

Measuring Valuation of Liquidity with Penalized Withdrawals*

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

Abstract

We use penalized withdrawals from retirement savings accounts as a revealed-preference tool to document three findings on American households' valuation of liquidity. First, local supply of credit explains over 30 percent of the nationwide differences in the valuation of liquidity across labor markets. Second, locations severely affected by the Great Recession displayed large increases in the valuation of liquidity, with spillovers in local credit tightening accounting for three-thirds of the effect. Third, Black households rely more on self-insurance from penalized withdrawals, consistent with lower access to formal credit markets. Overall, our findings imply sizable welfare gains from richer policy targeting.

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

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

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

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

1 Introduction

In a world with no borrowing constraints, households should perfectly smooth their marginal utility of consumption over time, where any remaining variation should be purely driven by aggregate shocks and permanent changes in households' income or consumption profiles. In practice, households are imperfectly insured as they do not have access to sufficient liquidity from formal or informal lending institutions (Parker 1999; Johnson et al. 2006). As a result, the degree to which marginal utility today differs from expected marginal utility tomorrow—or, simply, the *valuation of liquidity*—would typically vary across households at any point in time. Such variation implies resources are misallocated, lending the possibility of welfare gains from directing funds to households with higher valuations of liquidity or from helping them borrow in equilibrium from households who have lower valuations today.

To harvest these welfare gains, we need to empirically identify differences in households' valuation of liquidity. This is, however, a difficult task due to two major challenges. First, the valuation of liquidity is not directly observable and relying instead on consumption data is problematic beyond common issues of data availability and measurement. Indeed, detecting fluctuations in consumption is neither sufficient nor necessary to infer that a household is imperfectly insured: preferences themselves may change over time, possibly as a function of economic and life circumstances, and lead consumption to fluctuate. Second, the valuation of liquidity is an equilibrium object determined not only by household-specific shocks, which affect their demand for funds, but also by the available supply of credit to households. As a result, observing shocks to income, or even directly to demand for liquidity, is not sufficient to characterize the equilibrium valuation of liquidity which is the welfare object of interest. Doing that would further require detailed knowledge of the available supply of credit for *each* household in any given time since, beyond market-level economic conditions, the supply of credit is itself a function of household-specific factors (such as credit scores).

In this paper, we offer a revealed-preference approach that overcomes these challenges and allows us to characterize how the valuation of liquidity by American households varies across time and space. We do so by leveraging a simple insight: households who are willing to take up pricey borrowing reveal a high valuation of consumption today versus consumption tomorrow—that is, a high valuation of liquidity. The logic rests on the simple notion that observing a household purchasing a good at a given price—here, borrowing at a given interest rate—implies that the household values the good by at least as much. In practice, implementing our revealed-preference approach requires a credit product that entails two characteristics: (1) wide availability to households (to allow for a comprehensive analysis and for households to reveal their preferences); and (2) observable price (to serve as the

benchmark against which preferences are revealed). In addition, a product with uniform pricing has important value added: it allows comparable assessments across time (to assess effects of changing economic conditions), across space (to assess variation across geographic locations), and across observable types of households (to assess disparities in access to credit). We show that penalized withdrawals from retirement savings accounts serve as a common credit product that is close to this ideal. It is widely available to households and it has an observable and constant marginal price (the 10 percent penalty). Based on this idea, we use U.S. administrative tax records from 1999-2018 to characterize American households' valuation of liquidity and its variation, carefully paying attention to its equilibrium nature.

To guide the empirical analysis, we develop a simple analytical framework to formalize the idea that penalized withdrawals can be used as a revealed-preference tool to characterize the valuation of liquidity. Specifically, we use the now-standard two-asset heterogeneous agent model to provide a clear structural interpretation for the key empirical objects that we measure as follows. First, we show that the *frequency of penalized withdrawals* is a symptom of a household having a sufficiently high marginal valuation of liquidity in equilibrium. Second, we use the model to introduce the concept of *penalized liquidity*, which is the amount of liquidity that must be provided to a group of households to insure them enough to keep their marginal valuation of liquidity bounded by the penalty amount. In turn, penalized liquidity provides a simple money metric to quantify the degree to which households are formally under-insured. Finally, the model clarifies the importance of designing an empirical strategy that considers both demand and supply drivers of liquidity as the marginal valuation of liquidity is fundamentally an equilibrium object.

As in any revealed-preference approach, our analysis relies on households' ability to optimize since we use their behavior to back out their preferences. Reassuringly, we provide several pieces of evidence on penalized withdrawals supporting the view that households are indeed optimizing on this margin. Specifically, we find that households withdraw only infrequently and that typical penalized withdrawals are not linked to account closures (so that households are at an interior solution). As we discuss, these patterns are consistent with households optimally deciding to use their retirement savings accounts to access liquidity when needed and less with some leading behavioral interpretations (specifically, myopia, mental accounting, and narrow bracketing) driving the observed patterns in penalized withdrawals. Accordingly, the penalized amount withdrawn maps to our concept of penalized liquidity, providing us with a simple yet robust assessment of the degree to which households lack formal insurance.

We then proceed to our core contribution of empirically characterizing Americans' valuation of liquidity and its determinants. We begin by investigating how the equilibrium

valuation of liquidity at the household-level is affected by plausible shocks to their demand for funds. This serves as a first step in our analysis, where we empirically corroborate the underpinnings of our model that households use penalized withdrawals to mitigate liquidity needs. Previous work studied how leakages from retirement accounts increase after household shocks (Goodman et al. 2021), but here we focus on penalized distributions as guided by our theory as the key to reveal the valuation of liquidity. We study major financial household events—unemployment and large income declines—that could lead to sudden increases in the demand for liquidity, and we find that these adverse shocks clearly lead to sudden and persistent increases in penalized withdrawals. For example, a 30 percent decline in household income leads to a 9.5 percentage points (pp) jump in the propensity to make a penalized withdrawal and to an increase of approximately \$2,000 of needed penalized liquidity. We find that even households with plausibly large wealth holdings are imperfectly insured and see their valuation of liquidity spike as a result of income shocks, although, as expected, to a lesser extent. Importantly, using the recently-developed (and highly-validated) race imputation in the IRS tax records (Cronin et al. 2023, Fisher 2023), we provide novel evidence that Black households rely on penalized withdrawals to a higher extent than White households. As an example, following unemployment, their take-up of penalized withdrawals spikes by 35 percent more than it does for White households.

As we show in the model, perfect credit markets imply that the supply of funds is horizontal at the risk-free interest rate, so that any shift in demand should be purely absorbed by changes in the amount borrowed. Therefore, our event study findings that valuation changes when demand shifts further corroborate that households face a meaningfully inelastic supply of credit, which motivates the investigation of supply-side determinants of the valuation of liquidity. Household-level supply of liquidity could be determined both by the local environment to which a household is exposed (including formal institutions and social/informal support) and by household characteristics that govern access to credit (such as the household’s credit score). We leverage variation in the average annual share of households that make penalized withdrawals across commuting zones (CZs) to study the determinants of the supply of credit. Specifically, we leverage estimation strategies that rely on household migration across labor markets in the U.S in two steps.

In the first step, we employ a movers design (Finkelstein et al. 2016) that allows us to quantify the share of the total geographic differentials across CZs attributed to location itself. We find clear changes at the time of the move that then balance out with a high degree of persistence, which imply that permanent location characteristics strongly pass through to household withdrawals. Indeed, one of our key findings is that place effects can explain about a third of the overall spatial differences in penalized withdrawals across the country.

We provide a series of investigations that support this conclusion, by corroborating the validity of the movers design and by studying alternative candidate explanations (changing economic conditions, tax optimization, and learning) for which we find limited or no support.

In the second step, we estimate the location and household fixed effects with the AKM model (Abowd et al. 1999) that relies on similar identifying assumptions. We then explore correlates with regional differences using CZ-level social and economic characteristics, which allows us to shed light on potential drivers of the systematic heterogeneity across households and locations in the access to credit, i.e., in the supply of liquidity. Corroborating our credit supply interpretation, we find a strong correlation with the location component (and not with the household component) when considering measures of local credit insecurity and median home values (as high home values can provide collateral). We then also investigate correlations with the percent of Black residents in a given community, where we find a persistent relationship with the household fixed effects, suggesting that households in Black communities indeed have limited access to credit. Moreover, using the household fixed effects and the race imputation, we find that, regardless of their location and with similar economic circumstances, households with a Black primary-earner systematically rely on self-insurance via early withdrawals by 30 percent (2.9 pp) more than households with a White primary-earner (whose average rate is 9.5 pp). Importantly, this implies that the race differentials we have uncovered in the event studies are tied to Black households' inaccessibility to cheaper formal means of credit.

We conclude by applying our tool to a study of the Great Recession, a leading episode of severe worsening in the supply of credit that allows us to study the dynamic role of location in the valuation of liquidity. Argento et al. (2015) studied aggregate nationwide patterns of potential leakages from retirement accounts around the Great Recession. As we are interested in place effects, we leverage local variation based on the degree to which a CZ had been affected (Yagan 2019), and we again focus the analysis on penalized withdrawals as guided by our model. We find clear novel evidence that more affected commuting zones, as measured by unemployment shocks, have seen a larger increase in penalized withdrawals. We decompose the overall effect into a direct component (driven by a household's income and employment status) and an indirect component (e.g., through local credit market spillovers) by flexibly accounting for household-level economic circumstances. We find that about two-thirds of the overall effect can be attributed to the indirect spillover component, suggesting that the impact of aggregate local unemployment on the valuation of liquidity operates mainly through a credit crunch, that is, through a decrease in the local supply of liquidity.

Overall, our work introduces penalized withdrawals as a powerful yet overlooked tool that carries rich information on the valuation of liquidity among American households and its het-

erogeneity. We find that local supply of credit is a key determinant of household valuation of liquidity, identifying an important channel by which location shapes behavior and welfare. This result provides strong motivation for enriching the targeting of social insurance from primarily household circumstances to also locations over time, which could generate promising welfare improvements. We additionally uncover important racial disparities in liquidity valuation via access to credit that cannot be explained by geography or household-level circumstances, giving guidance to policies that aim to reduce racial inequalities. Underscoring its practical relevance, our analysis shows that we can use penalized withdrawals as a dynamic tool to monitor the valuation of liquidity across localities or sub-populations to guide the design of the large and growing government programs that provide social insurance.

Related Literature. Interest in the valuation of liquidity spans several fields, including public finance and macroeconomics. It encompasses the analysis of insurance and capital market inefficiencies, liquidity constraints and households’ ability to smooth marginal utility, and the optimal design of social insurance (see, e.g., Zeldes 1989; Parker 1999; Souleles 1999; Johnson et al. 2006; Card et al. 2007; Chetty and Finkelstein 2013). To this broad literature, we make three main contributions.

Our first contribution is to propose and validate a new tool to assess household valuation of liquidity—a crucial input for model calibrations and the design of optimal policies—that overcomes major identification and measurement challenges. First, consumption is hard to measure: data are usually partial, measurement requires accounting for the flow value of durable goods, and it necessitates the use of economies of scale in household production technology.¹ Second, a fundamental challenge in welfare evaluations is the need to translate observed household behaviors to normative values by estimating households’ preferences, in a way that flexibly allows for preferences to be heterogeneous and state-dependent (Landais and Spinnewijn 2021).² Our approach overcomes these two challenges by relying on a margin—the choice of making a penalized withdrawal—which is both readily measurable and directly reveals information on a household’s valuation of liquidity. As such, it inherently reflects underlying heterogeneity and is robust to any preference specification and state dependence, i.e., it freely allows household preferences to change in response to shocks or to changes in household conditions, such as parenthood and aging.

Our second contribution is to offer a novel comprehensive analysis of the anatomy of

¹Recognizing the significant challenges involved in analyzing consumption data, recent work has developed new methods based on studying labor supply behavior (e.g., Shimer and Werning 2007; Chetty 2008; Landais 2015; Hendren 2017; Fadlon and Nielsen 2019).

²See also discussions in, e.g., Finkelstein et al. (2009, 2013), Chetty and Finkelstein (2013), and Coyne et al. (2019).

the valuation of liquidity across the U.S., identifying the underlying driving forces of its variation. In doing so, we contribute to several sub-strands of the economic literature. First, we contribute to the growing influential work showing the key role of location in determining well-being in the U.S., from education, to earnings and intergenerational mobility, to health (e.g., Chetty and Hendren 2018a,b; Finkelstein et al. 2016, 2021; Card et al. 2023). We find that place effects can explain a third of the spatial gaps in Americans’ valuation of liquidity, providing novel evidence for how access to credit and liquidity—a key economic input into household well-being over the life cycle—is shaped by a location.³ Second, we contribute to the important work on racial disparities in economic outcomes in the U.S. (e.g., Bayer and Charles 2018; Chetty et al. 2020; Derenoncourt and Montialoux 2021; Derenoncourt et al. 2021; Bartscher et al. 2021). We provide evidence that, regardless of their location and with similar economic circumstances, Black households systematically display higher valuation of liquidity. Notably, this evidence provides new empirical support for the notion that Black families in the U.S. have more limited access to alternative channels of credit and are systematically underserved in the credit market.⁴ Third, we contribute to the extensive work on the Great Recession (e.g., Chodorow-Reich 2014; Chodorow-Reich et al. 2019; Yagan 2019). We offer novel insights into the effects of the Great Recession by providing new clear evidence on the important dynamics of the local availability of credit and valuation of liquidity and by showing a quantitatively large role for a market spillover effect.

Finally, our third contribution is to provide a new set of moments on the relationships between economic shocks and the valuation of liquidity that are informative for the emerging quantitative macro literature with heterogeneous agents (e.g., Krueger et al. 2016; Kaplan et al. 2018; Auclert 2019; Auclert et al. 2020; Laibson et al. 2021). Our results strongly support the notion that even wealthy households may be liquidity constrained (Kaplan et al. 2014), and they offer a new set of targeted moments and external validation tests for quantitative exercises. Targeting directly our moments on the valuation of liquidity—that is, on distortions in the Euler equation—has the key advantage over moments on consumption and income in that it is robust to different preference specifications and state dependence. This is even more important given the recent evidence that preference heterogeneity is central to properly account for the joint distribution of household level changes in income and consumption (Parker 2017; Aguiar et al. 2020).

³In this respect, Keys et al. (2020) analyze geographic variation in financial distress (focusing on collections, defaults, and bankruptcy), which could offer a look into some of the particular channels by which the variation in the valuation of liquidity in the U.S. that we analyze could be explained.

⁴Our findings are in line with recent results from Ganong et al. (2020), who study the consumption responses to typical labor income shocks and find higher elasticity for Black and Hispanic households, suggesting racial disparities in consumption smoothing.

Structure of the paper. Section 2 discusses the institutional details of penalized withdrawals and describes our data. Section 3 introduces the conceptual framework that formalizes the link between penalized withdrawals and household valuation of liquidity as well as guides the empirical analysis. We then turn to our main empirical analysis: Section 4 studies household-level events and the valuation of liquidity; Section 5 studies local market supply of liquidity and households’ access to credit based on spatial analysis; and Section 6 studies the evolution of local valuation of liquidity during the Great Recession. Section 7 discusses policy implications of our results, and Section 8 concludes.

2 Institutional Background, Data, and Preliminary Facts

We begin by describing how penalized withdrawals work institutionally, introducing our dataset, and explaining how we measure penalized withdrawals for the U.S. population. We then empirically describe the extent to which households rely on penalized withdrawals from retirement savings accounts as a source of short-term liquidity.

2.1 Institutional Setting

Many financial savings instruments require that money is held for a specified period of time or until a certain date. These include Health Savings Accounts (HSAs), Certificates of Deposits (CDs), and, most prominently, retirement savings accounts—either employer-sponsored 401(k)s or private Individual Retirement Accounts (IRAs). Within these retirement accounts, holders may withdraw funds “early” but must pay a penalty when doing so. Specifically, in the U.S., holders of retirement accounts must pay a penalty of 10 percent, above and beyond their income tax liability, for withdrawals that occur prior to age 59.5. The presence of this penalty provides us with the underlying basis of our approach as we will elaborate.

Some early withdrawals are excepted from tax penalties based on the reason for withdrawal. An exception is granted for the following events: account rollovers (e.g., across employers or from 401(k) to an IRA upon a job separation); permanent disability; death of account holder (allowing spouses to withdraw with no penalties); funds used for higher education; unreimbursed medical costs over 10 percent of the household’s adjusted gross income (AGI); first time home purchase; and separation from employment for those over age 55.⁵

To put penalized withdrawals in context, it is useful to describe the different ways in which U.S. households can access credit in the short run. Using data from the 2009 TNS

⁵For more details, see IRS website at <https://www.irs.gov/retirement-plans/plan-participant-employee/retirement-topics-tax-on-early-distributions>.

