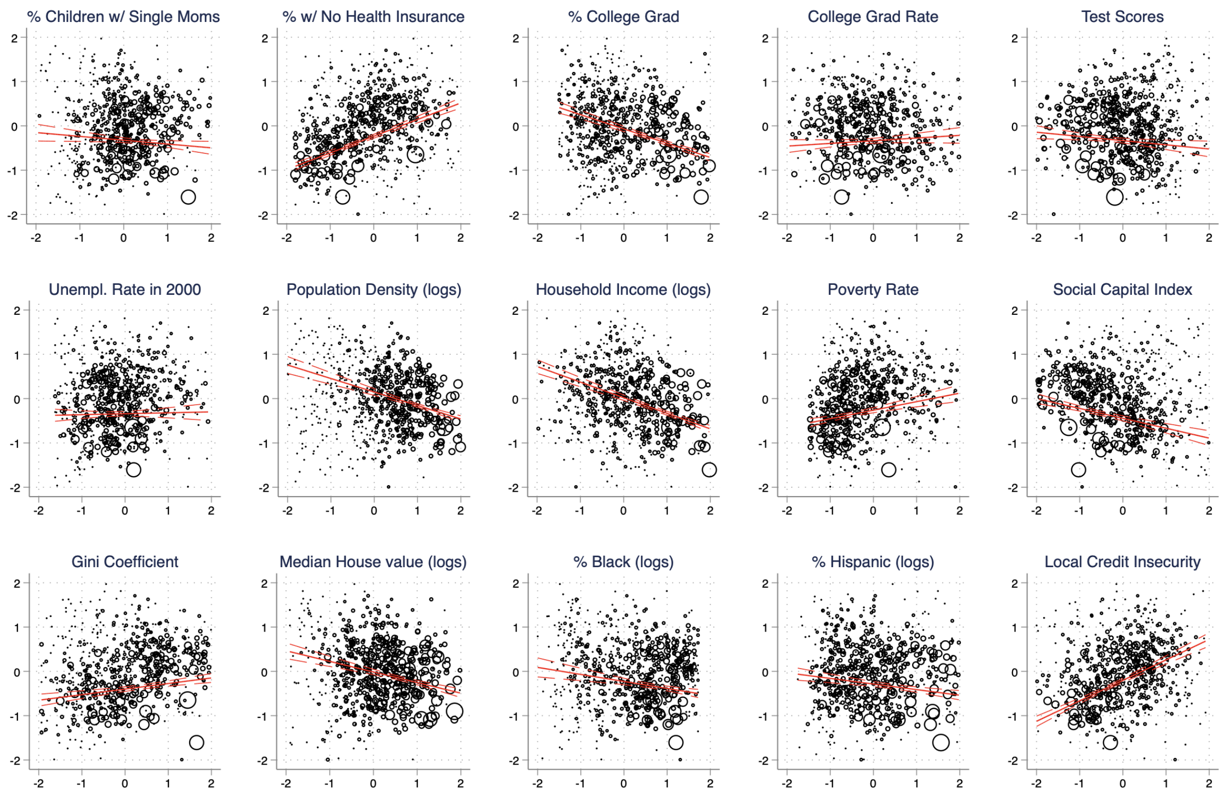

Notes: These figures display correlations of the location fixed effects, Γ_z, as estimated using equation (5), with CZ-level social and economic characteristics.

