

NBER WORKING PAPER SERIES

THE WEALTHY HAND-TO-MOUTH

Greg Kaplan
Giovanni L. Violante
Justin Weidner

Working Paper 20073
<http://www.nber.org/papers/w20073>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2014

We thank Yu Zhang for outstanding research assistance, and Mark Aguiar, Karen Pence, David Romer, Rob Shimer, David Weil, and Justin Wolfers for comments. This research is supported by grant no. 1127632 from the National Science Foundation. I received compensation from the Brookings Institution in conjunction with this research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Greg Kaplan, Giovanni L. Violante, and Justin Weidner. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Wealthy Hand-to-Mouth
Greg Kaplan, Giovanni L. Violante, and Justin Weidner
NBER Working Paper No. 20073
April 2014
JEL No. D1,D3,D9,E2,E6

ABSTRACT

The wealthy hand-to-mouth are households who hold little or no liquid wealth (cash, checking, and savings accounts), despite owning sizable amounts of illiquid assets (assets that carry a transaction cost, such as housing or retirement accounts). We use survey data on household portfolios for the U.S., Canada, Australia, the U.K., Germany, France, Italy, and Spain to document the share of such households across countries, their demographic characteristics, the composition of their balance sheets, and the persistence of hand-to-mouth status over the life cycle. The portfolio configuration of the wealthy hand-to-mouth suggests that these households may have a high marginal propensity to consume out of transitory income changes, a prediction for which we find empirical support in PSID data. We explain the implications of this group of consumers for macroeconomic modeling and fiscal policy analysis.

Greg Kaplan
Department of Economics
Princeton University
Fisher Hall
Princeton, NJ 08544
and NBER
gkaplan@princeton.edu

Justin Weidner
Department of Economics
Princeton University
Princeton, NJ, 08544
jweidner@princeton.edu

Giovanni L. Violante
Department of Economics
New York University
19 W. 4th Street
New York, NY 10012-1119
and NBER
glv2@nyu.edu

1 Introduction

The life-cycle permanent income hypothesis is a valuable organizing framework to analyze both household survey and aggregate time-series data on the joint dynamics of income and consumption. At the same time, economists have long recognized that certain aspects of these data are at odds with some of the model’s most salient predictions. This is true for both the standard version of the model (Friedman, 1957; Hall, 1978) and the more recent “buffer-stock” versions (Deaton 1991; Carroll 1997). In particular, both at the micro and macro level, it is common to estimate a large sensitivity of consumption with respect to transitory changes in income, whereas according to the theory these income dynamics should be smoothed.¹ It is also common to find that expected consumption growth is uncorrelated with the real interest rate, a result that implies a breakdown of the forward-looking Euler equation holding with equality, as long as the elasticity of intertemporal substitution is not zero.²

The most direct way to account for these facts is through the existence of a sizable share of *hand-to-mouth* (HtM) consumers in the population: consumers that spend all of their available resources in every pay-period. HtM consumers have a high marginal propensity to consume (MPC) out of transitory income changes, which can account for the high correlation between consumption and the transitory component of income growth, even for anticipated income shocks. Moreover, HtM consumers are not on their Euler equations, and thus they are a source of misalignment between movements in the interest rate and movements in aggregate consumption growth. The main challenge to this view asserts that standard measurements using micro data on household balance sheets conclude that the fraction of households with near zero *net worth*, and hence who consume all of their income each period, is too small to quantitatively reproduce the facts discussed above.

Measuring HtM behavior using data on net worth is consistent with the vast majority of heterogeneous-agent equilibrium macroeconomic models. These frameworks either feature either a single asset or two assets with different risk profiles, but the same degree of liquidity. Notable examples are the Bewley models featuring uninsurable idiosyncratic risk and credit constraints, in the tradition of Huggett (1996), Aiyagari

¹Some notable examples of micro-level evidence on excess sensitivity are Parker (1999), Souleles (1999), Shapiro and Slemrod (2003a, 2003b, 2009), Johnson, Parker, and Souleles (2006), Parker et al. (2013), and Broda and Parker (2014). See Jappelli and Pistaferri (2010), for a recent survey. Campbell and Mankiw (1989, 1990, 1991) provide evidence based on macroeconomic time-series.

²See, again, Campbell and Mankiw (1989, 1990, 1991), but also Attanasio and Weber (1993), and Ludvigson and Michaelides (2001).

(1994), Rios-Rull (1995), and Krusell and Smith (1998), and the spender-saver models that feature impatient and patient consumers with complete markets in the tradition of Campbell and Mankiw (1989). This latter class of models has been revived recently by Gali, Lopez-Salido, and Valles (2007), Eggertson and Krugman (2012), and Justiniano, Primiceri, and Tambalotti (2013), among others, to analyze macroeconomic dynamics around the Great Recession. Models by Krusell and Smith (1997) and Carroll, Slacalek, and Tokuoka (2014a, 2014b) combine the spender-saver insight of heterogeneity in patience with a standard one-asset incomplete-markets model.

In this paper, we argue that measurements of HtM behavior inspired by this class of models are misleading because they miss what we call the *wealthy hand-to-mouth* households. These are households who hold sizable amounts of wealth in illiquid assets (such as housing or retirement accounts), but very little or no liquid wealth. As a result, they consume all of their disposable income every period. Clearly, such households would not be picked up by standard measurements since they have positive—and often substantial—net worth.

To obtain a comprehensive measurement of HtM behavior with cross-sectional survey data about household portfolios, a far better strategy is to use a model with two assets, one liquid and one illiquid, as the guiding framework. The illiquid asset yields a higher return, but it can only be accessed by paying a transaction cost. Recent examples of this two-asset environment are Angeletos et al. (2001), Laibson et al. (2003), Chetty and Szeidl (2007), Alvarez, Guiso, and Lippi (2012), Huntley and Michelangeli (2014), and Kaplan and Violante (2014a, 2014b). Through the lens of this two-asset model, there are two types of HtM households. The poor hand-to-mouth (P-HtM) hold little or no liquid wealth and no illiquid wealth; the wealthy hand-to-mouth (W-HtM) also hold little or no liquid wealth, but do have significant amounts of illiquid assets on their balance sheet. Just like the P-HtM households, W-HtM households have large MPCs out of small transitory income fluctuations. However, we show W-HtM households are more similar to non HtM (N-HtM) households along many other important dimensions. As a result, the W-HtM cannot be fully assimilated into either group. Rather, they are best represented as a third, separate class of households.

This paper investigates W-HtM behavior theoretically and empirically, and examines this peculiar but sizable group’s implications for macroeconomic modeling and policy analysis.

First, we ask why households with significant wealth would optimally choose to con-

sume all of their income every period, instead of using their wealth to smooth shocks. To answer this question, in Section 2 we develop a stylized model based on Kaplan and Violante (2014a). The model reveals that, under certain parameter configurations, optimal portfolio composition has positive amounts of illiquid wealth and zero liquid wealth. Such wealthy HtM households are better off bearing the welfare loss from income fluctuations rather than smoothing consumption. This is because the latter option requires holding large balances of cash and foregoing the high return on the illiquid asset (and, therefore, the associated higher level of long-run consumption). This explanation is consistent with calculations by Browning and Crossley (2001) who showed that, in a plausibly parameterized life-cycle buffer stock model, the utility loss from setting consumption equal to income, instead of fully optimizing, is second order. Cochrane (1989) and Krusell and Smith (1996) perform similar calculations in a representative agent environment. The model also provides useful guidance for our empirical strategy. In Section 3 we outline this strategy in detail and explain how we approach measurement issues.

Next, we ask how large the share of W-HtM households is in the population, what their demographic characteristics are, relative to the other two groups, how their balance sheets compare with that of the N-HtM, and how persistent their HtM status is over their life cycle. This empirical analysis is based on cross-sectional survey data on household portfolios for eight countries: the U.S., Canada, Australia, the U.K., Germany, France, Italy and Spain. We describe these data in Section 4. When the literature on household portfolios has previously examined these data, its emphasis has been on the allocation between risky and safe assets (see Guiso, Haliassos, and Jappelli (2002) for a thorough cross-country comparison). Instead, our focus is on the liquidity characteristics of the portfolio. In Section 5, we study U.S. data, for which we have several repeated cross-sections between 1989 and 2010, as well as a two-year panel for 2007-2009. In Section 6, we present a comparative cross-country analysis with survey data from 2010 and surrounding years.

The analysis of U.S. data leads to six main findings. First, between 25 and 40 percent of U.S. households are HtM, with our preferred estimate at one-third of the population. We find that one-third of HtM households are poor HtM and two-thirds are wealthy HtM, meaning the W-HtM represent the vast majority of this group, and would be missed by measurement of HtM behavior based on net worth. Third, households appear to be most frequently P-HtM at young ages, whereas the age profile of the W-HtM is hump-shaped and peaks around age 40. Fourth, the W-HtM typically hold sizable

amounts of illiquid wealth: for example, the median at age 40 is around \$50,000. Fifth, W-HtM households appear very similar to the unconstrained N-HtM in terms of their age-profile of income and the shares of illiquid wealth held in housing and retirement account. Finally, we determine W-HtM status is slightly more transient than P-HtM status.

Some interesting findings emerge from a comparison of the U.S. economy with the other countries we study. In all of the other countries, W-HtM households are a much greater share of the population than P-HtM households, even more so than in the United States. However, the total fraction of HtM households varies significantly across countries. Like in the U.S., HtM households represent more than 30 percent of the population in Canada, U.K., and Germany, but represent 20 percent or less of the population in Australia, France, Italy, and Spain. For the euro area countries, we observe that holdings of consumer debt are minimal, suggesting that the substantial liquid wealth seen, even among the income-poor, may act as a buffer stock that substitutes for expensive and limited access to credit.

In Section 7 we show that a household's HtM status has strong predictive power for its consumption response to transitory shocks. We apply the identification strategy from Blundell, Pistaferri, and Preston (2008) to U.S. income and consumption panel data to measure the MPC out of transitory income shocks for each type of household. We find that W-HtM and P-HtM households have significantly stronger responses than N-HtM households. In contrast, when we split households into HtM groups based on net worth only, we do not find a significant difference in the consumption responses of those two groups.

In Section 8, we argue that the W-HtM deserve their own separate status, in the cast of characters populating macroeconomic models. We use our empirical estimates of the share of households in each HtM group, together with simulated MPCs from three alternative structural models of consumption behavior, to show that the W-HtM cannot be assimilated to either the P-HtM or the N-HtM. We highlight four areas where frameworks that do not explicitly model W-HtM households provide misguided intuition about the effects of fiscal policy: the degree of MPC non-linearity with respect to the transfer size, the asymmetry of the consumption response with respect to equal-size income windfall and losses, the optimal phasing-out of stimulus payments with income for maximizing the impact on aggregate consumption, and the extent of cross-country dispersion in consumption responses to a fiscal transfer. Section 9 summarizes and concludes the paper.

2 Wealthy hand-to-mouth behavior: a simple model

We start by analyzing a simple three-period model in order to illustrate the determinants of hand-to-mouth behavior. In this section, we keep the presentation to a bare minimum. Appendix A contains a more thorough analysis of the problem. The model is also useful to determine how to detect a household's HtM status in the data and, as such, it provides guidance for our measurement exercise.

Household problem. Consider a household that lives for three periods, $t = 0, 1, 2$, but consumes only in the last two periods. Preferences over consumption at $t = 1, 2$ are given by

$$v_0 = u(c_1) + u(c_2), \quad (1)$$

with no discounting between periods, and with $u' > 0, u'' < 0$. The variable c_t denotes nondurable consumption at date t .

In period 0, the household has an initial endowment ω and makes a portfolio allocation decision. Two assets are available as saving instruments. An illiquid asset a pays off a gross return R before the consumption decision in period 2, but cannot be accessed at the time of the consumption decision in period 1. A liquid asset m can be accessed before the consumption decision in both periods, but pays a return $1 < R$. For now, we do not allow the agent to borrow, that is, to take a negative position in the liquid asset, but we later relax this assumption.

After the initial portfolio allocation decision, households receive income y_1 and make their consumption and liquid saving decision at $t = 1$. In the last period $t = 2$, they receive income y_2 and consume this amount, their liquid savings from $t = 1$, and their savings allocated to the illiquid asset at $t = 0$, plus the accrued capital income. Therefore, the only two decisions to analyze are the initial portfolio allocation decision and the consumption/saving decision at $t = 1$. Finally, note that since the income path (y_1, y_2) is known at $t = 0$, there is no uncertainty.

Our characterization of hand-to-mouth behavior concerns the asset position at the time of the $t = 1$ consumption decision. We define a household as not hand-to-mouth (N-HtM) if, after consuming at $t = 1$, it holds a positive amount of liquid assets, i.e., $m_2 > 0$ and $a \geq 0$. As is clear from (1), this household will choose $c_1 = c_2$. We define a household as poor hand-to-mouth (P-HtM) if, after consuming at $t = 1$, it does not hold any liquid or illiquid assets: $m_2 = 0$ and $a = 0$. We define a household as wealthy hand-to-mouth (W-HtM) if, after consuming at $t = 1$, it holds a positive

amount of illiquid assets but no liquid assets: $m_2 = 0$ and $a > 0$. Therefore, the $t = 1$ consumption/saving decision determines whether an agent is HtM, and the initial portfolio allocation at $t = 0$ determines whether a HtM agent is poor or wealthy HtM. For both HtM households, $c_1 < c_2$.