Global Economic Crisis survey, Lusardi et al. (2011) examine households' ability to come up with \$2,000 within 30 days if the need arises. They find that penalized withdrawals are indeed perceived by many households as a relevant liquidity tool.⁶

2.2 Data

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

Data sources and sample construction. We use U.S. administrative tax records based on a 10 percent random sample of U.S. tax filers from 1999-2018 aggregated to the household level. Specifically, we select 10 percent of individuals based on the last 4 digits of their Social Security Number (SSN). We then pull tax records for taxpayers who report those SSNs on Form 1040 (income tax return) for either the primary filer or the spouse. In cases where spouses indicate that they are married filing separately, we combine their data to build a single household return comparable to those married filing jointly. Once we have constructed these households, we create a consistent panel for them throughout our data's time range. We enrich the data from income tax returns with data from information returns filed by third parties (e.g., Form W-2, Form 1099-R, etc.).

We restrict our sample to households who have an individual in the age range 45-59 as a primary filer and who have a retirement account. The age range condition allows us to focus on prime-age households who likely have retirement accounts and for whom this tool is more relevant. We identify households as having a retirement account in a given year if up to that year (within our sample period of 20 years) they report making a contribution to a 401(k) or an IRA account on Form W-2 or Form 5498, or if they have outstanding balances in IRA accounts as reported on Form 5498. Our core sample consists of approximately 10.5 million households.

Variable definitions. The key outcome we study is penalized withdrawals from retirement savings accounts (401(k)s/IRAs) prior to age 59.5 with a penalty of 10 percent. We observe whether households took a distribution based on Form 1099-R (Box 1). We know whether or not the distribution was subject to the 10 percent penalty based on the distribution code reported on Form 1099-R (Box 7).

While there are several codes that correspond with penalized withdrawals, it is possible that a distribution coded as penalized on Form 1099-R is not ultimately penalized. For exam-

⁶The 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 possession (18.8), liquidate retirement investments even if penalty is required (11.1), pawn assets (7.7), friends (7.4), unsecured loan (7.1), home equity line of credit (HELOC)/second mortgage (4.3), payday/payroll advance loan (3.6), liquidate investments (2.3), sell home (0.4).

ple, account administrators might not know the reason for the account owner’s withdrawal, and, without any additional input from the account owner, they might mark a distribution as penalized. However, if the account owner qualified for an exception but did not report this exception to the fund manager, they would not necessarily be liable for the 10 percent penalty. There are also some exceptions that are correctly reported with the code “no known exception” but may not be subject to the 10 percent penalty, including withdrawals for unreimbursed medical expenses. As we explain in detail in Appendix A, we can account for these cases using Form 5329. Taxpayers who receive a Form 1099-R indicating the 10 percent penalty on early distributions are able to claim an exception on Form 5329 if one applies but was not accounted for on the original Form 1099-R.

For the household’s economic circumstances, we use Form 1040. We define the household’s overall income as the household-level Adjusted Gross Income (AGI) minus the amount of penalized withdrawals. AGI includes earnings, capital income, retirement income, and taxable Social Security benefits. Labor supply outcomes are based on earnings, where we define employment as having positive earnings in a given year.

Additional information we use includes the following. We gather annually-reported location information based on the address provided on Form 1040. We extract information on outstanding IRA balances from Form 5498, which indicates the fair market value of all IRA accounts (Box 5). This value includes all investments in the account at year end on December 31, where account trustees and custodians are responsible for ensuring that all IRA assets (including those not traded on established markets or not having a readily determinable market value) are valued annually at their fair market value. We use Form 1099-G (unemployment benefits) to define unemployment events. We extract capital income from Schedule D of Form 1040 when we study variation by household capital income, and we extract employer ID (EIN) from Form W-2 to identify job switch events.

Finally, we use the race and Hispanic origin imputations that leverage administrative tax data using the methodology developed and described in Fisher (2023). This methodology uses information on a taxpayer’s name, location at a given time, family characteristics, and income variables to predict race and ethnicity. Dummy variables for race and Hispanic origin are then created based on which estimated probability is highest for each taxpayer.⁷

⁷Imputations based on first and last name and geography are widely used in social sciences and in economics more specifically (recent applications include, e.g., Baron et al. 2023 and Hofstra et al. 2020). The methodology we use further improves on previous methods by including tax variables in the imputation procedure. This procedure has been tested with promising results, especially with respect to imputing the probabilities of being Black or Hispanic (see, e.g., Cronin et al. 2023 and Costello et al. 2024).

2.3 Key Facts on Penalized Withdrawals

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

Fact 1: Most households have retirement accounts. Appendix Figure D.1 shows the prevalence of retirement savings accounts across U.S. 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).⁸

Fact 2: Penalized withdrawals are widely used but infrequently. Next, panels (c) and (d) of Appendix Figure D.1 show that penalized withdrawals are widely used by households throughout the age and income distributions. Almost 10 percent of households within our age group make a penalized withdrawal in any given year. Penalized withdrawals are prevalent across the age distribution, but they fall, as expected, after age 55, when separation from employers becomes an expected 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

⁸The 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 U.S. We use data from waves 7-14, which cover the years 2004-2018, and focus on households with primary respondents between the ages of 45-59 for whom we can identify an account type (DC or DB) or whether the household reported not having an account. Among these households, we calculate that 14,392 have at least one defined contribution account, which amounts to a share of 84.14%. We note that the HRS is a representative sample of overall households in the U.S., whereas we focus on tax filers and thereby exclude non-filers who have less resources and could be expected to have accounts at lower rates. Indeed, in Appendix Figure D.3 we find that overall prevalence rates shift downwards moderately when non-filers are included, with an average account prevalence rate of 83.8% over ages 45-59.

are not concentrated among a few households, but are a prevalent liquidity tool across the whole population. Panel (e) of Appendix Figure D.1 shows that almost half of all households observed for 15 consecutive years in our sample take a penalized withdrawal in at least one year. Moreover, the typical household withdraws infrequently, consistent with the hypothesis that households use penalized withdrawals as a tool to access liquidity when the need arises. Finally, panel (f) of Appendix Figure D.1 shows among households who made a withdrawal in some period, the distribution of subsequent years within our data frame that the household made additional withdrawals.⁹ The figure displays a large mass at zero, consistent with penalized withdrawals reflecting temporary financial constraints that require short-run liquidity.

Fact 3: Withdrawn amounts are sizable, yet accounts are not fully depleted.

Panel (a) of Appendix Figure D.2 shows the CDF of the dollar amounts of penalized distributions. The typical withdrawal is approximately \$5,000. Importantly, penalized withdrawals are usually not associated with an account closure and they deplete only a relatively small fraction of the available funds. Here, we leverage the fact that the data include outstanding balances for IRA accounts. We look at households who have an IRA account at time $t - 1$ and who make a penalized withdrawal from an IRA account between periods $t - 1$ and t . Panel (b) of Appendix Figure D.2 shows that the share of households who deplete funds is consistently below half throughout the account balance distribution and that it is much lower, as expected, among households with non-trivial amounts in their accounts. Second, in panel (c) of Appendix Figure D.2 we plot the CDF of the ratio of penalized IRA distributions out of balances for households that do not fully deplete their accounts: the median withdrawal depletes approximately 25 percent of outstanding IRA balances. Overall, the evidence shows that most households are within an interior solution with respect to their withdrawal decision margin.¹⁰ This evidence is consistent with the interpretation that penalized distributions are the 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 for the use of penalized withdrawals as a revealed preference tool.

⁹We 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).

¹⁰This 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.

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. 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.

Summary. Taken together, these four facts provide evidence that households use penalized withdrawals as a mean to mitigate short-run needs for liquidity. This evidence accordingly motivates the use of penalized withdrawals as a revealed-preference tool to characterize the needs and valuation of liquidity across American households. In Appendix B, we further address two potential concerns with our approach. First, 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. Accordingly, we complement our data with information on premature withdrawals among American families from the HRS.¹¹ Second, any revealed-preference approach relies on the assumption that agents are maximizing choices on the margin investigated. Importantly, on top of having described how the preliminary evidence we provided here is consistent with this notion, we further discuss in detail in the Appendix B how this evidence is, reassuringly, inconsistent with leading alternative behavioral explanations—in particular, narrow bracketing (e.g., Thaler 1999), mental accounting (e.g., Read et al. 1999), or myopia/present bias (e.g., Laibson 1997; O’Donoghue and Rabin 1999)¹²

3 Conceptual Framework

We next develop a simple conceptual framework with two goals. The first goal is to formalize the idea that penalized withdrawals can be used to measure households’ valuation of liquidity. This will provide the mapping between withdrawal behavior and valuation with an explicit layout of the underlying assumptions. The second goal is to motivate our empirical analysis by illustrating how the valuation of liquidity is an equilibrium object determined by

¹¹Despite small samples, the key benefit from doing so is that households are asked to provide the reasons they withdrew funds prematurely. In Appendix B, we first show that the average amounts of withdrawals, once we focus on a comparable sample, are aligned with our administrative data. Then, we show that households report using early withdrawals to finance concurrent expenditure needs or repay outstanding debt. These results thus corroborate the indirect evidence provided from the tax data that early withdrawals are a signal of liquidity needs.

¹²Of course, while the evidence is inconsistent with these behavioral explanations governing the results, they could still naturally play a role.

both the local supply of credit and households' demand for liquidity.

3.1 Model Setup

Household i lives in region z and chooses consumption over the life cycle. The household earns income in each period, $y_{i,t}$, which can be used for consumption, $c_{i,t}$, or saved in either a liquid asset or a retirement savings account for future consumption. The household also receives an additional share φ of earnings directly deposited into the retirement savings account. Finally, in each period the household experiences a liquidity shock, $\varepsilon_{i,t}$, drawn from a distribution $F(\varepsilon)$.¹³ This shock captures consumption needs that affect the disposable income—e.g. an unexpected health bill.¹⁴

To finance consumption beyond current income flows, the household can borrow liquid assets in the financial market by paying the risk-free rate r and a premium $\rho_{i,z}(b_{i,t})$, which is household specific (i), location specific (z), and increases in the amount borrowed at time t , $b_{i,t}$.¹⁵ Households can also withdraw from the retirement savings account, but this entails the possible payment of a penalty.¹⁶ Specifically, if the withdrawal is done before a statutory retirement age (denoted as time t^*), there is a marginal penalty rate τ , so that only $1 - \tau$ dollars are available for consumption of each dollar withdrawn. Due to the penalty, τ , we refer to the retirement savings account as the illiquid account.

We denote the balances in the liquid and illiquid accounts at the beginning of period t by $a_{i,t}$ and $k_{i,t}$, and we let $\Delta a_{i,t}$ and $\Delta k_{i,t}$ represent the net flows across time periods within these accounts. The total value of borrowed liquid funds, defined by $b_{i,t}$, is equal to $-\Delta a_{i,t}$ if the household already had no liquid assets at time $t - 1$ ($a_{i,t-1} < 0$) and it is equal to the max between 0 and $-a_{i,t}$ otherwise.

We let the flow utility, $u(c_{i,t}; h_{i,t})$, be indexed by state vector $h_{i,t}$ that is both household-specific and time-specific. This vector could capture, for example, the whole history of shocks until time t (excluding the current shock $\varepsilon_{i,t}$) and any other characteristic that may change the utility from consumption, such as marital status or fertility choices. This flexibility allows for state dependence in preferences, hence permitting that the demand for liquidity could be driven by household shocks that directly affect preferences for consumption (e.g.,

¹³ $F(\varepsilon)$ is constant over time and across individuals.

¹⁴The last assumption is useful for constructing a tight mapping to the empirical specifications that is provided in Lemma 4.

¹⁵The cost $\rho_{i,z}(b_{i,t})$ should be interpreted as a perceived shadow cost of funds, which captures the expected optimal borrowing choices across alternative sources of funds that are available for and known to the household. Therefore, it is a reduced-form measure of the local household-specific supply of credit as perceived by household i .

¹⁶In the model, we abstract from tax optimization considerations for withdrawing. This model simplification is justified by our empirical analysis in the later sections, in which we show that controlling empirically for tax motives has only a very modest effect on withdrawal behavior.

severe health shocks) rather than indirectly via available income.¹⁷

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

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

subject to

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

where β is the discount factor. It is important to emphasize that we index the value functions both by time and by the household state vector, $h_{i,t}$, since the value of the problem varies both across time and across households even conditional on the liquid and illiquid asset stocks $(a_{i,t-1}, k_{i,t-1})$.

3.2 Valuation of Liquidity and Penalized Withdrawals

We next define our main object of interest and characterize how it behaves in our setting. Recall that we aim to assess the value that a household assigns in equilibrium to moving a dollar from tomorrow to today—i.e., the household's valuation of liquidity.

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

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

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

To build intuition on $\theta_{i,t}(c_{i,t}; h_{i,t})$, Lemma 1 shows its value in different benchmark scenarios. Proofs for all Lemmas are provided in Appendix C.

¹⁷This level of flexibility highlights a strength of our approach: we directly reveal valuation of liquidity from household behavior without having to rely on structural assumptions that would map behaviors to preferences.

Lemma 1: Benchmarks for the Equilibrium Valuation of Liquidity. *The equilibrium valuation of liquidity for household i at time t and for any history $h_{i,t}$ satisfies the following:*

1. *If credit markets are perfect—i.e., $\rho_{i,z}(b) = 0$ for all b —then $\theta_{i,t}(c_{i,t}; h_{i,t}) = 1$.*
2. *If the household saves in the liquid asset—i.e., $\Delta a_{i,t} > 0$ —then $\theta_{i,t}(c_{i,t}; h_{i,t}) = 1$.*
3. *If the household borrows an amount b from the liquid asset—i.e., $b > 0$ —then $\theta_{i,t}(c_{i,t}; h_{i,t}) = \frac{1}{1-\rho'_{i,z}(b)} = 1 + \frac{\rho'_{i,z}(b)}{1-\rho'_{i,z}(b)}$.*

A few comments are in order. First, note in Part 1 of Lemma 1 that, even in the presence of perfect credit markets, the household's consumption may fluctuate over time. This could happen, for example, as a function of changes in household circumstances, consumption needs, or preferences. Yet, the valuation of liquidity is equal to 1 since the household's Euler equation must be undistorted as households can save and borrow with no limits at the same interest rate r . This observation illustrates the challenge in common analyses of fluctuations in consumption to recover information on fluctuations in marginal valuation, which we overcome with our revealed-preference approach. Second, Part 2 of Lemma 1 shows that households who save also have valuation of liquidity equal to 1, suggesting that, just as in the case of perfect credit markets, they are perfectly smoothing marginal utility over time in expectation at time t .¹⁸ Finally, Part 3 of Lemma 1 indicates that, if a household borrows from the liquid asset, the valuation of liquidity in equilibrium is a function of the shadow cost of funds. It shows that if we could observe for each household both their borrowing/savings behavior and their shadow cost of an additional unit of borrowing, $\rho'_{i,z}(b)$, we would be able to directly pin down their valuation of liquidity from the data. Intuitively, if you are willing to borrow at a high interest rate, your marginal valuation of funds today must be as large. In practice, however, researchers cannot directly observe $\rho'_{i,z}(b)$, as this value must take into account all the available means of credit that each household could have access. This requires knowing the particular interest rates each household would face in the credit market (on any possible form of credit, such as different credit cards and bank loans), and it more broadly requires knowing the shadow value of the household's informal means of credit (such as borrowing from relatives).

The first contribution of this paper is to offer an indirect revealed-preference approach that conveys important information on our object of interest, $\theta_{i,t}(c_t; h_{i,t})$. We bypass the

¹⁸It is important to note that this result in expectations does not imply that the realized ex-post marginal utility is constant over time. For example, in the presence of precautionary savings due to prudence (Carroll 1997), $\theta_{i,t}(c_t; h_{i,t})$ would be equal to 1 as long as a household saves, yet consumption (and, accordingly, the realized marginal utility) may fluctuate over time as a function of household income shocks.

identification hurdles we highlighted by relying on an existing widely-available credit product—penalized withdrawals from retirement savings accounts—whose marginal price is uniform. In turn, it allows us to conduct comparisons across households, time, and space. We formalize this result in Lemma 2 below.

Lemma 2: Equilibrium Valuation of Liquidity at Withdrawal. *Consider a household i that makes at time $t = t^* - 1$ a penalized withdrawal from the illiquid account without fully exhausting it, i.e. $\Delta k_{i,t} > 0$ and $k_{i,t} > 0$.¹⁹ The equilibrium valuation of liquidity for this household satisfies*

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

The intuition for this result is straightforward as it considers the last period before the statutory “retirement” age. At this stage, a household can freely withdraw funds from the illiquid account next period. However, if the household still decides to withdraw today paying a penalty of τ , then we can infer that the household does not have access to cheaper credit and that it values a marginal dollar today by as much as the penalty for withdrawing.

This simple insight applies more broadly but needs to be refined. In general, if a household makes a penalized withdrawal, we can infer that their valuation of liquidity is strictly larger than 1, but it does not have to be equal to $1 + \frac{\tau}{1 - \tau}$. As we explain formally in Appendix C, observing a penalized withdrawal implies that a household is willing to pay a price $1 + \frac{\tau}{1 - \tau}$ to move funds from the illiquid account to the liquid account (whereas $\theta_{i,t}(c_t; h_{i,t})$ is the ratio of valuations of liquid dollars). This price would equal the valuation of liquidity if in the next period the value of a dollar in the liquid and illiquid accounts are identical. Importantly, this condition is satisfied as long as the household is not expected to make any more penalized withdrawal until the retirement age t^* ; a condition the documented empirical withdrawal patterns in Section 2.3 are strongly consistent with.²⁰ In the general case within our model, illiquid dollars are less valuable than liquid dollars, which implies that $1 \leq \theta_{i,t}(c_{i,t}; h_{i,t}) \leq 1 + \frac{\tau}{1 - \tau}$ for households that are withdrawing at time $t < t^* - 1$ and that do not fully deplete their retirement savings account ($\Delta k_{i,t} > 0$).

