# Penalized Withdrawals and Households' Valuation of Liquidity*

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

January 2026

**Abstract**

We introduce penalized early withdrawals from retirement savings accounts as a revealed-preference tool to infer households' valuation of liquidity. The approach leverages a well-defined price of liquidity—the statutory early-withdrawal penalty—and allows us to bound marginal liquidity valuations without imposing structure on preferences or requiring direct data on credit access. Using US administrative tax records from 1999-2018, we document three main findings. First, persistent geographic differences account for over 30 percent of the nationwide differentials in liquidity valuation across labor markets, pointing to an important role for local credit conditions. Second, during the Great Recession, liquidity valuations rose sharply in harder-hit areas, with spillovers from local credit-market tightening explaining roughly two-thirds of the overall effect. Third, Black households rely more heavily on penalized withdrawals than otherwise similar White households, even conditional on income and location, consistent with systematically more limited access to lower-cost liquidity. Together, these results show that penalized withdrawals provide a scalable and informative lens on household liquidity constraints, revealing both the equilibrium determinants of liquidity valuation and persistent disparities in access to credit.

*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.

†US 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. Deviations from this benchmark would reflect only aggregate shocks or permanent changes in income or consumption. In reality, however, households face limited access to liquidity from both formal and informal credit sources (Parker 1999; Johnson et al. 2006), resulting in the partial insurance of shocks documented by Blundell, Pistaferri, and Preston (2008). Consequently, the extent to which marginal utility today exceeds expected marginal utility tomorrow—the "valuation of liquidity"—can vary substantially across households. Such variation implies differences in the marginal value of resources across agents and points to potential welfare gains from reallocating liquidity toward households with higher valuations.

Measuring differences in households' valuation of liquidity is empirically challenging for two fundamental reasons. First, the utility-based valuation of liquidity is not directly observable, even in settings with rich consumption data.[1] Linking observed consumption fluctuations to marginal utility requires strong assumptions about preference stability, which are particularly difficult to justify in environments characterized by major life events—such as health shocks—that plausibly alter households' circumstances and preferences.[2]

Second, the valuation of liquidity is inherently an equilibrium object. A household's willingness to pay for an additional dollar of liquidity depends not only on its own circumstances, but also on the prevailing credit conditions. As a result, consumption or income shocks alone provide an incomplete picture of liquidity valuation, as they do not capture contemporaneous variation in credit supply. Fully capturing the valuation of liquidity would therefore require not only information on consumption and income shocks, but also granular, time-varying data on the credit a household could access—data that, to our knowledge, are unavailable at a national scale.

This paper introduces a revealed-preference approach that sidesteps these challenges by using observed prices and choices to bound households' valuation of liquidity. The core idea is simple: when a household chooses to incur a known cost to access liquidity, it reveals that its marginal valuation of liquidity weakly exceeds that cost; conversely, when a household with access to that option chooses not to incur the cost, its valuation must lie weakly below it. Taken together, withdrawals and non-withdrawals bracket households' valuation of liquidity

[1] Recent advances include methods that infer consumption from budget residuals using 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).

relative to a known price. We operationalize this logic using penalized early withdrawals from retirement accounts, which are widely accessible and carry an explicit and observable marginal price—a statutory early-withdrawal penalty of 10 percent. Based on this insight, we use US administrative tax records from 1999-2018 to study households' valuation of liquidity and its variation across time and space.

We begin by developing a simple theoretical framework that links 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. In the model, observing a penalized withdrawal corresponds to situations in which a household's marginal valuation of liquidity exceeds the early-withdrawal penalty, while non-withdrawal indicates that it remains below it. We then characterize the amount of liquidity that would be sufficient to keep a household's marginal valuation of liquidity weakly below the penalty. We refer to this amount as "penalized liquidity," which provides a money-metric measure of under-insurance. The framework also highlights the importance of empirical strategies that account for both household- and market-level determinants, since liquidity valuation is shaped by equilibrium forces.

Motivated by this framework, we organize the empirical analysis in three steps.

In the first step, we examine how household-level factors shape the valuation of liquidity. Using within-household variation over time, we study how life events that increase liquidity needs—such as unemployment or large income declines—affect the likelihood and intensity of penalized withdrawals.[3] We find that adverse shocks are followed by substantial increases in reliance on penalized withdrawals. For example, unemployment is associated with a 10.4 percentage point (pp) increase in the probability of withdrawing and approximately $1,600 in additional withdrawn funds. Even households with substantial wealth exhibit meaningful, though attenuated, responses, consistent with the wealthy hand-to-mouth framework (Kaplan et al. 2014). Using 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 manage liquidity needs. Following unemployment, take-up among Black households increases by about 35 percent more than among White households, despite experiencing lower earnings reductions upon the event.

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

[3] Some prior work studied leakages from retirement accounts following household shocks (e.g., Goodman et al. 2021), but our approach focuses on penalized withdrawals as the object that reveals the valuation of liquidity. We also complement studies that examine withdrawal behavior around age 59.5, when the early-withdrawal penalty disappears (e.g., Goda et al. 2018; Rong 2023; Stuart and Bryant 2024).

of stress in credit markets.[4] We find that households in more affected areas exhibit larger increases in penalized withdrawals than households in less affected areas. Decomposing the effect, we show that roughly two-thirds reflects channels beyond direct household-level shocks to income and employment, consistent with tightening local credit conditions.

In the third and final step, we bring together household- and market-level factors in a unified empirical framework. Using a mover design within a two-way fixed effects model, we estimate the relative contribution of each set of factors to variation in liquidity valuation. We find that local place effects account for roughly one-third of the spatial differences in penalized withdrawals. Leveraging the timing and nature of moves, we assess and rule out several alternative explanations—including changes in households' economic conditions, tax optimization, and learning—as primary drivers of this result. We then estimate and interpret both household and location fixed effects. Location effects are strongly correlated with proxies for local credit conditions, such as measures of credit insecurity and median home values, which may capture access to collateral. Household effects, in contrast, are closely related to race: even after controlling for geography and income, households with a Black primary earner are 30 percent (2.9 percentage points) more likely than households with a White primary earner to rely on penalized withdrawals.

Taken together, our findings establish penalized withdrawals as a previously under-utilized empirical window into households' valuation of liquidity in the United States. We show that local credit conditions play a central role in shaping these valuations, highlighting the importance of place in determining access to liquidity beyond what is captured by income alone. We also document persistent racial differences in reliance on penalized withdrawals that are not explained by income or location. Finally, by quantifying when and where households are willing—or unwilling—to pay to convert illiquid wealth into liquid resources, our results provide new empirical moments that can inform models of household behavior under liquidity constraints with a minimal set of standard primitive assumptions.

**Related Literature.** This paper relates to several strands of work in public finance and macroeconomics that study household insurance, liquidity constraints, and consumption smoothing. A central theme in this literature is that households face imperfect insurance against idiosyncratic shocks due to capital market imperfections, which limits their ability to smooth marginal utility over time (e.g., Zeldes 1989; Parker 1999; Souleles 1999; Johnson et al. 2006; Card et al. 2007; Blundell et al. 2008; Chetty and Finkelstein 2013). More recent work emphasizes that the liquidity composition of wealth, rather than total wealth alone,

[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.

plays a key role in shaping household behavior and policy transmission, both in theory and in quasi-experimental settings (e.g., Cui and Sterk 2021; Kreiner et al. 2019). This paper contributes to this broad literature in three distinct ways.

*First*, we introduce a revealed-preference approach to studying households' valuation of liquidity. A long-standing challenge in prior work is mapping observed behavior to marginal utilities in the presence of heterogeneity and potential state dependence in preferences (see, e.g., Landais and Spinnewijn 2021). A second challenge is that liquidity valuation is an equilibrium object, shaped not only by households' needs but also by the credit conditions they face—information that is rarely observed at scale. We address both challenges by focusing on a well-defined and widely observed household margin: the decision to make a penalized early withdrawal from a retirement account. Within our framework, this choice provides transparent bounds on the valuation of liquidity, inferred directly from observed prices and choices rather than indirectly via consumption or income dynamics.

*Second*, we provide a comprehensive empirical characterization of the valuation of liquidity across US households and locations. In doing so, we contribute to a growing literature that emphasizes the role of place in shaping economic outcomes, including education, earnings, health, and mobility (e.g., Chetty and Hendren 2018a,b; Finkelstein et al. 2016, 2021; Card et al. 2025). We show that place effects account for roughly one-third of the spatial differences in households' valuation of liquidity, demonstrating that access to liquidity is shaped by local conditions over and above income differences.[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). Even after accounting for income and geography, households with a Black primary earner rely on penalized withdrawals at substantially higher rates, consistent with persistent differences in access to affordable credit.[6]

*Third*, we deliver new empirical moments that discipline how economic shocks translate into households' valuation of liquidity. These moments are directly relevant for quantitative heterogeneous-agent models that study insurance, liquidity constraints, and policy transmission (e.g., Krueger et al. 2016; Kaplan et al. 2018; Auclert 2019; Auclert et al. 2020; Laibson et al. 2021). By focusing on valuations revealed through costly liquidity choices—rather than on income or consumption responses—our moments are robust to preference heterogeneity and state dependence, and therefore provide complementary calibration targets in environments where standard Euler-equation-based approaches may be misspecified (e.g., Kaplan

[5] Keys et al. (2020) document substantial geographic variation in financial distress (e.g., collections, defaults, and bankruptcy), shedding light on mechanisms that may underlie place effects in liquidity valuation.

[6] Relatedly, Ganong et al. (2020) document higher consumption elasticities to income shocks for Black and Hispanic households, consistent with differences in liquidity access and smoothing capacity.

et al. 2014; Parker 2017; Aguiar et al. 2020).

**Structure of the paper.** Section 2 describes the institutional features of penalized early withdrawals, introduces the administrative tax data, and presents motivating descriptive evidence. Section 3 develops a conceptual framework that links penalized withdrawals to households' valuation of liquidity and clarifies the economic interpretation of our empirical objects. The empirical analysis proceeds in three parts. Section 4 studies how household-level shocks affect the valuation of liquidity, exploiting within-household variation over time. Section 5 examines how market-level shocks shape liquidity valuation by exploiting spatial variation in exposure to the Great Recession. Section 6 then brings household- and market-level forces together using fixed-effects and mover designs to quantify the relative importance of place and household characteristics in determining liquidity valuation. Section 7 discusses the conceptual foundations of our empirical approach and its broader relevance and policy implications. Section 8 concludes.

# 2 Background, Data, and Motivating Facts

We describe the institutional details governing penalized early withdrawals from retirement accounts and how we measure them in administrative data, and then present baseline facts that motivate their use as a revealed-preference object.

## 2.1 Institutional Setting

A variety of savings vehicles in the US impose restrictions on early access to funds, most prominently employer-sponsored 401(k) plans and 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, creating a well-defined and observable cost of accessing liquidity prior to the statutory age threshold.

Some early withdrawals are exempt from the penalty, including rollovers, permanent disability, death, qualified education expenses, large medical bills, or first-time home purchases.[7] In our analysis, we restrict attention to distributions that are explicitly coded as penalized and are not associated with these statutory exceptions.

Penalized withdrawals are one of several tools households may use when facing short-term liquidity needs. Survey evidence from Lusardi et al. (2011) shows that 11 percent of households report they would draw on retirement savings—even if doing so requires paying a

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

penalty—when faced with a $2,000 emergency, ranking above several high-cost alternatives.[8]

## 2.2 Data

We provide a brief description of our data here. Additional details are provided in Appendix A.

**Main data sources and sample construction.** We use US administrative tax records from 1999 to 2018, based on a 10 percent random sample of Social Security Numbers (SSNs) for computational reasons given the large data. We link each SSN to individual income 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, including Form W-2 (wage income), Form 1099-R (retirement distributions), and Form 5498 (retirement account balances and contributions).

Our main analysis focuses on households with a primary filer aged 45-59 for whom we observe evidence of holding a retirement account. We classify a household as holding a retirement account if contributions to a 401(k) or IRA or positive IRA balances are reported in Form W-2 or Form 5498. This definition yields a core sample of approximately 10.5 million households.

**Variable definitions.** Our key outcome variable is penalized withdrawals from retirement savings accounts, defined as 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 using Form 1099-R, with Box 1 capturing the distribution amount and Box 7 indicating the type of distribution based on its code. Several Box 7 codes correspond to penalized withdrawals.

In some cases, distributions coded as penalized may ultimately qualify for statutory exceptions. Financial plan administrators may lack complete information about the reason for a withdrawal and may therefore assign a penalty code by default. Taxpayers can subsequently claim eligible exceptions (e.g., large medical expenses or higher education costs) using Form 5329. We use information from Form 5329 to adjust our classification and ensure that our measure captures only withdrawals subject to the 10 percent penalty. Appendix A provides full details on the classification rules and corrections.

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

[8] The different tools and the share of households who expect to use each tool (given in parentheses) are: savings (52.4), family (29.6), work more (22.9), credit cards (20.9), sell possessions (18.8), liquidate retirement investments even if a penalty is required (11.1), pawn assets (7.7), friends (7.4), unsecured loan (7.1), home equity line of credit (HELOC) or second mortgage (4.3).

[9] Specifically, when spouses file separately, we combine their records to construct a single household return comparable to those 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 identify job separations or switches using employer identifiers (EINs) reported on Form W-2. We use Form 1099-G to identify unemployment events and Schedule D of Form 1040 to measure capital income.

We obtain information on IRA balances from Form 5498, which reports the fair market value of all IRA holdings at year-end (Box 5), as well as information on IRA contributions and rollovers.

We determine household location using mailing address information reported annually on Form 1040, allowing us to track households over time and study geographic variation in economic conditions and withdrawal behavior.

Finally, we incorporate administrative imputations of race and Hispanic origin using the methodology of Fisher (2023). Households are assigned to a race or 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 conduct selected robustness checks, we also use data from the Health and Retirement Study (HRS), a nationally representative longitudinal survey of US households. We use HRS waves 7-14 (2004-2018), focusing on households aged 45-59, to compare patterns of retirement account ownership and withdrawal behavior with those observed in the administrative data.

## 2.3 Baseline Facts on Penalized Withdrawals

We document a set of empirical facts describing how US households use penalized early withdrawals from retirement savings accounts. Appendix B provides the full technical details of the analysis we summarize here.[11]

**Prevalence of retirement accounts.** Among households with a primary filer aged 45-59, nearly 90 percent hold at least one retirement account during the sample period. Using data from the Health and Retirement Study (HRS), we corroborate this estimate: 84.1 percent of households in the same age range report holding a defined-contribution account. The high prevalence reflects the fact that our analysis is done at the household level as the economic decision unit, following the rich work on consumption, insurance, and risk sharing that establishes this perspective (see, e.g., Blundell, Pistaferri, and Saporta-Eksten 2016,

[10]See Cronin et al. (2023) and Costello et al. (2024) for validation exercises and applications.

[11]In the same appendix, and in Section 7, we discuss potential behavioral interpretations of penalized withdrawals and how such interpretations would affect our reading of the empirical patterns.

Pruitt and Turner 2020, and Fadlon and Nielsen 2021).[12]

**Frequency of withdrawals.** Penalized withdrawals are widely used across households throughout the age and income distributions, with roughly 10 percent of households making a penalized withdrawal in a given year. At the same time, most households withdraw only occasionally rather than repeatedly across years, indicating episodic use consistent with responses to short-term liquidity needs.