Solution. We begin with the initial portfolio allocation decision at $t = 0$:

$$\begin{aligned}
 v_0 &= \max_{m_1, a} u(c_1) + u(c_2) \\
 &\quad s.t. \\
 a + m_1 &= \omega \\
 c_1 + m_2 &= y_1 + m_1 \\
 c_2 &= y_2 + m_2 + Ra \\
 m_1 &\geq 0, a \geq 0
 \end{aligned}$$

where the first line is the resource constraint in the portfolio choice; the second and third lines are the budget constraints at $t = 1$ and $t = 2$; and the final line collects the inequality constraints on the choice variables. The first order condition of this problem with respect to a gives

$$u'(c_1) \left[1 + \frac{\partial m_2}{\partial a} \right] \geq u'(c_2) \left[R + \frac{\partial m_2}{\partial a} \right], \quad (2)$$

where the inequality is strict when $a = 0$. The derivative $\partial m_2 / \partial a$ reflects the dependence of the liquid savings decision at $t = 1$ on the amount held in illiquid assets. The resulting initial portfolio allocation implicitly determines the endowment points $(y_1 + \omega - a, y_2 + Ra)$ immediately prior to the consumption/saving decision at $t = 1$.

We now turn to this consumption saving decision at $t = 1$, given the predetermined amount invested in liquid wealth $m_1 = \omega - a$:

$$\begin{aligned}
 v_1(a) &= \max_{c_1, m_2} u(c_1) + u(c_2) \\
 &\quad s.t. \\
 c_1 + m_2 &= y_1 + \omega - a \\
 c_2 &= y_2 + m_2 + Ra \\
 m_2 &\geq 0
 \end{aligned}$$

where the first and second lines are the budget constraints at $t = 1$ and $t = 2$, and the

third line imposes the nonnegativity constraint on the choice variable. The first-order condition of this problem is:

$$u'(c_1) \geq u'(c_2), \quad (3)$$

where the strict inequality holds whenever the constraint binds and $m_2 = 0$. For example, when y_1 is high enough relative to y_2 , the agent wants to save some of its income into period 2, and $m_2 > 0$. In contrast, when y_1 is low enough relative to y_2 , the agent would, ideally, like to borrow and it is constrained at $m_2 = 0$. This “short-run” Euler equation in (3) states that, at $t = 1$, the relative price of consumption between $t = 1$ and $t = 2$ is equal to one, the return on the liquid asset.

Combining (3) with (2) yields

$$u'(c_1) \geq Ru'(c_2). \quad (4)$$

This is because $u'(c_1) = u'(c_2)$ when m_2 is interior, and because m_2 is unaffected by a marginal change in a when the household is at a constraint. This long-run Euler equation in (4) states that, from the agent’s viewpoint at $t = 0$, the relative price of consuming at $t = 1$ versus $t = 2$ is R . Comparing (4) and (3), the intertemporal trade-off appears to change between $t = 0$ and $t = 1$ because the illiquid asset is available as a saving instrument only at $t = 0$.

The “short-run” Euler equation (3) implies

$$m_2 = \max \left\{ \frac{y_1 + \omega - y_2 - (1 + R)a}{2}, 0 \right\}. \quad (5)$$

Since we are interested in characterizing HtM behavior, we focus on the case where $m_2 = 0$. Equation (5) reveals that a sufficient condition for this case is $y_2 \geq y_1 + \omega$: for a given initial endowment, income in period 2 is so large, relative to period 1, that even when the total endowment ω is saved into the liquid asset, the household still desires to consume more at $t = 1$.

To make further progress on the solution, we assume that u is in the constant elasticity of substitution (CES) class with elasticity of intertemporal substitution σ . Then, the long-run Euler equation (4) gives

$$a = \max \left\{ \frac{R^\sigma (y_1 + \omega) - y_2}{R + R^\sigma}, 0 \right\}, \quad (6)$$

From (6), we conclude that the household is W-HtM when

$$R > \left(\frac{y_2}{y_1 + \omega} \right)^{\frac{1}{\sigma}} \quad (7)$$

and is P-HtM when the opposite (weak) inequality holds.

It is useful to explain the role of the model's parameters in determining W-HtM behavior. A high relative return R makes the illiquid asset more attractive by raising its effective return, thereby inducing the agent to tolerate wider consumption differences across periods in order to achieve a higher overall consumption level. Steep income growth y_2/y_1 reduces the appeal of the illiquid asset as a saving instrument, as the income path already guarantees high consumption later in life. The higher the elasticity of intertemporal substitution σ , the more the household is willing to absorb a jump in consumption across periods, and so the more likely it is to save into the illiquid asset even if y_1 is low relative to y_2 .³

Since the model is deterministic, W-HtM households choose to invest in the illiquid asset at $t = 0$, even though they know with certainty that they will be constrained in the next period. By acting this way, they consume even less at $t = 1$ and make themselves even more constrained. Put differently, the shadow value of an additional unit of income at $t = 1$ is higher for the W-HtM than for the P-HtM. If we let this multiplier be λ , for a P-HtM agent $\lambda = u'(y_1 + \omega) - u'(y_2)$ and for a W-HtM agent $\lambda = u'(y_1 + \omega - a) - u'(y_2 + Ra)$, which is larger. Nevertheless, this choice is optimal because the welfare gain from the rise in the overall level of lifetime consumption more than compensates for the welfare loss from the consumption gap between $t = 1$ and $t = 2$.

MPC out of a transitory shock. Suppose that after the initial portfolio allocation decision, but before the consumption decision at $t = 1$, the household receives an unexpected income shock, such as a transfer τ from the government. What is the household's MPC out of this transfer? A N-HtM household has an MPC of exactly one-half, since there is no discounting and it smooths the payment equally across the two periods. If the transfer is small enough not to throw the agent off its kink ($m_2 = 0$), then the HtM household's MPC out of the transfer will be 1. This occurs as long as $\tau \leq y_2 - (y_1 + \omega) + (1 + R)a$. This condition is weaker for a W-HtM than for a P-HtM

³Equation (7) reveals that the model is homothetic in y_1, y_2 , and ω . In this sense, a high-income household is as likely to be W-HtM as a low-income one, as long as the life-cycle slope of their income profiles is the same.

because, as explained above, the former household is more constrained.⁴ Finally, note that all these results carry over to the case of an anticipated transfer, as long as the transfer is small enough that it does not change HtM status at $t = 1$.

Taking stock. Our two-period model is an extremely stylized environment that is useful to describe how W-HtM behavior can arise as a result of giving up gains from additional consumption smoothing in exchange for the opportunity of investing in a high-return asset that yields higher levels of average lifetime consumption. This insight survives in more general environments. We briefly discuss five extensions.

First, for some illiquid assets like housing or large durables such as vehicles, the most significant component of their return is the service flow they provide to the owner. At the same time, they have a consumption commitment component, meaning periodic expenditures that cannot be avoided such as maintenance and repair. Consider a version of our model with the following in period $t = 1$. The illiquid asset yields a utility flow ϕa proportional to the stock, and these services are perfect substitutes with c_1 (housing can be rented out and thus transformed into c_1); and the illiquid-asset's owner must incur expenditures κa . Then, the counterpart of condition (7) is one where R is simply replaced by $R/(1 - \kappa + \phi)$, the effective return on the illiquid asset.

Second, when the agent can access unsecured credit, there is a second kink in the budget constraint at the credit limit; this is in addition to the kink at $m_2 = 0$. The model in Appendix A shows that in this case, households can be W-HtM and P-HtM either at the zero kink or at the credit limit.

Third, as we show in Kaplan and Violante (2014a), in the presence of income uncertainty, a W-HtM prefers bearing the welfare loss from income fluctuations rather than holding the large balances of cash required for consumption smoothing. Saving in the liquid asset means foregoing the high return on the illiquid asset and the associated higher level of long-run consumption. This explanation is reminiscent of calculations made by Cochrane (1989), Krusell and Smith (1996), and Browning and Crossley (2001) who demonstrated that in several different contexts the utility loss from setting consumption equal to income, instead of fully optimizing, can be second order.

Fourth, in the model the illiquid asset is inaccessible in the intermediate period. In a

⁴In fact, Kaplan and Violante (2014a) show that, in a richer life-cycle version of this two-asset model with uninsurable income risk, the average MPC out of transitory income shocks among W-HtM households is larger than the average MPC among P-HtM households. We return to this point in Section 8.

more general environment where the illiquid asset can be accessed by paying a fixed transaction cost, the household may decide to deposit an unexpected positive windfall into the illiquid account, or to smooth a negative shock by withdrawing from the illiquid account. This behavior could potentially alter the model’s implications for the MPC of W-HtM agents. In Kaplan and Violante (2014a), we show that this is the case only if the shock is large relative to the transaction cost. We return on this point in Section 8.

Finally, in our two-period model, we have abstracted from discounting, but it is easy to see that with geometric discounting between periods, all the qualitative conclusions remain intact. Hyperbolic discounting introduces an additional reason to save in illiquid assets, since illiquidity protects quasi-hyperbolic households from future consumption splurges (see Angeletos et al., 2001; Laibson et al., 2003), and therefore makes it even easier to generate W-HtM behavior.

3 Identifying hand-to-mouth households in the data

The stylized model in Section 2 illustrates that there are two types of HtM households: poor hand-to-mouth (P-HtM) who do not hold any illiquid wealth, and wealthy hand-to-mouth (W-HtM) who own positive amounts of illiquid wealth. For each type of HtM household, there are two kinks in the intertemporal budget constraint where MPCs out of small income changes can be large: zero liquid assets and the unsecured credit limit.⁵ According to the theory, a household is HtM at the zero kink in period t if it consumes all its cash-on-hand for the period, and carries zero liquid wealth between t and $t + 1$. Similarly, a household is HtM at the credit limit if, at the end of period t , it has borrowed up to the limit.

Given the theoretical definition of HtM status, ideally, we would observe balances of liquid wealth at the end of the pay-period — the period that starts at income receipt and ends just before the next income receipt. Unfortunately, surveys either report average balances over the period, or balances at a random point in time (the interview date). As a result, HtM status will be measured with error.

To fully understand this issue, think about a continuous-time generalization of the

⁵The unsecured credit limit is always a hard constraint. The zero liquid asset position is a hard constraint for the subset of households who do not have access to credit, and a kink for virtually all others, since the interest rates on credit cards and other non-collateralized loans are typically much larger than the return on liquid assets.

model in Section 2 where income is paid discretely at the beginning of the period as liquid wealth, but consumption occurs continuously –and is constant– over the period. Then, given the timing mismatch between the discrete income payment and the continuous consumption expenditures, one expects to observe positive (or above credit limit) balances of liquid wealth even for the HtM households: this makes their identification especially challenging. In Appendix B, we lay out this enriched version of the model.

We now describe our identification strategy —which builds upon the one used in Kaplan and Violante (2014a)— starting with the case where liquid balances observed from the survey are an average over the period.

Average balances. Let y_{it} denote the income of household i in pay-period t , let a_{it} denote holdings of illiquid wealth, and let m_{it} denote average balances of liquid wealth over the pay period.

Panel (a) of Figure 1 depicts a graphical representation of the dynamics of income and average cash-on-hand m_{it} over a pay period for a HtM household who starts and ends the period at the zero kink. Its liquid balances peak at y_{it} , when income is paid into the liquid account at the beginning of the pay period, and are depleted constantly until they reach zero at $t + 1$. Average balances over the period are equal to half income.

A conservative criterion to identify HtM agents on the zero kink in the data is therefore to count those survey households whose average liquid wealth balances are positive (to capture the fact they are not borrowing), but are equal to or less than half their earnings per pay-period, where “half” is due to the assumption that resources are consumed at a constant rate. Specifically, a household is P-HtM at the zero kink if

$$a_{it} \leq 0, \quad \text{and} \quad 0 \leq m_{it} \leq \frac{y_{it}}{2} \tag{8}$$

and W-HtM at the zero kink if

$$a_{it} > 0, \quad \text{and} \quad 0 \leq m_{it} \leq \frac{y_{it}}{2}. \tag{9}$$

The case $a_{it} < 0$ is very rare in survey data. It occurs when housing equity is negative because a decline in house prices has pushed the market value of the house below the residual value of the mortgage. We include these households among the P-HtM because, even though they own illiquid assets, they effectively have no means of using them to smooth consumption and, as such, these households are more similar to the P-HtM.

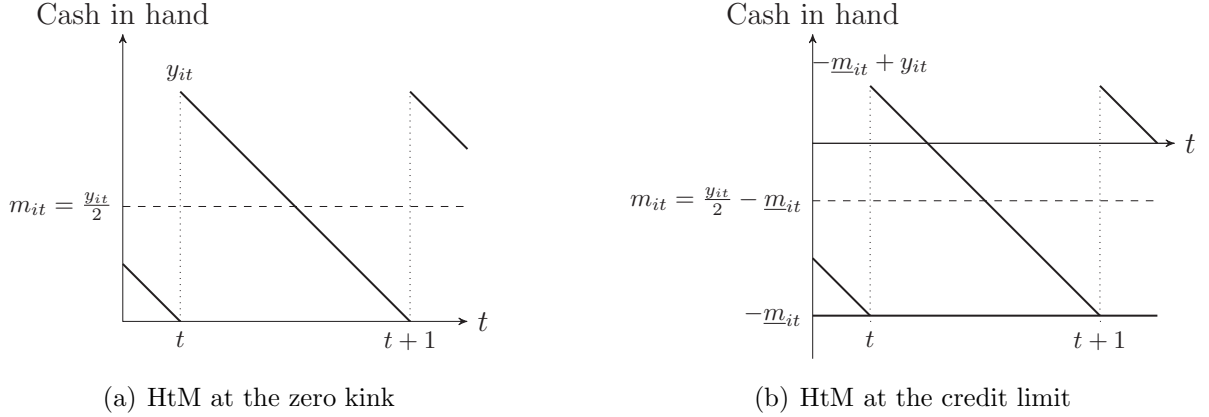


Figure 1: Illustration of two cases of HtM behavior

This estimator of the number of HtM households provides a *lower bound* because, although all N-HtM households would always hold average liquid balances above half their earnings, some HtM households may also hold, on average, liquid balances above half their earnings. For example, a household that starts the period with positive liquid savings, in addition to its earnings, and ends the period with zero liquid savings is HtM, but its average liquid balance is above half its earnings, and so it would not be counted as HtM by this criterion. Appendix B makes this point formally.