This last result highlights another key aspect of our setting: as long as a household has

¹⁹The assumption that $k_{i,t} > 0$ guarantees that the Euler equation is satisfied with equality, hence that $\theta_{i,t}(c_{i,t}; h_{i,t}) = 1 + \frac{\tau}{1 - \tau}$. If the household fully depletes the retirement savings account, which we have shown to be rare in the data, then we would get $\theta_{i,t}(c_{i,t}; h_{i,t}) > 1 + \frac{\tau}{1 - \tau}$. Throughout the characterization we focus on the typical empirically relevant case in which $k_{i,t} > 0$ holds.

²⁰Consistent with this observation, our results are almost identical when we repeat our main analysis focusing on a restricted sample of households in the age range 55-59, who are thus “just” before the statutory “retirement” age of 59.5. See Appendix Figure D.15.

funds available in their retirement account, their valuation of liquidity must be bounded. Since penalized withdrawals provide households with a mean to self-insure against liquidity shocks, access to funds without using them implies that the excess valuation of a dollar is limited by its price (i.e., the penalty). Lemma 3 formalizes this point.

Lemma 3: Valuation of Liquidity with Access to Penalized Withdrawals. *Consider a household i at time t . For any history $h_{i,t}$ such that $k_{i,t} > 0$, we have that*

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

with strict inequality if $\Delta k_{i,t} = 0$.

While penalized withdrawals allow households to self-insure, they also lead to retirement savings leakages which may undermine old-age financial security. It is thus natural to define a notion of penalized liquidity, which is the total leakage from the illiquid account that allows a subset of households to self-insure. This is a measurable empirical object, and it provides an intuitive measure of how well the formal credit market works.

Definition 2: Penalized Liquidity. *Consider a set of N households denoted by \mathcal{S} . Their average “penalized liquidity” from time t to time t' , $\Lambda_{t,t'}(\mathcal{S})$, is given by the average of the sum of their penalized withdrawals from the illiquid account:*

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

Lemma 2 and Lemma 3 taken together provide the foundations on which the empirical analysis is grounded. Observing households making penalized withdrawals gives information on their equilibrium valuation of liquidity (Lemma 2), but also not observing households making penalized withdrawals when they have available funds implies that their valuation of liquidity is relatively low (Lemma 3). The frequency and magnitude of penalized distributions, as captured by $\Lambda_{t,t'}(\mathcal{S})$, therefore measure how well-insured a given set of households is. Specifically, it answers the question: how much liquidity needs to be injected for a given set of households \mathcal{S} to keep them sufficiently well-insured so that their valuation of liquidity is bounded by $1 + \frac{\tau}{1 - \tau}$? Evaluating how $\Lambda_{t,t'}(\mathcal{S})$ varies across different sets of households provides a natural way to compare their differential access to financial markets and other potentially informal means of insurance.

3.3 Liquidity Demand and Supply

The valuation of liquidity is an equilibrium object which depends both on the household's demand for funds and on the supply of credit in the local market. Accordingly, we now describe how either shocks to demand or supply could lead to increases in households' valuation of liquidity and trigger a withdrawal. This characterization provides a structural interpretation of the empirical analysis.

To proceed, we define the demand and supply curves. We express them as the inverse relationship between the quantity of funds that are either demanded by households or supplied to households at a given "price," which in our setting is the effective marginal interest rate. Accordingly, we define the supply function to already take into account the household's maximization decision to access credit from the cheapest source. For this reason, it is convenient to focus on the last period before t^* , so that the marginal cost of a penalized withdrawal is simply $\frac{1}{1-\tau}$. We also focus on the empirically relevant set of households that have available illiquid funds, i.e. $k_{i,t} > 0$, at all times.

A household's demand for liquidity states the marginal interest rate, $\mathbb{D}_i(\bar{b})$, at which the household demands to borrow (or save) an amount of funds equal to $\bar{b} \equiv [\Delta k_{i,t} + \Delta a_{i,t} + \tau \Delta k_{i,t} \mathbb{I}_{(\Delta k_{i,t} < 0)} \mathbb{I}_{(t < t^*)} + \rho_{i,z}(b) \mathbb{I}_{(b > 0)}] - (1+r) a_{i,t-1}$ to finance consumption. Note that \bar{b} includes funds from any source, both liquid and illiquid, and thus we use a different notation relative to b , which captures only funds borrowed from the liquid account.

Definition 3: Demand for Liquidity. *The demand for liquidity is given by a function $\mathbb{D}_i(\bar{b})$ which solves the equation*

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

where $x_{i,t} \equiv (1-\varphi)y_{i,t} + (1+r)a_{i,t-1} - \varepsilon_{i,t}$ is the cash-on-hand for household i at time t , net of the liquidity shock.

By definition, \bar{b} is the amount of funds borrowed to satisfy consumption, which is just $c_{i,t} = x_{i,t} + \bar{b}$. The demand for liquidity can be thought of as the level of consumption that households would choose if they borrow at a specific interest rate given by $\mathbb{D}_i(\bar{b})$. Accordingly, it is a decreasing curve which provides the marginal interest rate at which a certain amount of funds is demanded. It inherits the decreasing slope from the fact that the equilibrium valuation of liquidity is itself decreasing in the level of consumption c for any utility function that satisfies decreasing marginal utility.

Definition 4: Supply of Liquidity. *The supply of liquidity is given by a function $\mathbb{S}_i(\bar{b})$ which satisfies*

$$\mathbb{S}_i(\bar{b}) \equiv \begin{cases} 1+r & \text{if } \bar{b} \leq 0 \\ \frac{1+r}{1-\rho'_{i,z}(\bar{b})} & \text{if } \bar{b} > 0 \text{ and } \rho'_{i,z}(\bar{b}) < \tau. \\ \frac{1+r}{1-\tau} & \text{if } \Delta k_{it} < 0 \text{ (and thus } \rho'_{i,z}(\bar{b}) = \tau) \end{cases}$$

The supply of liquidity provides the marginal interest rate the household must pay in order to borrow when the household wishes to access funds \bar{b} (from either the liquid and illiquid accounts).²¹ Due to the way we defined the supply and demand curves, at the equilibrium \bar{b} it has to be that $\mathbb{S}_i(\bar{b}) = \mathbb{D}_i(\bar{b})$. Moreover, again by definition, at the equilibrium amount of borrowing \bar{b} we have that $\theta_{i,t}(c_{i,t}; h_{i,t}) = \frac{\mathbb{D}_i(\bar{b})}{1+r}$. Hence, the valuation of liquidity is equal to 1 if and only if the shadow cost of capital at the demanded amount of funds \bar{b} is equal to the risk-free interest rate.

The upper panels of Figure 1 (panels (a)-(c)) illustrate the demand and supply curves and the resulting equilibrium valuation of liquidity in three cases. In panel (a), we consider perfect credit markets: the supply of funds is horizontal at the risk-free interest rate, which implies that any shift in demand is purely absorbed by changes in the amount borrowed. In this case, liquidity needs are perfectly insured and the Euler equation is undistorted. Panel (b), instead, considers the case of imperfect credit markets for a household that does not have access to a retirement savings account, i.e., $k_{i,t} = 0$. The supply curve is upward sloping, due to the convex cost $\rho_{i,z}(b)$. Borrowing funds is costly, which leads the household to borrow less, and have a higher valuation of liquidity in equilibrium, θ_1 . Lastly, panel (c) introduces the possibility of making a penalized withdrawal from a retirement account, and thus corresponds to the general case from Definition 4. As long as households have funds in the illiquid account, they can access illiquid funds paying the penalty τ , thus facing marginal cost of liquidity given by $\frac{1+r}{1-\tau}$. In the example shown in the figure, the possibility to make a penalized withdrawal decreases the equilibrium valuation of liquidity ($\theta_2 < \theta_1$), which captures their ability to self-insure, and it correspondingly also decreases the amount borrowed from the liquid account ($b_{3,a} < b_2$). The additional funds for self-insurance, i.e., the penalized liquidity (captured by $b_{3,b} - b_{3,a}$), are withdrawn at a penalty τ .

Next, we define the probability that a household makes a penalized withdrawal—which is the central object empirical analysis—and we show how it is affected by shocks either to the demand for liquidity or to the supply of funds.

²¹The definition of $\mathbb{S}_i(\bar{b})$ spans all the support of \bar{b} since we are assuming that $k_{i,t} > 0$.

Lemma 4: Probability of Making a Penalized Withdrawal. *There exists a threshold value of the liquidity shock, $\bar{\varepsilon}_{i,t}(h_{i,t})$, such that a household i with history $h_{i,t}$ makes a penalized withdrawal at time t if and only if $\varepsilon_{i,t} \geq \bar{\varepsilon}_{i,t}(h_{i,t})$. With history $h_{i,t}$, the probability that the household makes a penalized withdrawal, $\mathcal{P}(h_{i,t})$, is therefore given by*

$$\mathcal{P}(h_{i,t}) = 1 - F(\bar{\varepsilon}_{i,t}).$$

The lower panels of Figure 1 (panels (d)-(f)) illustrate how shifts to either the supply or demand curves could trigger a penalized withdrawal. Panel (d) first characterizes a household that borrows from liquid funds. Then, in panel (e), we subject the household to an upward shift of the credit supply schedule—i.e., a credit crunch.²² We model a credit crunch as an upward rotation in the additional cost of borrowing $\rho_{i,t}(b)$: if there are tighter conditions in the financial market, the household needs to pay a higher cost for each dollar borrowed. Of course, a credit crunch does not affect the availability of funds in the retirement savings account. As a result, the household cuts drastically the borrowing from the credit market (from b_1 to $b_{2,a}$) and partially compensates for this reduction by making a penalized withdrawal (of the amount $b_{2,b} - b_{2,a}$). The credit crunch leads to an increase in the valuation of liquidity, but access to the illiquid account bounds the effect via the self-insurance mechanism of penalized withdrawals.

Finally, in panel (f) we consider an increase in demand, for example, due to an unexpected decline in income, $y_{i,t}$. The overall amount of borrowing increases (from b_1 to $b_{3,b}$), triggering a penalized withdrawal (with penalized liquidity of an amount $b_{3,b} - b_{3,a}$). Facing a larger demand for funds, the household relies on their illiquid account to avoid being subject to a large increase in the cost of funds. Once again, observing a withdrawal from the illiquid account signals a relatively higher valuation of liquidity, but the possibility to make penalized distributions from retirement savings account allows households to self-insure and bounds fluctuations in valuation of liquidity in equilibrium.

The figure also highlights how, for both supply and demand shocks, the amount withdrawn—i.e., the “penalized liquidity” we defined above—is a convenient measure for the size of the shock. This discussion is summarized in Lemma 5 below.

Lemma 5: Determinant of Penalized Withdrawals. *Consider shifters, $\varphi_{i,\mathbb{D}}$ and $\varphi_{i,\mathbb{S}}$, to the demand and supply curves such that: $\mathbb{D}_i(\bar{b}) = \varphi_{i,\mathbb{D}} + \mathbb{D}(\bar{b})$ and the cost of funds is $\rho_i(b) = \varphi_{i,\mathbb{S}}\rho(b)$. For any household i and history $h_{i,t}$, the probability of making a penalized*

²²For simplicity, we are assuming in this illustration that the household has access to a large amount of illiquid funds $k_{i,t}$, so that they cover the whole support of borrowing b shown in the figure.

withdrawal, the penalized liquidity, and the expected equilibrium valuation of liquidity (with expectations taken over the shock $\varepsilon_{i,t}$) increase in both $\varphi_{i,\mathbb{D}}$ and $\varphi_{i,\mathbb{S}}$.

From Model to Data. Before turning to the data, we summarize how this section has set the foundation for the empirical analysis. We note that we are not going to bring the model structurally to the data. Rather, the model provides the foundation for the empirical regressions, and it maps them onto objects that have a clear theoretical meaning. This is done in the following way.

Equation (1) and Lemma 5 lead and motivate our empirical analysis. Specifically, equation (1) motivates us to study the choice of penalized withdrawals which provides a signal that reveals a high valuation of liquidity. Lemma 5 shows that a high valuation of liquidity could be driven by either demand forces or supply forces, and it motivates us to leverage the empirical variation to unpack both drivers. Accordingly, *Section 4* will first study household level events. By focusing on variation within households over time, we identify life-cycle events that affect the *demand* for liquidity.²³ This maps to \mathbb{D}_i at time t and primarily uses variation in $y_{i,t}$. Then, *Section 5* will instead study the determinants of the local *supply* of credit, which maps to \mathbb{S}_i and is governed by $\rho_{i,z}(b)$. It will unpack these supply-side determinants into components that are specific to locations (Γ_z) and components that are specific to households (α_i). Postulating that the shadow cost of capital is a function of a household component and a location component—specifically, $\rho_{i,z} = \alpha_i + \Gamma_z$ —a simple movers design that includes a rich set of controls to account for variations in the demand for liquidity (to the extent possible) can unpack the local supply of credit. We will additionally study how these components vary as a function of observable characteristics. Finally, *Section 6* will consider an extended version of the model, in which we conceptually allow the location component of the supply of credit to vary over time, i.e., $\Gamma_{z,t}$. We will study how $\Gamma_{z,t}$ has been dynamically affected by the Great Recession and accordingly translated to households' equilibrium valuation of liquidity. We now turn to our empirical analysis.

4 Household Events and Valuation of Liquidity

In this section, we study how household-level adverse economic events lead to changes in households' valuation of liquidity. This traces out how shifts in the demand for liquidity lead to movement along the supply function and to changes in the equilibrium valuation of liquidity (as illustrated in panels (a)-(c) of Figure 1. This section offers an empirical

²³For robustness, we later verify in Appendix Figure D.5 that the results remain almost identical when we conduct this analysis conditional on households who stay in the same commuting zone around the specific events that we study.

verification of our model by corroborating that shocks indeed lead to take-up of penalized withdrawals, and it quantifies the degree to which household-level shocks shift the valuation of liquidity. It also provides benchmark estimates at the household level against which we will compare the market-level shocks that we study later.

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

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

where $y_{i,t}$ is either an indicator for a penalized withdrawal for household i at time t , or the amount withdrawn in dollars (including zeros); $x_{i,t}$ is a full set of age fixed effects for the primary-filer and (cyclical) calendar year fixed effects; and α_i are household fixed effects.²⁴ We let $r(i, t)$ denote the year relative to the event timing for household i at time t , so that $I_r = \mathbb{I}_{r(i,t)=r}$ represent a set of relative time indicators. We take the baseline year to be -2 to capture changes in trends that could happen toward the realization of the event.²⁵ We plot β_r around different events to trace the evolution of households’ withdrawal behavior, where we are interested in capturing behavioral responses to the realization of the event as well as in anticipation of the event to evaluate the full dynamics of the valuation of liquidity. We use this empirical framework to study the dynamics of household penalized withdrawals around unemployment and income changes.

4.1 Unemployment Event

We define an unemployment event as the first period we observe at least one of the household members receiving unemployment benefits.

Results. Figure 2 plots the event study coefficients β_r , estimated when the outcome is either an indicator for making a penalized withdrawal (in Panel (a)) or the amount withdrawn (in Panel (b)). As the event approaches, we see an increase in penalized withdrawals that is then followed by a large spike at the year of the event. Through the lens of our model,

²⁴In the samples on which we run these regressions, we include all households to help with identification of non-event coefficients and we accordingly add to $x_{i,t}$ a dummy for households who do not experience an event.