**Size of withdrawals.** The median withdrawal amount is approximately $5,000. Most withdrawals involve partial rather than full liquidation of retirement balances, consistent with interior choices along the withdrawal margin.

**Link to income shocks.** Penalized withdrawals are strongly associated with income declines. Among withdrawing households, nearly 60 percent experience a year-over-year income loss, and they are more than twice as likely as non-withdrawing households to experience a decline of 50 percent or more.

These descriptive patterns establish some necessary conditions for our revealed-preference framework that we develop in the next section: penalized withdrawals are episodic, sizable, and systematically linked to adverse economic conditions.

# 3 Conceptual Framework

We develop a simple framework that serves two purposes. First, it formalizes how penalized withdrawals from retirement accounts reveal households' valuation of liquidity, making explicit the assumptions under which observed withdrawal behavior maps into marginal valuations. Second, it clarifies why liquidity valuation is an equilibrium object, shaped jointly by households' liquidity needs and the credit conditions they face, thereby motivating our empirical designs.

## 3.1 Model Setup

We consider a household $i$ in region $z$ that makes consumption and portfolio decisions over the life cycle. In each period $t$, the household receives income $y_{i,t}$ and chooses consumption

[12] Because our analysis is based on tax filers, it excludes the lowest-income non-filers, who are less likely to hold retirement accounts. When incorporating non-filers, overall prevalence declines somewhat to an average of 83.8 percent (Appendix Figure D.3).

$c_{i,t}$, liquid saving or borrowing, and net contributions to (or withdrawals from) an illiquid retirement account. To reflect the typical institutional setting, a fixed share $\varphi$ of earnings is automatically deposited into retirement savings.

Each period, the household faces an expenditure shock $\varepsilon_{i,t}$ drawn from $F(\varepsilon)$, capturing unexpected liquidity needs that reduce resources available for discretionary consumption (e.g., medical expenses or earnings interruptions following health shocks; Dobkin et al., 2018; Fadlon and Nielsen, 2021).

To fund consumption beyond current income, the household may borrow at an interest cost equal to the risk-free rate $r$ plus a premium $\rho_{i,z}(b_{i,t})$, which depends on household $i$ and region $z$ and increases with the amount borrowed $b_{i,t}$. We interpret $\rho_{i,z}(b_{i,t})$ as a reduced-form shadow cost of borrowing that captures both credit constraints and effective access to all available borrowing opportunities.

Households may also withdraw from retirement savings, but withdrawals made before the statutory retirement age $t^*$ incur a marginal penalty rate $\tau$, so that only $1-\tau$ dollars are available for consumption per dollar withdrawn. Prior to $t^*$, retirement savings are therefore imperfectly liquid at a known marginal cost.

We denote balances in liquid and illiquid accounts at the start of period $t$ by $a_{i,t}$ and $k_{i,t}$, respectively, and let $\Delta a_{i,t}$ and $\Delta k_{i,t}$ denote net flows during period $t$ (so withdrawals from the illiquid account correspond to $\Delta k_{i,t} < 0$). If the household begins the period with zero or negative liquid assets, any further reduction in liquid wealth must be financed through borrowing, 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 $h_{i,t}$ that may include histories of shocks and factors affecting consumption preferences, such as marital status, fertility, and health. Allowing for state-dependent preferences ensures that liquidity demand can respond directly to changes in circumstances rather than only through income fluctuations. Importantly, the flexibility of our approach does not require mapping observed consumption choices into preferences in order to recover households' valuation of liquidity.

We let $V_t\left(a_{i,t-1}, k_{i,t-1}; h_{i,t}\right)$ be the value of the problem, so that

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

subject to

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

where $\beta$ is the discount factor. Notably, the value function is indexed by both time and the household state vector $h_{i,t}$, allowing the problem's value to vary across households and periods even conditional on 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, which captures the marginal rate at which a household trades off a liquid dollar today for a liquid dollar tomorrow. We denote this valuation by $\theta_{i,t}$. Intuitively, $\theta_{i,t}$ represents the marginal shadow price associated with smoothing consumption over time. Penalized withdrawals then serve as a diagnostic device: observing a household incur the statutory penalty reveals that the cost of alternative sources of liquidity exceeds this threshold. Importantly, this interpretation is independent of the household's broader portfolio optimization problem. It allows us to identify a key moment of the distribution of liquidity valuations without requiring that penalized withdrawals be the household's first—or even primary—response to a liquidity shock.

**Definition 1: Equilibrium Valuation of Liquidity.** *The equilibrium valuation of liquidity for household $i$ at time $t$ is defined as*

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

The empirical analysis uses withdrawal behavior to obtain revealed-preference information about $\theta_{i,t}(c_{i,t}; h_{i,t})$. We formalize this link in the following lemmas, which characterize what can be inferred about households' valuation of liquidity from observed penalized withdrawals.

**Lemma 1: Household's Valuation of Liquidity at Withdrawal.** *Consider a household $i$ that makes a penalized withdrawal from its retirement account at time $t = t^* - 1$ while* *retaining a positive balance.*[13]

*The household's valuation of liquidity satisfies*

$$\theta_{i,t} = \frac{1}{1-\tau}, \qquad \text{with } \frac{1}{1-\tau} \geq \frac{1}{1-\rho'_{i,z}(b)}. \tag{1}$$

*Proof.* See Appendix C. □

The intuition is straightforward. At $t = t^* - 1$, the household anticipates unrestricted access to retirement assets in the following period. Choosing to pay the penalty $\tau$ to withdraw today therefore reveals that a liquid dollar today is valued at least $\frac{1}{1-\tau}$ times a liquid dollar tomorrow. Moreover, withdrawing rather than borrowing implies that the marginal cost of liquid credit exceeds the effective cost of accessing illiquid funds.

**Lemma 2: Household's Valuation of Liquidity at Withdrawal Away from Retirement.** *Consider a household $i$ that makes a penalized withdrawal from its retirement account at time $t < t^* - 1$ while retaining a positive balance. Then the household's valuation of liquidity satisfies*

$$\theta_{i,t} = \frac{1}{1-\tau} \Lambda_{i,t}, \tag{2}$$

*where*

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

*captures the marginal value of an illiquid dollar relative to a liquid dollar in the following period and satisfies*

$$1 - \tau \leq \Lambda_{i,t} \leq 1.$$

*Moreover:*

1. $\Lambda_{i,t} = 1 - \tau$ if the household expects that a marginal dollar held in the retirement account at $t + 1$ will be withdrawn before retirement with certainty;

2. $\Lambda_{i,t} = 1$ if the household expects that, from $t + 1$ until retirement, liquidity will not be valuable at the margin—so that an additional dollar carried in liquid form provides no advantage over carrying it in the retirement account.

[13] The assumption that $k_{i,t} > 0$ guarantees that the Euler equation is satisfied with equality. We show in Appendix B that this is the empirically typical case. In the atypical case in which the household fully depletes the retirement account, the implication is $\theta_{i,t} > \frac{1}{1-\tau}$.

*Proof.* See Appendix C. □

Lemma 2 clarifies how Lemma 1 generalizes to earlier periods. A penalized withdrawal at $t < t^* - 1$ typically implies $\theta_{i,t} > 1$, but it need not imply $\theta_{i,t} = \frac{1}{1-\tau}$. The reason is that the withdrawal trades off a liquid dollar today against an illiquid dollar tomorrow, whereas $\theta_{i,t}$ captures the trade-off between liquid dollars across time. The wedge between the two is summarized by $\Lambda_{i,t}$, which reflects expectations about the future usefulness of liquidity. When households do not expect to make additional penalized withdrawals before retirement and do not expect to borrow at a positive marginal cost $\rho'(b) > 0$, illiquid and liquid dollars are locally equivalent and $\Lambda_{i,t} = 1$, so Lemma 2 collapses to Lemma 1. When instead households expect to fully withdraw their retirement balances before retirement, $\Lambda_{i,t} = 1 - \tau$ and, at that point in time, the household's Euler equation is not distorted. The evidence showing households withdraw only episodically and that they typically do not deplete their accounts suggests that the first case is empirically more relevant.

**Lemma 3: Household's Valuation of Liquidity at Non-Withdrawal.** *Consider a household $i$ with positive retirement balances that does not make a penalized withdrawal at time $t$. The household's valuation of liquidity satisfies*

$$\theta_{i,t} \;\leq\; \frac{1}{1-\tau}.$$

*Proof.* See Appendix C. □

Lemma 3 highlights that revealed-preference information comes not only from households that make penalized withdrawals, but also from those that do not. A penalized withdrawal provides access to liquidity at a known marginal cost. When a household with positive retirement balances chooses not to withdraw, it reveals that the marginal value of liquidity does not justify paying this cost. As a result, non-withdrawal places an upper bound on the household's valuation of liquidity. Taken together, withdrawals and non-withdrawals bracket the distribution of liquidity valuations in the population.

At any point in time, we can compute the share of households with given characteristics who make penalized withdrawals. This share is also the probability that an individual household with those characteristics makes such a withdrawal, providing a natural measure of the prevalence of high liquidity valuation within that group. This leads us to the following definition.

**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*

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

*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 capture the intensive margin of self-insurance, we define an empirical measure—penalized liquidity—which quantifies the amount of liquidity households extract through 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.*

Together, the probability of withdrawals, $\mathcal{P}_t(\mathscr{I})$, and the amount of withdrawals, $\Lambda_{t,t'}(\mathscr{I})$, provide empirically tractable revealed-preference moments that allow us to infer information about households' valuation of liquidity. The extensive margin of withdrawal captures the fraction of households with a high valuation of liquidity at a given point in time. The intensive margin captures the amount of liquidity that would need to be provided to a group of households who have access to penalized withdrawals to keep their overall valuation of liquidity bounded. Comparing these objects across groups further allows us to assess systematic differences in exposure to binding liquidity constraints and access to alternative insurance mechanisms.

## 3.3 Liquidity Demand and Supply

The valuation of liquidity is an equilibrium object jointly determined by households' demand for funds and the supply of credit they face. Shocks to either side of the market can therefore raise the equilibrium valuation of liquidity and trigger a penalized withdrawal.

To organize this intuition, we represent liquidity demand and supply as inverse schedules that map quantities of funds into an effective marginal interest rate. For expositional clarity, we focus on the final period before $t^*$, when the marginal cost of accessing retirement savings simplifies to $\frac{1}{1-\tau}$.

A household's (inverse) demand for liquidity specifies the marginal interest rate, $\mathbb{D}_i(\bar{b})$, at which the household is willing to finance an amount $\bar{b}$ of consumption given their available sources of funds.

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

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

*where* $x_{i,t} \equiv (1-\varphi)y_{i,t} + (1+r)a_{i,t-1} - \varepsilon_{i,t}$ *denotes household* $i$*'s cash-on-hand at time* $t$*.*

By definition, $\bar{b}$ is the total amount of funds used to finance consumption, since $c_{i,t} = x_{i,t}+\bar{b}$. The function $\mathbb{D}_i(\bar{b})$ therefore traces the marginal interest rate at which the household is willing to finance additional consumption. Under standard preferences with diminishing marginal utility, this demand is downward sloping.

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

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

The function $\mathbb{S}_{i,z}(\bar{b})$ summarizes the marginal cost of accessing funds of amount $\bar{b}$, combining liquid borrowing and penalized withdrawals from illiquid savings. When borrowing costs rise to the penalty rate, households optimally substitute toward penalized withdrawals.

In equilibrium, the household chooses $\bar{b}$ such that $\mathbb{S}_{i,z}(\bar{b}) = \mathbb{D}_i(\bar{b})$. At this point, the valuation of liquidity satisfies

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

so that $\theta_{i,t} = 1$ if and only if the shadow cost of liquidity equals the risk-free rate.

**Shocks to Demand and Supply of Liquidity and Penalized Withdrawals.** Figure 1 illustrates how shocks to liquidity demand or credit supply affect the equilibrium valuation of liquidity and can trigger penalized withdrawals.

In the upper panels (a)-(c), we illustrate equilibrium liquidity valuation under three environments. Panel (a) depicts perfect credit markets, where a flat supply curve at the risk-free rate fully insures liquidity needs and leaves the Euler equation undistorted. Panel (b) introduces borrowing frictions for a household without retirement savings ($k_{i,t} = 0$): an upward-sloping supply curve driven by the convex borrowing cost $\rho_{i,z}(b)$ raises the equilibrium valuation of liquidity to $\theta_1$. Panel (c) adds access to penalized withdrawals. When illiquid assets are available, households can obtain liquidity at marginal cost $\frac{1+r}{1-\tau}$, lowering the equilibrium valuation to $\theta_2 < \theta_1$ and replacing part of liquid borrowing with penalized withdrawals, measured by $b_{3,b} - b_{3,a}$.

The lower panels (d)-(f) show how shocks to credit supply or liquidity demand trigger penalized withdrawals. Panel (d) depicts a baseline in which the household borrows only from liquid sources. In panel (e), a credit tightening—modeled as an upward shift in $\rho_{i,z}(b)$—raises the marginal cost of borrowing, leading the household to substitute away from liquid credit toward penalized withdrawals while partially stabilizing liquidity valuation. Panel (f) illustrates a demand shock, such as an unexpected decline in income $y_{i,t}$, which increases total borrowing and again induces a penalized withdrawal. In both cases, access to illiquid savings dampens the response of liquidity valuation to shocks.

These scenarios illustrate that penalized withdrawals respond systematically to both demand and supply shocks, with larger withdrawals reflecting larger liquidity needs and a greater residual demand for self-insurance.

**From Model to Data.** We conclude this section by outlining how the framework guides our empirical analysis. Lemmas 1-3 imply that observed withdrawal behavior has a direct revealed-preference interpretation. Penalized withdrawals identify households whose valuation of liquidity meets or exceeds the statutory penalty threshold, while non-withdrawals among households with positive retirement balances reveal that valuations remain below it. Accordingly, changes in withdrawal take-up trace movements in the prevalence of high liquidity valuation, while changes in withdrawn amounts capture the intensity with which households self-insure once the threshold is crossed.

Guided by this interpretation, Sections 4 and 5 study how transitory shocks on each side of the market affect the equilibrium valuation of liquidity. Section 4 focuses on household-level events that raise liquidity demand and examines how these events shift households across the penalty threshold and alter the amount of penalized liquidity extracted. Section 5 studies shocks to the supply of credit, exploiting variation induced by the Great Recession to trace analogous movements in liquidity valuation driven by local credit conditions. Section 6 then turns to the permanent components of liquidity valuation, using a movers design to separate the role of household-specific factors from location-specific factors, inclusive of both local credit markets and broader economic conditions that shape access to liquidity.

# 4 Valuation of Liquidity after Household-Level Events

In this section, we study how adverse household-level economic events affect households' valuation of liquidity. We show that negative shocks raise the share of households whose valuation crosses the statutory penalty threshold, triggering penalized withdrawals, and we quantify the implied increase in the prevalence of high liquidity valuation. These estimates further provide household-level benchmarks that we later compare to the effects of market-level shocks.

**Estimating Equation.** We estimate the following event-study specification:

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

where $y_{i,t}$ denotes the withdrawal outcome of household $i$ at time $t$, $x_{i,t}$ includes age fixed effects for the primary filer and (cyclical) calendar-year fixed effects, and $\alpha_i$ are household fixed effects that absorb all time-invariant household characteristics.[14] Time relative to the event period is denoted by $r(i,t)$, and $I_r = \mathbb{I}\{r(i,t) = r\}$ are relative-time indicators. We use year $r = -2$ as the omitted baseline to allow for potential changes in trends prior to the realization of the event.[15] We estimate dynamic effects over a 16-year window from $r = -5$ to $r = 10$.