Figure D.13: Correlations with Households Fixed Effects

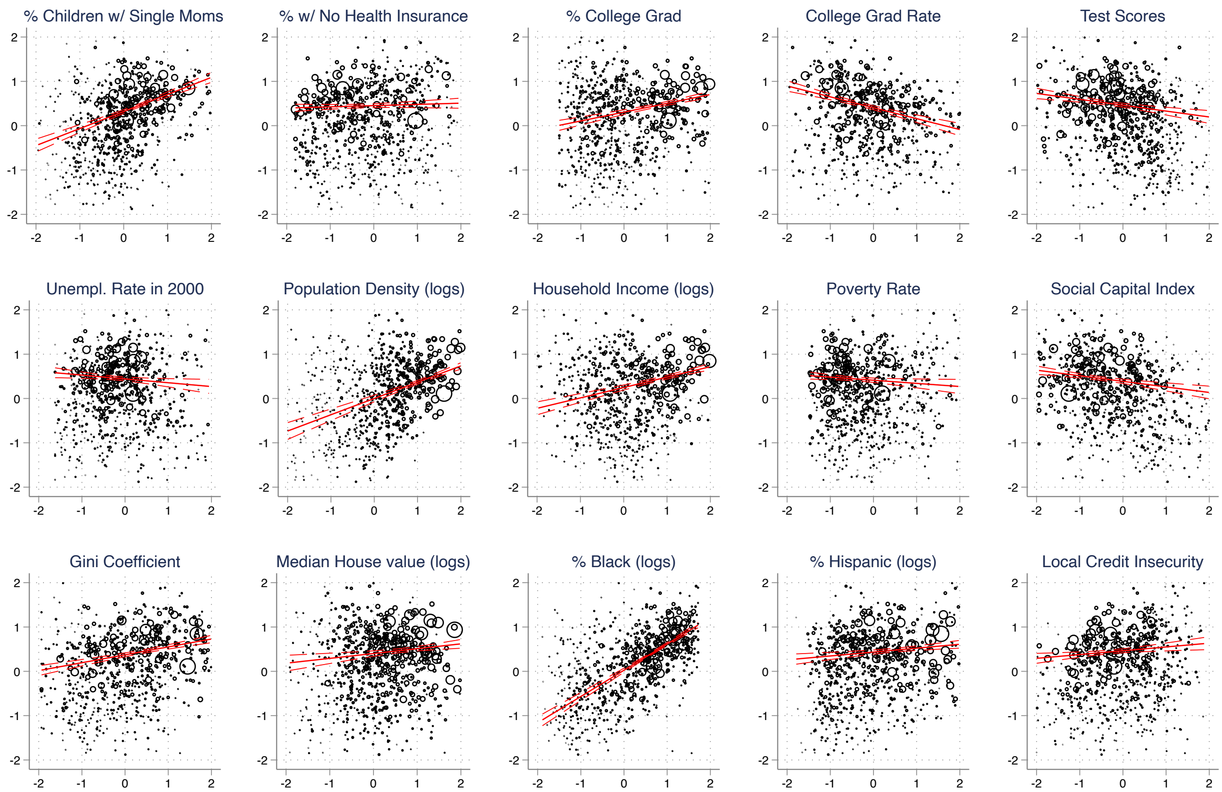

Notes: These figures display correlations of the household fixed effects, α_i, as estimated using equation (5) and collapsed at the CZ level, with CZ-level social and economic characteristics.