Next, consider a HtM household at the credit limit $-\underline{m}_{it} < 0$. This is a household that consumes all its cash-on-hand for the period, as well as all its available credit. For consistency with the strategy above, we propose to count a household as P-HtM at the credit limit if

$$a_{it} \leq 0, \quad m_{it} \leq 0 \quad \text{and} \quad m_{it} \leq \frac{y_{it}}{2} - \underline{m}_{it}, \quad (10)$$

and to count it as W-HtM at the credit limit if

$$a_{it} > 0, \quad m_{it} \leq 0 \quad \text{and} \quad m_{it} \leq \frac{y_{it}}{2} - \underline{m}_{it}. \quad (11)$$

Panel (b) of Figure 1 depicts the dynamics of income and average cash-on-hand m_{it} over a pay period for a HtM household that starts and ends the period at the credit limit. It is easy to see that this criterion is also conservative: a household that starts the period at t with liquid wealth above its credit limit and ends the period at $t+1$ having exhausted all its borrowing capacity, would carry an average balance above the limit, and would therefore escape our criterion based on equations (10) and (11).

Balances at a point in time. Some surveys report balances of liquid wealth at

the interview date, which can be thought of as a random point during the pay-period. Is it still true in this case that our estimator based on criteria (8)-(11) provides a lower bound on the fraction of HtM households? In Appendix B we show that we would always miss some truly HtM households. However, we may mistake a non-HtM households for HtM if its end-of-period liquid balances are less than one-half of its income away from zero or from the credit limit if it is borrowing. For a bi-weekly pay-period, this means that the only problematic households are those with one week or less of income in excess of their kink —households which, for practical purposes, one may want to identify as HtM anyway.

Consumption commitments. Recent literature has emphasized the existence of pre-committed consumption expenditures —expenditures that a household is committed to incur every pay-period, unless it pays a transaction cost (either monetary or in terms of time) to modify its previous commitments (see, for example, Chetty and Szeidl, 2007; Shore and Sinai, 2010). These expenditures include rent, mortgages or other loan payments, utility bills, fees for school, gym, or clubs, and alimony. The key feature of committed expenditures is that they are bulk expenditures incurred at a point in time that discretely deplete a household’s balance of liquid wealth.

How does the presence of such expenditures affect our identification strategy? Let \bar{c}_{it} be the amount of committed expenditures for household i at date t . If \bar{c}_{it} is incurred at the beginning of a pay period, the criterion to identify a HtM household (say, at the zero kink) should be amended as $m_{it} \leq (y_{it} - \bar{c}_{it}) / 2$, while if it is incurred at the end of the period, the criterion should be $m_{it} - \bar{c}_{it} \leq y_{it}/2$. In the first case, our baseline measurement overestimates HtM status, and in the second case it underestimates it. Instead, if committed expenditures are incurred smoothly over the period or are paid in the middle of the pay period, then the criterion should be $m_{it} - \bar{c}_{it}/2 \leq (y_{it} - \bar{c}_{it}) / 2$ which is the same as our baseline measurement. We verify the robustness of our estimates with respect to those consumption commitments that we can measure in our survey data by using these alternative assumptions about the timing of expenditures.

Definition of HtM in terms of net worth. For comparison with net-worth based theories of HtM behavior, we also compute the fraction of HtM agents in terms of net worth. Let $n_{it} = a_{it} + m_{it}$ be the net worth of agent i in period t . Then, a household is HtM in net worth (HtM-NW) if

$$0 \leq n_{it} \leq \frac{y_{it}}{2} \quad \text{or,} \quad n_{it} \leq 0 \quad \text{and} \quad n_{it} \leq \frac{y_{it}}{2} - \underline{m}_{it} \quad (12)$$

Direct survey questions. Finally, whenever the data allow, we also use direct survey questions as alternate estimates of the fraction of HtM households. These questions typically ask whether expenditures over the last month have exceeded income, abstracting from purchases of large durable goods such as housing or cars, and whether the household usually spends more than its income. Counts of HtM households derived from these questions provide a useful check on the reliability of our identification strategy based on reported liquid wealth and income.

4 Survey data on household portfolios

The countries included in our study are the U.S., Canada, Australia, the U.K., and the four largest economies in the euro area: Germany, France, Italy, and Spain. Data for the first four countries come from their own separate surveys, the U.S. Survey of Consumer Finances (SCF), the Canadian Survey of Financial Security (SFS), the Household, Income and Labour Dynamics in Australia (HILDA) survey, and the U.K. Wealth and Assets Survey (WAS). Data for the euro area countries come from the Household Finance and Consumption Survey (HFCS), a joint project administered by all of the central banks of the Eurosystem. Appendix C contains a detailed description of all these cross-sectional surveys.

In order to categorize a household as W-HtM, P-HtM, or N-HtM, we need information on its labor income and on the amounts of assets and liabilities held in various categories of its balance sheet. In the rest of this section, we discuss sample selection and comparability across surveys. Next, we present some descriptive statistics on the asset and liability distribution across countries.

4.1 Sample selection and data comparability

Each individual survey is tailored to its own country and, as such, the questions asked and the definitions of particular asset classes vary across surveys. Our main goal is to be as consistent as possible in selecting the sample, and in defining income, liquid, and illiquid wealth across surveys.

Sample selection. In all surveys, we restrict our analysis to households in which

	U.S.	Canada	Australia	U.K.	Germany	France	Italy	Spain
Survey	SCF	SFS	HILDA	WAS	HFCS	HFCS	HFCS	HFCS
Years	1989-2010	2005	2010	2008-10	2008-10	2008-10	2008-10	2008-10
Initial sample size	35513	5267	7317	18510	3565	15006	7951	6197
<u>Exclusions</u>								
Not age 22-79	2098	373	782	1655	246	1428	846	559
Negative income	9	10	0	0	0	0	0	0
All inc. from self empl.	4334	—	202	334	228	890	721	658
Final sample size	29072	4884	6333	18176	3091	12688	6384	4980

Table 1: Summary information on the survey data used. Self employment income is not provided in the SFS for Canada.

the head is between 22 and 79 years of age, and drop households only if their income is negative or if all of their income originates from self-employment.⁶ Table 1 summarizes the survey years we use for each country, the sample selection, and the final sample sizes. Since all these surveys oversample the rich, we always use weights to construct sample statistics.

Income. In choosing our definition of income, we make an attempt to include all labor income plus government transfers that are regular inflows of liquid wealth. We exclude interests, dividends, and other capital income because they are realized more infrequently. Income in the SCF is gross wages and salaries, self-employment income, regular private transfers such as child support and alimony, public transfers such as unemployment benefits, food stamps, and Social Security Income (SSI), and regular income from other sources excluding investment income. Income in the Canadian SFS is after-tax total income. There is no distinction between labor, capital, and self-employment income. In the HILDA, income is wages and salaries, self-employment income, regular private transfers such as child support and alimony, and public benefits such as the Australian Government Parenting Payment. For the U.K. WAS, we define income as net employee earnings, net self-employment income, plus any public benefits such as the Jobseeker’s Allowance and Maternity Allowance. Income in the HFCS is gross income from wages, salaries, and self-employment, unemployment benefits, regular private transfers such as child support and alimony, and regular public transfers.⁷

⁶The only exception to our age range is for the U.K. WAS which provides ages in 5 year age bins, so we include households with heads between 20 and 79 years of age.

⁷The reference period for the income questions differs between surveys. For income variables in the SCF, the survey asks for annual income in the previous year. For example, the 2010 SCF uses 2009 as its reference period for income. The income reference period differs by country in the HFCS. France

The main discrepancy in income measurement across surveys is that income in Canada is reported after taxes, whereas all other countries survey gross income before taxes. For most households, except the self-employed, taxes are withheld at the source and hence the amount paid into the liquid account —and available for spending— is net of taxes. Thus, using income before taxes does somewhat overstate the fraction of HtM households by inflating the liquid wealth threshold. Whenever possible, we verify the robustness of our results to an adjustment for the individual tax liability.

Liquid wealth. In the U.S. SCF, we consider liquid assets to be checking, saving, money market and call accounts plus directly held mutual funds, stocks, corporate bonds and government bonds. Liquid assets in the Canadian SFS are deposits in financial institutions plus holdings in mutual funds, other investment funds, stocks and bonds. In the Australian HILDA, liquid assets include balances in bank accounts, equity investments, and cash investments (bonds). In the U.K. WAS, liquid assets include bank accounts, Individual Savings Accounts (ISAs), and holdings of shares, corporate bonds, and government bonds.⁸ For the Euro area HFCS, liquid assets are cash, sight (also called current, draft, or checking) accounts, mutual fund holdings, shares in publicly traded companies, and corporate or government bond holdings.

The main shortcoming in the definition of liquid wealth is the absence of information on cash holdings. To address this problem, we resort to an imputation procedure based on data from the 2010 Survey of Consumer Payment Choice, administered by the Federal Reserve Bank of Boston (see Foster et al., 2013). We compute the ratio of average cash holdings measured in this survey to the median value of checking, saving, money market and call accounts from the 2010 SCF. We then inflate the value of each household’s checking, saving, money market and call accounts by this ratio in all surveys.⁹

We define liquid debt in the SCF as the sum of all credit card balances that accrue

and Germany both use 2009 as a reference period, Spain uses 2007, and Italy uses 2010. Wave Two of the WAS (2008-2010) asks questions regarding the “usual” amounts for monthly income and benefits. The 2005 SFS uses 2004 as its reference period, and gave its respondents the option of skipping the income questions and using linked data from their 2004 tax return. Wave Ten of the HILDA uses the 2009-2010 financial year which runs from July 1, 2009 to June 30, 2010 for its reference period for income.

⁸ISAs are accounts designed for the purpose of saving with a favorable tax status. A broad range of asset categories, including cash, can be held in ISAs. There are no restrictions to how much and when funds can be withdrawn.

⁹Average cash holdings, excluding large-value holdings in 2010 was \$138. Median checking, saving, money market and call accounts in the 2010 SCF is \$2500, making the ratio about 5.5%. In the HFCS, information on cash holdings is available for Spain from a non-core module. We check the median ratio of cash to sight accounts and find it to be about 5% in Spain.

interest, after the most recent payment. Liquid debt in the SFS is credit card and installment debt. Liquid debt in the HILDA is credit card debt. In the U.K. WAS, liquid debt is credit card debt, plus any balances on store cards, hire purchases, and mail orders. In the HFCS, liquid debts are considered to be the balance on credit cards, after the most recent payment, which accrue interest, and any balances on credit lines or bank overdrafts which also accrue interest.

The measure of liquid wealth that we use to compute HtM status is net liquid wealth, or liquid assets minus liquid debt. We also examine a narrower definition of net liquid wealth that excludes directly held mutual funds, stocks, and bonds from liquid assets, and a broader one that includes outstanding debt in home-equity lines of credit as liquid debt.

Illiquid wealth. Net illiquid wealth in the SCF includes the value of housing, residential and non-residential real estate net of mortgages and home equity loans, private retirement accounts (such as 401(k)s, IRAs, thrift accounts, and future pensions), cash value of life insurance policies, certificates of deposit, and saving bonds. Illiquid wealth in the Canadian SFS is the value of the principal residence and other real estate investment less mortgages on the properties and lines of credit that use property as collateral. It also includes retirement savings such as Registered Retirement Savings Plans, Registered Retirement Income Funds, employer pension plans, and other retirement funds. In the HILDA, illiquid wealth is net equity in home and other real-estate properties plus life insurance policies and superannuation (government-supported, compulsory private retirement funds).¹⁰ In the U.K. WAS, we take illiquid wealth to include the value of the main residence, other houses, and land net of mortgages and land debt, plus occupational and personal pensions, insurance products, and National Savings products. The definition of net illiquid wealth in the HFCS is the value of the household's main residence and other properties net of mortgages and unsecured loans specifically taken out to purchase the home, plus occupational and voluntary pension plans, cash value of life insurance policies, certificate of deposits, and saving bonds.

We also explore broader definitions of illiquid wealth that include the value of businesses for the self-employed, the resale value of vehicles net of the loans taken out to purchase them, and other non-financial wealth not included in our baseline, such as antiques,

¹⁰Superannuation has some features of private retirement accounts, such as 401(k) accounts in the U.S., which we include in illiquid wealth, and some features of public pensions (the compulsory nature of a minimum contribution) which we exclude from illiquid wealth. Because of this ambiguity, we also offer a sensitivity analysis where we exclude superannuation wealth from illiquid assets.

artwork, jewels, and gold.¹¹ Changing the definition of illiquid wealth affects only the split between poor and wealthy HtM, but not the total number of HtM households.

Reference period. The reference period for the liquid and illiquid wealth questions varies across surveys. In the SCF, it is the interview date for most assets, but for some, such as checking and saving accounts, when the respondent was unsure, the interview could prompt for an average balance over the month. The Canadian SFS asks for information on assets and debts for “a time as close as possible to the date of the interview.” Both the WAS and HILDA ask for current balances or values of assets and liabilities. In the HCFS, France, Germany, and Spain use the interview date, and Italy uses December 31, 2010.

4.2 Descriptive statistics

Table 2 reports some basic descriptive statistics on household income, liquid and illiquid wealth holdings, and portfolio composition, for each country in the sample.

In all countries, the typical household portfolio structure is rather simple. It comprises a small amount of liquid wealth in the form of bank accounts, some housing equity, and a private retirement account. In particular, the median holdings of other financial assets such as directly held stocks, bonds, mutual funds, and life insurance are zero everywhere. This is a well known fact in the empirical study of household portfolios (see Guiso, Halassios, and Jappelli, 2002).

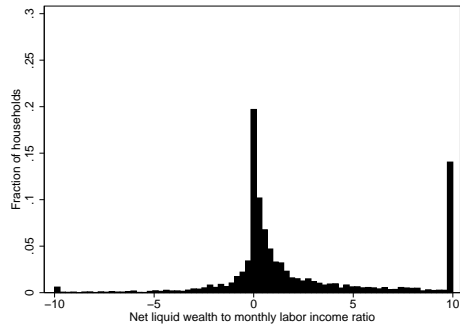
There are, however, some interesting cross-country differences. First, the median net worth to median income ratio varies a lot across countries: from just above one in Germany and the U.S. to over six in the U.K., Italy, and Spain. With respect to net liquid wealth, consumer credit appears a lot less frequent in the Euro area: less than 10 percent of households have credit card debt in France, Italy, and Spain, compared to 30 to 40 percent in the Anglo-Saxon countries. Figure 2, which plots the distribution of net liquid wealth to monthly income for the eight countries, reinforces this observation.