²⁵We choose -2 as baseline year because the year -1 coefficient can incorporate anticipation but also potential effects of the onset of an event, given the annual frequency of the data at the end of a calendar year and the defined timing of the event. For example, households who experience an event of a large decline in income (which we take to be at least 20 percent) between the end of period -1 and period 0 would be assigned a “large income decline” event at 0 , but the process of a decline in income could have already (and likely) began throughout year -1 .

these patterns imply an increase in the valuation of liquidity, which in turn maps to underinsurance of unemployment shocks.²⁶

To consider magnitudes, panel (a) shows that the share of households with sufficiently high valuation of liquidity to trigger a penalized withdrawal doubles at the onset of the event: at baseline (in $t = -2$) households make penalized withdrawals at a rate of 9.9 percentage points (pp), which increases at the unemployment event (in $t = 0$) by 10.4 pp. Panel (b) shows that households make additional penalized withdrawals of an average of approximately \$1,600 in that same year. This amount exactly maps to the concept of *penalized liquidity* we defined in the theoretical framework; that is, it is the average amount of liquidity injection that is needed to keep the marginal valuation of liquidity at or below the withdrawal penalty. Two observations are useful to interpret the magnitude of the penalized liquidity. Comparing the withdrawn amounts to the decline in household income around the event, we find that penalized withdrawals compensate on average for less than 8 percent of the average income decline (which is approximately \$20,900 at the onset of the event, see panel (a) of Appendix Figure D.4). This suggests that, on average when including zeros of non-withdrawing households, the households in our sample are relatively well-insured. However, the relatively small average masks substantial heterogeneity when we consider comparing households that are induced to withdraw to those who do not. Dividing the point estimates of the effects on amounts and on take-up at the year of the event, we see that the typical household that makes a penalized withdrawal takes out about \$19,000 from their retirement account.²⁷

Heterogeneous Effects. Next, we explore how different types of households vary in the degree to which they are insured against unemployment by studying the extent to which they rely on withdrawals as self-insurance upon the event. Doing so allows us to shed new light on the household-level determinants of the valuation of liquidity and to further validate our approach by exploring how our results relate to prior work. It is also directly informative

²⁶The findings go hand-in-hand with the important literature on the effects of unemployment on earnings and consumption, which has shown large declines in consumption in the short run with lingering effects on earnings in the long run (See, e.g., Sullivan and Von Wachter 2009; Kolsrud et al. 2018; Schmieder et al. 2018; Ganong and Noel 2019; Gerard and Naritomi 2021.). Unlike these assessments of income or consumption, our investigation of the valuation of liquidity is robust to the possibility that preferences are themselves affected by employment status. This could be the case, for example, if employment leads to different consumption needs, such as the classic substitution to cooking meals at home while unemployed and the reduction in time and monetary costs involved in commuting. Indeed, a key advantage of our framework is that it freely allows for state dependence in preferences for any shock we would consider.

²⁷This back-of-the-envelope calculation should be interpreted with caution since it assumes that the unemployment event does not affect the amounts withdrawn for those households that would have taken a penalized contribution even in the absence of the event.

for targeting households with a higher valuation of liquidity along observable/measurable dimensions, such as age, wealth, and location. It is useful to emphasize, however, that the heterogeneity analysis is only correlational. We use the frequency of withdrawals that has a direct interpretation, whereas amounts may also reflect other aspects such as differences in average household income.

We first study how the effects of unemployment may differ by race. We plot the event study where we split households by the imputed race of the primary earner. Panel (c) of Figure 2 provides plots for households with primary earners that are either Black or White. We find that withdrawal frequencies are significantly higher among Black households at the onset of unemployment, despite the fact their income decline is lower (as shown in panel (b) of Appendix Figure D.4). That is, Black households are much more likely to rely on self-insurance from retirement accounts when they experience an unemployment event, which in turn signals more limited access to cheaper ways of securing liquidity among Black families.

In panels (d)-(e) of Figure 2 we look at how withdrawals upon the unemployment event vary by household characteristics, focusing on the age of the primary filer and household capital income.²⁸ We find that, prior to age 55 when job separations become eligible for non-penalized withdrawals, there is a gradient with respect to age in withdrawal probability upon unemployment.²⁹ This gradient is consistent with the idea that older households, who had more time to accumulate a buffer stock of savings, may be more resilient to shocks and have lower liquidity needs.³⁰ When studying responses by household capital income as a measure for non-housing wealth, we split households into those with negative, zero, or positive capital income, and bin households with positive capital income into four groups. In the non-negative range, the evidence is closely consistent with our revealed-preference interpretation of making penalized withdrawals: households with access to alternative financial means that can provide liquidity have lower increases in the valuation of liquidity as they have lower remaining residual risk.³¹ While the differences across households are sizable, it is worth

²⁸We present heterogeneity results from a regression that simultaneously includes these categories, as well as additional household level controls: home ownership (an indicator based on claimed property tax or mortgage interest deductions from Schedule A and mortgage information from Form 1098), a dummy for whether the primary filer is married, the number of dependents, and the average household income based on information from all years within our data range.

²⁹As our analysis is at the household level, it is possible to have withdrawals with penalty after age 55 of the primary filer if the secondary filer could be younger than 55. In Appendix Figure D.6 we report the unemployment event study estimates for observations with primary filers younger than 55 to account for the change in withdrawal rules at that age.

³⁰Another possible interpretation in line with our conceptual framework is that the gradient is driven by inter-temporal substitution as households are approaching an age in which they can withdraw without paying a penalty.

³¹The somewhat lower levels of withdrawals at the region of negative capital income could be reflective of the notion that these households still have better access to capital markets as compared to households that

noticing that even households in the top quartile of capital income display a meaningful increase in penalized withdrawals when experiencing the event. Considering that those households have an average capital income close to \$40,000 per year, this result strongly corroborates the notion that even wealthy households might be liquidity constrained (e.g., Kaplan et al. 2014).

4.2 Income Changes

We next look at income changes. First, we look at large income losses as an event, which we define 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 (a) and (b) of Figure 3 display the event study coefficients for frequency and amounts of withdrawals. We again find spikes at the time of the event, where the magnitudes imply that households, on average, need approximately \$2,000 in penalized liquidity to keep their valuation of liquidity within bounds. As for unemployment, Black households display a greater increase in withdrawals upon large income losses relative to White households (see panel (e) of Figure 3).

We then refine the analysis of variation in income changes by studying households' withdrawals as a function of the deviation of their income flow from their average income across our data period. We split households by whether a member of the household switched jobs that year because job changes themselves, as displayed in Appendix Figure D.8, lead to increased take-up.³² Panels (c) and (d) of Figure 3 reveal a clear pattern. First, there is a strong gradient with respect to income losses, so that larger income losses lead to a higher frequency of withdrawals. This supports our model where penalized withdrawals are used as means of short-run self-insurance. Second, we find stark asymmetry around zero, where behavior completely flattens in the income increases domain. This is consistent with the self-insurance hypothesis, and it rules out alternative explanations. Specifically, it is inconsistent with the notion of strategic withdrawals for tax purposes as driving the penalized withdrawals behavior. In that case—that is, if households withdrew funds with penalty differentially as their current marginal tax rate changes—we would expect to observe some

are unlikely involved in capital markets altogether (those with zero capital income).

³²This could be driven by several factors such as increases in the valuation of liquidity in the transitional period, as well as alternative considerations such as salience or simply choosing to cash out if the balance is negligible. Recall that we exclude account rollovers from employer-sponsored accounts, which are just mechanical transfers of funds and could be common upon job separation. That said, upon job separation, low balances below a certain threshold can be automatically paid out in cash to the departing employee, with thresholds of \$5,000 prior to 2005 and \$1,000 thereafter (accordingly amounts between \$1,000 and \$5,000 can be automatically rolled over into an IRA post-2005.) To account for negligible balances and these automatic passive penalized distributions, Appendix Figure D.7 replicates our event study analyses but where the outcome variables are indicators for taking penalized withdrawals that are higher than given thresholds.

degree of a gradient in the entire income changes domain. Third, even households experiencing positive income changes make non-negligible penalized distributions, with average withdrawal amounts close to 1 percent of average income. This suggests that, as long as households are maximizing on the margin (as the evidence pointed to in Section 2.3), the equilibrium valuation of liquidity is not only driven by income shocks, but it is also driven by changing consumption needs through expenditure shocks (such as unobservable health shocks or child-related expenses, which were captured by the shock ε_{it} in our model). As a result, even perfectly insuring households against negative income shocks would still be insufficient to achieve marginal utility smoothing over time.

5 Local Supply of Credit and Valuation of Liquidity

Section 4 has shown that households are only imperfectly able to self-insure, so that they face a limited supply of credit (as shown in Figure 1). Therefore, it is natural and important to investigate the determinants of the local supply of credit and how it varies across households and locations.

In an ideal case, we would want to know the full schedule of credit supply and the corresponding shadow costs of capital for each and every household, which would require data on all household-specific formal and informal available credit instruments. With such data, it would then be possible to study how credit access varies across household characteristics and locations, providing information on how to optimally target liquidity injections across household types and their location of residence. To the best of our knowledge, such rich database is essentially impossible to build, as we rarely observe the full set of credit vehicles available to each household (e.g., informal lending among family members) and it is conceptually challenging to quantify the cost of borrowing. In this section, we overcome this challenge by leveraging our revealed-preference tool, combined with empirical designs that are able to dissect the determinants of the local supply of credit.

We focus on spatial variation as geographic localities capture the local financial and social credit environments to which households are exposed. We proceed in three related steps: (i) we show that there are large differences in penalized withdrawals across locations; (ii) we use a standard movers design to quantify the share of the spatial differences casually attributed to locations; and (iii) we study how the estimates of location effects and of average household effects within a location correlate with a battery of locality-level observables to offer insights on aspects that can potentially govern or mediate the valuation of liquidity.

Large Variation across Regions. Panel (a) of Figure 4 plots the average annual share of households that make a penalized withdrawal by commuting zones (CZs). We find large

differences across regions, with a mean of 9.8 pp and a standard deviation of 1.7 pp.³³ This variation could capture differences across locations in either the demand for liquidity (for example, due to a higher unemployment rate) or in the local supply of credit. The local supply of credit could itself be driven either by characteristics of the households who live in that location (such as average household credit score) or by “true” location effects due to the local environment to which a household is exposed. This environment includes traditional financial institutions (such as banks) as well as local social networks and informal support (such as religious organizations). These components can be decomposed using a standard statistical model with income controls and household and location fixed effects. This model, which we describe below, will guide our analysis in the entire section.

Statistical Model. We use the following statistical model for household behavior (adopted from Abowd et al. 1999 and Finkelstein et al. 2016 and adjusted to our setting):

$$(3) \quad y_{i,z,t} = \alpha_i + \Gamma_{z(i,t)} + x_{i,t}\lambda + \varepsilon_{i,t}.$$

In this specification, $y_{i,z,t}$ is an indicator equal to one if household i makes a penalized withdrawal in CZ z at time t ; α_i is a household fixed effect; $\Gamma_{z(i,t)}$ are location fixed effects determining the household’s outcome, where $z(i,t)$ indexes the location of household i in year t ; $x_{i,t}$ is a vector of potential time-varying controls, including indicators for age of primary filer, (cyclical) calendar year fixed effects, and household-level economic conditions. As is well known, this specification is identified off movers across CZs.

5.1 Movers Design Implementation

We next use a movers design to establish whether and to what degree the large spatial differentials we found are due to persistent characteristics of the local environment. Specifically, we directly analyze outcomes of households who have moved across CZs, and we use the difference in intensity of penalized withdrawal behavior across the household’s original location and new location as the source of variation. We discuss the identifying assumption and its validity below.

To proceed, we further develop the statistical model in equation (3) for households who switch locations in the following way (similar to Finkelstein et al. 2016). For household i ,

³³To address known biases in plug-in estimates of second moments due to sampling errors (see, e.g., Andrews et al. 2008), we estimate the standard deviations of location-level statistics based on a split-sample approach (as in, e.g., Finkelstein et al. 2021 and Card et al. 2023). Specifically, with a random sample split, we conduct the estimation on each subsample separately and assess the variation based on the covariance of estimates across subsamples.

whose location changed from z_0 to z_1 , we denote by Δ_i the difference in average propensity of taking penalized withdrawals between the destination CZ and the origin CZ: $\Delta_i \equiv y_{z_1} - y_{z_0}$, where $y_z \equiv E[y_{i,z,t}]$ is the average taken over all time periods and households in location z . Empirically, we include in these averages only households that are non-movers to attain “leave-out” means. Δ_i is the sum of the differences in the locations’ and households’ contributions to the observed share of withdrawals across households. We define $r(i, t)$ as the period relative to the household’s move, and we let $I_{r(i,t)>0}$ denote an indicator for time periods after the move. Lastly, we parameterize the model by letting the difference across locations that is attributable to location effects be $\theta \equiv \frac{\Gamma_{z_1} - \Gamma_{z_0}}{y_{z_1} - y_{z_0}}$. That is, the parameter θ captures the aggregate share of the overall differentials across CZs in making penalized withdrawals that is causally determined by location. With this parameterization, we then get the following equation for households who move:

$$y_{i,t} = \alpha_i + \Gamma_{z_0} + \theta I_{r(i,t)>0} \Delta_i + x_{i,t} \lambda + \varepsilon_{i,t},$$

where θ is our parameter of interest. It represents the average passthrough of the overall difference between the new and old location into households’ withdrawals in the years following the move.

Estimating Equation. A direct empirical analogue for the latter equation, which estimates the mean effects in the post-move years, takes a standard difference-in-differences form:

$$(4) \quad y_{i,t} = \mu_i + \theta \times Post_{i,t} \times \Delta_i + x_{i,t} \lambda + \varepsilon_{i,t},$$

where $\mu_i = \alpha_i + \Gamma_{z_0}$. Here, $Post_{i,t}$ is an indicator variable that equals 1 in the post-move years and equals 0 in the pre-move years.³⁴ We take an extended version of this equation to the data to allow for flexible dynamics by estimating the following event-study specification:

$$(5) \quad y_{i,t} = \mu_i + \sum_{r \neq -2} \beta_r \times I_r + \sum_{r \neq -2} \theta_r \times I_r \times \Delta_i + x_{i,t} \lambda + \varepsilon_{i,t},$$

where $I_r = \mathbb{I}_{r(i,t)=r}$ are indicators for time relative to the move. To be consistent with the previous section, our baseline period is taken to be two years prior to the move ($r = -2$). The event study specification in equation (5) allows us to test for parallel trends in the pre-move period (based on θ_r for $r < -1$) and to investigate dynamics in location effects in the post-move period (based on θ_r for $r > 0$). Robust standard errors are clustered at the

³⁴Note that the vector $x_{i,t}$ also includes the baseline variable $Post_{i,t}$ and that Δ_i is absorbed by the household fixed effect.

origin CZ level.

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

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

Interpretation and Robustness. Our interpretation of these results—guided by the conceptual framework from Section 3—is that, when households move to locations with worse local supply of credit, they have to rely more on penalized withdrawals for liquidity. We next provide a series of empirical checks that supports this conclusion. In particular, we investigate the validity of the movers design in identifying the location pass-through, we explore leading threats to identification, and we study the potential role of other explanations or mechanisms for the patterns we found. For these exercises we focus on our main specification that uses as dependent variable the indicator for a penalized withdrawal.

First, we check our identifying assumption of parallel trends: our design requires that households’ underlying trends in withdrawals do not systematically differ by Δ_i . The standard testable implication of this assumption is whether there are differential trends in the pre-move period across households with differential Δ_i . Reassuringly, even when we estimate

equation (5) for an extended window that runs from year -5 to $+10$ (on an unbalanced sample of households) we find that there are virtually no pre-trends (see panel (c) of Figure 5 as well as Appendix Figure D.11).

Another aspect to consider is that mover designs cannot account for shocks that simultaneously differ across households with varying treatment intensity Δ_i and align exactly with the timing of moves. We therefore ask: can the results be explained by differential changes to household economic conditions interacted with timing of move to differentially intense locations? Two pieces of evidence are inconsistent with this notion. Looking again at panel (c) of Figure 5, we see a high degree of persistence for up to 10 years in the estimates for the role of location, θ_r . This is in contrast to the effects of household-level shocks, which have been shown above to be transitory with clear dissipating dynamics. Similarly, Appendix Figure D.9 shows a comparable pattern of transitory dynamics for the move event itself, as captured by the “event study” coefficients β_r in the movers equation (5). These combined findings are hard to reconcile with patterns in passthroughs being driven by shocks aligned with the time of move.³⁵ We then directly account for household-specific economic conditions that could change around the move and, potentially, in a differential way across locations with varying degrees of withdrawal intensity. We run specifications that include a flexible set of (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. The results show that the estimates hardly change in terms of either dynamics or magnitudes. See panel (d) of Figure 5.

Second, it is could be possible in principle that the results are driven by households learning about withdrawals from peers when they move to a higher intensity location. As an initial observation, we note that this channel is inconsistent with the immediate jump upon the move and the limited dynamics thereafter (akin to discussion in Finkelstein et al. 2022). To further test the learning hypothesis directly, we focus on households who had already used this liquidity tool and made a penalized withdrawal in the pre-move periods. Albeit with less precision due to the additional constraint (and with increased noise in the longer horizon region where there are less households), panel (e) of Figure 5 shows the results are very similar, suggesting again that learning is not driving our findings.

Third, an additional explanation could be tax optimization, whereby households’ pe-

³⁵We note that the moderate decline in the estimates in the extended window of post-move years is attributable to attrition and return moves (see panel (a) of Appendix Figure D.10). They attenuate the persistence in the effects since we assign a household the same destination location for the entire post-move period, whether they subsequently moved or not, because these behaviors could be endogenous to the initial move. Panel (b) of Appendix Figure D.10 illustrates this point: when we scale the estimates by the share of movers that are still at the assigned destination, the dynamics flattens out.

nalized withdrawal behavior is governed by their marginal tax rate. We have already seen evidence inconsistent with this conjecture in the analysis of household events, where we find no gradient in the region of positive income changes. To further investigate it in the context of moves, we add controls for a location’s top marginal tax rate (that varies over state and time) flexibly interacted with time relative to the move. The small attenuation in estimates in panel (f) of Figure 5 suggests that tax motives might play at most a minor role.