Figure D.14: CDFs of Share of Black Households by Commuting Zones and ZIP Codes

(a) Raw Distributions

(b) ZIP Code Means Relative to CZ Means

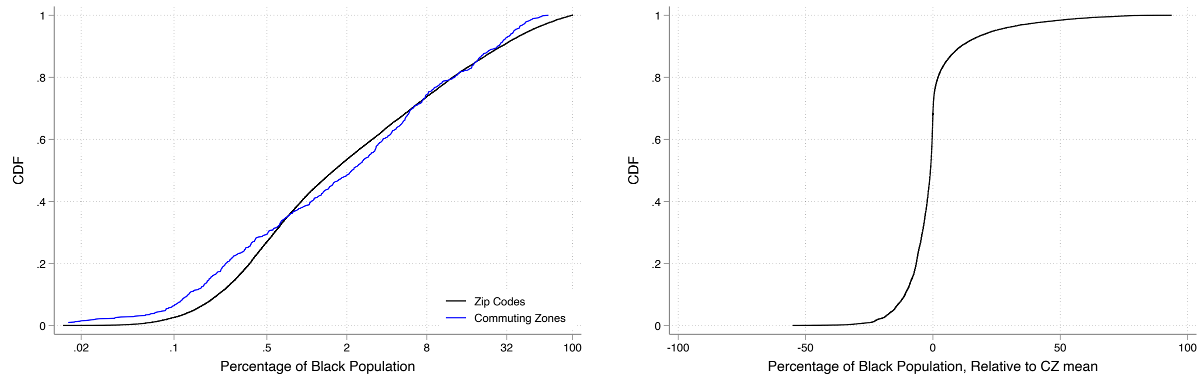

Notes: These figures display cumulative density functions (CDFs) for the share of Black households. Panel A displays CDFs across Commuting Zones (CZs) and across 5-digit ZIP Codes, and panel B displays the CDF across 5-digit ZIP Codes relative to the Commuting Zone means.

Figure D.15: Event Study: Large Income Losses (Primary Filers of Ages 55-59)

(a) Frequency of Withdrawals (b) Withdrawn Amounts

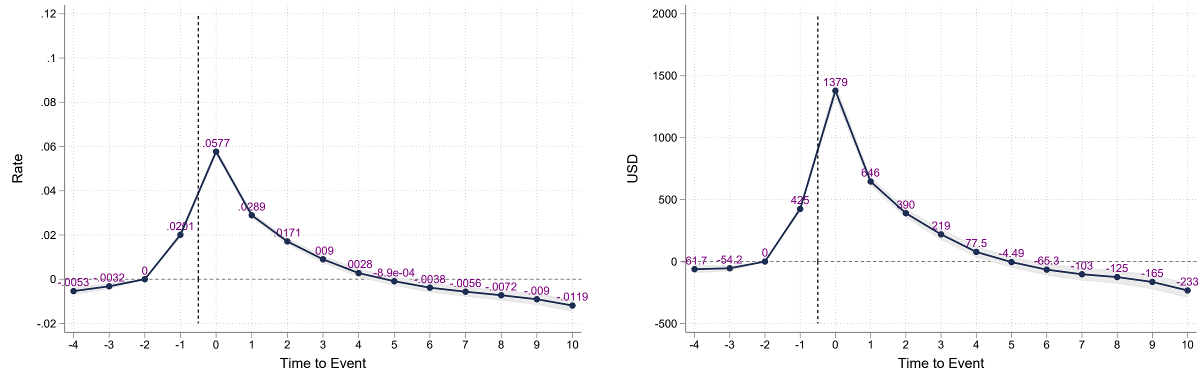

Figure D.16: Movers Analysis (Primary Filers of Ages 55-59)