Housing equity forms the majority of illiquid wealth for households in every country, with the exception of Germany where median housing wealth is zero, since only 48 percent of the population are homeowners. This homeownership rate is at least 10 percentage points less than in all other countries (see also Eymann and Börsch-Supan,

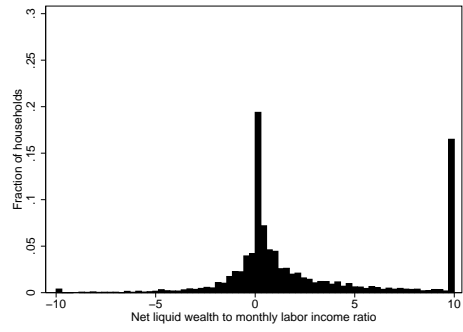
¹¹In our robustness checks with respect to business equity we include all households whose income is entirely from self-employment as long as they had non-negative income from their business.

	US		CA		AU		UK	
	Median	Frac. Pos.	Median	Frac. Pos.	Median	Frac. Pos.	Median	Frac. Pos.
Income (age 22-59)	47040	0.984	49905	1.000	79555	0.993	29340	0.979
Net Worth	56721	0.883	112418	0.877	380889	0.984	187157	0.880
Net liquid wealth	1714	0.750	2643	0.716	12139	0.880	2111	0.632
Cash, checking, saving, MM accounts	2640	0.923	2873	0.864	8709	0.978	2639	0.766
Directly held stocks	0	0.142	0	0.109	0	0.351	0	0.160
Directly held bonds	0	0.014	0	0.106	0	0.015	0	0.154
Revolving credit card debt	0	0.382	0	0.412	0	0.296	0	0.405
Net illiquid wealth	52000	0.761	100713	0.752	347500	0.939	174999	0.843
Housing net of mortgages	29000	0.629	64238	0.648	250000	0.714	81400	0.677
Retirement accounts	1508	0.526	871	0.518	61000	0.863	58560	0.766
Life insurance	0	0.186	0	0.033	0	0.064	0	0.110
	DE		FR		IT		ES	
	Median	Frac. Pos.	Median	Frac. Pos.	Median	Frac. Pos.	Median	Frac. Pos.
Income (age 22-59)	35444	0.994	31518	0.999	26116	0.987	26961	0.991
Net Worth	46798	0.949	108976	0.966	165420	0.919	178925	0.967
Net liquid wealth	1319	0.853	1453	0.925	5226	0.769	2685	0.890
Cash, checking, saving, MM accounts	1154	0.876	1255	0.953	4181	0.769	2261	0.908
Directly held stocks	0	0.110	0	0.151	0	0.043	0	0.106
Directly held bonds	0	0.050	0	0.015	0	0.146	0	0.014
Revolving credit card debt	0	0.225	0	0.076	0	0.049	0	0.086
Net illiquid wealth	39306	0.876	104214	0.922	148524	0.803	171161	0.885
Housing net of mortgages	0	0.476	86372	0.607	148524	0.716	162491	0.847
Retirement accounts	0	0.245	0	0.039	0	0.088	0	0.037
Life insurance	0	0.493	0	0.378	0	0.193	0	0.245

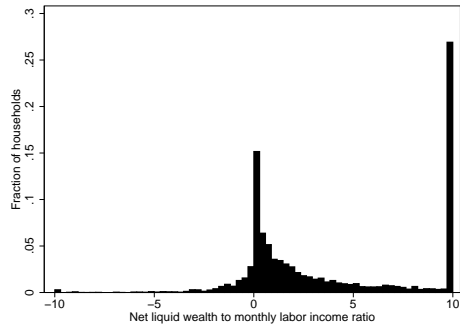
Table 2: Data for the U.S. are from the 2010 survey only. All figures are in local currency units. Data for Canada is adjusted to 2010 CA\$ using the Canadian CPI. From the Federal Reserve Board's G.5 release, the average exchange rates in the survey years are 1.2 CA\$, 1.1 AU\$, 0.6 British pounds, and 0.7 euros per U.S. dollar.



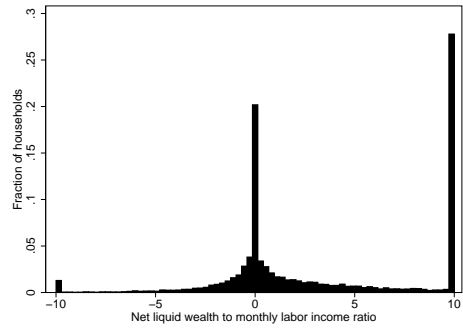
(a) United States



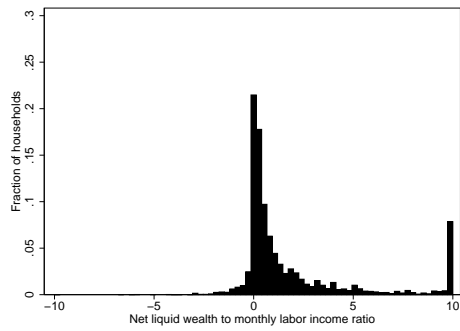
(b) Canada



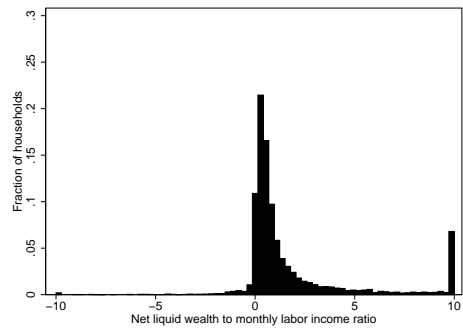
(c) Australia



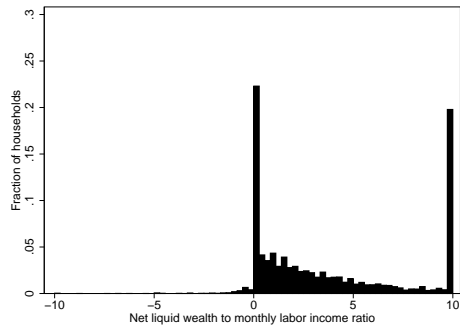
(d) United Kingdom



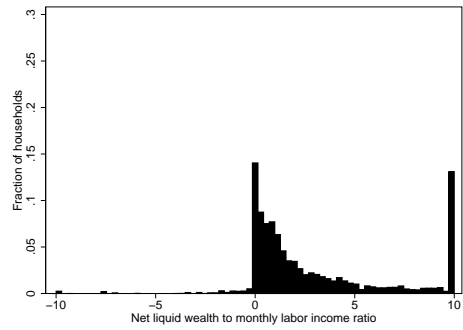
(e) Germany



(f) France



(g) Italy



(h) Spain

Figure 2: Distribution of liquid wealth to monthly income ratios by country.

2002). The median value of housing equity relative to median annual income is especially remarkable in Italy and Spain, where it exceeds six.

There are also large differences in the fraction of households with positive private retirement wealth: in the Anglo-Saxon countries, at least half of all households hold a personal retirement account, whereas in France, Italy, and Spain less than one in ten do. Surely, a big part of the explanation is in the generosity of the public pension system in these countries: according to the OECD, replacement rates for the median earner are between 60 and 70 percent in these countries, compared to 40 percent in the U.K. and the United States (see OECD, 2013). The size of private retirement wealth in Australia and the U.K. is astonishing. In Australia, this is partly due to the “superannuation” regulations that require all employers to generously contribute to tax-deferred retirement accounts on behalf of their employees.¹² In the U.K., the Pension Schemes Act of 1993 created tax-free employer-sponsored (defined benefits) occupational pensions and (defined contributions) personal pensions. The Pension Act of 2008 established that workers must choose to opt out of an employer’s occupational pension plan, rather than opt in (see Banks and Tanner (2002) for more details).

Finally, the proportion of households with life insurance in their portfolio is much higher in the euro area than in the Anglo-Saxon countries. We conjecture that solid intergenerational family ties, and a stronger precautionary savings motive linked to the lower female participation rate, may account for these differences.

5 United States

In this section, we report the main findings for the U.S., using data from the 1989-2010 waves of the SCF. We begin by estimating the fraction of HtM households and assessing the robustness of our estimates to a variety of aspects of the definition adopted in Section 3. We then analyze the key demographic characteristics of N-HtM, P-HtM, and W-HtM households, and we examine their portfolio composition in more detail.

¹²In the survey years, the compulsory minimum employer contribution rate was 9 percent of the employee salary.

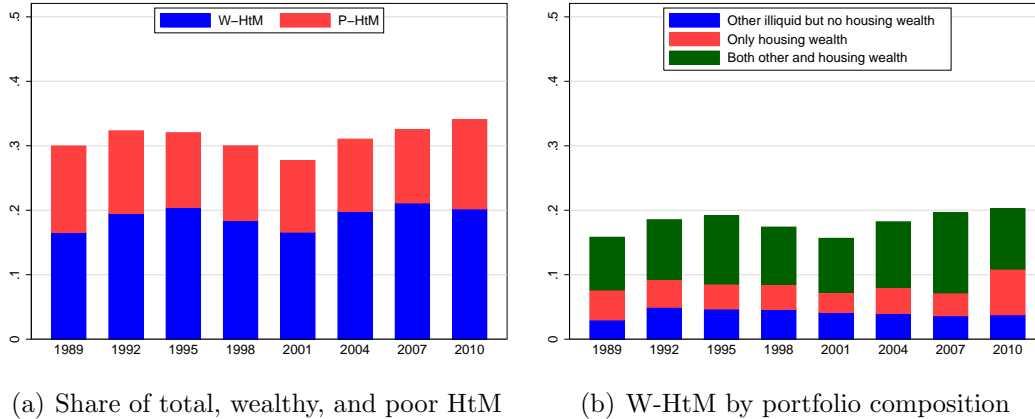


Figure 3: Time-series of fraction of HtM households in the U.S.

5.1 The share of HtM households

Our definition of HtM status is based on equations (8)-(12). Since the SCF does not report individual data on the frequency of pay, we need to make an assumption that applies to all households. Consumer Expenditure Survey (CEX) data from 1990-2010 reveal that 32 percent of respondents are paid weekly, 52 percent of respondents are paid bi-weekly, and the rest are paid = monthly or at lower frequencies.¹³ Based on these findings, in the benchmark analysis we set the pay-frequency to two weeks. In the benchmark, we also set the household credit limit to one month of income. The SCF asks respondents to report their credit limit, but most of the other surveys do not, and hence for comparability we choose a common limit.¹⁴

Panel (a) of Figure 3 plots the fraction of HtM households in the U.S. population over the period 1989-2010 and their split between wealthy and poor HtM. Our estimates indicate that, on average, 31 percent of U.S. households are HtM over this period. Of these, roughly one-third are poor HtM and two-thirds are wealthy HtM. This our paper's first main result: the vast majority of hand-to-mouth households own illiquid assets. Looking at changes over time across the two decades covered by our data, the fraction of HtM households remains fairly stable and the split between poor and wealthy does not change significantly. The first line of Table 3 reports that the share

¹³We thank Yiwei Zhang for providing us with these tabulations based on Zhang (2014).

¹⁴The choice of one month of income for the benchmark is consistent with the SCF self-reported limits. When setting the limit for households without credit cards to zero, the median self-reported limit to income ratio is 0.54 in 1989. It grows steadily to 1.7 in 2007 and then drops to 1.2 in 2010. This evolution of credit limits is even more remarkable when conditioning only on credit card holders (around 70 percent of the population): the median limit to income ratio rises from 1.2 in 1989 to 3.4 in 2007, and then drops to 2.8 in 2010.

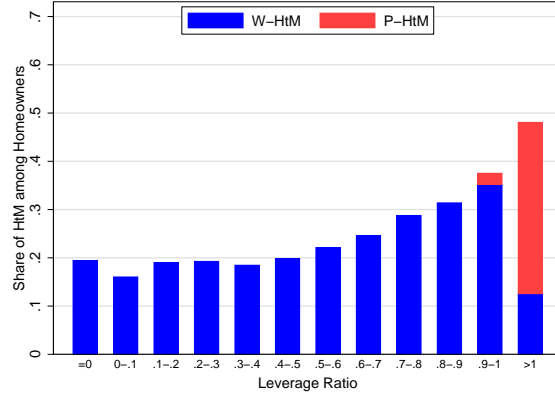


Figure 4: Share of HtM households among homeowners by leverage ratio. SCF 1989-2010.

of U.S. households that are HtM in terms of net worth is less than 14 percent. Thus looking at the wealth distribution through the eyes of net worth misses over half of the HtM households in the United States.¹⁵

Panel (b) explores the illiquid asset portfolio of the W-HtM households by plotting the share of W-HtM households that own housing, retirement wealth, or both. About one-half of W-HtM have both, about a third have positive housing but no retirement wealth, and a sixth have positive retirement wealth but no housing. A deeper look into the portfolio of HtM households reveals that, if we condition on homeownership, the leverage ratio is a strong predictor of HtM status. Figure 4 shows that the fraction of HtM households doubles from 20 to 40 percent as the leverage ratio rises towards one, as regular mortgage payments absorb a significant fraction of disposable income and leave households with little or no liquid savings.

5.1.1 Robustness

Figure 5 and Table 3 summarize our sensitivity analyses. Panel (a) of Figure 5 plots the shares of poor and wealthy HtM weighted by income. Not surprisingly, the total fraction of HtM households is smaller than its unweighted counterpart: HtM households represent roughly 20 percent of total U.S. income, since their income is below the U.S. average. When we weight by income, the W-HtM represent three-quarters of HtM households. Panel (b) plots HtM shares when the pay-period is set to a month instead of two weeks. The fraction of HtM households increases by 9 percentage points and

¹⁵HtM-NW are always more numerous than the P-HtM because there are some households with liquid wealth above the threshold, who are therefore not HtM, but with enough negative illiquid wealth (i.e., negative home equity) to push their net worth below the threshold.