**Interpretation through revealed preference.** Our main extensive-margin outcome is an indicator for making a penalized withdrawal. Lemmas 1-3 imply that this choice has a direct revealed-preference interpretation. A household that makes a penalized withdrawal reveals that its valuation of liquidity exceeds the statutory threshold, whereas a household with positive retirement balances that does not withdraw reveals that its valuation remains weakly below it. Accordingly, our event-study coefficients $\beta_r$ can be read as tracing changes in the prevalence of households whose valuation of liquidity exceeds $\frac{1}{1-\tau}$ around the event.

[14]We include all households in the estimation sample to improve identification of non-event coefficients and therefore add to $x_{i,t}$ an indicator for households that never experience the event. In addition, Appendix Figure D.5 replicates the analysis among households that remain in the same Commuting Zone around the event, with nearly identical results.

[15]Because the data are annual, the coefficient at $r = -1$ may reflect both anticipation and partial realization of the event.

## 4.1 Unemployment Event

We define an unemployment event as the first year in which at least one household member receives unemployment benefits. Figure 2 plots the event-study coefficients $\beta_r$ when the outcome is an indicator for making a penalized withdrawal (panel (a)) or the amount withdrawn (panel (b)). As the event approaches, penalized withdrawals increase sharply, with a pronounced spike in the event year. Interpreted through our framework, these patterns imply a substantial increase in households' valuation of liquidity, consistent with incomplete insurance against unemployment shocks.[16]

In terms of magnitudes, panel (a) shows that the share of households with sufficiently high valuation of liquidity to trigger a penalized withdrawal roughly doubles at the onset of unemployment, rising by 10.4 pp relative to a baseline of 9.9 at $r = -2$. Panel (b) shows that households withdraw an additional $1,600 on average in the event year. Interpreted through our model, this amount measures the liquidity households self-provide to partially insure themselves. Equivalently, it provides a benchmark for the amount of outside liquidity that would have been required—at a marginal cost below the penalty—to keep the household's valuation of liquidity bounded by $\frac{1}{1-\tau}$ in the event year.

To benchmark this magnitude, we compare withdrawn amounts to the decline in household income around the event. Penalized withdrawals offset, on average across all households, about 8 percent of the income decline (that is approximately $20,900 at onset; Appendix Figure D.4, panel (a)). Scaling withdrawal amounts by the increase in take-up, we find that households that withdraw extract roughly $19,000 in penalized liquidity in the event year.

**Heterogeneous Effects.** We next examine heterogeneity in withdrawal responses following unemployment. This analysis sheds light on systematic differences in liquidity valuation across households and helps identify groups with persistently higher reliance on costly self-insurance. We focus on withdrawal frequencies, which have a clearer interpretation than amounts, as the latter scale mechanically with income.

Panel (c) of Figure 2 presents 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 make penalized withdrawals

[16]These findings complement a large literature documenting sharp short-run consumption declines and persistent earnings losses following unemployment (see, e.g., Sullivan and Von Wachter 2009; Kolsrud et al. 2018; Schmieder et al. 2018; Ganong and Noel 2019; Gerard and Naritomi 2021). Unlike income- or consumption-based approaches, our framework allows for state-dependent preferences, such as consumption-leisure complementarities that lead to differences in types of consumption needs associated with unemployment. The classic example of this case is the substitution to cooking meals at home while unemployed and the corresponding reduction in time and monetary costs involved in commuting.

following unemployment, consistent with more limited access to alternative, lower-cost liquidity.

Panel (d) examines heterogeneity by capital income, which we use as a proxy for non-housing wealth. Households are grouped into six bins: negative capital income, zero capital income, and four equally-sized bins among those with positive capital income. Withdrawal responses decline with capital income, consistent with lower marginal valuation of liquidity among households with greater financial means as implied by the model. Nonetheless, even households in the top capital-income quartile—averaging nearly $40,000 in capital income—exhibit meaningful increases in withdrawals, consistent with the presence of liquidity constraints among relatively wealthy households (Kaplan et al., 2014).

## 4.2 Income Changes

We next study how changes in household income affect liquidity valuation. A large income loss event is defined as the first year in which total household income falls by more than 30 percent relative to the previous year.

Panels (a) and (b) of Figure 3 show event-study estimates for withdrawal frequency and amounts. Penalized withdrawals increase sharply in the event year, indicating that households require approximately $2,000 in penalized liquidity, on average, to keep marginal valuation within bounds. As in the unemployment analysis, Black households exhibit substantially larger increases in withdrawals than White households following large income losses (panel (e)).

To further characterize these responses, we relate withdrawals to deviations of household income from its longer-run average, distinguishing households by whether a member switches jobs in the event year. We do this because job changes themselves (as displayed in Appendix Figure D.7) could lead to increased take-up.[17]

Panels (c) and (d) of Figure 3 reveal three patterns. First, withdrawal frequency increases sharply with income losses, aligning closely with the model's predictions. Second, responses are strongly asymmetric: withdrawals rise with income losses but remain flat for income gains. The latter rules out tax strategic considerations as primary explanations for withdrawals, as those would have to display a gradient in income gains which change the house-

[17]Job changes could increase withdrawals due to higher liquidity valuation during transition periods or other factors such as cashing out small balances. We already exclude account rollovers from employer-sponsored, as they 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 penalized distributions, Appendix Figure D.6 replicates our event study analyses but where the outcome variables are indicators for taking penalized withdrawals that are higher than given thresholds.

hold's marginal tax rate. Third, even households experiencing income gains make penalized withdrawals, averaging to about 1 percent of annual income. This suggests that equilibrium liquidity valuation reflects not only income shocks but also expenditure-driven needs, such as health or child-related expenses. It further highlights the strength of our revealed-preference approach that accommodates multi-dimensional sources of liquidity needs beyond observable shocks to income. As a result, even full insurance against negative income shocks would not be sufficient to fully smooth marginal utility over time.

# 5 Valuation of Liquidity during the Great Recession

We next examine how large, market-wide economic shocks influence households' valuation of liquidity, focusing on the Great Recession.

**Estimating Equation.** We estimate the following "dose-response" event-study specification:

$$y_{i,z,t} = \Gamma_z + \alpha_i + x_{i,t}\lambda + \sum_{r=2000, r\neq 2006}^{2017} \beta_r I_r + \sum_{r=2000, r\neq 2006}^{2017} \theta_r I_r \times Treat_z + \varepsilon_{i,t}, \tag{4}$$

where $y_{i,z,t}$ is the outcome for household $i$ residing in Commuting Zone (CZ) $z$ in year $t$, $I_r$ are calendar-year indicators for 2000-2017 (with 2006 as the omitted baseline), $x_{i,t}$ includes age fixed effects for the primary filer and calendar-year fixed effects, $\Gamma_z$ are CZ fixed effects, and $\alpha_i$ are household fixed effects.

$Treat_z$ measures local exposure to the Great Recession and is defined as the change in the CZ-level unemployment rate from 2007 to 2009, following Yagan (2019). Our parameters of interest are $\theta_r$, which capture the differential response of households in locations experiencing a 1 pp larger increase in unemployment.

**Results.** Panels (a) and (b) of Figure 4 plot the estimates of $\theta_r$ for withdrawal frequency and withdrawal amounts, respectively. Households in Commuting Zones more severely affected by the Great Recession exhibit significantly larger increases in penalized withdrawals. The response peaks in 2009: a 1 pp larger local unemployment increase raises the probability of making a penalized withdrawal by 0.403 pp and increases penalized liquidity by \$63.4 per household (including zeros for non-withdrawing households).

Aggregating effects over the period 2007-2012, we estimate a cumulative increase of 1.47 pp in withdrawal incidence and \$251 in penalized liquidity. These patterns indicate a persistent elevation in households' valuation of liquidity in harder-hit locations, manifested as a sustained increase in the share of households triggered to make penalized withdrawals and in the amount of costly self-insurance undertaken.

To benchmark these magnitudes against household-level unemployment shocks, we normalize the market-level estimates by unemployment incidence. This normalization corresponds to a counterfactual 100 pp increase in local unemployment, analogous to an individual household entering unemployment. Under this scaling, the locality-level effect on withdrawal incidence is approximately four times larger than the household-level unemployment effect documented in Section 4 (40.3 pp versus 10.4 pp). Interpreted through our framework, this comparison implies that roughly one-quarter of the Great Recession effect reflects increased household liquidity demand, while the remaining three-quarters reflects market-level changes, inclusive of the deterioration in local credit supply. A similar conclusion emerges for withdrawal amounts: the normalized market-level effect ($6,340) substantially exceeds the household-level unemployment effect ($1,590).

Motivated by this decomposition, we explicitly separate the cumulative Great Recession effect into direct household-level and indirect market-level components. We augment equation (4) with flexible controls for household economic conditions—unemployment status, wage earnings, and gross income—fully interacted with calendar-year indicators. The resulting estimates are shown by the dashed lines in panels (a) and (b) of Figure 4. The indirect component accounts for approximately three-quarters of the total effect, consistent with a market's tightening credit and worsening economic conditions affecting households in distressed locations. Within our framework, these results indicate that economic conditions that go beyond those directly-linked to the household and relate to contractions in market credit supply were a dominant driver of the increase in households' equilibrium valuation of liquidity during the Great Recession.

In sum, the Great Recession illustrates how market-wide shocks raise households' valuation of liquidity through both direct income disruptions and indirect market-level economic and credit market spillovers. They shift a large mass of households from a region where liquidity valuation remains bounded below the penalty to one in which costly self-insurance through penalized withdrawals becomes optimal.

# 6 Unpacking Persistent Drivers of Valuation of Liquidity

In this final empirical analysis, we study permanent drivers of households' valuation of liquidity. Specifically, we ask whether persistent differences in penalized withdrawals—and hence in households' valuation of liquidity—across local markets are primarily attributable to household characteristics or to the locations in which households live. To address this question, we first use a movers design to quantify the relative importance of place versus household factors. We then relate the estimated location and household effects to local economic and financial characteristics, shedding light on the mechanisms underlying persistent differences in the valuation of liquidity.

As a starting point, we document substantial spatial heterogeneity in penalized withdrawals. Panel (a) of Figure 5 plots the average annual share of households making a penalized withdrawal by Commuting Zone (CZ), aggregated over 1999-2018. The mean withdrawal rate is 9.8 pp, with a cross-location standard deviation of 1.7 pp.[18]

## 6.1 Movers' Design

To quantify the causal contribution of location to differences in households' valuation of liquidity, we exploit variation from households that move across Commuting Zones. 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 and to their local environment.

We begin with a two-way fixed effects specification (Abowd et al., 1999):

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

where $y_{i,t}$ is an indicator for whether household $i$ makes a penalized withdrawal at time $t$, $\alpha_i$ captures persistent household characteristics, $\Gamma_{z(i,t)}$ are location fixed effects, and $x_{i,t}$ includes time-varying controls such as age indicators for the primary filer, (cyclical) calendar-year fixed effects, and household-level economic conditions. $z(i,t)$ indexes household $i$'s location in period $t$. Household fixed effects capture persistent traits affecting the valuation of liquidity, while location fixed effects capture persistent features of the local environment, including access to formal credit markets and informal insurance networks.

The logic of the move quasi-experiment is straightforward: if a household moves from a low- to a high-withdrawal area and increases its use of penalized withdrawals, it indicates that persistent destination features affect liquidity valuation. The design compares post-move outcomes for movers experiencing differential treatment intensities based on the gap in withdrawal behavior between their origin and destination. The identifying assumption is that these movers would have exhibited similar withdrawal patterns had the move not occurred, leading to standard parallel pre-trends necessary conditions and their respective tests we provide below.

[18]We address bias in plug-in estimates of second moments due to sampling error (e.g., Andrews et al. 2008) by estimating dispersion using a split-sample approach as in Finkelstein et al. (2021).

We formalize this intuition as Finkelstein et al. (2016) in the following way. Let $y_{z_j} \equiv E[y_{i,t} \mid z(i,t) = z_j]$ represent the location level average withdrawal rate. The share of the differences in withdrawal rates across location that is causally attributable to location—our key parameter of interest—is then captured by:

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

Empirically, we calculate expectations with means using non-movers only (i.e., "leave-out" means) to avoid mechanical correlations in the estimations below.

We define the difference in average withdrawal rates between household $i$'s destination CZ ($z_D$) and origin CZ ($z_O$) as the "treatment intensity" of the move by:

$$\Delta_i \equiv y_{z_D} - y_{z_O}.$$

We let $r(i,t)$ denote event time relative to the move and $I_{r(i,t)>0}$ indicate post-move periods. This parameterization leads to the following empirical specification for movers:

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

where $\mu_i = \alpha_i + \Gamma_{z_O}$ captures the household and origin-location fixed effect. In this specification, the same parameter $\theta$ behaves as the average "passthrough" of cross-location differences in withdrawal rates into movers' own behavior. That is, it captures how much of the observed CZ-level differences in penalized withdrawals is transmitted to movers. Estimation of this equation using the sample of movers identifies the share of observed differences in withdrawals across locations that is driven by location itself, $\theta$, simply as the regression coefficient on the "treatment intensity" variable, $\Delta_i$.

**Estimating Equation.** We estimate the full dynamic counterpart of equation (6) around the move using the following "dose-response" event-study specification:

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

where $I_r = \mathbb{I}\{r(i,t) = r\}$. The omitted category is two years prior to the move ($r = -2$), consistent with the earlier event-study analyses. Standard errors are clustered at the origin CZ level. Indexing $\theta$ by $r$ conveniently allows for testing parallel pre-trends in the pre-period as well as for potential dynamics in location effects in the post-period.

**Results.** In Figure 7, we first estimate equation (7) on a balanced sample of households observed from three years before to five years after a move. Panel (a) reports results for withdrawal take-up, and panel (b) of reports results for withdrawal amounts. We find no evidence of differential pre-trends across households moving to locations with different withdrawal intensities, supporting the identifying assumptions. At the time of the move, withdrawal behavior adjusts sharply and remains persistently elevated thereafter.

The estimated transmission of location-level differences into household withdrawal probability averages 0.34 in the post-move period (panel (a)), implying that approximately one-third of cross-location differentials in penalized withdrawals is attributable to place effects. Moving from a low- to a high-withdrawal location (from the 5th to the 95th percentile) increases required penalized liquidity by approximately \$340 per year (panel (b)); where, among households that adjust on the withdrawal extensive margin, this increase corresponds to average additional withdrawals of roughly \$16,900.

We then extend the analysis to a window from year $-5$ to $+10$ (on an unbalanced sample of households) with similar findings. Again, there are virtually no pre-trends in support of the design. As for the post-event period, we find a high degree of persistence for up to 10 years in the role of location. Moreover, the moderate decline in the estimates is attributable to "attrition" and return moves displayed in Appendix Figure D.8 (panel (a)). They attenuate the persistence since we assign a household the same destination for the entire post-move period, whether they subsequently moved or not, because these behaviors could be endogenous to the initial move itself. Panel (b) of Appendix Figure D.8 illustrates this point: when scaling the estimates by the share of movers who are still at the assigned destination, the dynamics flatten out.