5.2 Investigating Variation in Systematic Components

Within the large spatial variation in panel (a) of Figure 4, we have shown that location characteristics (Γ_z) explain approximately one third of the total variation. The remainder of the variation must be driven by spatial differences in households’ composition ($\alpha_i + x_{i,t}\lambda$). Recall that the passthrough estimates hardly change when we flexibly control for dynamic household economic conditions. This suggests that the remainder of the variation attributed to households’ composition is not the result of differences in their temporary demand for liquidity, captured by $x_{i,t}\lambda$; rather, it is likely a result of systematic household differences, captured by α_i .

Motivated by this evidence, we turn to investigate the potential drivers of these supply-side components. That is, we estimate equation (3) to provide us with separate estimates for: (1) the location fixed effects, Γ_z , which capture the local market-level supply of credit; and (2) the household fixed effects, α_i , which capture a household’s access to credit (irrespective of their location and of several household-level economic outcomes for which we control in $x_{i,t}$: unemployment, wage earnings, and gross income, with lagged, current, and lead values).

Panels (b) and (c) of Figure 4 first display maps of the estimated location effects (Γ_z) and the estimated household fixed effects (α_i) averaged within a CZ. There is meaningful geographic variation: the location effects have a standard variation of 1 pp and the CZ-level average household fixed effects have a standard deviation of 1.8 pp. We then investigate their correlations with a battery of CZ-level characteristics (taken from Chetty et al. 2016 unless noted otherwise). Panel (d) of Figure 4 reports all the normalized regression coefficients from a series of univariate OLS regressions, while Appendix Figures D.12 and D.13 include all the corresponding scatter plots. We focus here on discussing the characteristics which display particularly strong correlations with the estimates of our statistical model.

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

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

Household Effects. Household fixed effects most notably correlate with measures of racial composition: households who live in communities with a high share of Black residents are significantly more likely to make penalized withdrawals irrespective of their current location. How should we interpret this result? Withdrawals are a financial instrument that, conditional on having a retirement account, does not discern (or discriminate) across households of different social groups. Heterogeneity in relying on withdrawals across different types of households can therefore reveal their differential access to alternative means of credit. Households in Black communities reveal a high valuation of liquidity, suggesting they have more limited access to alternative channels of credit.

Through the lens of our model, the patterns indicate that households who reside in communities with a higher share of Black families have a systematically higher valuation of liquidity. As such, this result provides an important aspect that could help guide place-based policies aimed at increasing equity in access to credit, by providing a cursor for the direction in which liquidity should be injected. These disparities can be driven by the lack of access to credit by Black households or by the lack of access to credit by non-Black households who live in communities with a higher share of Black households. Three pieces of evidence, however, provide support that the lack of access to credit among Black families is the more likely explanation.

First, the location effects themselves are uncorrelated with the share of Black families in a community. It implies that, when randomly drawn from the population, a household who moves into an area with a high share of Black families would not see its own penalized withdrawals increasing. Instead, it means that limited access to credit follows Black families wherever they go.

Second, we repeat the entire analysis at finer geographic units; in particular, we run specification (3) at the 5-digit ZIP Code level and project α_i onto their associated ZIP Codes. We then use the resulting estimates to correlate the household fixed effects and the ZIP Code fixed effects with the share of Black households within the ZIP Code, controlling for CZ fixed effects. As shown in panel (d) of Figure 4, the point estimates are remarkably similar: the empirical correlations across CZs are very similar to those within CZs across

ZIP Codes.³⁶

Finally, we directly look at the correlation between our estimated household-level fixed effects and the primary earner’s imputed race. We provide this analysis in Table 1, which reports estimates based on underlying specifications of equation (3) that either exclude or include our flow economic controls. Indeed, we find (in column 2) that Black households systematically rely on penalized withdrawals more than White households. The mean of the household-specific fixed effects for Black households is 2.94 pp higher than the mean of the household-specific fixed effects for White households of 9.5 pp, that is, over 30 percent higher. The table also illustrates that the gap cannot be explained by household economic circumstances as captured by our flexible income and employment measures.

All in all, the findings strongly suggest that Black households experience strong limitations to accessing credit, which go above and beyond their financial circumstances (as captured by a battery of controls) and are irrespective of where they live (as isolated by the location fixed effects).³⁷

6 Valuation of Liquidity During the Great Recession

Localities can encompass both stable components, such as established institutions, and time-varying components, such as aggregate shocks and changing economic conditions. In Section 5, we have focused on characterizing the stable component of location in determining valuation of liquidity. In this final piece of our empirical analysis, we consider locations as an evolving entity, and we study their dynamic evolution during a leading episode that could have led to a severe worsening of local credit: the Great Recession.

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

$$(6) \quad y_{i,z,t} = \sum_{r \neq 2006, r=2000}^{r=2017} \beta_r \times I_r + \sum_{r \neq 2006, r=2000}^{r=2017} \theta_r \times I_r \times Treat_z + \Gamma_z + \alpha_i + x_{i,t} \lambda + \varepsilon_{i,t}.$$

As dependent variables, we consider either an indicator for a penalized withdrawal, or the total amount withdrawn in dollars (including the zeros). Then, turning to the independent variables, I_r are calendar year indicators running from years 2000 to 2017, where year 2006 is taken to be the baseline year; $x_{i,t}$ is a full set of primary-filer age fixed effects, (cycli-

³⁶We note that, due to spatial segregation, there is a lot of variation both across CZs and within CZs across ZIP Codes in the share of Black households (see Appendix Figure D.14).

³⁷We also note that similar patterns hold when we explore a location’s percent of children living with single mothers, who represent another classic example of economically-disadvantaged households. Our results suggest that single mothers may also have limited access to credit within their communities.

cal) calendar year fixed effects, and potential household-specific economic controls; Γ_z are commuting zone fixed effects; and α_i are household fixed effects. $Treat_z$ is the treatment intensity of location z in terms of unemployment shock. Specifically, we utilize the measures from Yagan (2019), who considers the change in a commuting zone’s unemployment rate between the years 2007 and 2009. Our parameters of interest are θ_r , which capture the relative change in behavior in a locality that is exposed to a 1 pp larger local unemployment shock.

Results. The solid lines in panels (a) and (b) of Figure 6 plot the θ_r estimates from equation (6) for either the frequency or the amounts of withdrawals. Locations that were about to be hit differentially by the Great Recession, were on similar trajectories prior to the event (in support of the research design). Then, commuting zones more severely affected by the Great Recession, as measured by unemployment increases, have seen a larger increase in penalized withdrawals. The response peaks around the height of the Great Recession with an increase in frequency of local penalized withdrawals of 0.403 pp per 1 pp in local unemployment shock in year 2009, which results in a total increase of \$63.4 in penalized liquidity per household. Calculating the cumulative effects from 2007-2012, we find an overall increase of 1.47 pp in the probability of taking a penalized withdrawal and a corresponding overall increase of \$251 in penalized liquidity.

To make the magnitudes comparable with the results from Figure 2, we need to normalize the coefficients to show the effects of a 100 pp increase in unemployment, so that it is comparable to the individual event of becoming unemployed. Interestingly, with this normalization we find that the flow effect on making a penalized withdrawal of a locality-level unemployment shock is about 4 times as large as the direct effect of a household-level unemployment shock that we have estimated (40.3 vs. 10.4 pp). This suggests that the effect of a 1 pp of local unemployment on the valuation of liquidity is about one-quarter due to increases in household demand ($0.258=10.4/40.3$) and about three-quarters due to decrease in local supply. We reach a very similar conclusion using the results for the amount withdrawn in (multiplying the point estimate by 100, we get \$6,340 to be compared with \$1,590).

In light of these patterns, we break down the cumulative impact of the Great Recession into: (1) a direct effect, through, e.g., the specific household’s income and employment; and (2) an indirect effect, through market-level spillovers. We do so by flexibly accounting for household-level economic circumstances in estimating equation (6), where we specifically add as controls unemployment, wage earnings, and gross income, with lagged, current, and lead values, as well as their interactions with calendar year dummies. The estimates are reported in the dashed lines in panels (a) and (b) of Figure 6. Consistent with the previous

argument, we find that the indirect impact accounts for approximately three-quarters of the overall effect of the Great Recession. Our findings are therefore closely consistent with a tightening of local credit conditions for all workers in distressed locations. Overall, the Great Recession provides a leading example of how evolving local circumstances can have an important role on American households' need for and valuation of liquidity, not only directly but also through meaningful market spillovers.

7 Policy Implications: Discussion

Defined-contribution retirement savings accounts have become the backbone of the American retirement system. Indeed, Siliciano and Wettstein (2021) find that most older households with heads born between 1920 and 1940 had access to a defined-benefit (DB) pension plan, and that this share had dropped rapidly with the youngest Baby Boomers, born in 1965, having almost no access to DB plans. Instead, as we have shown above, the vast majority of households have access to defined contribution (DC) accounts and keep positive balances within those accounts. Over time, an increasing share of employers have been automatically enrolling workers in these plans (requiring employees to make an active decision to opt out), and all workers can establish and contribute to IRAs on their own. Moreover, in recent years, many states have taken steps to increase participation in retirement savings plans; for example, California, Oregon, and Illinois have adopted automatic enrollment IRA programs (auto-IRAs), under which employers not offering an employer-sponsored account to any of their employees must facilitate payroll deductions from workers' paychecks to be transferred to state-facilitated IRAs (Bloomfield et al. 2023). More recently, the passing of the Securing a Strong Retirement Act of 2022 (the SECURE Act 2.0) provides multiple changes whose goal is to increase retirement savings, specifically by expanding automatic enrollment in employer-provided retirement plans, simplifying rules for small businesses and easing restrictions on how employees could jointly offer a plan, and helping those near retirement save more for longer.³⁸ Thus, access to retirement savings accounts—and with it access to penalized withdrawals as a short-run liquidity tool—is wide and expected to grow even further. In turn, our conceptual contribution of how we can learn from Americans' choices of making penalized withdrawals about their valuation of liquidity is widely relevant and practical.

With that in mind, a main policy takeaway from our findings is that there could be meaningful welfare improvements from enriching the targeting of social insurance policies and sections of the tax code that affect liquidity. Richer policies could depend on households'

³⁸See <https://www.congress.gov/bill/117th-congress/house-bill/2954/text>.

specific economic conditions (addressing their demand for liquidity), on locality (e.g., by improving a community’s access to credit), as well as on time and local economic conditions (e.g., by an inter-temporal reallocation of the same funds). Our work stresses that we can indeed use penalized withdrawals as a practical dynamic tool to monitor the evolving local valuation of liquidity and guide such richer targeting.

There is a variety of ways in which this targeting could happen through the social insurance system or the tax code. Most immediately, the tax penalty itself could become a function of household-level, location-level, and aggregate-level economic conditions. Indeed, tax penalties are already waived in the tax code for several qualified household-level events (such as spousal death) that are believed to increase households’ liquidity needs. Moreover, Congress has recognized premature withdrawals as a potential avenue for liquidity and has adjusted the penalty price in the wake of major events that caused shocks to liquidity among American taxpayers. Specifically, localized exceptions have been offered in the past for some natural disasters, including Hurricane Katrina and, most recently, in 2020 Congress waived penalties on withdrawals of up to \$100,000 from qualified retirement accounts for COVID-19-related purposes. Our analysis points to welfare gains from such systematic price adjustments to the cost of funds in savings accounts through the tax code. For example, the tax penalty may be especially burdensome on lower-income taxpayers who already face relatively higher prices in credit markets, and so policymakers could consider tailoring the penalty amount to a taxpayer’s income level, especially around events predictive of penalized withdrawals.

Additionally, the tool developed in this paper can serve to identify targets for other location-level incentives aimed at equalizing access to financial services across communities. For example, a program similar to Empowerment Zones (EZs), which allowed businesses in economically distressed areas to receive employment tax credits, could be implemented to specifically target the financial services sector in financially underdeveloped communities as identified by our findings where many households have high liquidity needs.

Finally, our findings provide a precursor for potential welfare gains from new financial products and the coming regulation of these markets. With the large spatial variation in available credit that we have uncovered, easy-access financial technology solutions (FinTech) have the potential to reach households in need of credit who live in financial deserts with limited traditional credit means, allowing for more equitable access to credit nationwide.

8 Conclusion

This paper introduces conceptually and validates empirically penalized withdrawals from retirement savings accounts as a novel robust tool that carries information on households’

valuation of liquidity. We use this tool to characterize the anatomy of the equilibrium valuation of liquidity among American families based on IRS tax records and offer several sets of findings. First, we find that the local supply of credit can explain over 30 percent of the nationwide differences in the valuation of liquidity across labor markets. Second, analyzing the Great Recession, we find that aggregate local labor market shocks lead to large increases in the valuation of liquidity, where spillovers in local credit tightening account for three-thirds of the overall effect. Third, while we show that the use of penalized withdrawals for liquidity needs is pervasive, we provide novel findings that Black households rely on self-insurance from penalized withdrawals to a larger extent, as compared to White households with similar economic conditions and regardless of where they live. These results provide novel evidence suggesting that Black American families are systematically underserved by formal credit markets and have limited access to cheaper means of securing liquidity throughout the country.

References

- Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High Wage Workers and High Wage Firms. *Econometrica* 67(2), 251–333.
- Aguiar, M. A., M. Bils, and C. Boar (2020). Who Are the Hand-to-Mouth? Technical report, National Bureau of Economic Research.
- Amador, M., I. Werning, and G.-M. Angeletos (2006). Commitment vs. Flexibility. *Econometrica* 74(2), 365–396.
- Andrews, M. J., L. Gill, T. Schank, and R. Upward (2008). High wage workers and low wage firms: negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 171(3), 673–697.
- Argento, R., V. L. Bryant, and J. Sabelhaus (2015). Early Withdrawals from Retirement Accounts During the Great Recession. *Contemporary Economic Policy* 33(1), 1–16.
- Auclert, A. (2019). Monetary Policy and the Redistribution Channel. *American Economic Review* 109(6), 2333–67.
- Auclert, A., M. Rognlie, and L. Straub (2020). Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model. Technical report, National Bureau of Economic Research.
- Baron, E. J., J. J. Doyle Jr, N. Emanuel, P. Hull, and J. P. Ryan (2023). Racial discrimination in child protection. 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.
- Bloomfield, A., K. M. Lee, J. Philbrick, and S. Slavov (2023). How do firms respond to state retirement plan mandates? Technical report, National Bureau of Economic Research.
- Card, D., R. Chetty, and A. Weber (2007). Cash-on-Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market. *Quarterly Journal of Economics* 122(4), 1511–1560.
- Card, D., J. Rothstein, and M. Yi (2023, August). Location, location, location. Working Paper 31587, National Bureau of Economic Research.
- Carroll, C. D. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *Quarterly Journal of Economics* 112(1), 1–55.

- Chetty, R. (2008). Moral Hazard vs. Liquidity and Optimal Unemployment Insurance. *Journal of Political Economy* 116(2), 173–234.
- Chetty, R. and A. Finkelstein (2013). Social Insurance: Connecting Theory to Data. In *Handbook of Public Economics*, Volume 5, pp. 111–193. Elsevier.
- Chetty, R. and N. Hendren (2018a). The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects. *Quarterly Journal of Economics* 133(3), 1107–1162.
- Chetty, R. and N. Hendren (2018b). The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates. *Quarterly Journal of Economics* 133(3), 1163–1228.
- Chetty, R., N. Hendren, M. R. Jones, and S. R. Porter (2020). Race and economic opportunity in the united states: An intergenerational perspective. *The Quarterly Journal of Economics* 135(2), 711–783.
- Chetty, R., M. Stepner, S. Abraham, S. Lin, B. Scuderi, N. Turner, A. Bergeron, and D. Cutler (2016). The Association Between Income and Life Expectancy in the United States, 2001-2014. *Journal of the American Medical Association* 315(16), 1750–1766.
- Chodorow-Reich, G. (2014). The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008-9 Financial Crisis. *Quarterly Journal of Economics* 129(1), 1–59.
- Chodorow-Reich, G., J. Coglianese, and L. Karabarbounis (2019). The Macro Effects of Unemployment Benefit Extensions: a Measurement Error Approach. *Quarterly Journal of Economics* 134(1), 227–279.
- 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 (2019). Market Inefficiency and Household Labor Supply: Evidence from Social Security’s Survivors Benefits. Technical report, National Bureau of Economic Research.
- 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.
- 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.
- 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.
- Feldstein, M. (1985). The Optimal Level of Social Security Benefits. *Quarterly Journal of Economics* 100(2), 303–320.
- Finkelstein, A., M. Gentzkow, D. Li, and H. L. Williams (2022). What drives risky prescription opioid use? evidence from migration. Technical report, National Bureau of Economic Research.
- Finkelstein, A., M. Gentzkow, and H. Williams (2016). Sources of Geographic Variation in Health Care: Evidence From Patient Migration. *Quarterly Journal of Economics* 131(4), 1681–1726.
- Finkelstein, A., M. Gentzkow, and H. Williams (2021). Place-Based Drivers of Mortality: Evidence from Migration. *American Economic Review* 111(8), 2697–2735.
- Finkelstein, A., E. F. Luttmer, and M. J. Notowidigdo (2009). Approaches to Estimating the Health State Dependence of the Utility Function. *American Economic Review* 99(2), 116–21.
- 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.
- 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., 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.
- Goodman, L., J. Mortenson, K. Mackie, and H. R. Schramm (2021). Leakage From Retirement Savings Accounts In The United States. *National Tax Journal* 74(3), 689–719.
- Hamdani, K., C. K. Mills, E. Reyes, and J. Battisto (2019). Unequal Access to Credit: The Hidden Impact of Credit Constraints. *Federal Reserve Bank of New York*.
- Hendren, N. (2017). Knowledge of Future Job Loss and Implications for Unemployment Insurance. *American Economic Review* 107(7), 1778–1823.
- Hofstra, B., V. V. Kulkarni, S. Munoz-Najar Galvez, B. He, D. Jurafsky, and D. A. McFarland (2020). The diversity–innovation paradox in science. *Proceedings of the National Academy of Sciences* 117(17), 9284–9291.
- 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.
- 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. (2015). Assessing the Welfare Effects of Unemployment Benefits Using the Regression Kink Design. *American Economic Journal: Economic Policy* 7(4), 243–78.
- 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.
- Read, D., G. Loewenstein, and M. Rabin (1999). Choice Bracketing. *Journal of Risk and Uncertainty* 19(1-3), 171–97.
- 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.