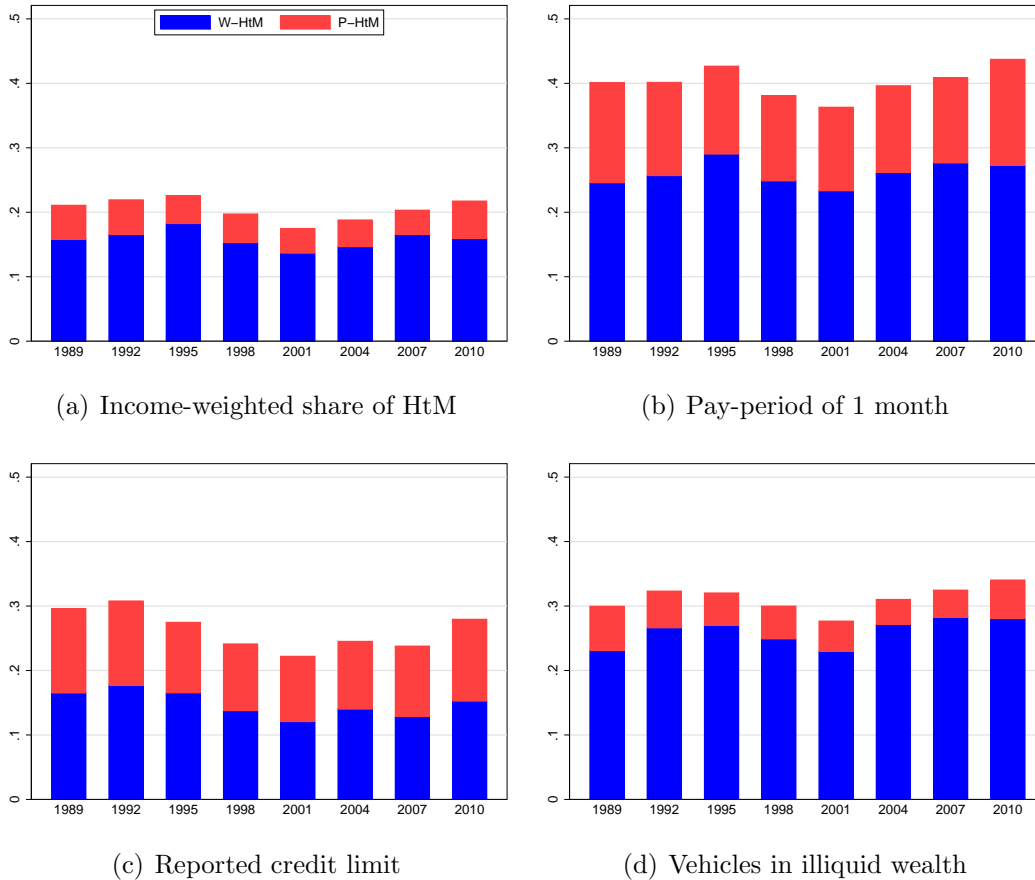


Figure 5: Time series of fraction of HtM households in the U.S., alternate definitions.

W-HtM account for most of the difference with the baseline. Symmetrically, the fourth line of Table 3 shows that, when the pay-period is set to one week, the share of W-HtM drops by 5 percentage points. In panel (c) of Figure 5, we verify the robustness of our estimates with respect to the tightness of the credit limit. When we use the self-reported credit limit in the SCF, the fraction of HtM households drops by 5 percentage points, with a lower number of W-HtM households accounting for all of the drop. Finally, panel (d) shows that by including vehicles as illiquid wealth, we move roughly half of the P-HtM into the W-HtM group but, by construction, the total share of HtM households in the population is unchanged.

Table 3 contains a number of other sensitivity analyses. We begin with direct questions on HtM status. The SCF contains a combination of sequential questions aimed at assessing whether “over the past year, [household] spending exceeded, or was about the same as, income, and such expenditures included purchases of a home or automobile or spending for any investments.”¹⁶ Based on this definition, the share of HtM households

¹⁶These questions (X7510, X7509, X7508) were included in the survey starting from 1992.

	P-HtM	W-HtM	N-HtM	HtM	HtM-NW
Baseline	0.121	0.192	0.688	0.312	0.137
In past year, $c > y$	0.130	0.309	0.561	0.439	—
Usually, $c > y$	0.089	0.156	0.756	0.244	—
Financially fragile households	0.173	0.331	0.497	0.503	0.209
Reported credit limit	0.114	0.147	0.738	0.262	0.126
1 year income credit limit	0.102	0.118	0.780	0.220	0.108
Weekly pay period	0.106	0.150	0.744	0.256	0.119
Monthly pay period	0.141	0.261	0.598	0.402	0.164
Higher illiquid wealth cutoff	0.131	0.181	0.688	0.312	0.137
Ret. acc. as liquid for 60+	0.121	0.183	0.696	0.304	0.137
Businesses as illiquid assets	0.114	0.193	0.693	0.307	0.129
Direct as illiquid assets	0.120	0.217	0.663	0.337	0.137
Other valuables as illiquid assets	0.117	0.196	0.688	0.312	0.132
Excludes cc puzzle households	0.163	0.183	0.654	0.346	0.177
HELOCs as liquid debt	0.120	0.181	0.699	0.301	0.135
Usual income	0.119	0.198	0.683	0.317	0.137
Disposable income - Reported	0.121	0.188	0.691	0.309	0.137
Disposable income - Single	0.120	0.187	0.693	0.307	0.136
Comm. cons. - beg. of period	0.102	0.166	0.732	0.268	0.116
Comm. cons. - end of period	0.149	0.272	0.579	0.421	0.174

Table 3: Robustness results for fraction HtM in each category in the SCF pooled 1989-2010. Financially fragile households includes those households within \$2,000 in liquid assets of their income threshold as HtM. Higher illiquid wealth cutoff requires households to have above \$1,000 in illiquid assets to be considered W-HtM. Ret. acc. as liquid for 60+ puts retirement accounts into liquid wealth for households above age 60. Businesses as illiquid assets drops the self employment income sample selection and adds business assets to illiquid wealth and self employment income to income. Direct as illiquid assets classifies directly held mutual funds, stocks, corporate and government bonds as illiquid assets. Disposable income subtracts federal income taxes estimated from NBER’s TAXSIM from income. Disposable income - Reported assumes that each household files their actual marital status and number of children as dependents. Disposable income - Single assumes that every household files as single with no dependents. Comm. cons. - beg. of period assumes the household’s committed consumption is incurred at the beginning of the period. Comm. cons. - end of period assumes the household incurs it at the end of the period.

is around 44 percent. W-HtM households account for two-thirds of the total, and fluctuations in this measure over time very closely follow those in the baseline definition of Figure 5(a). The third row of Table 3 also reports results for another sequence of direct questions in the SCF. The first question asks households “Which of the following statements comes closest to describing your saving habits?” We label a household as HtM if it responds “Don’t save - usually spend more than (or as much as) income.” Roughly 24 percent of households are HtM according to this definition.

It is reassuring that our baseline estimate of HtM households sits in between the counts

based on these two direct questions, since we interpret the first question as providing an upper bound and the second a lower bound. Our baseline calculations refer to the current HtM status for a household. In the first set of direct questions, households that spent more than their income over the past year because they dissaved or borrowed are not truly HtM, but they would be classified as such based on the first set of direct questions. Conversely, the second set of direct questions asks about the usual HtM status, and therefore those households who are, at the time of the survey, transitorily into a HtM status would answer negatively to the question. The cross-sectional correlation between our indicator of HtM status and the one provided by these two questions is about 0.3 for each.

Our estimates of HtM households are related to calculations of “financially fragile” households by Lusardi, Schneider, and Tufano (2011). Based on an ad-hoc survey, they document that a quarter of U.S. households report that they would *certainly* be unable to come up with \$2,000 in 30 days, and a similar fraction reports that they could *probably* not come up with the funds to deal with an ordinary financial shock of this size. These authors also emphasize that there are many solidly middle-class households in this last group. In line three of Table 3, we compute the fraction of households who are less than \$2,000 away from the liquid wealth thresholds for being defined as HtM. We find that 50 percent of households are “financially fragile” according to this definition. Of these, 17 percent have no illiquid assets, but 33 percent own housing and/or retirement wealth. The P-HtM could be mapped into the Lusardi et al. survey respondents who would certainly not come up with this amount, and the W-HtM into those that would probably be unable to cope. Overall, our estimates are in line with those of Lusardi et al., but they also suggest a more nuanced interpretation. Households in the second group have the means to deal with a shock of this size, for example, by using their illiquid wealth as collateral for a loan. They may choose not to do it because the transaction costs involved dominate the welfare gain from smoothing such small shock, but for larger shocks, they will choose to adjust and smooth consumption. We return to this shock-size asymmetry of behavior in Section 8.¹⁷

The other robustness checks in Table 3 are conducted with respect to the definition of illiquid wealth, debt, income, and the timing of consumption expenditures. Using a higher illiquid wealth threshold in the definition of W-HtM (\$1,000 instead of \$1) moves about 1 percentage point of households from W-HtM into P-HtM. Broadening the definition of illiquid wealth to include business equity, or directly held stocks and

¹⁷Karen Pence (2011) makes a similar point in her discussion of Lusardi et al.

bonds, or other valuables (artwork, antiques, jewels, etc.) has small effects relative to the baseline.¹⁸ Including all private retirement wealth as liquid wealth for households aged sixty and above reduces the share of W-HtM households by less than 1 percentage point.

Around one-quarter of U.S. households simultaneously have positive liquid assets above $y/2$ and some revolving credit card debt.¹⁹ One may worry that many of these households have net liquid wealth close to zero, and they would therefore be counted as HtM, even though they have slack in both liquid wealth and credit. In Table 3 we show that excluding this group does not affect our calculations much because the distribution of HtM status within this group is not too different from the population distribution. Home equity lines of credit (HELOCs) were virtually nonexistent before the year 2000, but in the last decade they became a more common instrument to extract liquidity from housing.²⁰ Changing the definition of liquid debt by including used up HELOCs—while simultaneously increasing the credit limit by the total available line of credit—decreases the fraction of HtM households, as expected, but by only 1 percentage point.

The SCF collects data on a household’s normal, or usual, income as well as on their actual income. This alternate definition of income has no effect on our calculations. Recall that our definition of income is gross income before taxes and tax credits. Through the NBER TAXSIM, we have constructed, household by household, a measure of after-tax income.²¹ Under this income measure, the total fraction of HtM households declines, but quantitatively this effect is very small. The reason is that, in the U.S., the effective average tax rate is very small at the low end of the income distribution (around zero), mainly because of the EITC: even in the middle quintile it is only 10 percent. Finally, as explained in Section 3, accounting for committed expenditures has an ambiguous effect on the share of HtM agents, depending on whether the expenditures occur mostly at the beginning or at the end of the pay-period. Table 3 shows that these two opposite timing assumptions bound the share of total HtM households between 27 and 42

¹⁸When we include business equity, we also include in our sample all those households whose labor income comes entirely from self-employment. These households are excluded from the baseline sample.

¹⁹In the household finance literature, this observation is called the credit card puzzle (Telyukova 2013).

²⁰The fraction of home-owners with HELOCs in 2001 was 7.1%, 12.9% in 2007, and 10.7% in 2010. The average HELOC limit in 2001 was \$11,087, in 2007 it was \$18,984, and in 2010 it was \$19,070. The average percent of the HELOC used was 27.5% in 2001, 31.0% in 2007, and 31.6% in 2010.

²¹The variables we used in TAXSIM are year, marital status, the number of children, and the breakdown of income into its parts (wages, UI benefits, etc.). We deducted federal taxes from gross income. We assumed each household files their actual marital status and claims all their children living in the household as dependents. As an upper bound, we have also computed the case where they all file as single without dependents.

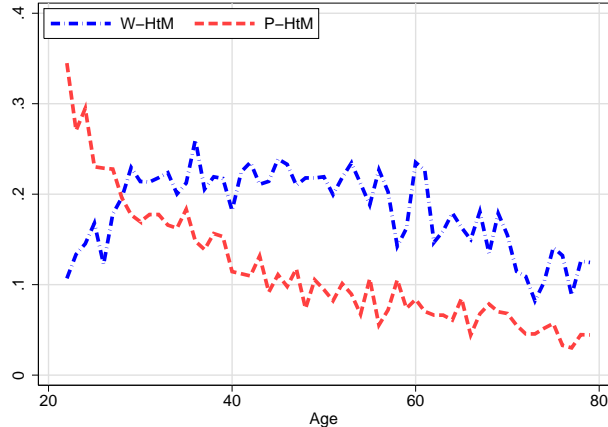


Figure 6: Age profile of fraction of HtM households in the U.S., pooled 1989-2010.

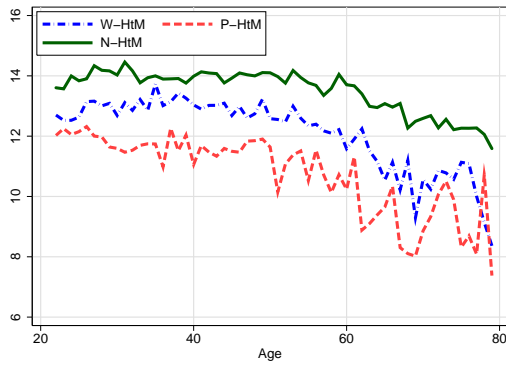
percent.