Guided by the framework in Section 3, we interpret these results as evidence that persistent local conditions—including market-wide economic factors and credit access—causally and substantially affect households' valuation of liquidity. We subject this interpretation to a series of robustness checks, summarized in Appendix Figure D.9. Panel (a) shows the results are not explained by changes in household economic conditions following a move; including a flexible set of controls for unemployment, wages, and gross income (with their lags, leads, and interactions with relative time) leaves the estimates virtually unchanged. Panel (b) indicates the results are not driven by movers learning about withdrawals from new peers; results are similar for households who had previously made penalized withdrawals before moving. Finally, panel (c) suggests tax optimization cannot be a primary explanation, as flexibly controlling for location-specific top marginal tax rates yields only modest attenuation. Together, these analyses support a locality's broader economic and credit conditions as primary drivers of persistent spatial differences in liquidity valuation.

### 6.2 Drivers of Location and Household Effects

Lastly, we investigate the factors underlying the estimated location effects ($\Gamma_z$) and household effects ($\alpha_i$). We estimate equation (5) including flexible controls for household economic conditions—unemployment status, wage earnings, and gross income with lagged, current, and lead values.

Panels (b) and (c) of Figure 5 display the estimated location effects and CZ-level averages of household fixed effects. Both exhibit substantial geographic variation, with standard deviations of approximately 1 pp and 1.8 pp, respectively. We relate these estimates to CZ characteristics using univariate regressions; Figure 6 reports normalized coefficients, with scatter plots in Appendix Figures D.10 and D.11.

**Location Effects.** Locations with higher values of the Credit Insecurity Index developed by the Federal Reserve Bank of New York (Hamdani et al., 2019) exhibit significantly higher reliance on penalized withdrawals, consistent with more limited access to alternative credit. We also find that areas with higher median home values display lower withdrawal propensities, consistent with collateral availability easing credit constraints. Both patterns align closely with the local credit channel of the location effects in withdrawals.

**Household Effects.** Household fixed effects are most strongly correlated with race. Households with Black primary earners rely substantially more on penalized withdrawals, even after controlling for income, employment, and location. Location effects are uncorrelated with the share of Black residents, while household fixed effects remain strongly correlated with race even within Commuting Zones and ZIP Codes (Figure 6). Table 1 shows that Black households' fixed effects exceed those of White households by over 30 percent, a gap that persists after conditioning for household economic circumstances. These findings suggest persistent racial disparities in access to credit that operate independently of income and geography.

## 7 Discussion

This section reflects on the conceptual foundations of our empirical approach and discusses its broader relevance and policy implications. We first assess the validity of using penalized withdrawals as a revealed-preference measure to bracket households' valuation of liquidity, then discuss the empirical scope of this tool, and finally highlight implications for policy design.

**Penalized Withdrawals as a Revealed-Preference Measure.** Penalized withdrawals from retirement accounts provide a tractable revealed-preference approach to studying households' valuation of liquidity. When a household chooses to incur a statutory 10 percent penalty to access funds early, it reveals that the marginal value of liquidity exceeds the marginal cost imposed by the penalty. The absence of withdrawals among households with positive retirement balances is equally informative, indicating situations in which liquidity valuation remains below this threshold. Observed withdrawal behavior therefore delivers a transparent revealed-preference bound on households' willingness to pay for liquidity.

Our approach builds on a long tradition in public finance and welfare analysis that infers preferences from responses to observed prices and incentives. A key identifying requirement is that households actively choose whether to withdraw. While inertia and default effects are well documented in the *accumulation* phase of retirement saving, the evidence points to more active and responsive behavior in the *decumulation* phase. As we document in Section 2.3, penalized withdrawals are episodic, typically partial rather than exhaustive, and strongly concentrated around adverse economic events. These patterns are difficult to reconcile with passive or mechanical behavior and instead support an optimizing interpretation.

The structure of the early-withdrawal penalty itself further strengthens this interpretation in comparison to passive behavior or time inconsistency. The immediate cost is incurred contemporaneously with the liquidity benefit, making the marginal tradeoff salient and limiting the role of present-biased preferences that complicate accumulation decisions. Such complications arise primarily in the savings phase, where the cost of consuming today is borne only far in the future at retirement. While behavioral frictions may still be present, their primary implication would be to introduce internalities in a full welfare analysis rather than to undermine the information content of observed withdrawal choices. Penalized withdrawals therefore remain informative about households' underlying liquidity needs and constraints even in such cases. We offer a detailed discussion of the possibility of behavioral interpretations in Appendix B.

**Scope and Empirical Reach.** Defined-contribution retirement accounts have become the dominant retirement savings vehicle in the US, with participation widespread across the income distribution. As a result, penalized withdrawals are not only conceptually meaningful but also empirically scalable. The margin we study applies to a large share of US households and can be measured precisely using administrative data. While other financial instruments, such as credit cards or HELOCs, often provide lower-cost liquidity, our measure specifically captures households for whom such alternatives are either exhausted or unavailable. Observing a penalized withdrawal thus identifies a situation of financial distress in which the shadow price of liquidity has cleared all lower-interest options available to the household. In this sense, the 10 percent penalty acts as a structural "cutoff" that filters out transitory or low-valuation liquidity needs, leaving a clean signal of deep credit-market binding.

As retirement balances grow and plan coverage continues to expand through automatic enrollment and recent policy initiatives, this margin is likely to become increasingly even more relevant for understanding how households manage liquidity in the presence of risk and incomplete insurance.

**Policy Implications.** Our findings have several implications for policy design. First, they suggest scope for refining the tax treatment of early withdrawals. Existing exemptions already recognize that liquidity needs vary across circumstances, and a more systematic approach—conditioning penalties on observable household characteristics, local credit conditions, or macroeconomic indicators—could improve insurance while preserving the commitment value of retirement accounts. Second, the large and persistent geographic differences we document point to a role for place-based policies aimed at improving access to lower-cost credit in underserved areas. Finally, emerging financial technologies may help bridge liquidity gaps in regions with limited traditional credit supply, potentially reducing households' reliance on costly self-insurance through retirement withdrawals.

Overall, penalized withdrawals offer a distinctive window into liquidity constraints in the US household sector. By quantifying how the value of liquidity varies across people and places, our approach provides a foundation for more targeted and effective policy interventions.

# 8 Conclusion

This paper introduces and empirically validates penalized withdrawals from retirement savings accounts as a revealed-preference tool for inferring households' valuation of liquidity. Using population-level tax records, we characterize the equilibrium distribution of liquidity valuation among American families and document three main findings.

First, we show that location effects account for over 30 percent of nationwide differences in liquidity valuation across labor markets. Second, analyzing the Great Recession, we find that market-level shocks generate large increases in liquidity valuation, with spillovers from local economic conditions and credit tightening accounting for the majority of the overall effect. Third, while reliance on penalized withdrawals for liquidity needs is widespread, Black households depend on this form of costly self-insurance to a significantly greater extent than White households with similar economic conditions, regardless of where they live. This final finding provides novel evidence that Black American families face systematically more limited access to lower-cost sources of liquidity throughout the country.

Taken together, our results demonstrate that penalized withdrawals offer a powerful lens for studying liquidity constraints and highlight the central roles of place, credit-market conditions, and racial disparities in shaping households' access to liquidity in the United States.

# 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 (2020). Optimal Illiquidity. Technical report, National Bureau of Economic Research.

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review 98*(5), 1887–1921.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review 106*(2), 387–435.

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 (2025). Location, location, location. *American Economic Journal: Applied Economics 17*(1), 297–336.

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. 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.

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 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 Eco-*

*nomics 100*(2), 303–320.

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., 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.

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*.

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.

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.

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.

Pruitt, S. and N. Turner (2020). Earnings risk in the household: Evidence from millions of us tax returns. *American Economic Review: Insights 2*(2), 237–254.

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.

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

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

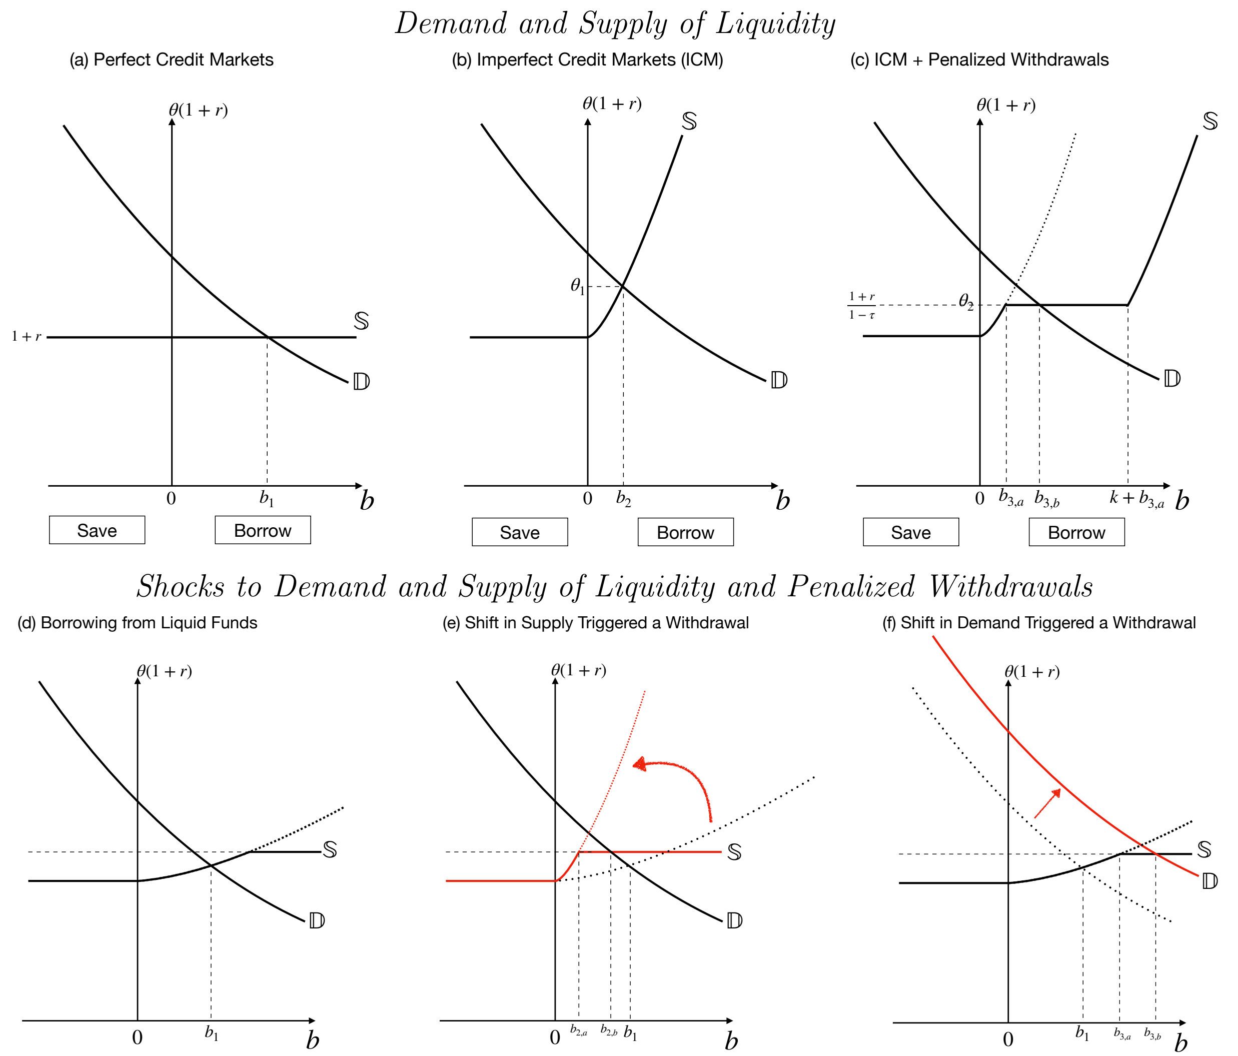


Notes: These figures illustrate how the supply and demand for liquidity, as defined in the main text, determine the equilibrium valuation of liquidity. The top three panels consider three different financial markets: perfect market (left), imperfect (middle), imperfect with the possibility to make a 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

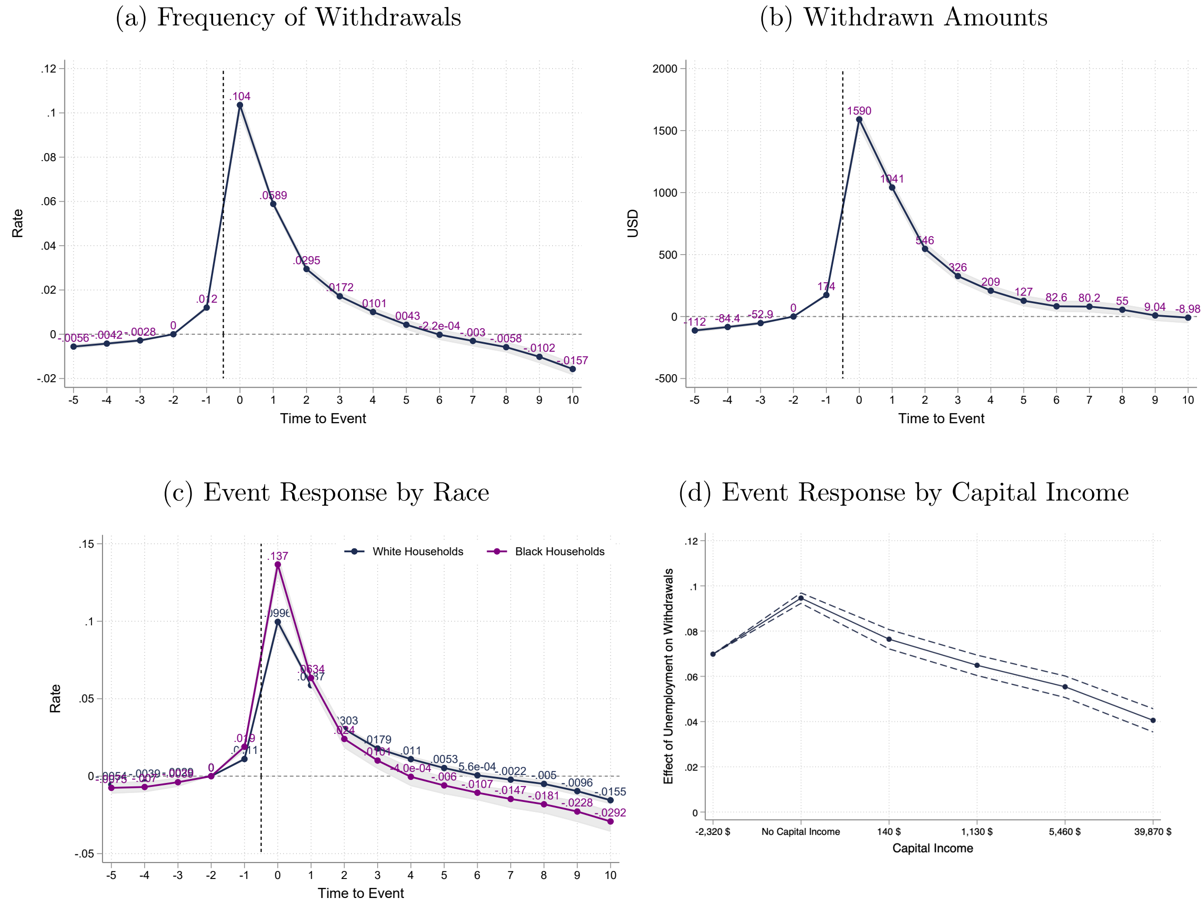


Notes: This figure studies penalized withdrawals around the event of household unemployment, defined as the first period we observe at least one of the household members receiving unemployment benefits. Panels (a) and (b) plot the event study coefficients from specification (3) when the outcome variables 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