- Shimer, R. and I. Werning (2007). Reservation Wages and Unemployment Insurance. *Quarterly Journal of Economics* 122(3), 1145–1185.
- Siliciano, R. L. and G. Wettstein (2021). Can the drawdown patterns of earlier cohorts help predict boomers behavior?
- Souleles, N. S. (1999). The Response of Household Consumption to Income Tax Refunds. *American Economic Review* 89(4), 947–958.
- 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

Demand and Supply of Liquidity

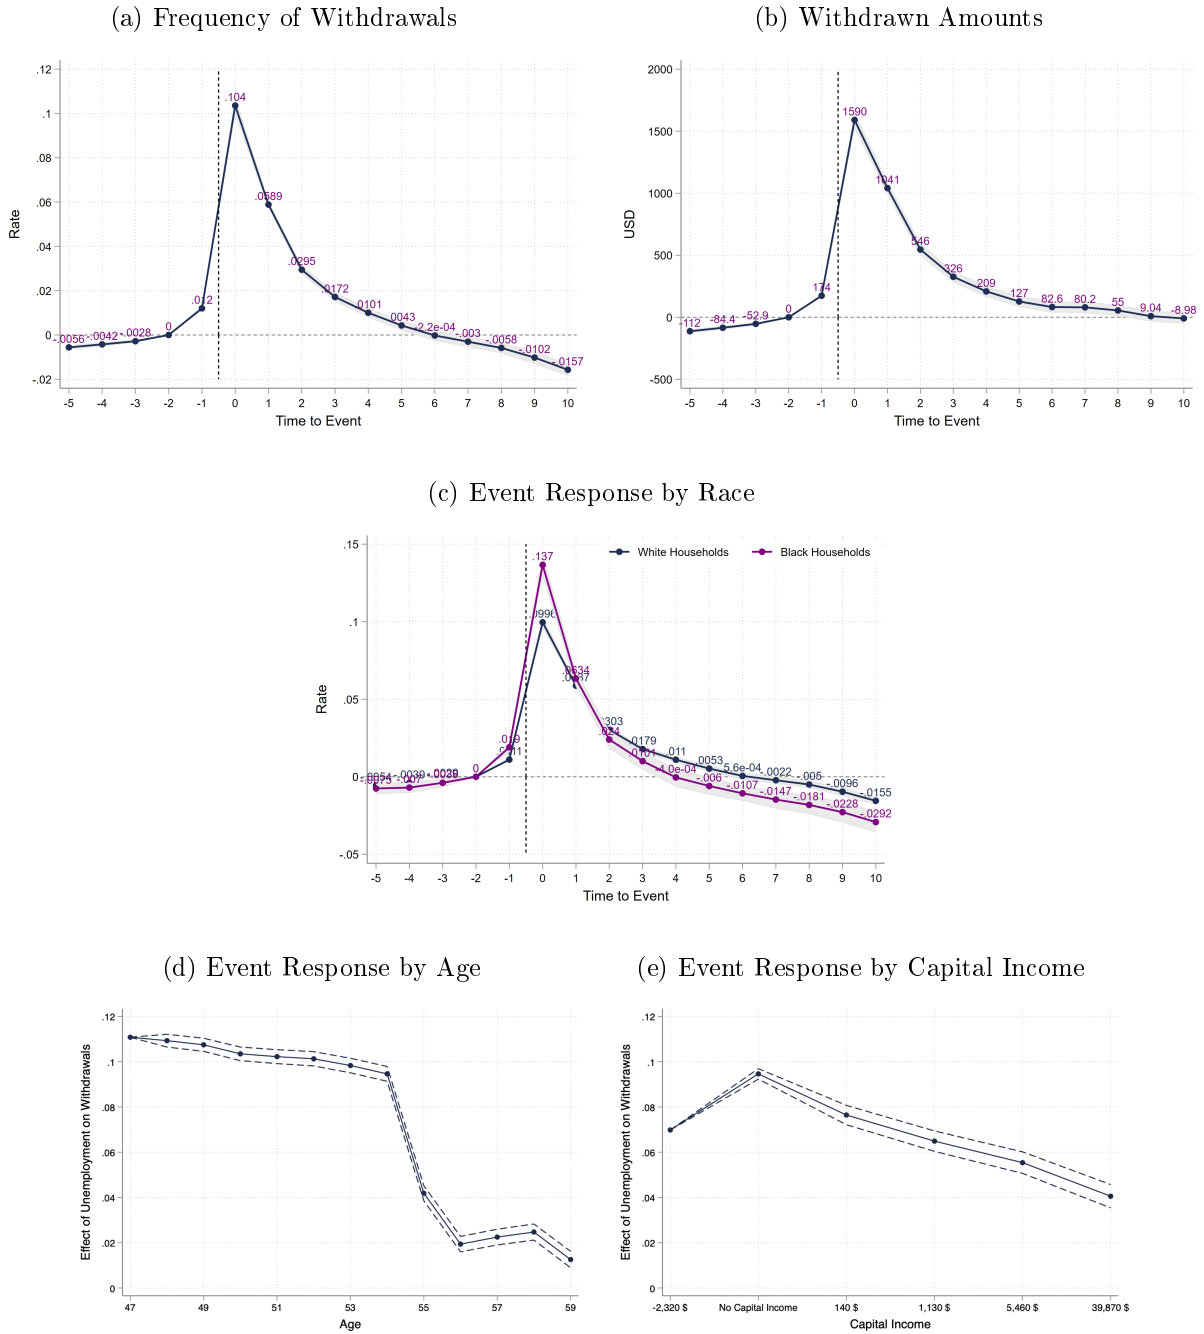


Shocks to Demand and Supply of Liquidity and Penalized Withdrawals



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 (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 (2) when the outcome variables are take-up and amounts of penalized withdrawals, respectively. Panels C-E 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 the primary filer's age. Panel E plots how the point estimates at time 0 vary as a function of household capital income. We split households into those with negative, zero, or positive capital income, and we bin the ones with positive capital income into four groups. On the x-axis, we then plot the average capital income within the corresponding bin. In panels D-E we present heterogeneity results from a regression that simultaneously includes these categories, as well as additional household-level controls: home ownership (an indicator based on property tax payments from Schedule A and mortgage debt from Form 1098), a dummy for whether the primary filer is married, the number of dependents, and the average household income based on information from all years within our data range.

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



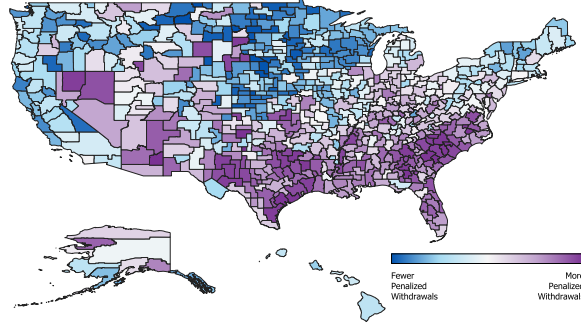
(e) Responses to Large Income Loss by Race



Notes: This figure studies penalized withdrawals around changes in household income. Panels A and B plot the event study coefficients from specification (2) 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 20 percent (relative to a previous year). Panels C-D study households' take-up and amounts of withdrawals as a function of the deviation of their income flow from their average income across our data period. We split households by whether a member of the household switched jobs that year because job changes themselves, as displayed in Appendix Figure D.8, lead to increased take-up. Panel E plots the event study of take-up of penalized withdrawals, split by whether the household's primary filer is Black or White.

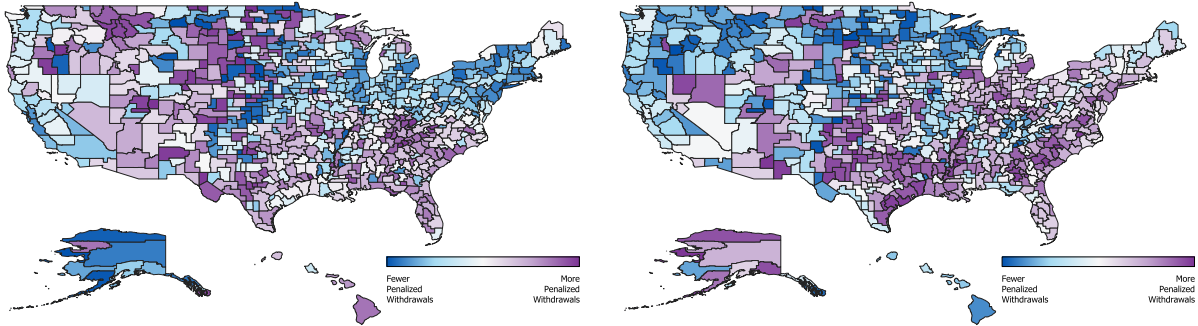
Figure 4: Geography of Withdrawals and Valuation of Liquidity

(a) Overall Variation

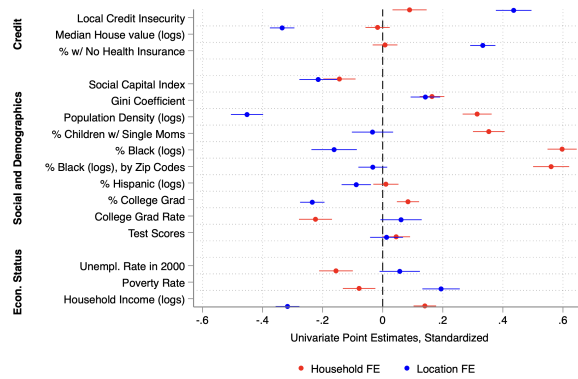


(b) Location Fixed Effects

(c) Household Fixed Effects

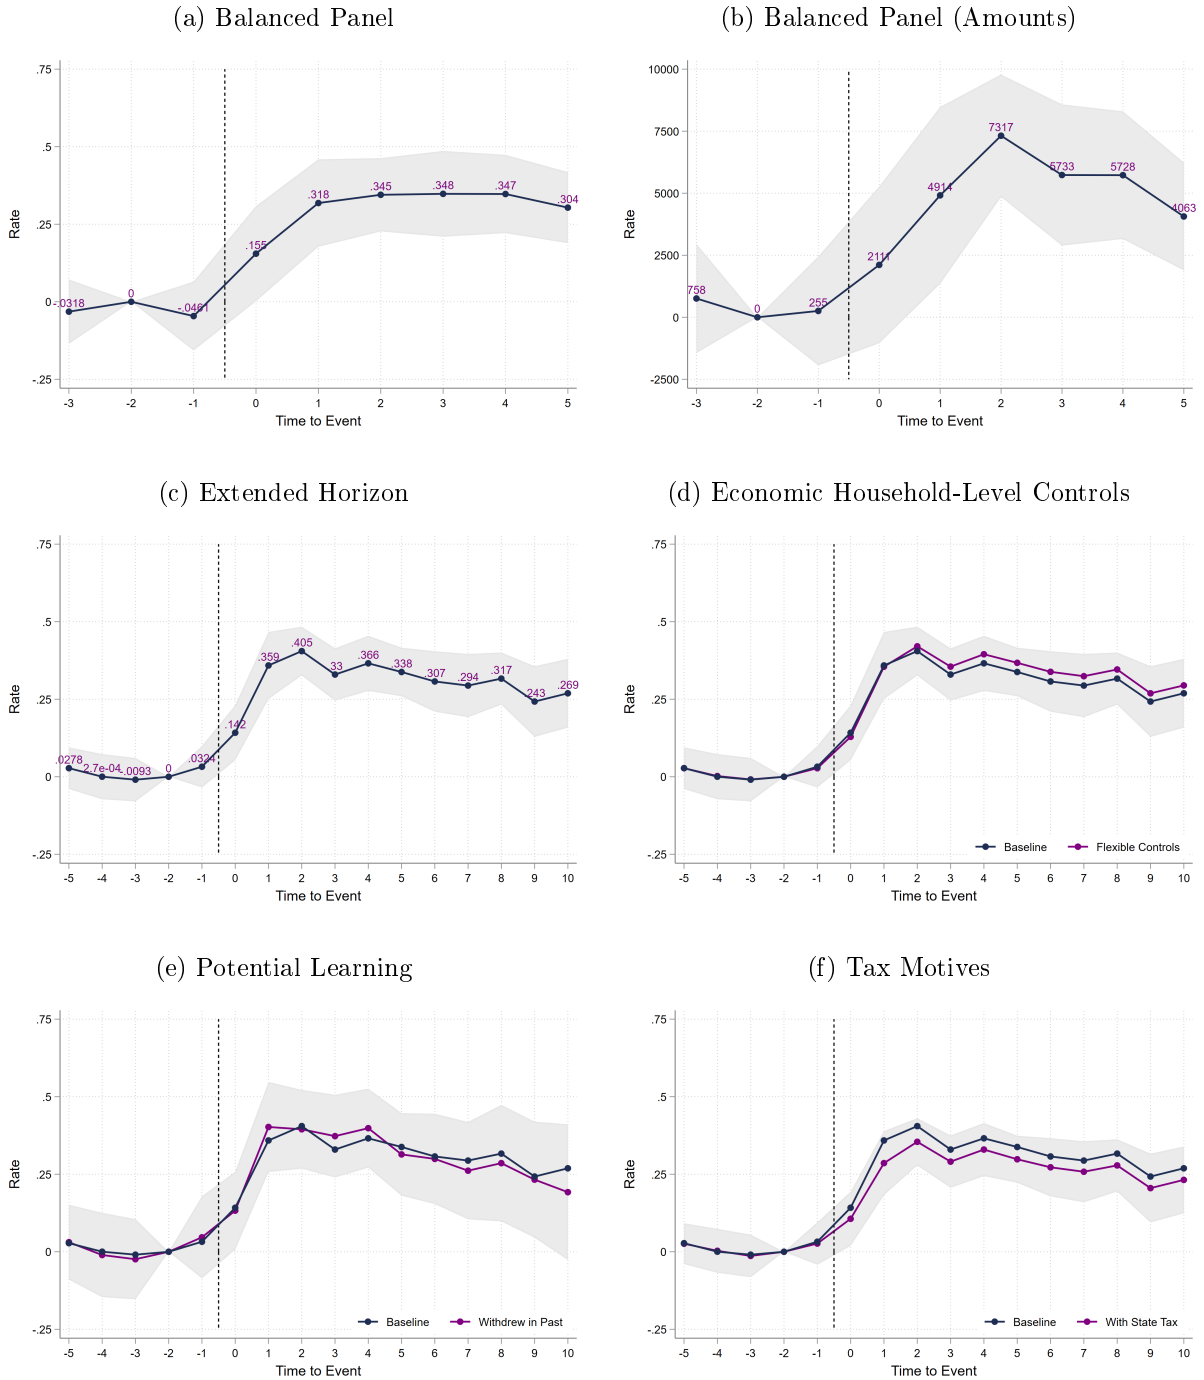


(d) Withdrawals and Locality Characteristics



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 (3) with household-level economic controls, panel B plots a map of the location fixed effects, Γ_z , and panel C plots a map of the household fixed effects, α_i , collapsed at the CZ level. The economic controls include unemployment, wage earnings, and gross income, with lagged, current, and lead values. Finally, panel D uses these estimates to display correlations of the regional differences across CZs with CZ-level social and economic characteristics. We display correlations of these characteristics separately for the location fixed effects, Γ_z , and for the household fixed effects, α_i , collapsed at the CZ level. We also study the correlation of the household fixed effect with a race indicator for the primary filer that assumes the value 1 for Black and the value 0 for White.