5.2 The demographics, portfolio composition, and status persistence of HtM groups

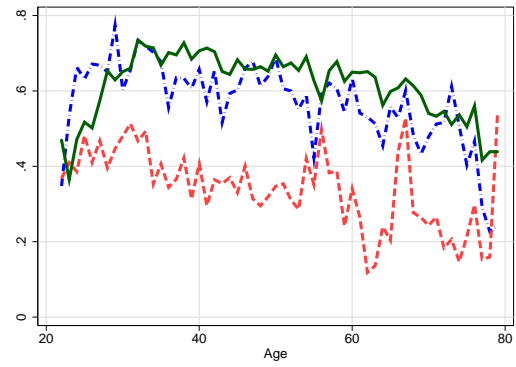
Demographics. We now turn to the demographic characteristics of the three groups of HtM households. Figure 6 plots the share of the population that is W-HtM and P-HtM by age.²² The bulk of P-HtM behavior is observed in the early stages of the life-cycle. The fraction of P-HtM households drops sharply until age 30, and keeps falling steadily over the life cycle until reaching roughly 5 percent in retirement. The age profile of the fraction of W-HtM households is instead markedly hump shaped: it peaks at around age 40, when over 20 percent of U.S. households are W-HtM, and remains above 10 percent throughout the life cycle. Accordingly, the share of N-HtM individuals increases steadily from 50 percent at age 22 to 80 percent in retirement.

The first three panels of Figure 7 report some demographic characteristics of the three HtM groups by age. N-HtM households have on average one more year of education than the W-HtM who, in turn, have one more year of education than the P-HtM. In terms of marital status, N-HtM and W-HtM households are indistinguishable, whereas the figure shows that the P-HtM households are 30 percent less likely to be married. In contrast, P-HtM and W-HtM are both more likely to have children than are N-HtM households.

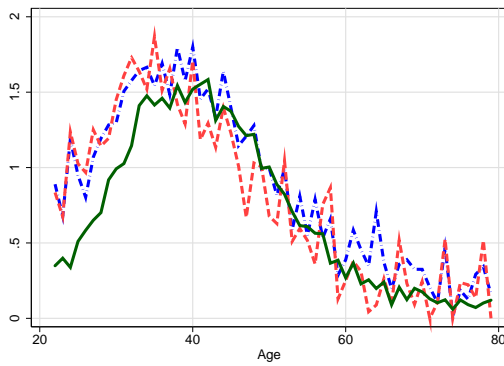
²²These plots are based on pooled data from all surveys and do not control for time or cohort effects. We verified that age profiles are similar in both cases, but become more noisy, and hence we present the raw data.



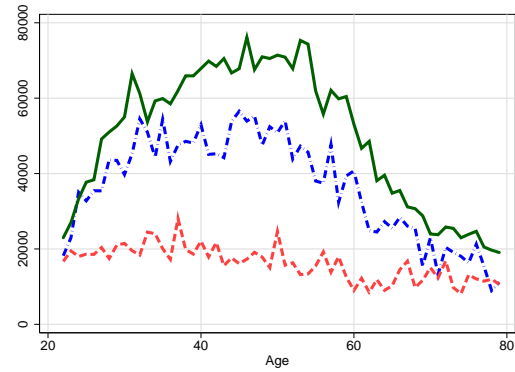
(a) Average years of education



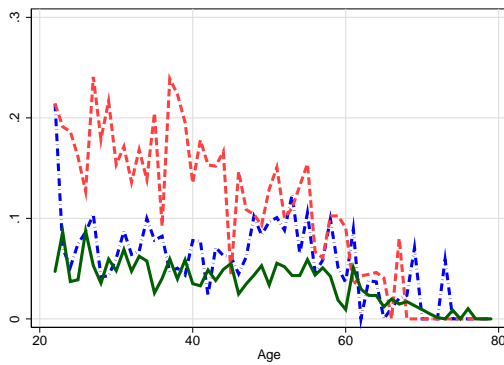
(b) Fraction married



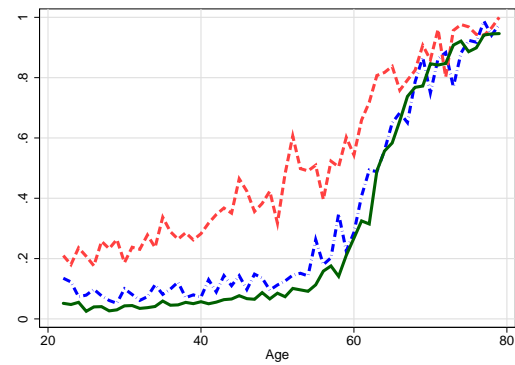
(c) Average number of children



(d) Median income



(e) Frac. with at least one unemp. member



(f) Frac. of income from gov. benefits

Figure 7: Age profile of demographic characteristics of the HtM in the U.S., pooled 1989-2010.

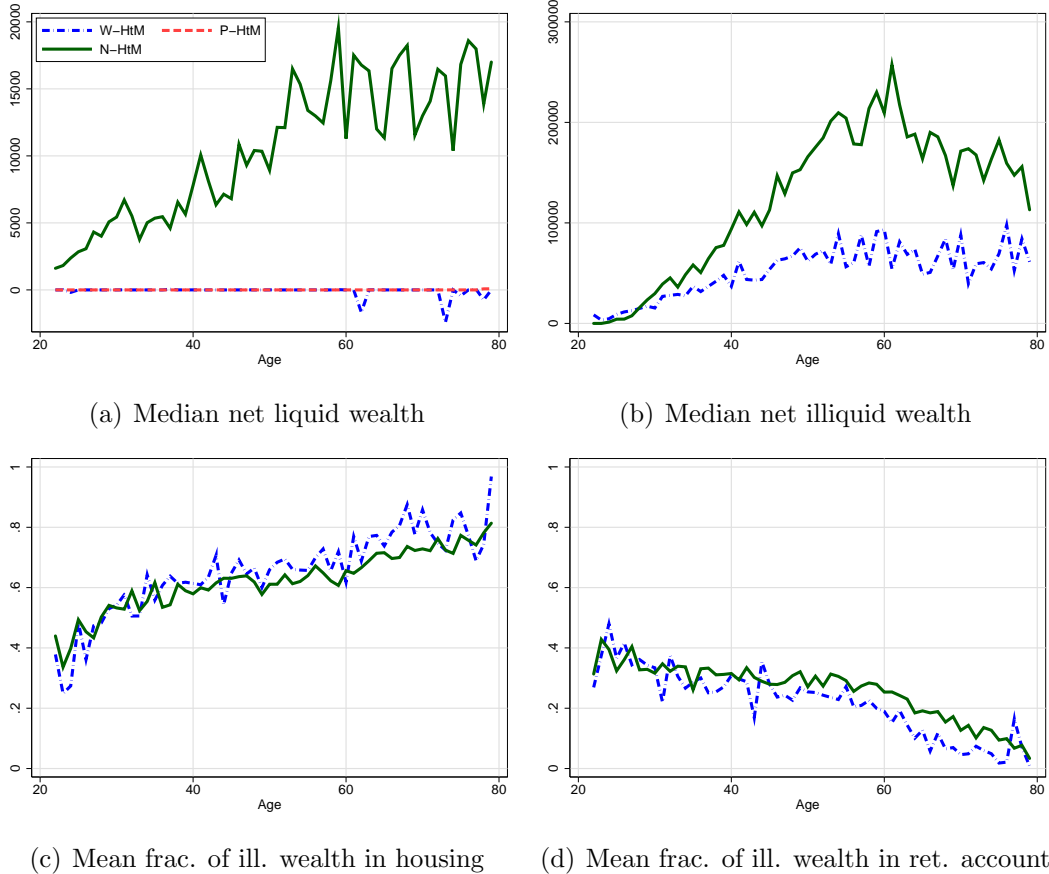


Figure 8: Age profile of the portfolio composition of the HtM in the U.S., pooled 1989-2010. To reduce the sensitivity to outliers, means are computed after trimming the overall top and bottom 0.1 percent of that statistic’s distribution.

Figure 7(d) shows that P-HtM households are income-poor, with median annual income around \$20,000 (in \$2010) during the working years, while the N-HtM are high-income households whose median earnings are \$70,000 at their life-cycle peak. The most surprising finding is that the W-HtM look a lot like the N-HtM in terms of their income path. The same conclusion holds for the incidence of unemployment and for the likelihood of receiving welfare benefits, which are both much lower for N-HtM and W-HtM households than for the P-HtM.

Portfolio composition. Figure 8 digs deeper into the balance-sheet composition of the three groups of HtM households. Panel (a) shows that median net liquid wealth holdings are zero at virtually every age for both the P-HtM and the W-HtM. Median net liquid wealth for N-HtM households grows steadily from about \$2,500 at age 25 until retirement, where it levels off at roughly \$15,000.²³ Panel (b) reveals that the

²³Recall, though, the overall median net liquid wealth across the whole population is less than

07 → 09	P	W	N
P	0.548	0.127	0.326
W	0.101	0.455	0.444
N	0.055	0.129	0.816
Ergodic	0.126	0.191	0.683

Table 4: Transition matrix for the 2007-2009 panel of the SCF. Fraction of households with the row HtM status in 2007 and the column HtM status in 2009. The last row reports the implied ergodic distribution.

W-HtM hold significant amounts of illiquid wealth: for example, median holdings at age 40 exceed \$50,000. Hence, W-HtM households are not just P-HtM households with small amounts of savings in less liquid assets. The next two panels of Figure 8 articulate this observation further. Panels (c) and (d) plot age profiles of the average fraction of illiquid wealth held in housing and retirement accounts for W-HtM and N-HtM households. The conclusion is striking: the lines are on top of each other, indicating that the portfolio allocation of these two groups is nearly identical.

Persistence. How persistent is a household’s HtM status? We answer this question by exploiting the 2007-2009 panel component of the SCF. Table 4 reports the 2-year transition matrix across the three HtM statuses for U.S. households. The diagonal elements of the matrix reveal that N-HtM status is by far the most persistent, and W-HtM status the most transient of the three. These transition probabilities imply that the expected length of HtM status is around 3.5 years for the W-HtM, 4.5 years for the P-HtM, and 11 years for the N-HtM.

6 Cross-country evidence

The previous section showed that around 30 percent of households in the United States are HtM, one-third of which are P-HtM and two-thirds of which are W-HtM. In this section we use household portfolio data from seven other developed economies to assess whether the prevalence of W-HtM households is a common feature of the wealth distribution across countries and, if so, whether the characteristics of W-HtM in terms of demographics, income, and balance sheets are similar to those in the U.S.

As discussed in Section 4, we focus our attention on three other Anglo-Saxon countries

\$2,000 (Table 2) and hence, even among the N-HtM, there are households with small amounts of liquid wealth.

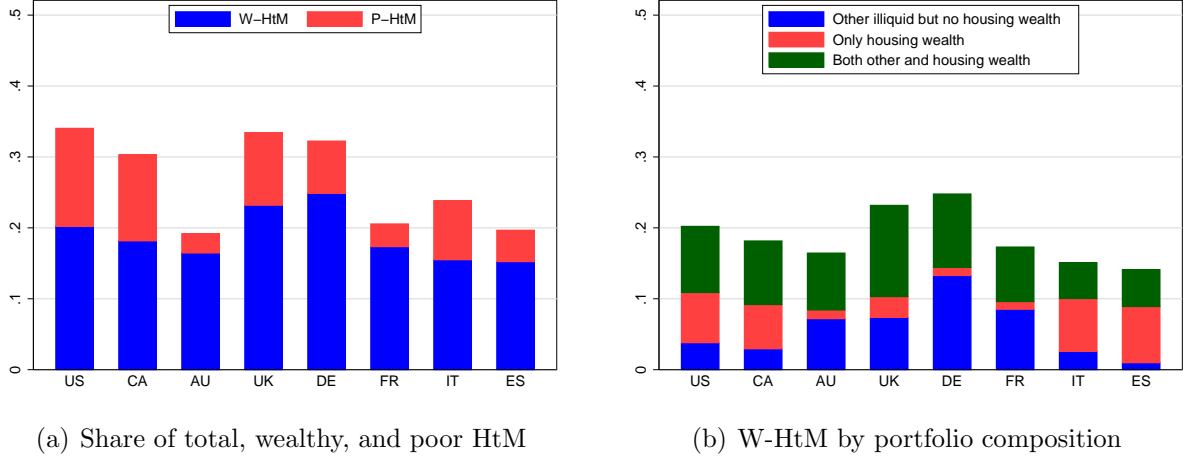


Figure 9: Fraction of HtM households across countries

—Canada, Australia, and the U.K.— and the four largest euro-area economies, Germany, France, Italy, and Spain. While data is available for more than one point in time for most of these countries, in order to keep the discussion manageable we focus on the most recent single cross-section in each country. For Australia and the European countries this is 2010, for the U.K. it is 2009, and for Canada it is 2005. For the sake of comparability, we use only the 2010 wave of the SCF for the United States.

Figure 9(a) shows the fraction of poor and wealthy HtM households in each country. There is a striking similarity in the overall fraction of HtM households, as well as in their breakdown between poor and wealthy, between the U.S., Canada, and the U.K. These three countries have a large share of HtM households, exceeding 30 percent. Australia is an outlier among the Anglo-Saxon countries in two ways: first, the total fraction of HtM is roughly half the fraction in the U.S., the U.K., and Canada; second, 90 percent of HtM households in Australia are W-HtM. Among the euro-area countries, France, Italy, and Spain have smaller shares of HtM households than the U.S., U.K., and Canada —around 20 percent— whereas in Germany this share is closer to 30 percent. Even for the Euro area countries, the fraction of W-HtM among the HtM households exceeds $2/3$. For all eight countries, Figure 9(a) shows there are more W-HtM households than P-HtM. Thus a wide-spread feature of international household portfolios is that a complete characterization of the fraction of the population that is likely to exhibit HtM behavior requires going beyond those with just low net worth.

Figure 9(b) reveals that there are significant differences in the portfolio composition for the W-HtM across countries. In Italy and Spain, virtually all the W-HtM own some housing wealth. Homeowners are also dominant among the group of W-HtM in the

U.S. and Canada. In contrast, around half of the W-HtM in Australia, Germany, and Canada have no housing wealth. Rather, the majority of their illiquid assets are held in private retirement accounts. Table D1 in the Appendix provides more information on the cross-country portfolio composition.