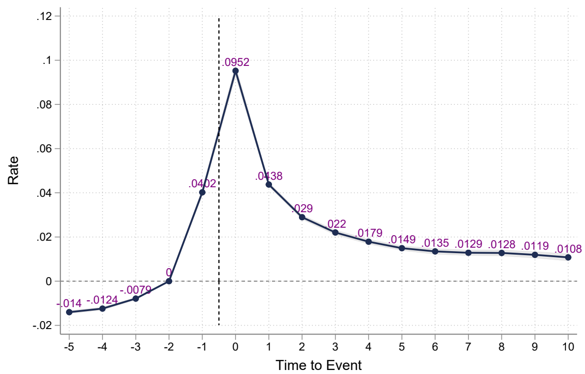


(b) Withdrawn Amounts after Large Income Loss

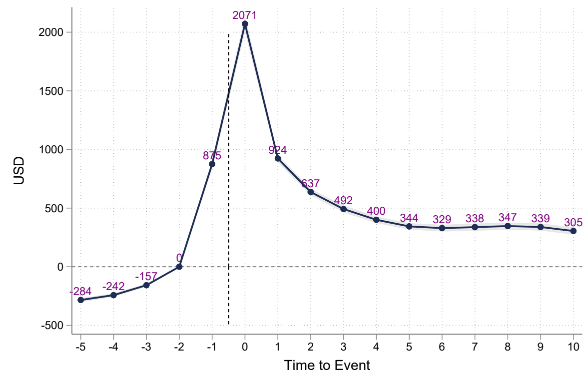


(c) Take-Up as a Function of Income Changes

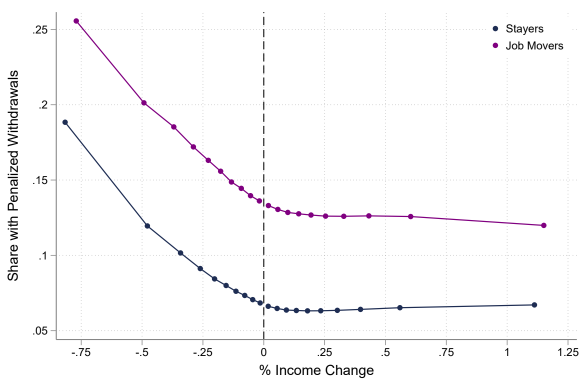


(d) Amounts as a Function of Income Changes

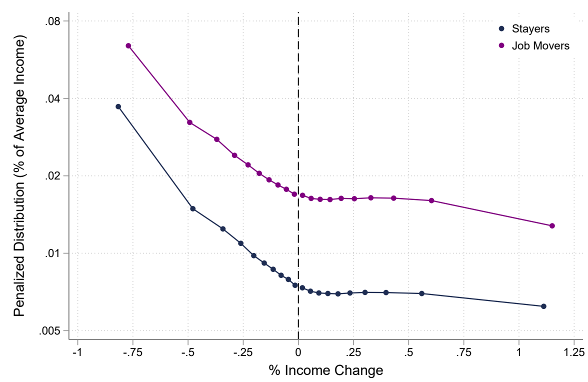


(e) Responses to Large Income Loss by Race

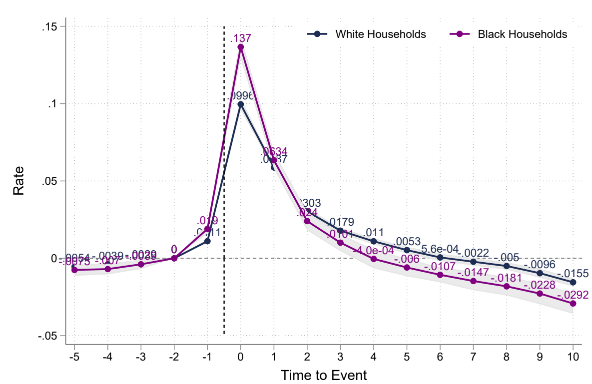


Notes: This figure studies penalized withdrawals around changes in household income. Panels (a) and (b) plot the event study coefficients from specification (3) when the outcome variables are take-up and amounts of penalized withdrawals, respectively, around the event of a large income loss. Large income 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.7, can 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

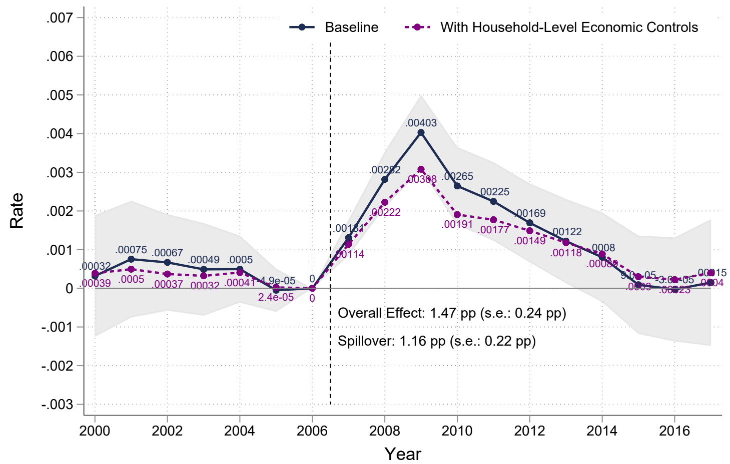


(b) Withdrawn Amounts

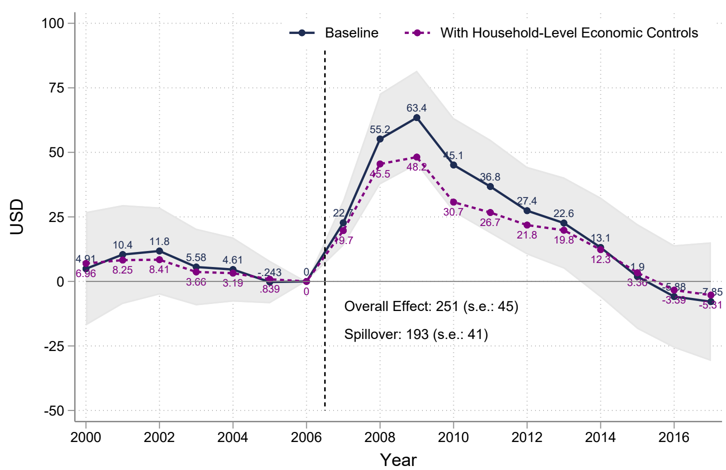


Notes: This figure displays estimates of the effect of the Great Recession on penalized withdrawals using equation (4). It provides estimates for the relative change in behavior in a locality that was exposed to a 1 percentage point larger local unemployment shock. Panel (a) analyzes the frequency of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

Figure 5: Geography of Withdrawals and Valuation of Liquidity

(a) Overall Variation

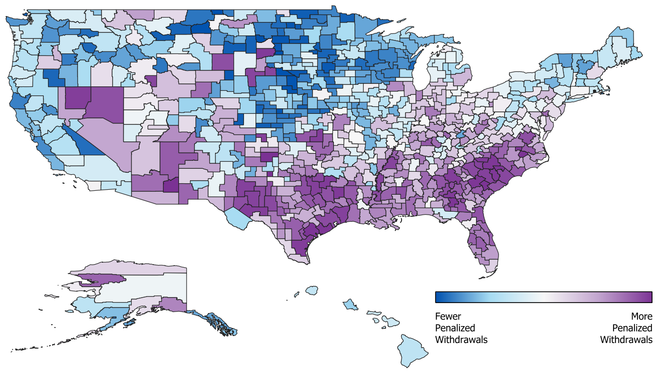


(b) Location Fixed Effects

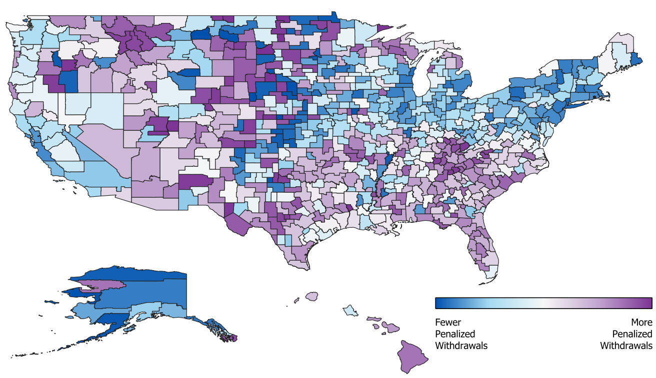


(c) Household Fixed Effects

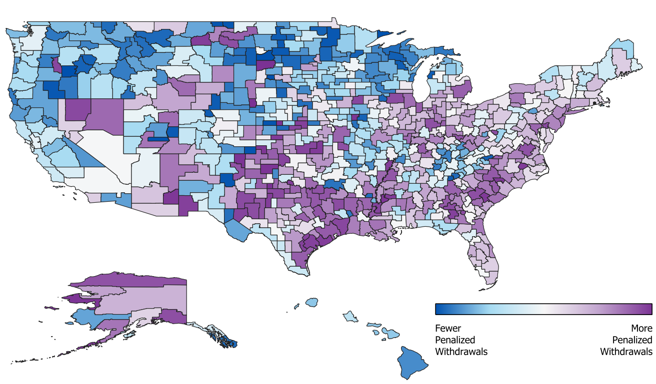


Notes: This figure studies the geography of penalized withdrawals and the valuation of liquidity. Panel (a) plots a map of the average annual share of households that have made a penalized withdrawal by Commuting Zones (CZs). Then, based on estimation of equation (5) with household-level economic controls, panel (b) plots a map of the location fixed effects, $\mathit{\Gamma}_z$, and panel (c) plots a map of the household fixed effects, $\alpha_i$, 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

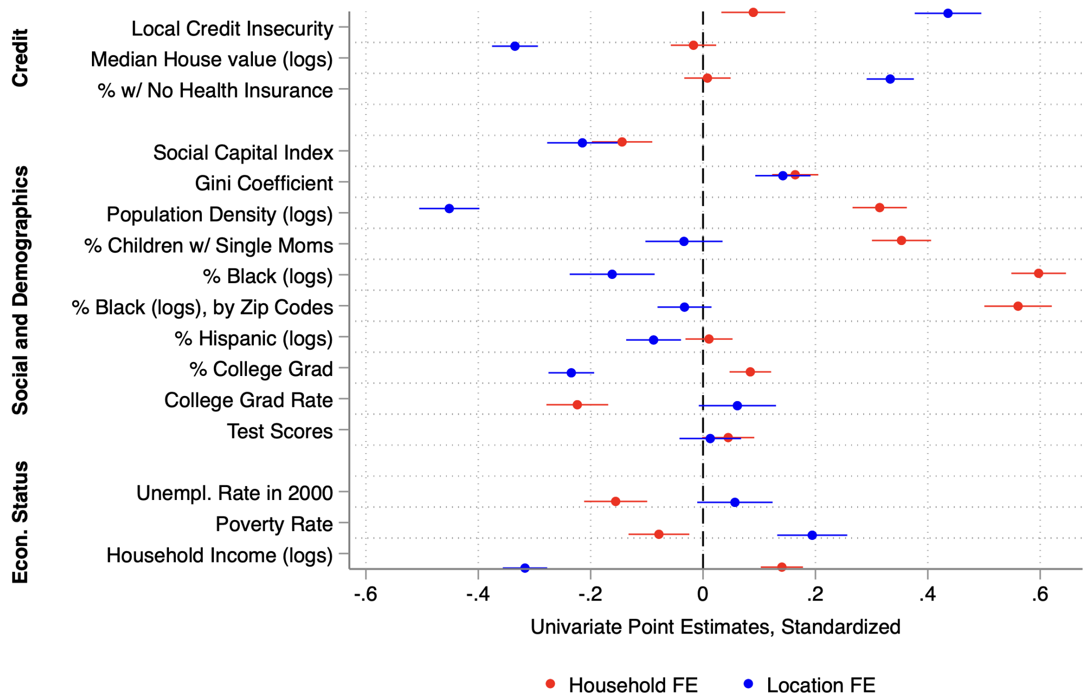


Notes: This figure displays correlations of the regional differences across CZs with CZ-level social and economic characteristics. We display correlations of these characteristics separately for the location fixed effects, $\Gamma_z$, and for the household fixed effects, $\alpha_i$, collapsed at the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

Figure 7: Movers Analysis

(a) Balanced Panel - Withdrawal Take-Up

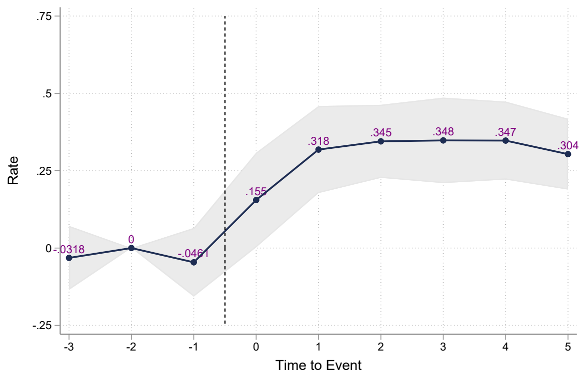


(b) Balanced Panel - Withdrawan Amounts

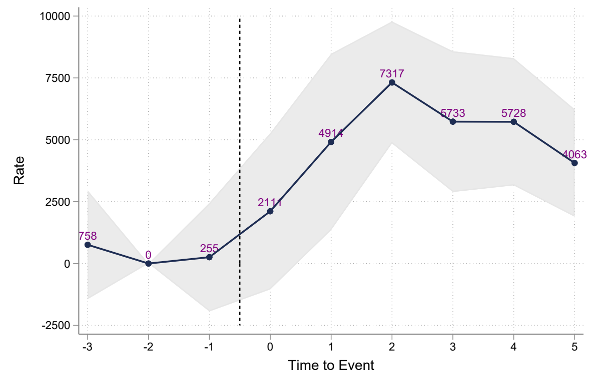


(c) Extended Horizon - Withdrawal Take-Up

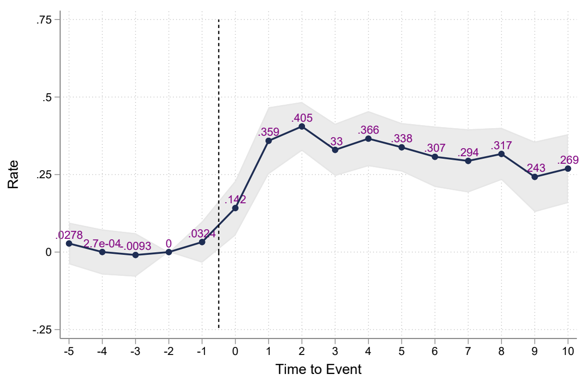


(d) Extended Horizon - Withdrawan Amounts

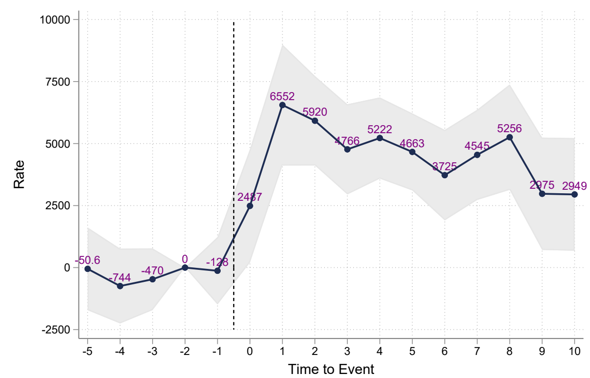


Notes: These figures display estimates for the share of spatial differentials in withdrawals that can be attributed to location, using the movers design specification of equation (7). Panels (a) and (b) show the estimates from a balanced panel of households we observe in the window that spans years [-3, +5] around the move for the outcomes take-up and amounts of penalized withdrawals, respectively. Panels (c) and (d) show the estimates from an unbalanced panel of households in an extended window that spans years [-5, +10] around the move.