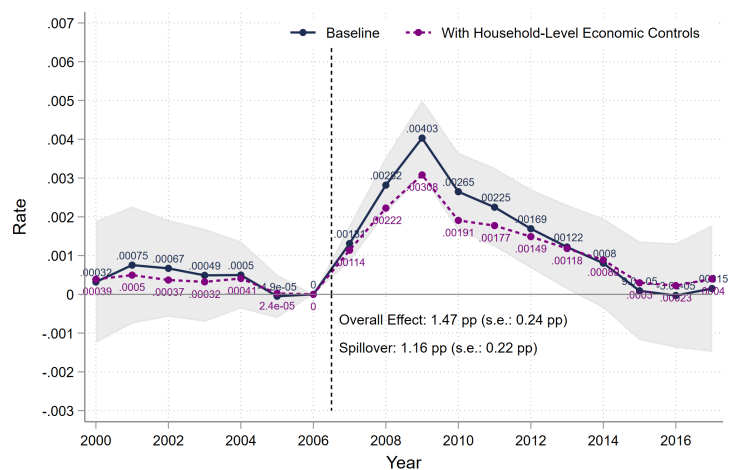
Figure 5: Movers Analysis




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

Figure 6: Penalized Withdrawals and Local Unemployment during the Great Recession

(a) Frequency



(b) Amounts



Notes: This figure displays estimates of the effect of the Great Recession on penalized withdrawals using equation (6). 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.

Table 1: Penalized Withdrawals and Race

	Making a Penalized Withdrawal	
	(1)	(2)
Black	0.0314 (0.0019)	0.0294 (0.0018)
Constant	0.1018 (0.0009)	0.0950 (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 estimated household-level fixed effect and the primary earner's race from specifications of equation (3) 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

We start by randomly selecting 10 percent of U.S. individuals. This selection is based on the last 2 digits of an individual's social security number. While social security numbers historically have not been completely random, the last 4 digits have always been randomly assigned. For these individuals, we collect date of birth and possible date of death from Social Security Administration administrative data.

We next build a sample of U.S. 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 range from tax years 1999 through 2018. We specifically focus on "prime age" households with primary filers aged 45-59.

We collect income information and location information from Form 1040 using administrative tax data. We measure a move in residence by checking whether the ZIP Code associated with Form 1040 changes from the previous year. We supplement this information with pension/retirement information present on Form W-2; information on retirement distributions found on Form 1099-R, including distribution amounts and codes associated with the type of withdrawal; and information from Form 1099-G regarding unemployment compensation.

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

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

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

Form 5329 (see below).

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

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

Variables 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 reported on Form 5498.

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

subscribes to a defined contribution plan with the employer, we at least know that most of the taxpayers identified as having access to retirement funds by our instrument participated in a retirement plan with an employer that offered a defined contribution plan.

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

Variables related to demographics and economic circumstances. We do not directly observe unemployment. However, Form 1099-G reports unemployment insurance (UI) payments made to individuals. As such, we assume any individual who receives UI payments has observed unemployment in that year.³⁹

While the administrative tax data do not explicitly show job changes or timing of those changes, we can 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.

To impute race and Hispanic origin, we use 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 and origin are then created based on which estimated probability is highest for each taxpayer.⁴⁰

We say that a taxpayer moves if the address reported on their tax return places them in a different commuting zone than in the year prior. Note that this omits local moves within the same commuting zone.

We define a large negative income shock as a deviation of 20 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.

³⁹Note 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.

⁴⁰Note 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: Implications

Taken together, the four facts we describe in Section 2.3 provide evidence that households use penalized withdrawals as a mean to mitigate short-run needs for liquidity. This evidence thus motivates us to use penalized withdrawals as a revealed-preference tool to characterize the needs and valuation of liquidity across American households. Yet, we address two potential concerns with our approach: first, 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; second, any revealed-preference approach rely on the assumption that agents are maximizing choices on the margin.

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

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

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

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

(a) Distribution of Amounts of Balances and Withdrawals

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

(b) Use of Withdrawals

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

Notes: These tables display summary statistics on defined contribution accounts from the Health and Retirement Study (HRS). We use HRS data from waves 7-14, which cover the years 2004-2018. The sample is restricted to respondents who are between the ages of 45-59.5. We focus on the 14,392 households who have defined contribution pension plans, who represent a population share of 84.14%. For these households, the average and median ages are 59.3 and 58.2 years old, respectively. The first line in Table 1 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. The second row of Table 1 displays the distribution of withdrawn amounts, but we note it includes only 222 observations for which there were non-missing positive values. Table 2 displays the distribution of the usages of the withdrawn amounts, counting 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 in Section 2.3 are closely consistent with various predictions of a model by 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, 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).⁴¹ In our context, the main concern is that the observed behavior could be generated purely by behavioral biases, such as narrow bracketing (e.g., Thaler 1999), mental accounting

⁴¹In fact, one traditional rationale for government intervention in retirement savings (particularly in the form of Social Security) has been that some individuals lack the foresight to save for their retirement years (Diamond 1977; Feldstein 1985).

(e.g., Read et al. 1999), or myopia/present bias (e.g., Laibson 1997; O’Donoghue and Rabin 1999), and may not convey information on the underlying valuation of liquidity. Reassuringly, as we next discuss, the evidence presented in Section 2.3 is not consistent with these interpretations.⁴²

We first consider *narrow bracketing*, whereby households do not integrate their entire portfolios into their decision making. The fact that most households withdraw sizable amounts and that the penalized withdrawals are only infrequently linked to the closure of a specific account 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 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, Figure D.1e 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).

⁴²Of course, while the evidence is inconsistent with these behavioral explanations governing the results, they could still naturally play a role.

C Proofs for Section 3 and an Additional Result

We first characterize the solution of the model to prove the Lemmas of Section 3. We then offer a generalization of Lemma 2 to the case in which households are further from the statutory retirement age.

Proofs of Lemmas from Section 3.

We start from the recursive formulation of the problem

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

subject to

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

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

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

$$(7) \quad \{b_{i,t} = 0\} : \quad u'(c_{i,t}; h_{i,t}) (\beta (1 + r))^{-1} = E_t \left[\frac{\partial V_{t+1}(a_{i,t+1}, k_{i,t+1}; h_{i,t+1})}{\partial a_{i,t+1}} \right]$$

$$(8) \quad \{\Delta a_{i,t} < 0, b_{i,t} > 0\} : \quad u'(c_{i,t}; h_{i,t}) (\beta (1 + r))^{-1} = \frac{E_t \left[\frac{\partial V_{t+1}(a_{i,t+1}, k_{i,t+1}; h_{i,t+1})}{\partial a_{i,t+1}} \right]}{1 - \rho'_{i,z}(b)}$$

$$(9) \quad \{\Delta k_{i,t} < 0\} : \quad u'(c_{i,t}; h_{i,t}) (\beta(1+r))^{-1} \geq \frac{E_t \left[\frac{\partial V_{t+1}(a_{i,t+1}, k_{i,t+1}; h_{i,t+1})}{\partial k_{i,t+1}} \right]}{1 - \tau}.$$

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

Lemma 1 and Lemma 2 are derived directly from the first order conditions, the definition of $\theta_{i,t}(c_{i,t}; h_{i,t})$, and the argument we just made on the relationship between $\frac{\partial V_{t+1}}{\partial k_{i,t+1}}$ and $\frac{\partial V_{t+1}}{\partial a_{i,t+1}}$, which are identical under the assumption of Lemma 2.

Lemma 3 is also derived directly from the first order conditions. As long as the household has funds in the illiquid account (i.e. $k_{i,t+1} > 0$ as assumed by Lemma 3), then either $\theta_{i,t}(c_{i,t}; h_{i,t}) < \frac{1}{1-\tau}$ and the household does not make a penalized withdrawal, or the household makes a penalized withdrawal but remains in an interior solution, hence satisfying the first order condition (9) with equality. In this latter case, $\theta_{i,t}(c_{i,t}; h_{i,t}) = \frac{1}{1-\tau} \pi(h_{i,t})$, where

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

is the relative value of an illiquid dollar, which has to be weakly smaller than one as discussed above.

To prove Lemma 4, we proceed in steps. First, notice that a trivial perturbation argument shows that in equilibrium the following conditions must hold:⁴³

$$(11) \quad \{b_{i,t} > 0, \Delta k_{i,t} < 0\} : \quad \frac{E_t \left[\frac{\partial V_{t+1}(a_{i,t+1}, k_{i,t+1}; h_{i,t+1})}{\partial a_{i,t+1}} \right]}{1 - \rho'_{i,z}(b)} = \frac{E_t \left[\frac{\partial V_{t+1}(a_{i,t+1}, k_{i,t+1}; h_{i,t+1})}{\partial k_{i,t+1}} \right]}{1 - \tau}$$

$$(12) \quad \{b_{i,t} > 0, \Delta k_{i,t} = 0\} : \quad \frac{E_t \left[\frac{\partial V_{t+1}(a_{i,t+1}, k_{i,t+1}; h_{i,t+1})}{\partial a_{i,t+1}} \right]}{1 - \rho'_{i,z}(b)} \leq \frac{E_t \left[\frac{\partial V_{t+1}(a_{i,t+1}, k_{i,t+1}; h_{i,t+1})}{\partial k_{i,t+1}} \right]}{1 - \tau}$$

$$(13) \quad \{b_{i,t} = 0, \Delta k_{i,t} < 0\} : \quad \frac{E_t \left[\frac{\partial V_{t+1}(a_{i,t+1}, k_{i,t+1}; h_{i,t+1})}{\partial a_{i,t+1}} \right]}{1 - \rho'_{i,z}(b)} \geq \frac{E_t \left[\frac{\partial V_{t+1}(a_{i,t+1}, k_{i,t+1}; h_{i,t+1})}{\partial k_{i,t+1}} \right]}{1 - \tau}.$$

Next, notice that keeping $\Delta a_{i,t}$ and $\Delta k_{i,t}$ fixed, $c_{i,t}$ is decreasing in $\varepsilon_{i,t}$, hence $u'(c_{i,t}; h_{i,t})$ is increasing in $\varepsilon_{i,t}$ (due to the curvature of the utility function). As a result, the higher is $\varepsilon_{i,t}$, the more the household needs to decrease asset accumulation for the Euler equations to hold. In turn, these observations imply that if $\varepsilon_{i,t}$ is large enough, the household would be induced to make a penalized withdrawal. To see why this is the case, consider a value of $\varepsilon_{i,t}$ sufficiently low that, even without dissaving any liquid asset (i.e., $b_{i,t} = 0$) the Euler equation (7) is satisfied. Then, consider

⁴³The first equality has to hold since otherwise the household would deviate and either borrow more or less and adjust the penalized withdrawals accordingly. Similarly, the other two inequalities must hold for the household to not want to withdraw any amount or to not want to borrow.

an increase in $\varepsilon_{i,t}$. The household needs to increase consumption today, which hence decreases the liquid assets tomorrow to bring back balance (where we note that $E_t \left[\frac{\partial V_{t+1}(a_{i,t+1}, k_{i,t+1}; h_{i,t+1})}{\partial a_{i,t+1}} \right]$ is decreasing in $a_{i,t+1}$ due, again, to the curvature of the utility function). Eventually, if $\varepsilon_{i,t}$ is sufficiently large, the household would run out of liquid assets and would therefore start borrowing. At this point, the Euler equation (8) applies. Following the same reasoning, if we keep increasing $\varepsilon_{i,t}$, the household would need to increase their borrowing. Eventually, equation (12) would not be satisfied, and thus the household would be induced to make a penalized withdrawal. This value of liquidity shock is $\bar{\varepsilon}_{i,t}(h_{i,t})$. We have thus proved Lemma 4.

Finally, we turn to Lemma 5, which follows an identical argument. Consider first an increase in the demand for liquidity, $\varphi_{i,\mathbb{D}}$, which for example could capture a decline in income leading to a higher valuation of liquidity for the same amount of borrowing (recall the definition of $\mathbb{D}(\bar{b})$). An increase in $\varphi_{i,\mathbb{D}}$ would imply that the same argument discussed above unfolds for lower values of $\varepsilon_{i,t}$. As a result, $\bar{\varepsilon}_{i,t}(h_{i,t})$ would be lower and penalized withdrawals will become more likely. Next, consider a decline in the supply of liquidity, i.e., an increase in $\varphi_{i,\mathbb{S}}$. A higher $\varphi_{i,\mathbb{S}}$ leads equation (12) to be violated at lower values of borrowing, thus once again reducing $\bar{\varepsilon}_{i,t}(h_{i,t})$ and leading to more frequent penalized withdrawals. The reason is mechanical: the denominator $\rho'_{i,z}(b)$ is more steeply increasing in b .

Generalization of Lemma 2.

Lemma 2 considered the case in which households find themselves just before the statutory retirement age. This assumption is convenient because it implies that the marginal value of liquid and illiquid funds tomorrow are identical. In general, however, illiquid funds are less valuable, and thus we are only able to bound the valuation of liquidity at a time of a withdrawal.

Lemma 2b: Withdrawals and Equilibrium Valuation of Liquidity. *If a household withdraws from the illiquid account at time $t < t^*$, then:*

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

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

We have already discussed above that $\pi(h_{i,t}) \leq 1$ since illiquid dollars are less valuable than liquid ones (unless no withdrawal is made before t^*). Next, consider the lower bound of $\pi(h_{i,t})$. To prove the lower bound we proceed by contradiction. Assume that the household is maximizing and that $\pi(h_{i,t}) < (1 - \tau)$. Then, build an alternative strategy by withdrawing one dollar from the illiquid account, paying the penalty τ , and transferring $(1 - \tau)$ dollars into the liquid account.

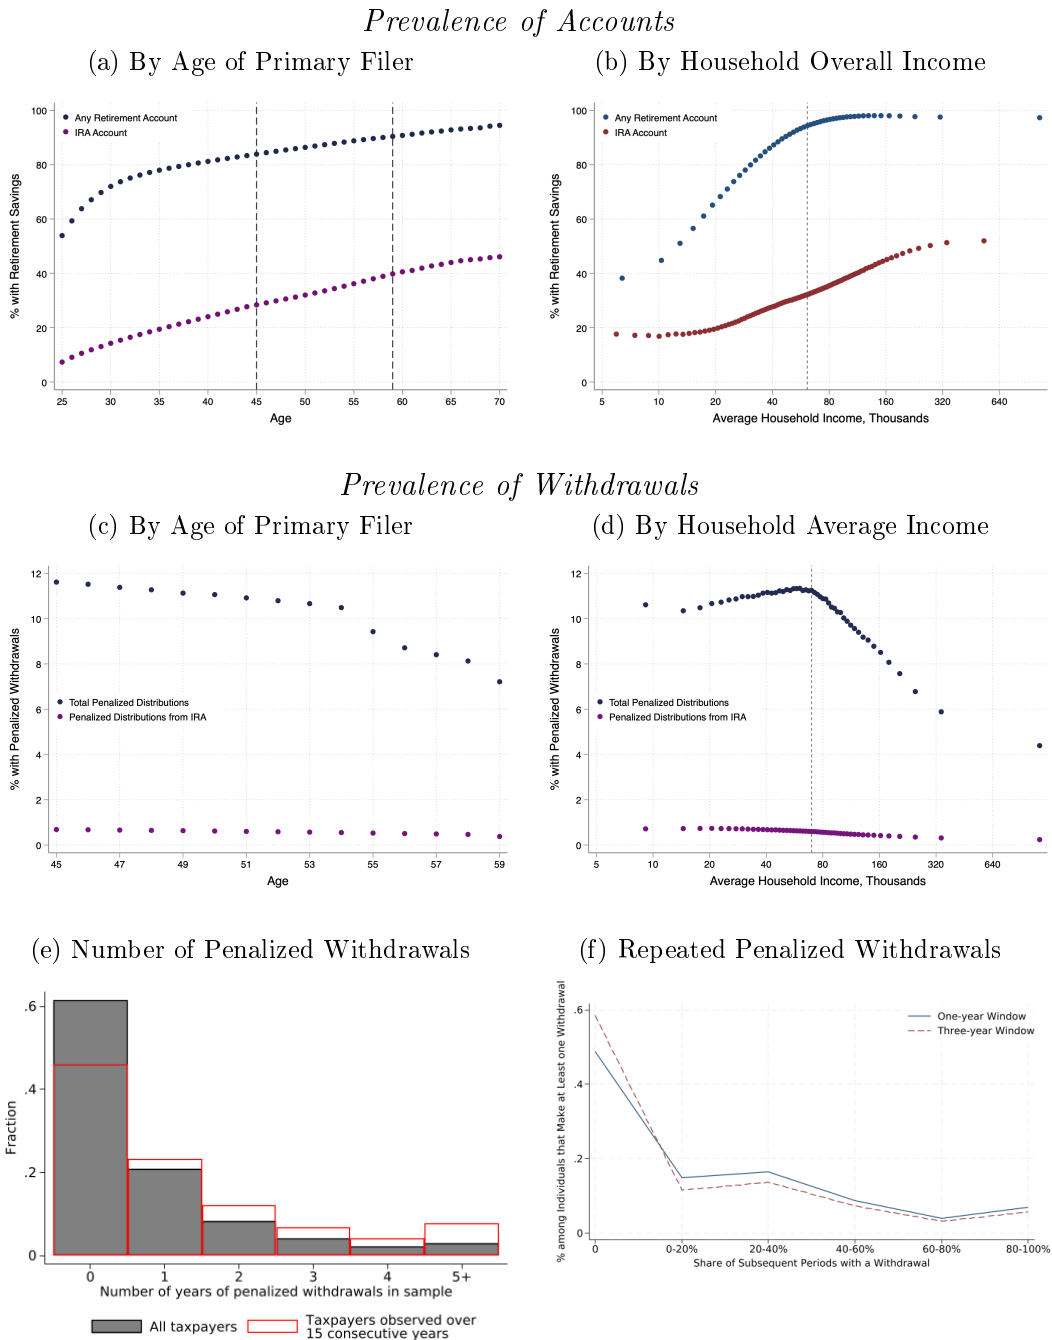
This deviation generates a total change in the household's value of the problem that is given by:

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

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

D Appendix Figures

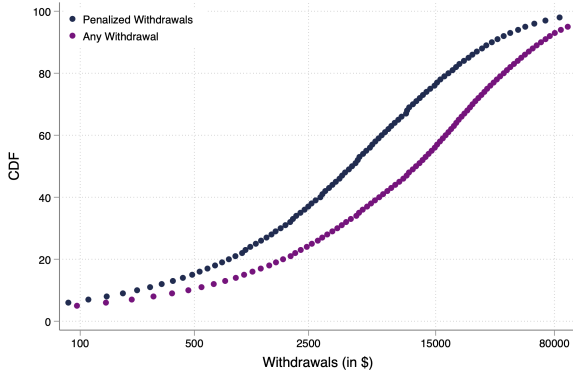
Figure D.1: Prevalence of Retirement Savings Accounts and Penalized Withdrawals



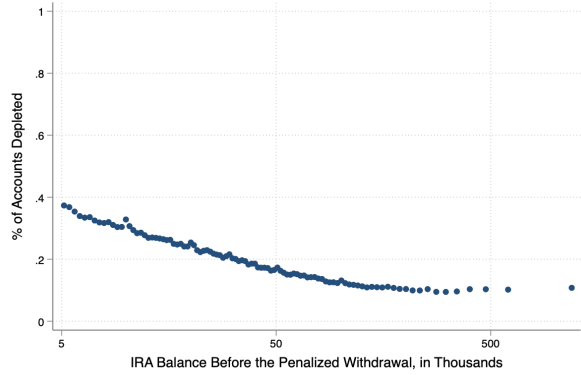
Notes: These figures illustrate the prevalence of retirement accounts and penalized withdrawals. We identify households as having accounts in a given year if up to that year within our sample period of 20 years they have made a contribution to 401(k)/IRA accounts or have balances in IRA accounts. The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

Figure D.2: Statistics on Penalized Withdrawals


(a) CDF of Withdrawals




(b) Share of IRA Accounts Fully Depleted



(c) Share of IRA Balances Withdrawn

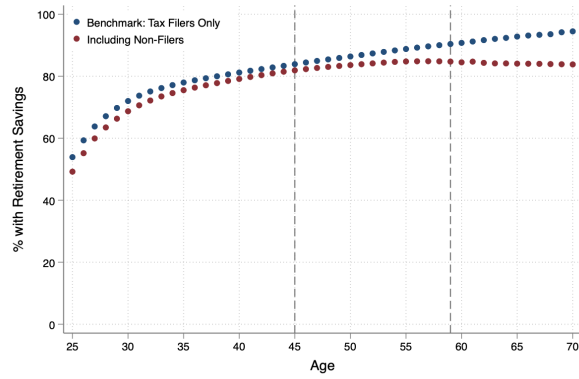


(d) Penalized Withdrawals and Income Changes



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


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



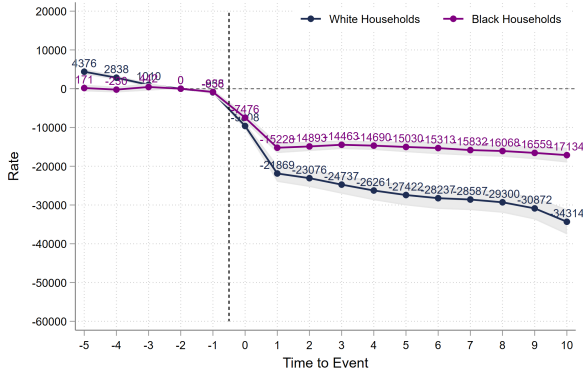
Notes: This figure illustrates the prevalence of retirement accounts by age. We compare our benchmark sample of households that file tax returns to a more inclusive sample, which combines in households that do not file tax returns. The figure shows that households that do not file tax returns are less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

Figure D.4: Unemployment Event: Adjusted Gross Income


(a) All Households



(b) By Race




(c) By Race in Percent




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

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




(a) Unemployment Event



(b) Large Income Loss

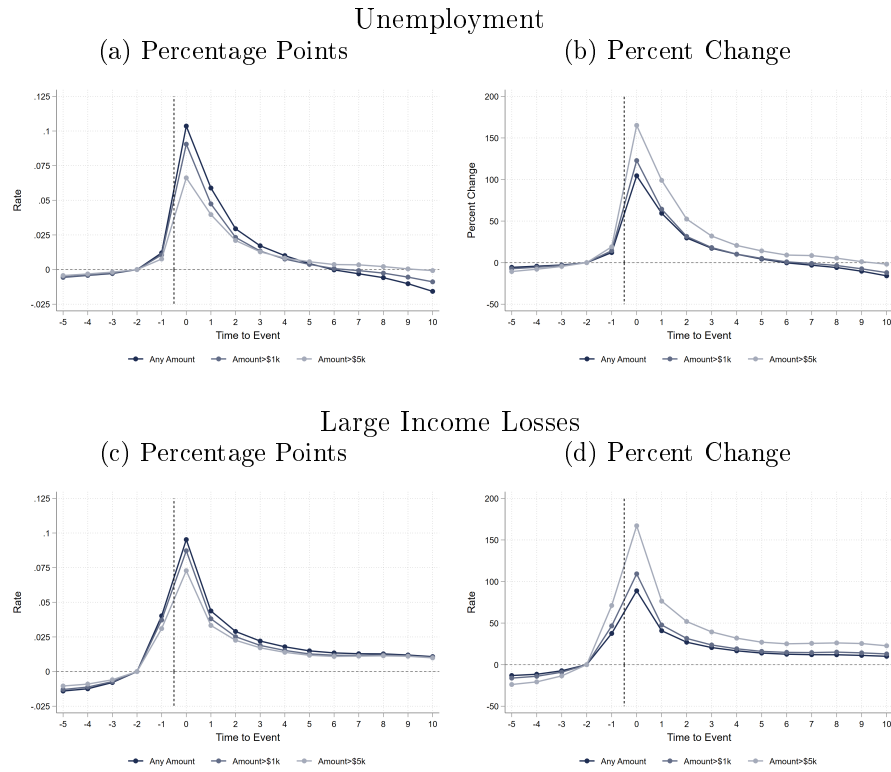
Notes: This figure plots the event study coefficients from specification (2) for the event of unemployment, as defined by the first period we observe at least one of the household members receiving unemployment benefits. We compare the overall sample to a restricted sample in which we include only households that do not change their commuting zone around the event. Specifically, we only consider households that are in the same commuting zone (CZ) in periods -1 and 1.

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



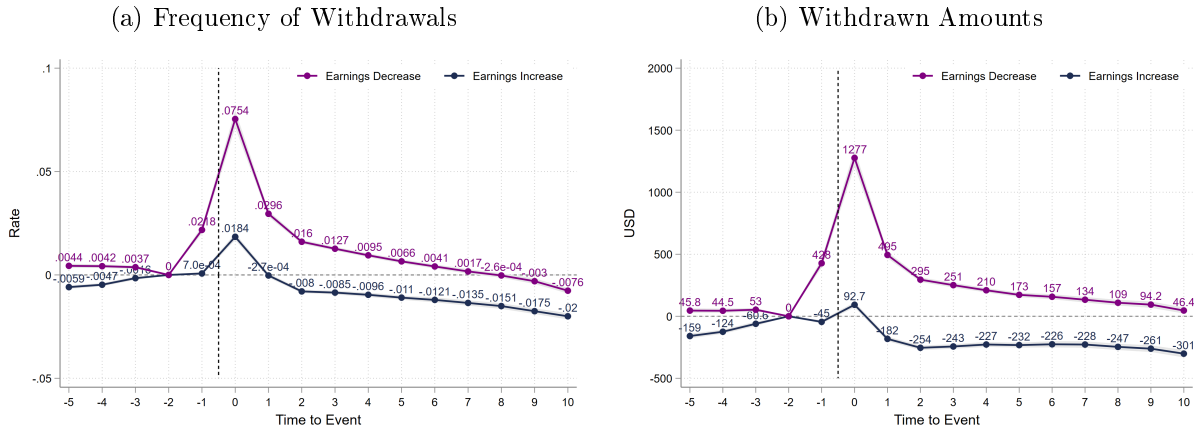
Notes: This figure plots the event study coefficients from specification (2) for the event of unemployment, as defined by the first period we observe at least one of the household members receiving unemployment benefits. We include observations of primary filers younger than 55.

Figure D.7: Event Studies by Amount Withdrawn



Notes: This figure plots the event study coefficients from specification (2) 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 at period $t = -2$.


Figure D.8: Event Study of Job Switch




Notes: This figure studies penalized withdrawals around the event of a job switch using specification (2). It focuses on the sample of households for whom we see a change in employer from period $t - 1$ to period t without an episode of being on unemployment benefits. We then split households by whether the employee experienced an earnings increase or an earnings decrease upon the switch. Panel A studies frequency of withdrawals, and Panel B studies withdrawal amounts.

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

(a) Balanced Panel



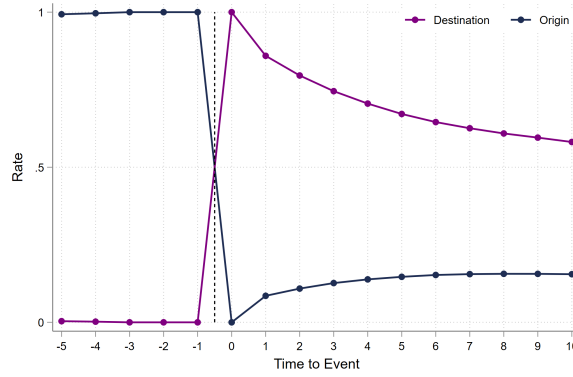
(b) Extended Horizon




Notes: These figures display estimates for the event study coefficients of a move (β_r) from the estimation of equation (5).

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

(a) Dynamics




(b) Passthrough Scaled by Movers Still in Destination



Notes: This figure provides additional analyses for the movers design. Panel A displays indicators for a household's geographic location around the move. In the movers design, we assign a household the same destination location for the entire post-move period. In this figure, we display indicator variables for whether, in a given period, the household remains in the assigned destination unit and whether the household returns to the assigned origin unit. Panel B scales the estimates for the movers analysis from panel C of Figure 5 by the share of movers still at the assigned destination.


Figure D.11: Movers Analysis–Extended Horizon

(a) Withdrawal Amounts



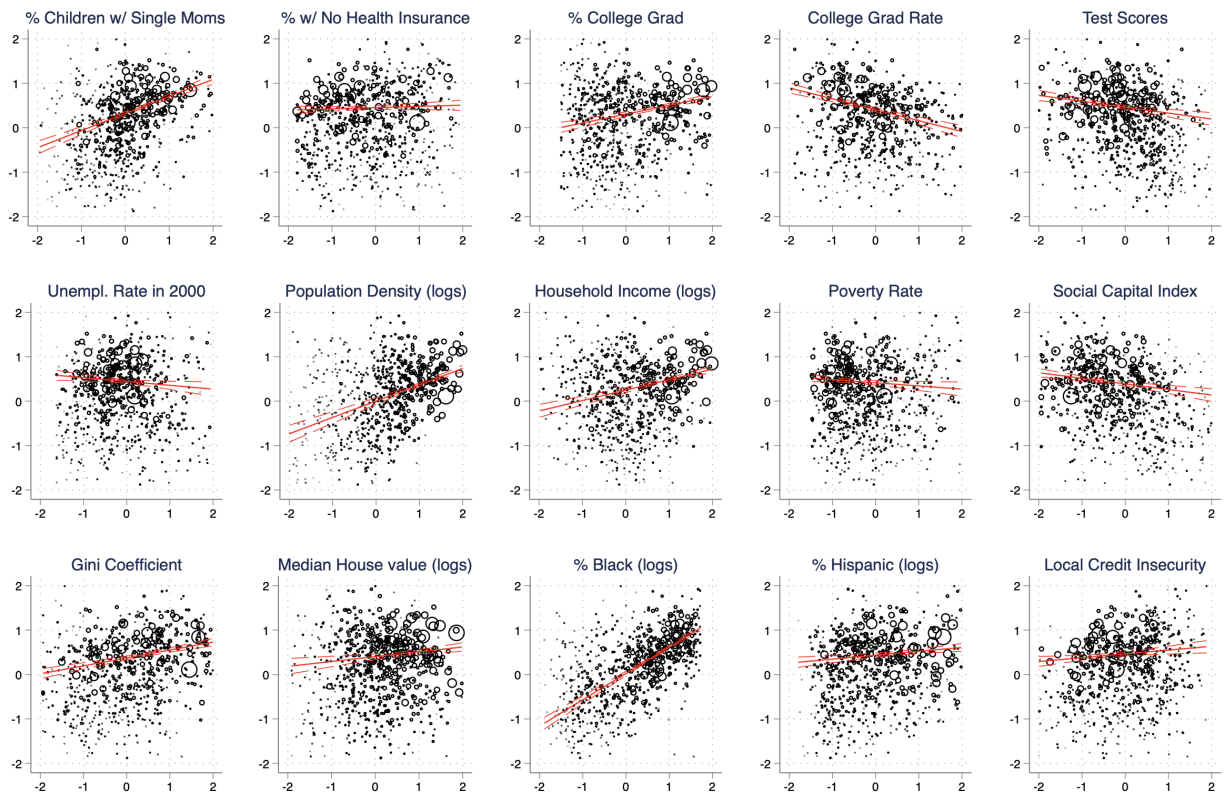
Notes: This figure displays estimates for the share of spatial differentials in withdrawals that can be attributed to location, using the movers design specification of equation (5). We show the estimates from an unbalanced panel of households on an extended time window that spans the years [-5, +10]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

Figure D.12: Correlations with Location Fixed Effects



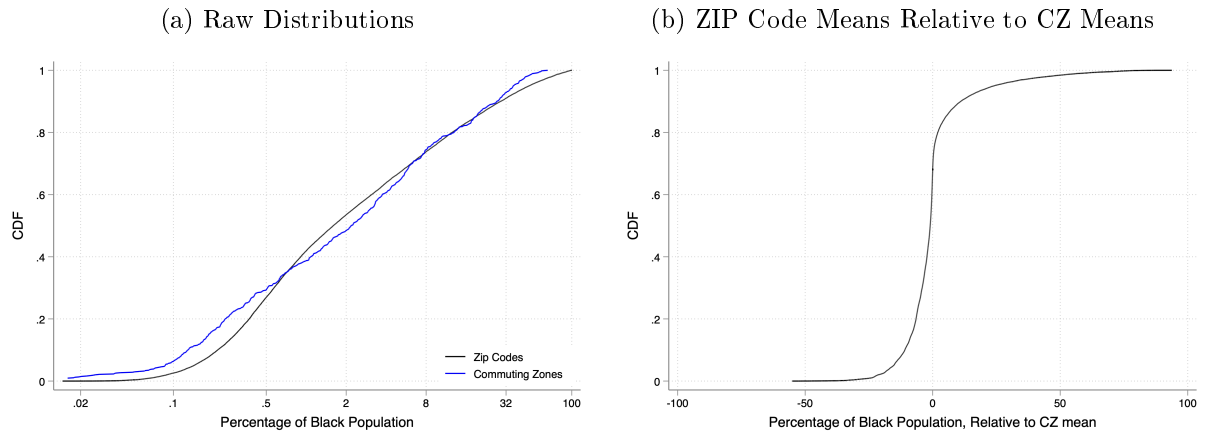
Notes: These figures display correlations of the location fixed effects, Γ_z , as estimated using equation (3), with CZ-level social and economic characteristics.

Figure D.13: Correlations with Households Fixed Effects



Notes: These figures display correlations of the household fixed effects, α_i , as estimated using equation (3) and collapsed at the CZ level, with CZ-level social and economic characteristics.

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




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


Figure D.15: Primary Filers of Ages 55-59

Event Study: Large Income Losses

(a) Withdrawal Frequency




(b) Withdrawal Amounts




Movers Analysis


(c) Withdrawal Frequency: Baseline



(d) Withdrawal Frequency: Flexible Controls



(e) Withdrawal Amounts: Baseline



(f) Withdrawal Amounts: Flexible Controls






Figure D.15: Primary Filers of Ages 55-59—*Continued*

Great Recession



(g) Withdrawal Frequency



(h) Withdrawal Amounts

Notes: In this figure, 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, the movers analysis, and the analysis of the Great Recession. 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.