What explains the fact that the euro area countries have a smaller fraction of HtM households than the U.S.? In the euro area countries, households hold more liquid wealth relative to their income compared to the United States. As is clear from Figure 2, this fact can be partially attributed to differences in liquid debt. The fraction of P-HtM households in the euro area countries with negative liquid wealth is 2 to 4 times smaller than in the Anglo-Saxon countries (see Table D1). Presumably, lower access to unsecured credit in Europe implies that there are more incentives for households to hold large balances of liquid wealth for transaction and precautionary reasons. For example, Vandone (2009) documents that, in 2006, the total value of consumer credit amounted to 25 percent of disposable income in the U.K., 15 percent in Germany and Spain, 12 percent in France, and only 10 percent in Italy.

Australia is the country with the largest share of W-HtM, among HtM households. Table D1 reveals that this can be traced to the very high share of the population that owns private retirement wealth. As explained in Section 4, the high ownership rate of retirement accounts in Australia is largely due to the superannuation regulations. When we exclude superannuation accounts as a component of wealth, the fraction of P-HtM in Australia rises from 3 to 9 percent, and the fraction of W-HtM drops accordingly.

Age profiles. Age profiles of the fraction of poor and wealthy HtM households in each country are shown in Figure 10. For most countries, the fraction of P-HtM households declines monotonically with age. The exceptions are Australia and France, where the age profiles of the P-HtM are flat. There are some marked differences in the age profiles of the W-HtM that can be explained by differences in portfolio holdings across countries. In countries where housing wealth is a substantial part of household portfolios, such as the U.S., Canada, and the U.K., the age profile is hump shaped with a peak in the early 40s. In contrast, in Australia and Germany, where a high fraction of W-HtM households hold retirement accounts, the share of W-HtM decreases with age.

An important caveat to these results is that because we infer age profiles from a single cross-section, we necessarily confound age, cohort, and time effects. This could explain,

for example, why in Spain the share of W-HtM falls steadily with age. This pattern may reflect time effects, since 25-35 year-olds have faced much harsher economic conditions upon entry into the labor market than earlier cohorts.²⁴

6.1 Robustness

Table 5 contains an extensive sensitivity analysis on our definitions of P-HtM and W-HtM households that parallels in Table 3.

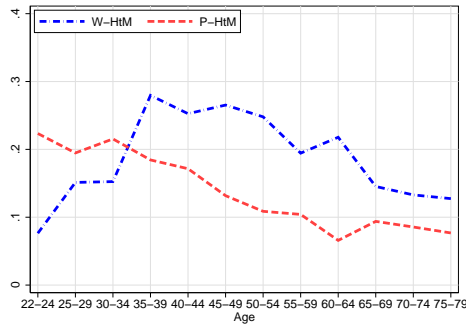
Questions on whether household spending exceeded income in the past year are present in all surveys. Similarly to the U.S., we find larger shares of both P-HtM and W-HtM households when we use these direct questions to measure the incidence of HtM behavior. The difference is especially marked for Italy and Spain where, according to this criterion, over 60 percent of households —and hence three times the baseline estimate— are HtM. Extending the credit limit from one month of income to one year of income has a substantial effect for the Anglo-Saxon countries, but virtually no impact for the euro area countries. This finding is in line with the empirical distribution of liquid assets documented in Figure 2, which showed that households with negative net liquid wealth are extremely rare in the euro area countries.²⁵

The fraction of “financially fragile” households (those with liquid balances lower than the threshold plus 2,000 local currency units) is only 10-15 percentage points larger than the share of HtM in the Anglo-Saxon countries, but in most of the euro area countries it is 30 percentage points larger. This result is consistent with the distributions of liquid wealth reported in Figure 2 showing that in euro area countries there is a large mass of households just to the right of the threshold.

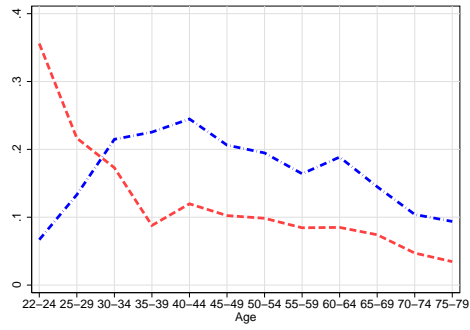
Shortening the pay-period to a week (or extending it to a month), from the bi-weekly baseline, has a small impact on the fraction of P-HtM households, but decreases (increases, respectively) the fraction of W-HtM households by 5 percentage points on average. Including vehicles as illiquid wealth shifts HtM households from poor to wealthy in every country, but to a lesser extent than in the United States. In two countries, Canada and Italy, including other non-financial assets (valuables, collectibles, jewels,

²⁴Figure D1 in the Appendix shows age-income profiles for each country by HtM status and confirms our findings from Section 5.2. The age-income profile for W-HtM households is much more similar to the profile of the N-HtM than to the profile for P-HtM. The only two exceptions are Italy and Spain, where the age-income paths for all three groups are very similar.

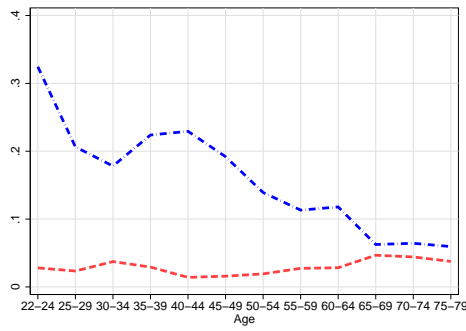
²⁵Recall that, based on the definitions of Section 3, changing the credit limit affects HtM status only for households with negative liquid debt.



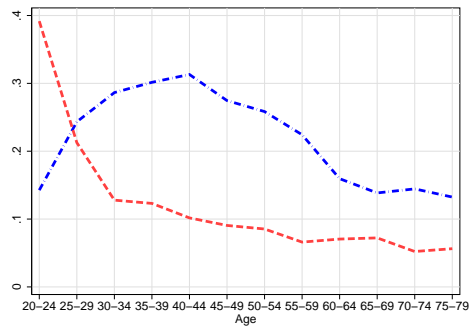
(a) United States



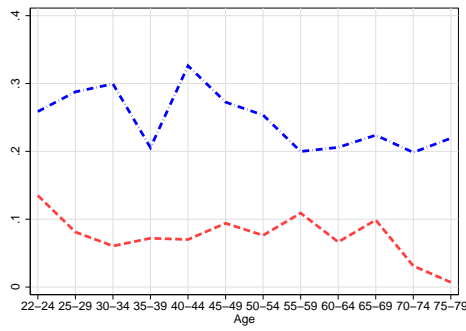
(b) Canada



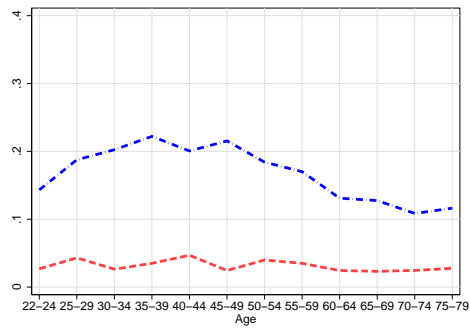
(c) Australia



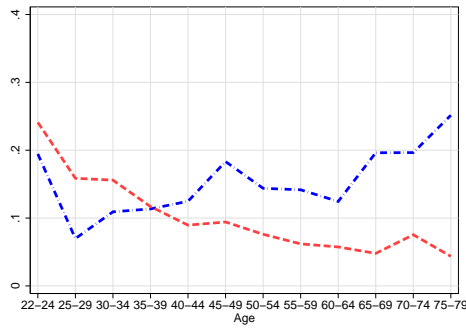
(d) United Kingdom



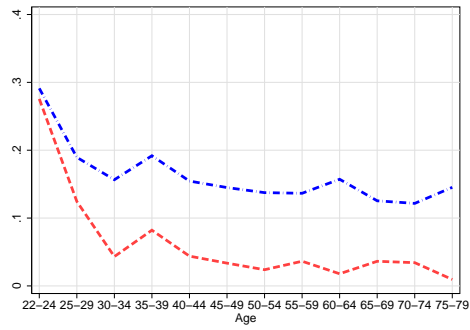
(e) Germany



(f) France



(g) Italy



(h) Spain

Figure 10: Age profile of fraction of HtM households by country.

P-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Baseline	0.138	0.121	0.027	0.103	0.074	0.032	0.083	0.044
In past year, $c > y$	0.157	0.181	0.020	0.092	0.090	—	0.156	0.091
Financially fragile households	0.198	0.190	0.042	0.139	0.110	0.070	0.117	0.092
1 year income credit limit	0.116	0.090	0.024	0.078	0.070	0.030	0.083	0.040
Weekly pay period	0.119	0.105	0.022	0.098	0.058	0.021	0.080	0.036
Monthly pay period	0.165	0.149	0.033	0.111	0.086	0.048	0.091	0.061
Vehicles as illiquid assets	0.060	0.081	0.012	0.065	0.052	0.002	0.028	0.024
Ret. acc. as liquid for 60+	0.138	0.122	0.027	0.103	0.074	0.032	0.083	0.044
Businesses as illiquid assets	0.132	0.115	0.027	0.102	0.071	0.031	0.076	0.043
Direct as illiquid assets	0.137	0.120	0.027	0.102	0.074	0.032	0.083	0.045
Other valuables as illiquid assets	0.134	0.008	0.025	0.099	0.071	—	0.034	0.044
Excludes cc puzzle households	0.174	0.146	0.034	0.124	0.078	0.032	0.086	0.046
HELOCs as liquid debt	0.135	0.127	—	0.103	0.074	0.032	0.083	0.044
Disposable income	0.137	—	—	0.103	—	—	0.080	—
Comm. cons. - beg. of period	0.116	—	—	—	0.066	0.025	0.076	0.036
Comm. cons. - end of period	0.175	—	—	—	0.092	0.050	0.090	0.064
W-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Baseline	0.202	0.182	0.165	0.232	0.248	0.173	0.155	0.152
In past year, $c > y$	0.327	0.409	0.189	0.250	0.392	—	0.474	0.596
Financially fragile households	0.337	0.305	0.261	0.363	0.523	0.585	0.257	0.404
1 year income credit limit	0.130	0.098	0.117	0.135	0.229	0.157	0.147	0.141
Weekly pay period	0.155	0.147	0.116	0.211	0.161	0.087	0.142	0.119
Monthly pay period	0.273	0.247	0.231	0.276	0.370	0.354	0.188	0.220
Vehicles as illiquid assets	0.281	0.223	0.180	0.269	0.270	0.204	0.211	0.173
Ret. acc. as liquid for 60+	0.187	0.161	0.153	0.196	0.245	0.173	0.154	0.152
Businesses as illiquid assets	0.206	0.188	0.166	0.232	0.251	0.173	0.158	0.154
Direct as illiquid assets	0.220	0.215	0.195	0.246	0.303	0.198	0.165	0.162
Other valuables as illiquid assets	0.207	0.295	0.167	0.235	0.252	—	0.204	0.153
Excludes cc puzzle households	0.192	0.179	0.151	0.247	0.236	0.166	0.157	0.148
HELOCs as liquid debt	0.192	0.107	—	0.154	0.238	0.166	0.147	0.140
Disposable income	0.200	—	—	0.237	—	—	0.149	—
Comm. cons. - beg. of period	0.173	—	—	—	0.219	0.127	0.148	0.138
Comm. cons. - end of period	0.284	—	—	—	0.344	0.336	0.173	0.199

Table 5: Robustness results for fraction P-HtM and W-HtM in each category. Financially fragile households includes those households within 2,000 local currency units in liquid assets of their income threshold as HtM. Ret. acc. as liquid for 60+ puts retirement accounts into liquid wealth for households above age 60. Vehicles as illiquid assets includes the value of other valuables for France as the value of vehicles combined with other valuables. Businesses as illiquid assets drops the self employment income sample selection and adds business assets to illiquid wealth and self employment income to labor income. Direct as illiquid assets classifies directly held mutual funds, stocks, corporate and government bonds as illiquid assets. Disposable income removes taxes from gross income. Taxes for the U.S. are estimated from NBER’s TAXSIM assuming all households file as single with no dependents. Comm. cons. - beg. (end) of period assumes households incur consumption commitments at the beginning (end) of the pay period.

etc.) in the definition of illiquid wealth shifts 12 and 5 percent of households from poor to wealthy HtM, respectively.²⁶ Including HELOCs among liquid debt has no effect, except in Canada, where the share of HtM increases by 8 percentage points.

Our baseline measure of income is income after transfers but before taxes, except for Canada where it is disposable income. For three countries, the U.S., the U.K., and Italy, we can analyze the effect of netting taxes at the source for every household.²⁷ In all these three countries, the effect of this correction is minor.

7 The consumption response of the wealthy hand-to-mouth to transitory income shocks

In the previous sections we documented a sizable presence of W-HtM households across a number of countries, but our survey data did not allow us to investigate the consumption behavior of this group of households. In this section we show evidence that, as predicted by the theory presented in Section 2, these households have a large MPC out of transitory income shocks. We use data from the Panel Study of Income Dynamics (PSID) to estimate the consumption response to transitory changes in income using the methodology proposed by Blundell, Pistaferri, and Preston (2008, hereafter BPP), and further examined in Kaplan and Violante (2010). The novelties of our empirical analysis, relative to BPP, are that we use a more recent sample period with enriched data and, most importantly, we estimate transmission coefficients of income shocks to consumption separately for different types of HtM households.

Data source. Estimating the consumption response to income shocks for households with different types of HtM status requires a longitudinal dataset with information on income, consumption, and wealth at the household level. Starting from the 1999 wave, the PSID contains the necessary data. The PSID started collecting information on a sample of roughly 5,000 households in 1968. Thereafter, both the original families and their split-offs (children of the original family forming a family of their own) have been followed. The survey was annual until 1996 and became biennial start-

²⁶There are differences in this question across surveys. The SCF and the HFCS ask about the single most valuable asset not previously mentioned. In HILDA, they ask about collectibles. In the Canadian SFS, valuables are meant to include also the content of the principal residence. In light of this, the result for Canada is not surprising.