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 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 or 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.

**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.[19] 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.[20]

---

[19]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.

[20]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 our model 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 percent. We note that the HRS is a representative sample of overall households in the US, 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 percent over ages 45-59.

**Fact 2: Penalized withdrawals are widely used and are episodic.** 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 only episodically, 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 time frame that the household made additional withdrawals.[21] 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 amount 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.[22] 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.

[21] 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).

[22] 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 consistent with households using 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.

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) of Appendix Table B.1 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 Withdrawan 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 percent. 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. 2020; Fadlon and Laibson 2021).[23] 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* (e.g., Read

[23] 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).

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 Proofs for Lemmas 1-3

Throughout, we consider the model in Section 3.1. Fix a household $i$ in region $z$ and suppress indices when there is no risk of confusion. Let $u'(c_t; h_t)$ denote the marginal utility of consumption in period $t$. We write

$$V_t \equiv V_t(a_{t-1}, k_{t-1}; h_t), \qquad V_{t+1,a} \equiv \frac{\partial V_{t+1}(a_t, k_t; h_{t+1})}{\partial a_{t+1}}, \qquad V_{t+1,k} \equiv \frac{\partial V_{t+1}(a_t, k_t; h_{t+1})}{\partial k_{t+1}}.$$

Whenever we refer to a *penalized withdrawal* at time $t$, we mean $\Delta k_t < 0$ and $t < t^*$, so that an extra dollar withdrawn delivers only $(1-\tau)$ dollars of consumption.

## Preliminaries: Optimality conditions for $\Delta a_t$ and $\Delta k_t$

Consider an interior choice of liquid saving/borrowing, $\Delta a_t$, and an interior choice of retirement flows, $\Delta k_t$, away from kinks (so that $b_t > 0$ implies differentiability of $\rho(\cdot)$ at $b_t$ and $\Delta k_t \neq 0$ implies differentiability of the penalty term). Using the budget constraint

$$c_t = (1-\varphi)y_t - \varepsilon_t - \Delta k_t - \Delta a_t + \tau\, \Delta k_t\, \mathbb{I}_{(\Delta k_t<0)} \mathbb{I}_{(t<t^*)} - \rho_{i,z}(b_t) \mathbb{I}_{(b_t>0)},$$

we have, for a penalized withdrawal ($\Delta k_t < 0$ and $t < t^*$),

$$\frac{\partial c_t}{\partial \Delta k_t} = -(1-\tau), \qquad \frac{\partial k_t}{\partial \Delta k_t} = (1+r),$$

where $k_t = (1+r)\big(k_{t-1} + \Delta k_t + \varphi y_t\big)$.

Similarly, when borrowing is used at the margin ($b_t > 0$ and locally $b_t$ increases one-for-one with $-\Delta a_t$), we have

$$\frac{\partial c_t}{\partial \Delta a_t} = -1 + \rho'_{i,z}(b_t), \qquad \frac{\partial a_t}{\partial \Delta a_t} = (1+r),$$

where $a_t = (1+r)\big(a_{t-1} + \Delta a_t\big)$.

The first-order condition for a penalized withdrawal ($\Delta k_t < 0$ and $t < t^*$) is therefore

$$u'(c_t; h_t)\,(1-\tau) = \beta(1+r)\, E_t[V_{t+1,k}]\,. \tag{8}$$

When borrowing is used at the margin ($b_t > 0$), the first-order condition for $\Delta a_t$ is

$$u'(c_t; h_t)\,\big(1 - \rho'_{i,z}(b_t)\big) = \beta(1+r)\, E_t[V_{t+1,a}]\,. \tag{9}$$

Finally, recall the definition of the equilibrium valuation of liquidity:

$$\theta_t \equiv \frac{u'(c_t; h_t)}{\beta(1+r) E_t[V_{t+1,a}]}, \qquad \Lambda_t \equiv \frac{E_t[V_{t+1,k}]}{E_t[V_{t+1,a}]}.$$

## Proof of Lemma 2

Lemma 2 states that for a household that makes a penalized withdrawal at $t < t^* - 1$ while retaining a positive balance, the household's valuation of liquidity satisfies

$$\theta_t = \frac{1}{1-\tau}\Lambda_t, \qquad \text{and} \qquad 1-\tau \leq \Lambda_t \leq 1,$$

together with the two special cases.

**Step 1: Derive** (2)**.** For a penalized withdrawal ($\Delta k_t < 0$ and $t < t^*$) with an interior choice along the withdrawal margin (which is ensured by $k_t > 0$), the first-order condition (8) holds with equality:

$$u'(c_t; h_t)\,(1-\tau) = \beta(1+r)\,E_t[V_{t+1,k}].$$

Divide both sides by $\beta(1+r)E_t[V_{t+1,a}]$ to obtain

$$\frac{u'(c_t; h_t)}{\beta(1+r)E_t[V_{t+1,a}]} = \frac{1}{1-\tau} \cdot \frac{E_t[V_{t+1,k}]}{E_t[V_{t+1,a}]},$$

which is exactly

$$\theta_t = \frac{1}{1-\tau}\Lambda_t.$$

**Step 2: Show $1-\tau \leq \Lambda_t \leq 1$ for $t+1 < t^*$.** Fix period $t+1 < t^*$. Consider the continuation problem at $t+1$ starting from some state $(a_t, k_t; h_{t+1})$ and consider a marginal increase in next period's retirement wealth, $k_{t+1}$. An extra illiquid dollar at $t+1$ cannot be more valuable than an extra liquid dollar, because liquid wealth weakly expands the feasible set relative to illiquid wealth (liquid wealth can always be carried or used directly, while retirement wealth is subject to the penalty when accessed). Formally, letting $V_{t+1}$ denote the value function at $t+1$,

$$\frac{\partial V_{t+1}}{\partial k_{t+1}} \leq \frac{\partial V_{t+1}}{\partial a_{t+1}} \quad \Rightarrow \quad E_t[V_{t+1,k}] \leq E_t[V_{t+1,a}],$$

hence $\Lambda_t \leq 1$.

For the lower bound, note that a marginal dollar held in the retirement account at $t+1 < t^*$ can always be converted into $(1-\tau)$ dollars of liquid resources immediately via a (marginal) penalized withdrawal. This option implies that an extra retirement dollar must be worth at least $(1-\tau)$ times an extra liquid dollar:

$$\frac{\partial V_{t+1}}{\partial k_{t+1}} \geq (1-\tau)\frac{\partial V_{t+1}}{\partial a_{t+1}} \quad \Rightarrow \quad E_t[V_{t+1,k}] \geq (1-\tau)E_t[V_{t+1,a}],$$

hence $\Lambda_t \geq 1-\tau$.

**Step 3: Special cases.** (i) *Lower bound.* If the household expects that a marginal dollar in the retirement account at $t+1$ will be withdrawn before retirement with certainty, then the marginal retirement dollar is effectively converted into liquid consumption only net of the penalty. In that case, the option described above is exercised with probability one at the margin and the lower bound binds:

$$E_t[V_{t+1,k}] = (1-\tau)E_t[V_{t+1,a}] \quad \Rightarrow \quad \Lambda_t = 1-\tau.$$

(ii) *Upper bound.* If the household expects that, from $t+1$ until retirement, liquidity will not be valuable at the margin (so that carrying a marginal dollar in liquid form provides no advantage over carrying it in the retirement account), then liquid and illiquid wealth are locally equivalent and the upper bound binds:

$$E_t[V_{t+1,k}] = E_t[V_{t+1,a}] \quad \Rightarrow \quad \Lambda_t = 1.$$

This completes the proof. □

## Proof of Lemma 1

Lemma 1 is the special case in which a household makes a penalized withdrawal at time $t$ while retaining a positive balance and the valuation of liquidity satisfies

$$\theta_t = \frac{1}{1-\tau}, \qquad \text{with} \qquad \frac{1}{1-\tau} \geq \frac{1}{1-\rho'_{i,z}(b)}.$$

**Step 1: $\theta_t = \frac{1}{1-\tau}$.** Lemma 1 corresponds to the case in which the next period treats liquid and illiquid wealth symmetrically at the margin. In the institutional environment, this holds at $t = t^* - 1$: at $t+1 = t^*$ retirement balances become freely accessible, so a marginal dollar in $k_{t+1}$ is locally equivalent to a marginal dollar in $a_{t+1}$. Therefore,

$$E_t[V_{t+1,k}] = E_t[V_{t+1,a}] \quad \Rightarrow \quad \Lambda_t = 1.$$

Applying Lemma 2 then yields

$$\theta_t = \frac{1}{1-\tau}\Lambda_t = \frac{1}{1-\tau}.$$

**Step 2: $\frac{1}{1-\tau} \geq \frac{1}{1-\rho'_{i,z}(b)}$.** Suppose borrowing is used at the margin at time $t$ (so that $b_t > 0$ and (9) holds). Dividing (9) by $\beta(1+r)E_t[V_{t+1,a}]$ gives

$$\theta_t = \frac{1}{1-\rho'_{i,z}(b_t)}.$$

Since the household withdraws rather than borrowing more, the marginal cost of obtaining liquidity through borrowing cannot be lower than the marginal cost of obtaining liquidity through a penalized withdrawal. Equivalently, at the optimum we must have

$$\frac{1}{1-\tau} \geq \frac{1}{1-\rho'_{i,z}(b_t)}.$$

(If the household is not borrowing at the margin, this inequality holds trivially because the right-hand side is not the relevant marginal cost of liquidity.) This completes the proof. $\square$

## Proof of Lemma 3

Lemma 3 states that if a household has positive retirement balances and does not make a penalized withdrawal at time $t$, then

$$\theta_t \leq \frac{1}{1-\tau}.$$

Consider a household with $k_t > 0$ and $t < t^*$. Suppose the household does not make a penalized withdrawal at $t$, i.e., it chooses a point on the withdrawal margin with $\Delta k_t \geq 0$ (or, more generally, it does not choose $\Delta k_t < 0$). Then a small deviation toward a penalized withdrawal, $\Delta k_t \downarrow 0$, must not increase value. The Kuhn–Tucker condition for the (one-sided) deviation in the direction of a penalized withdrawal implies

$$u'(c_t; h_t)\,(1-\tau) \;\leq\; \beta(1+r)\,E_t[V_{t+1,k}].$$

Divide both sides by $\beta(1+r)E_t[V_{t+1,a}]$ to obtain

$$\theta_t \;\leq\; \frac{1}{1-\tau}\Lambda_t.$$

Finally, for $t+1 < t^*$ we have $\Lambda_t \leq 1$ from the bounds established in the proof of Lemma 2, hence

$$\theta_t \leq \frac{1}{1-\tau}.$$

This completes the proof. $\square$

# D Appendix Figures

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

*Prevalence of Accounts*

(a) By Age of Primary Filer

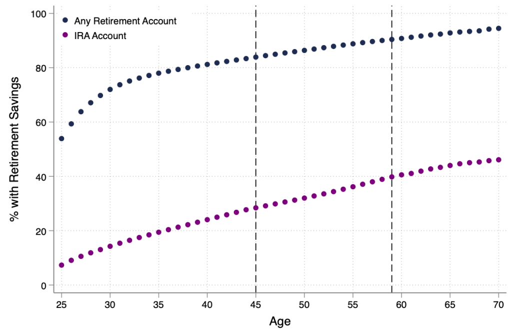


(b) By Household Overall Income

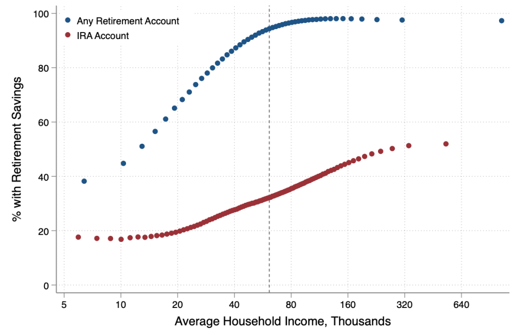


*Prevalence of Withdrawals*

(c) By Age of Primary Filer

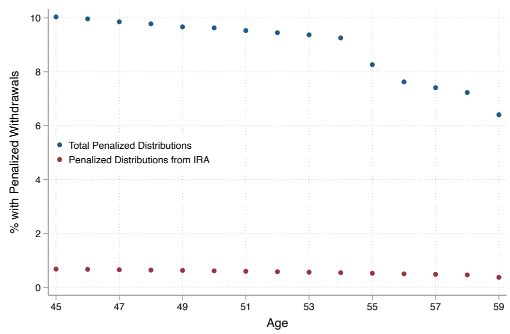


(d) By Household Average Income

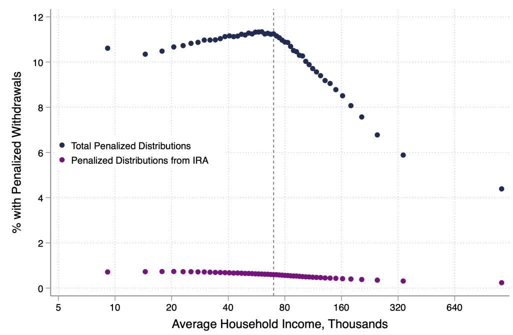


(e) Number of Penalized Withdrawals

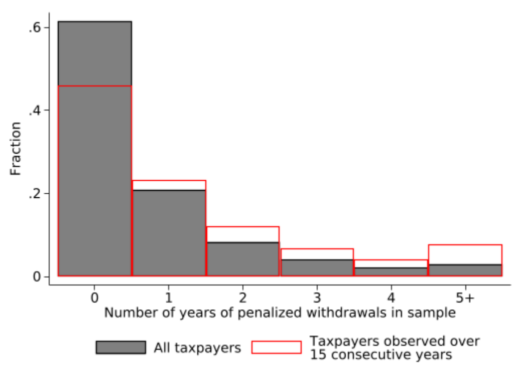


(f) Repeated Penalized Withdrawals

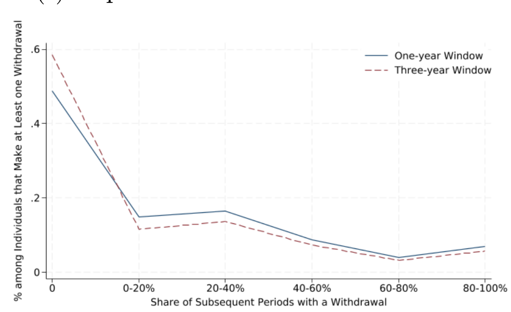


Notes: These figures illustrate the prevalence of retirement accounts and penalized withdrawals. We identify households as having accounts in a given year if up to that year within our sample period of 20 years they have made a contribution to 401(k)/IRA accounts or have balances in IRA accounts. The 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) or 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