(a) Withdrawal Frequency: Baseline (b) Withdrawal Frequency: Flexible Controls

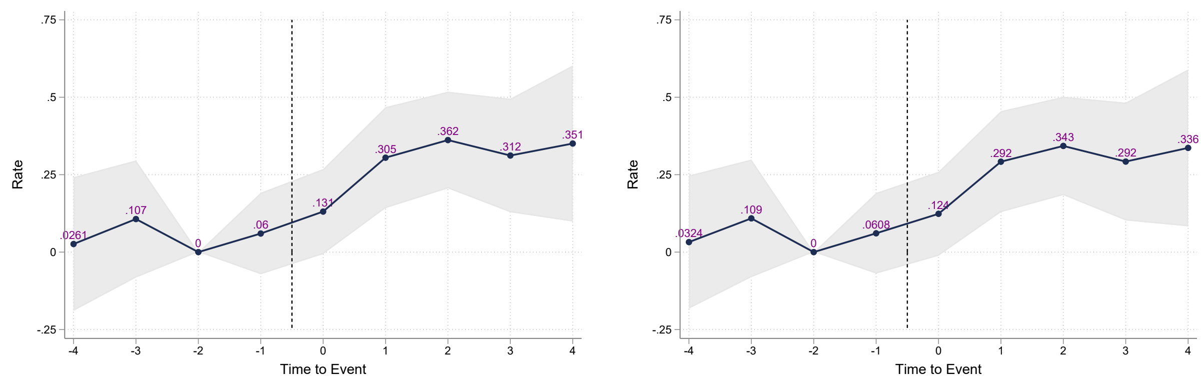

(c) Withdrawn Amounts: Baseline (d) Withdrawn Amounts: Flexible Controls

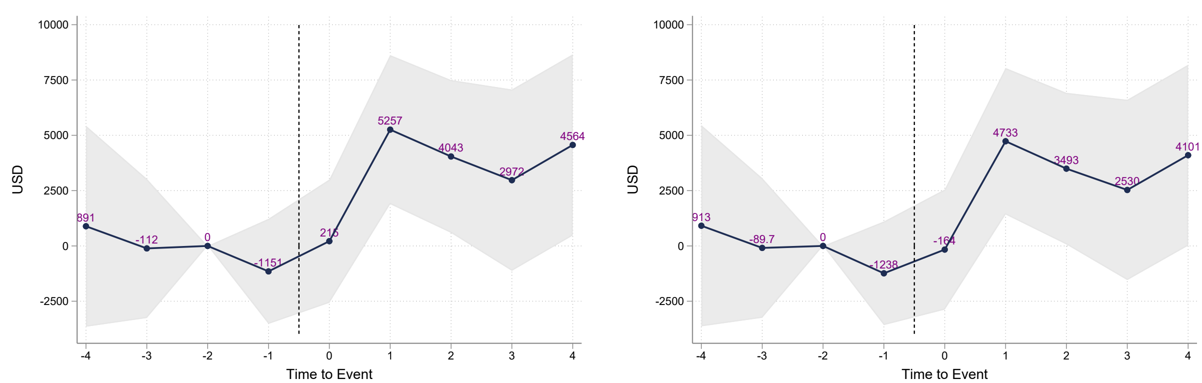

Notes: In these figures, we repeat the main analysis that pertains to households with a primary filer of ages 45-59 but when we constrain the sample to the age range 55-59 to focus on households near the statutory age of 59.5 when withdrawals become non-penalized. We provide the event study of a large income loss and the movers analysis. Note that we do not include the event study of an unemployment event, since withdrawals are already non-penalized for individuals over 55 who separate from their employer.

Figure D.17: Withdrawal Behavior During the Great Recession (Primary Filers of Ages 55-59)

(a) Frequency of Withdrawals

(b) Withdrawn Amounts

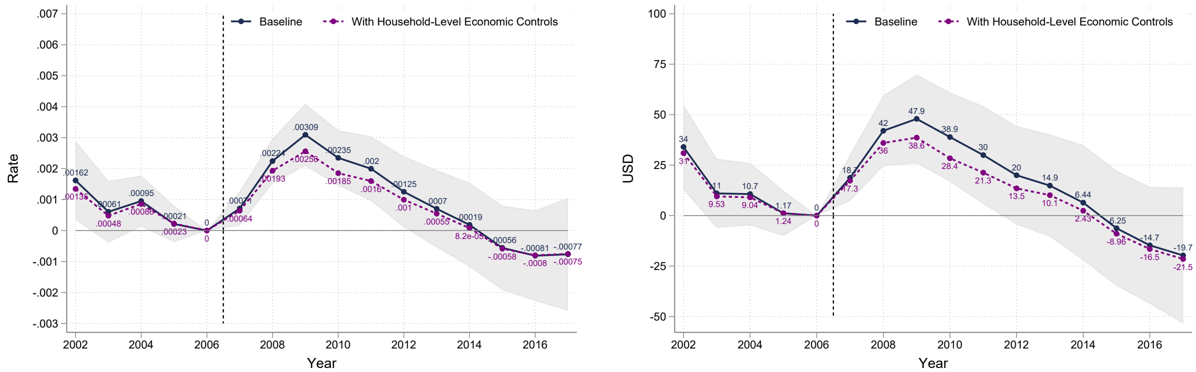

Notes: In this figure, we repeat the main analysis of the withdrawal behavior during the Great Recessions, when we constrain the sample to the age range 55-59 to focus on households near the statutory age of 59.5 when withdrawals become non-penalized.