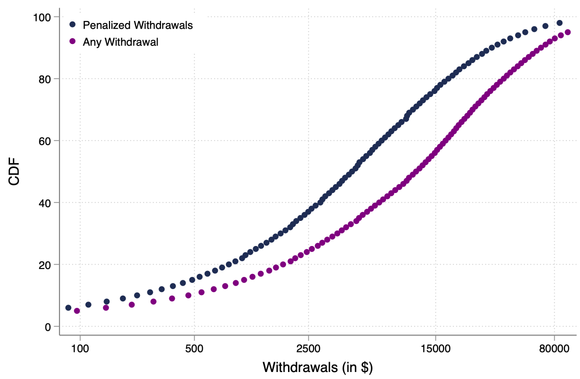


(b) Share of IRA Accounts Fully Depleted

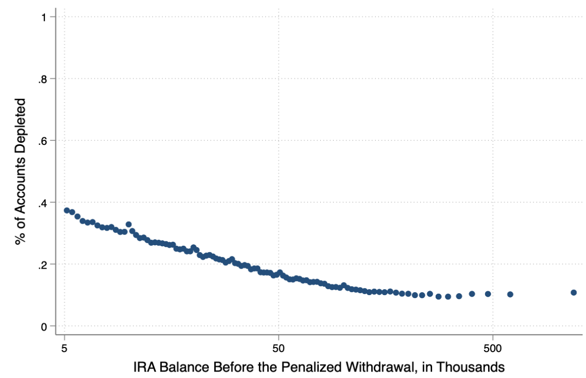


(c) Share of IRA Balances Withdrawn

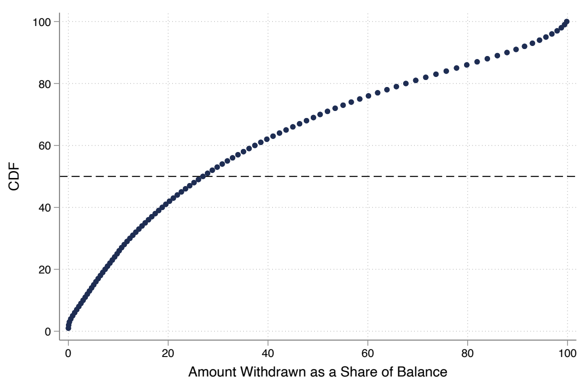


(d) Penalized Withdrawals and Income Changes

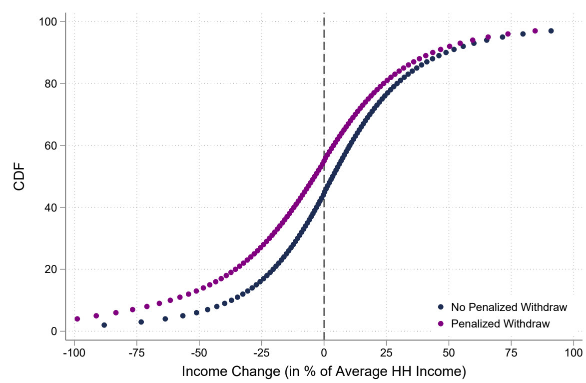


Notes: These figures provide different statistics regarding the behavior of penalized withdrawals. Panel (a) shows the overall CDF of amounts of penalized withdrawals and compares it with the overall CDF of amounts of withdrawals of any kind. Panels (b) and (c) focus on households who have an IRA account 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

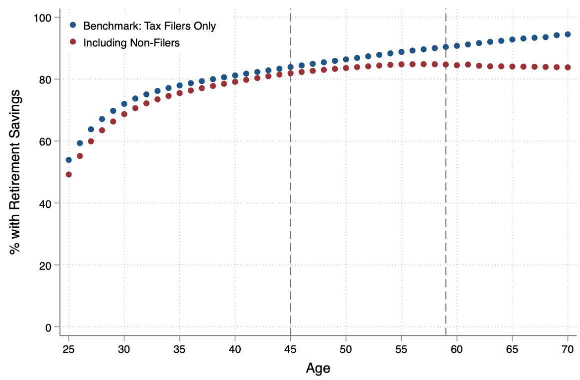


Notes: This figure illustrates the prevalence of retirement accounts by age. We compare our benchmark sample of households that file tax returns to a more inclusive sample, which combines in households that do not file tax returns. The figure shows that households that do not file tax returns are less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

Figure D.4: Unemployment Event—Adjusted Gross Income

(a) All Households

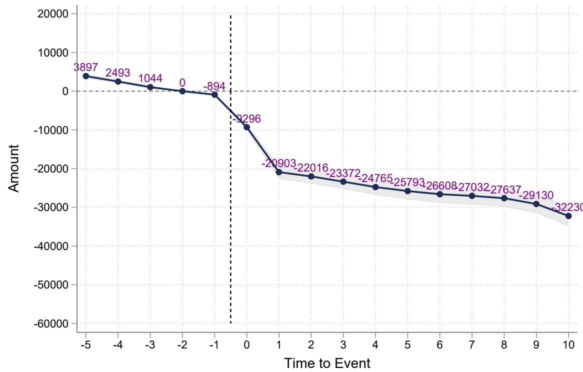


(b) By Race

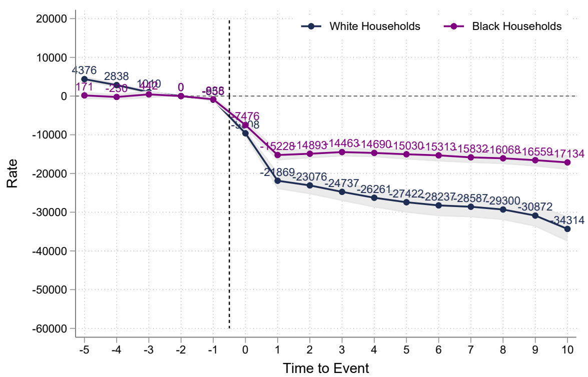


(c) By Race in Percent Change

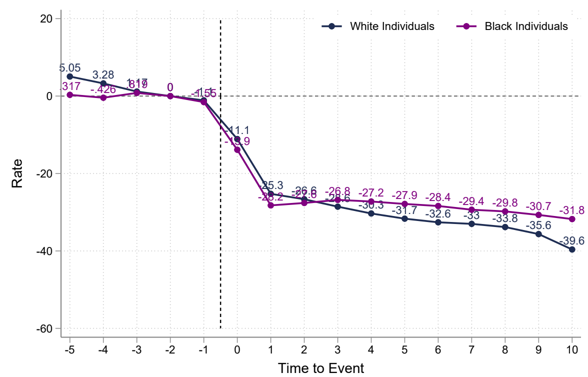


Notes: This figure studies households' adjusted gross income (AGI) around the event of unemployment, defined as the first period we observe at least one of the household members receiving unemployment benefits. Panel (a) plots the event study coefficients from specification (3) for the entire sample. 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: Event Studies for Households Who Stay in the Same CZ

(a) Unemployment

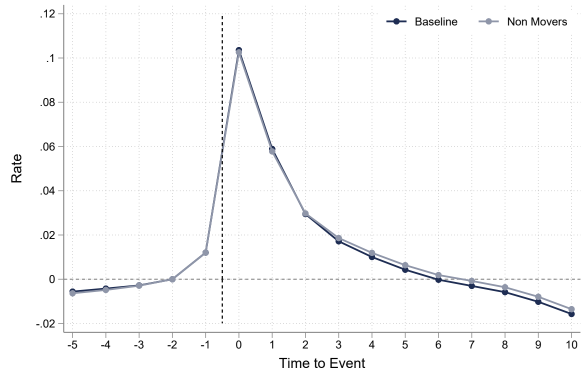


(b) Large Income Loss

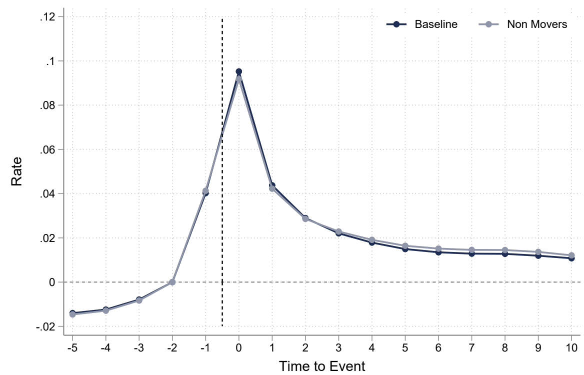


Notes: This figure plots the event study coefficients from specification (3) for the events of unemployment in panel (a) and large income losses in panel (b). We compare the overall sample to the sample of households that do not change their Commuting Zone around the event. Specifically, we consider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

Figure D.6: Event Studies by Amount Withdrawn

Unemployment

(a) Percentage Points

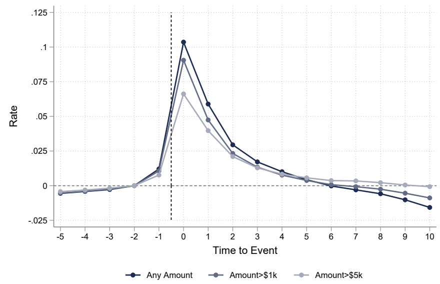


(b) Percent Change

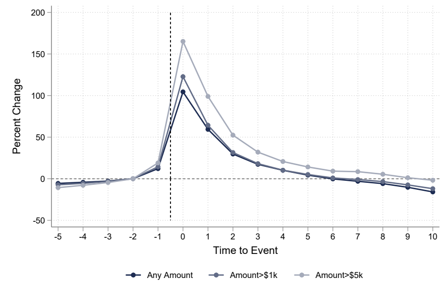


Large Income Losses

(c) Percentage Points

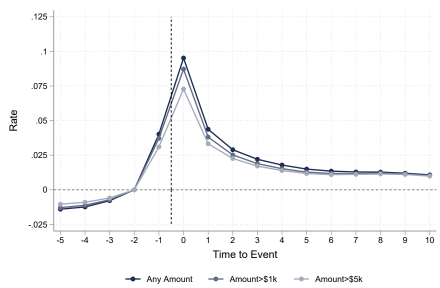


(d) Percent Change

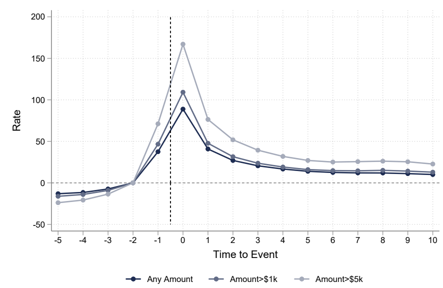


Notes: This figure plots the event study coefficients from specification (3) for the events of unemployment and large income loss. We study indicators for making penalized withdrawals of different amount thresholds: any amount, more than $1,000, and more than $5,000. For each event, the left panel reports 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.7: Event Study of Job Switch

(a) Frequency of Withdrawals

(b) Withdrawn Amounts

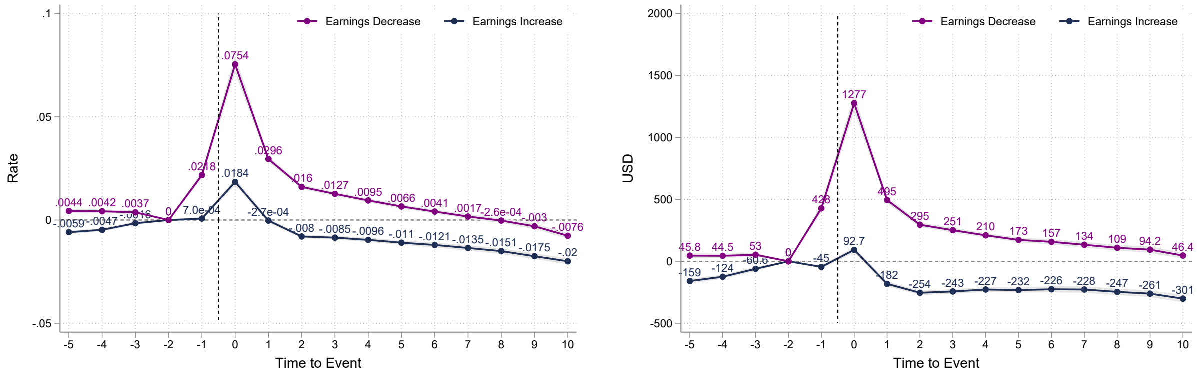


Notes: This figure studies penalized withdrawals around the event of a job switch using specification (3). It focuses on the sample of households for whom we see a change in employer from period $t-1$ to period $t$ without an episode of being on unemployment benefits. We then split households by whether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

Figure D.8: Movers Design—Attrition and Return Moves

(a) Dynamics

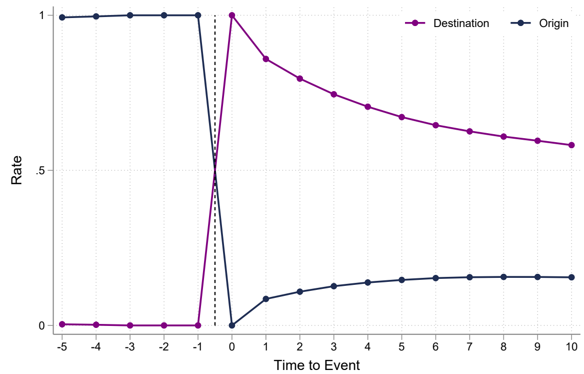


(b) Passthrough Scaled by Movers Still in Destination

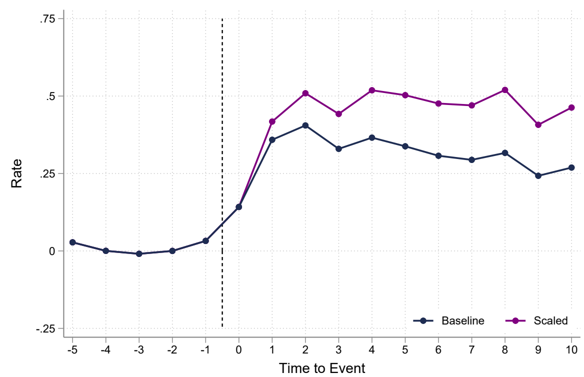


Notes: This figure provides additional analyses for the movers design. Panel (a) displays indicators for a household's geographic location around the move. In the movers design, we assign a household the same destination location for the entire post-move period. In this figure, we display indicator variables 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.9: Movers Analysis

(a) Economic Household-Level Controls

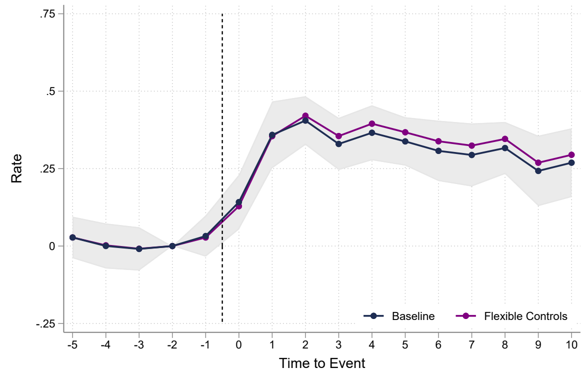


(b) Potential Learning

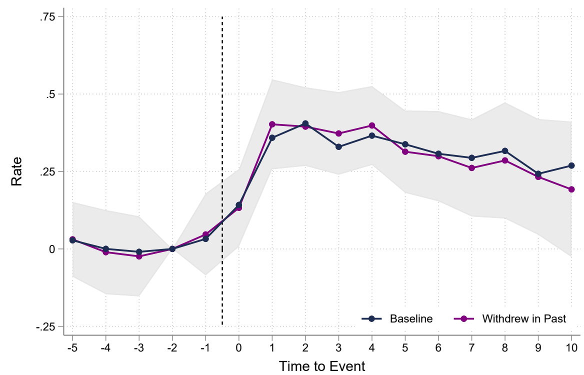


(c) Tax Motives

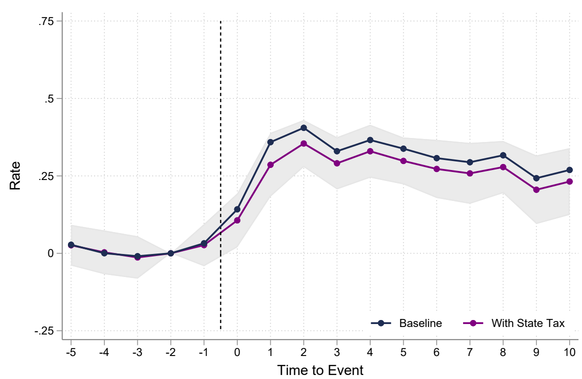


Notes: These figures display estimates for the share of spatial differentials in withdrawals that can be attributed to location, using the movers design specification of equation (7). The panels provide a series of robustness investigations. Panel (a) runs a specification that includes flexible (endogenous) 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 (b) 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 (c) 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.

Figure D.10: Correlations with Location Fixed Effects

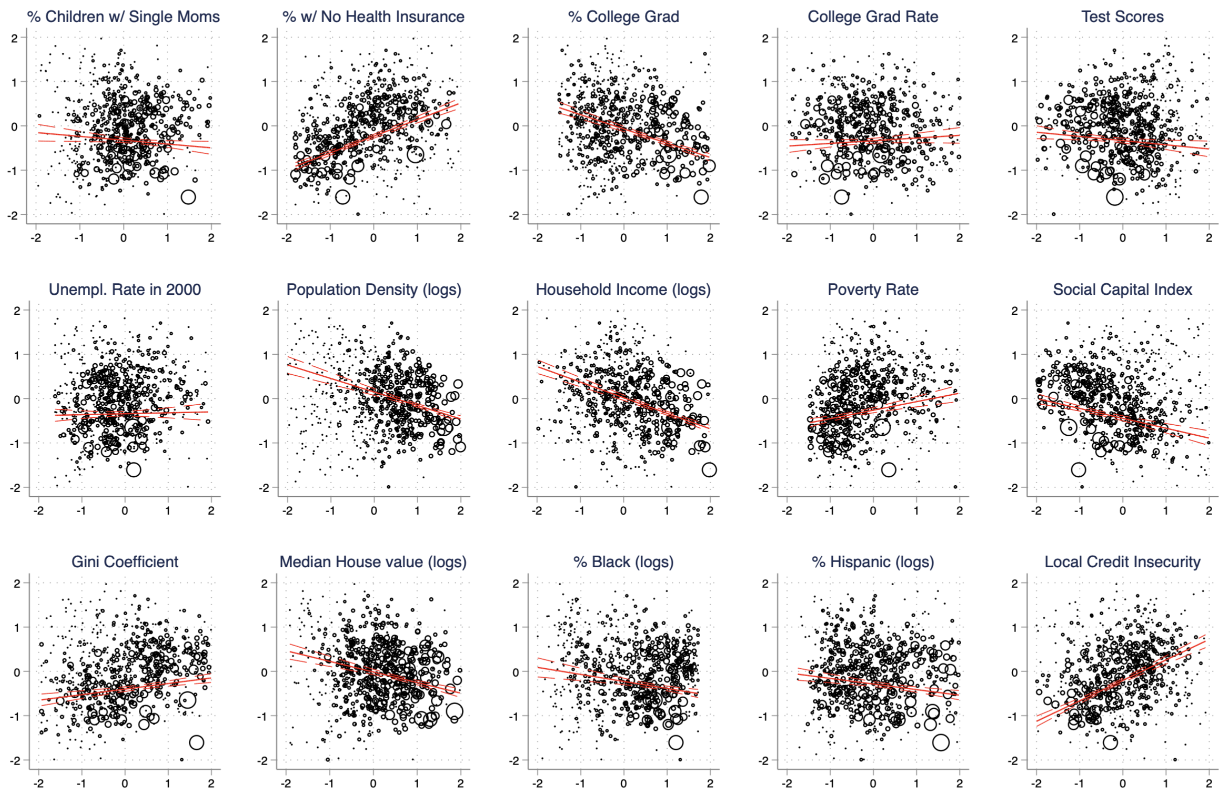


Notes: These figures display correlations of the location fixed effects, $\Gamma_z$, as estimated using equation (5), with CZ-level social and economic characteristics.

Figure D.11: Correlations with Households Fixed Effects

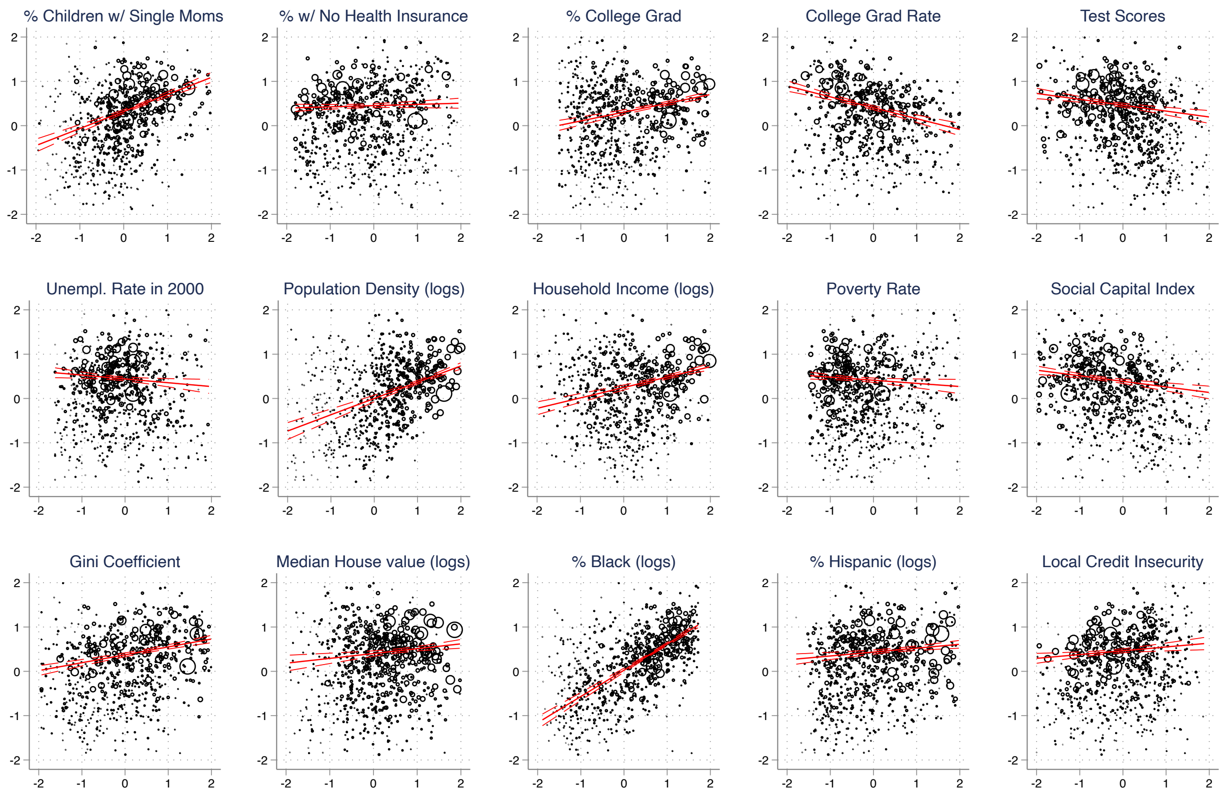


Notes: These figures display correlations of the household fixed effects, $\alpha_i$, as estimated using equation (5) and collapsed at the CZ level, with CZ-level social and economic characteristics.

Figure D.12: Event Study—Large Income Losses (Primary Filers of Ages 55-59)

(a) Frequency of Withdrawals

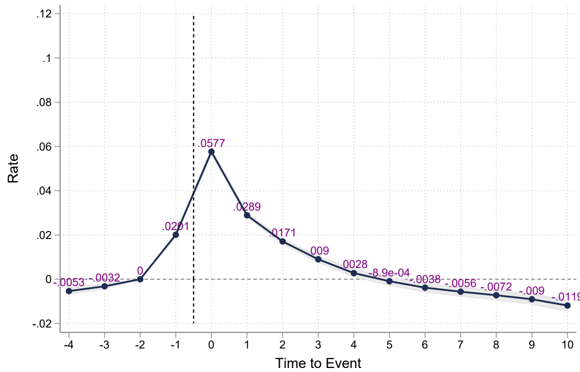


(b) Withdrawn Amounts

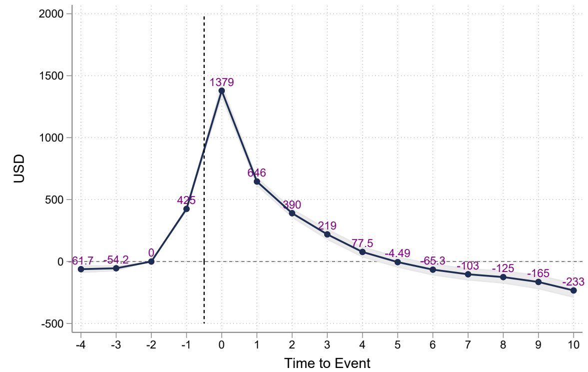


Figure D.13: Movers Analysis (Primary Filers of Ages 55-59)

(a) Withdrawal Frequency: Baseline

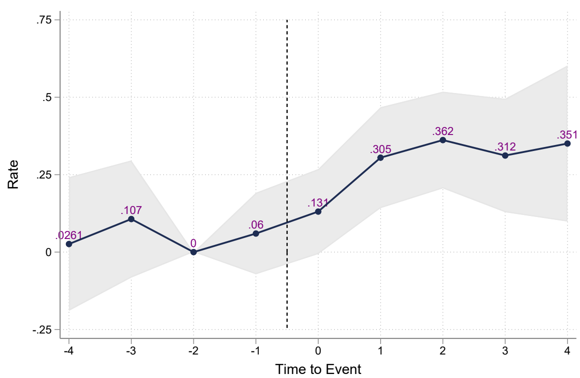


(b) Withdrawal Frequency: Flexible Controls

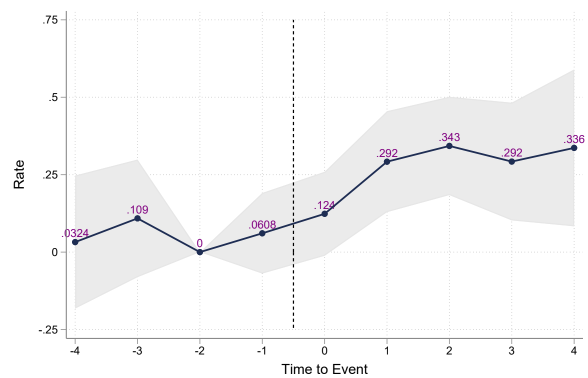


(c) Withdrawn Amounts: Baseline

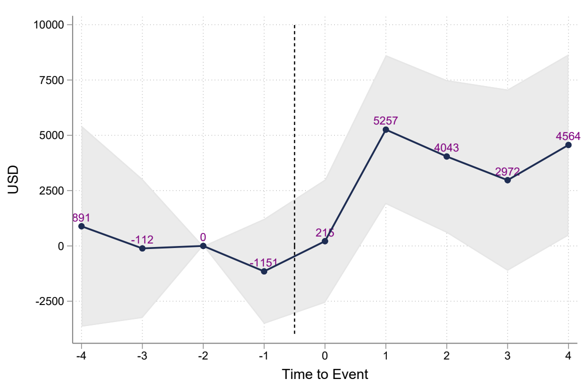


(d) Withdrawn Amounts: Flexible Controls

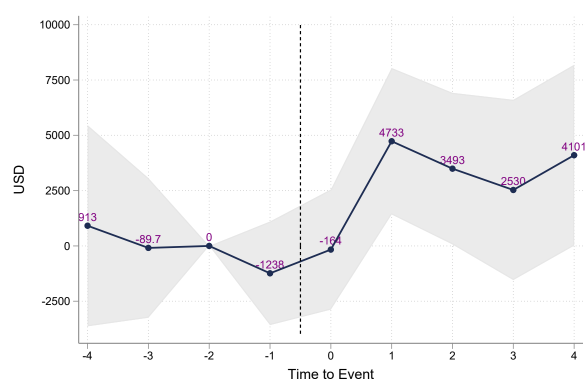


Notes: In these figures, we repeat the main analysis that pertains to households with a primary filer of ages 45-59 but when we constrain the sample to the age range 55-59 to focus on households near the statutory age of 59.5 when withdrawals become non-penalized. We provide the event study of a large 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.14: Withdrawal Behavior During the Great Recession (Primary Filers of Ages 55-59)

(a) Frequency of Withdrawals

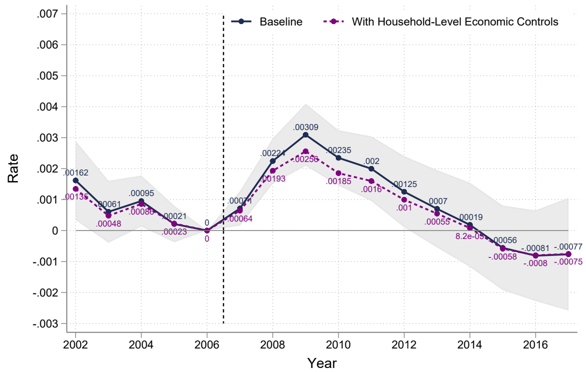


(b) Withdrawn Amounts

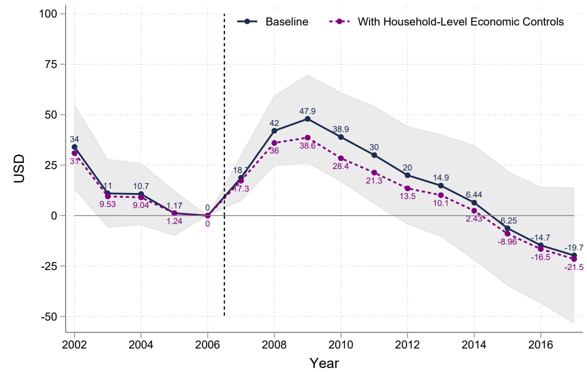


Notes: In this figure, we repeat the main analysis of the withdrawal behavior during the Great Recessions, when we constrain the sample to the age range 55-59 to focus on households near the statutory age of 59.5 when withdrawals become non-penalized.