²⁷For the U.S., we resort to an imputation based on TAXSIM as explained in Section 5.1.1. The U.K. and Italian surveys ask households about their tax liabilities.

ing in 1997. In 1999 the survey augmented the consumption information available to researchers so that it now covers over 70 percent of all consumption items available in the Consumer Expenditure Survey (CEX), and also asked a set of additional questions on the household balance sheet in every wave.²⁸

Sample selection. We start with the PSID Core Sample and drop households with missing information on race, education, or state of residence, and those whose income grows more than 500 percent, falls by more than 80 percent, or is below \$100. We drop households who have top-coded income or consumption. We also drop households that appear in the sample fewer than three consecutive times, because identification of the coefficients of interest requires a minimum of three periods. In our baseline calculations, we keep households where the head is 25-55 years old. Our final sample has 39,772 observations over the pooled years 1999-2011 (seven sample years).

Definitions. The construction of our consumption measure follows Blundell, Pistaferri, and Saporta-Eksten (2014). We include food at home and food away from home, utilities, gasoline, car maintenance, public transportation, child care, health expenditures, and education. Our definition of household income is labor earnings of the households plus government transfers. Liquid assets in the PSID include the value of checking and savings accounts, money market funds, certificates of deposit, savings bonds, and Treasury bills plus directly held shares of stock in publicly held corporations, mutual funds, or investment trusts. Before 2011, liquid debt is the value of debts other than mortgages, such as credit cards, student loans, medical or legal bills, and personal loans. In 2011, liquid debt includes only credit card debt. Net liquid wealth is liquid assets minus liquid debt. Net illiquid wealth is the value of home equity plus the net value of other real estate plus the value of private annuities or IRAs and the value of other investments in trusts or estates, bond funds, and life insurance policies.²⁹ Net worth is the sum of net illiquid and net liquid wealth. Given these definitions of income and wealth, the HtM status indicators are constructed exactly as outlined in Section 3, where the pay-period is assumed to be two weeks, and the credit limit is one month of income. In our PSID sample, 25 percent of households

²⁸Until 1999, the Wealth Files supplemented the annual survey every five years. Starting in 1999, they became biannual, like the survey itself. In 2009 and 2011, the wealth questions were enriched further with the Housing, Mortgage Distress, and Wealth Data Supplements.

²⁹The two main discrepancies with the SCF definitions are that we do not attempt a cash imputation, and both CDs and saving bonds are in liquid, instead of illiquid, wealth. Since these two saving instruments are not common, we do not expect this discrepancy to affect our results. For example, if we classify CDs and saving bonds as liquid wealth in the 2010 SCF, the fraction of HtM drops by only 1 percentage point.

are W-HtM, roughly in line with the SCF estimates, but the share of the P-HtM is 21 percent, and hence is almost twice as large as its counterpart in the SCF.

The BPP methodology. We refer the reader to BPP and to Kaplan and Violante (2010) for a thorough description of the methodology. Here, we only sketch the key steps. As in BPP, we first regress log income and log consumption expenditures on year and cohort dummies, education, race, family structure, employment, geographic variables, and interactions of year dummies with education, race, employment, and region. We then construct the first-differenced residuals of log consumption Δc_{it} and log income Δy_{it} . Recall that, since the survey is biannual, a period is two years. The income process y_{it} is represented as an error component model which comprises orthogonal permanent and i.i.d. components. Hence, income growth is given by

$$\Delta y_{it} = \eta_{it} + \Delta \varepsilon_{it}, \quad (13)$$

where η_{it} is the permanent shock and ε_{it} is the transitory shock. This is a common income process in the empirical labor literature, at least since MaCurdy (1982) and Abowd and Card (1989) who showed that this specification is parsimonious and fits income data well. The BPP estimator of the transmission coefficient of transitory income shocks to consumption, the MPC, is given by

$$\widehat{MPC}_t = \frac{\text{cov}(\Delta c_{it}, \Delta y_{i,t+1})}{\text{cov}(\Delta y_{i,t}, \Delta y_{i,t+1})}. \quad (14)$$

The true marginal propensity to consume out of a transitory shock is defined as

$$MPC_t = \frac{\text{cov}(\Delta c_{it}, \varepsilon_{it})}{\text{var}(\varepsilon_{it})} \quad (15)$$

The estimator in (14) is a consistent estimator of (15) if the household has no foresight, or no advanced information, about future shocks, i.e.:

$$\text{cov}(\Delta c_{it}, \eta_{i,t+1}) = \text{cov}(\Delta c_{it}, \varepsilon_{i,t+1}) = 0, \quad (16)$$

The estimator is implemented by an IV regression of Δc_{it} on Δy_{it} , instrumented by $\Delta y_{i,t+1}$. Note that $\Delta y_{i,t+1}$ is correlated with the transitory shock at t , but not with the permanent one. Kaplan and Violante (2010) show that the presence of tight borrowing constraints does not bias the estimate of the transmission coefficient for transitory shocks—an important finding in light of the fact that we are interested in the differential response of HtM households, who may be close to a constraint, and N-HtM

	P-HtM	W-HtM	N-HtM	HtM-NW	N-HtM-NW
Baseline	0.243*** (0.065)	0.301*** (0.048)	0.127*** (0.036)	0.229*** (0.054)	0.201*** (0.030)
Pre-tax earnings	0.131*** (0.043)	0.223*** (0.035)	0.122*** (0.027)	0.143*** (0.036)	0.164*** (0.023)
Include food stamps	0.217*** (0.059)	0.264*** (0.045)	0.105*** (0.035)	0.203*** (0.050)	0.171*** (0.029)
Cont. married households	0.095 (0.194)	0.193** (0.079)	0.079* (0.043)	-0.048 (0.129)	0.157*** (0.042)
Stable marital status	0.239*** (0.085)	0.282*** (0.054)	0.110*** (0.038)	0.190*** (0.070)	0.195*** (0.033)
Households with male heads	0.186** (0.080)	0.193*** (0.058)	0.073* (0.040)	0.150** (0.064)	0.129*** (0.035)
Monthly income	0.229*** (0.068)	0.288*** (0.053)	0.159*** (0.034)	0.236*** (0.057)	0.199*** (0.030)

Table 6: MPC out of transitory income shocks for different types of HtM households. Bootstrapped standard errors based on 250 replications in parenthesis. Pre-tax earnings: transfers excluded. Include food stamps: food stamps are included among transfers. Cont married households: sample restricted to continuously married households. Stable marital status: sample restricted to households with no change in marital status. Households with male heads: households with female heads (mostly single) excluded from the sample. Monthly earnings: pay-period set to one month instead of two weeks. Levels of Significance: *** = 0.01, ** = 0.05, * = 0.10.

households.

Results. Table 6 summarizes our results. In our baseline specification, the MPC of the W-HtM group is the highest, around 30 percent. In other words, in the first two years, the W-HtM households consume 30 percent of an unexpected change in income whose effect entirely dissipates within the period. The point estimate of the MPC for the P-HtM is 24 percent, and for the N-HtM is less than 13 percent. Given the well known measurement error present in survey data, especially for consumption expenditures, and the small sample size, it is not surprising that these estimates are somewhat imprecise. However, the difference between the MPC for the W-HtM and the N-HtM is statistically significant.

When the sample is split between HtM and N-HtM based on net worth, the estimated transmission coefficients are very similar across the two groups. The group of HtM-NW is essentially the same as the P-HtM, and in fact their estimated MPCs are similar. However, among the N-HtM-NW there are also many W-HtM households that artificially inflate the estimate of the MPC. Based on this household classification, one would conclude that there is no evidence of a differential response of consumption to income shocks for households with different HtM status. A classification based on liquid and

illiquid wealth, instead, finds economically significant differences.

The remaining rows in Table 6 offer a robustness analysis with respect to the definition of income and consumption, household composition, and the assumed pay-period. The ranking of MPCs between wealthy, poor, and non-HtM is always as in the baseline specification, and as predicted by the theory, the gap between HtM households based on the net worth criterion is always very small or is not statistically significant.

Our key finding that the consumption of the W-HtM displays excess sensitivity to transitory income shocks is in line with some recent findings. Misra and Surico (2013) expand on the research of Johnson, Parker, and Souleles (2006) and Parker, Souleles, Johnson, and McLelland (2013) on the U.S. fiscal stimulus payment episodes of 2001 and 2008. They conclude that, for both stimulus programs, the largest propensity to consume out of the tax rebate is found among households who own real estate but have high levels of mortgage debt. Cloyne and Surico (2013) exploit a long span of expenditure survey data for the U.K. and a narrative measure of exogenous income tax changes. They also find that homeowners with high leverage ratios exhibit large and persistent consumption responses to tax shocks. Baker (2013) combines several novel sources of household data on consumption expenditures, income, and household balance sheets to investigate the comovement of income and consumption, at the micro level, around the Great Recession. He finds that expenditures of highly-indebted households with illiquid assets are especially sensitive to income fluctuations. Overall, this body of work confirms our finding in Figure 4 that highly-leveraged homeowners are likely to be W-HtM and, hence, to have large MPC out of income shocks.

8 Implications for fiscal policy

What does the existence of W-HtM households, together with their large propensities to consume out of transitory income shocks, imply for how one should think about fiscal policy? In this section we use a series of policy simulations from three alternative models to argue that W-HtM households should be modelled as a separate group: ignoring them leads to a distorted view of the effects of fiscal stimulus policies on aggregate consumption.

The first model that we use is the two-asset incomplete markets model from Kaplan and Violante (2014a, 2014b, KV thereafter). We label this model SIM-2, since it extends the standard incomplete markets (SIM) life-cycle economy by adding a second illiquid

asset that pays a higher return —through both a financial component and a housing services component— but is subject to a transaction cost. For the reasons explained in Section 2, the illiquidity due to the transaction cost means that the model generates households of all three HtM types. The version of the model we use here does not allow borrowing and has a transaction cost of \$1,000.³⁰

The second model, which we label SIM-1, is a standard one-asset incomplete markets life cycle model. The version that we adopt is the same as in KV, but with the transaction cost set to zero, and recalibrated to data on net worth alone, rather than data on illiquid and liquid assets separately. Since this is a one-asset model, it generates only P-HtM and N-HtM households, and has no W-HtM households.

The third model, which we label SP-S, is a spender-saver model in the spirit of Campbell and Mankiw (1989) and, more recently, Gali, Lopez-Salido, and Valles (2007), Eggertson and Krugman (2012), and Justiniano, Primiceri, and Tambalotti (2013). In the SP-S model, some households (the savers) act as forward-looking optimizing consumers who can save in a single risk-free asset. The remaining households (the spenders) follow the rule-of-thumb consumption policy of consuming all their income in every period. This class of models is typically calibrated so that the distinction between the spenders and savers is based on their holdings of liquid wealth rather than net worth. Thus, in the SP-S model, the W-HtM and the P-HtM households are lumped together and considered to be the spenders, while the N-HtM households are considered to be the savers.

To summarize, SIM-2 is a two-asset economy in which the W-HtM households are explicitly modeled as a distinct group. SIM-1 is a “net-worth economy” where the W-HtM households are treated as if they were N-HtM households. Compared to SIM-2, SIM-1 greatly understates the fraction of HtM households. SP-S is a “liquid-wealth economy” where both the W-HtM and the P-HtM are treated identically as HtM households with an MPC that is always equal to 1. Thus, compared to SIM2, SP-S has the correct number of HtM households, but it greatly overstates their MPC.

From each of these three models, we simulate a cohort of households. For each household, we compute the quarterly consumption response to a one-time unexpected cash windfall, or cash loss, of different amounts (\$50, \$500, \$2,000). We then divide the sim-

³⁰We refer the reader to KV for a full description of the model, its calibration, and a comparison of the predictions of the model with life-cycle data, and with the aggregate consumption response to the 2001 and 2008 fiscal stimulus payments as estimated by Johnson, Parker and Souleles (2006), and Parker et al. (2013), respectively.

	SIM-2			SIM-1		SP-S	
	P-HtM	W-HtM	N-HtM	HtM	N-HtM	HtM	N-HtM
Average	0.35	0.44	0.06	0.14	0.02	1.00	0.02
Low Income	0.34	0.37	0.16	0.15	0.04	1.00	0.04
Middle Income	0.38	0.44	0.09	0.11	0.02	1.00	0.02
High Income	0.31	0.52	-0.02	0.12	0.01	1.00	0.01
Age \leq 40	0.38	0.42	0.08	0.16	0.02	1.00	0.02
Age 40-60	0.30	0.42	0.01	0.11	0.01	1.00	0.01
Age $>$ 60	0.39	0.51	0.13	0.04	0.04	1.00	0.04

Table 7: Quarterly MPC out of an unexpected \$500 transfer for the aggregate economy, and for various subgroups of the population, using group composition from the 2010 SCF. SIM-2: two-asset, life-cycle, incomplete-market model. SIM-1: one-asset, life-cycle, incomplete-market model. SP-S: spender-saver model.

ulated cohort into twenty-seven bins, based on three income terciles, three age classes (22-40, 40-60, 60+) and the three HtM groups. For each of these bins we compute the average consumption response from the model. To obtain an aggregate response of the economy, we need shares of the population in each of these twenty-seven groups. For this last step, we use our cross-sectional survey data from Sections 5 and 6.

Table 7 reports the quarterly average MPCs out of a \$500 windfall in the three models, for the HtM groups and for some sub-groups defined by income and age, using group shares from the 2010 SCF. In the SIM-2 model, MPCs are very small for all N-HtM households, except for those who are income-poor and old. For high-income households who are N-HtM, the average MPC is slightly negative. The intuition for this finding is discussed in detail in KV. It arises because the receipt of a \$500 windfall may trigger a household who has already accumulated lots of liquid wealth, and is close to its planned date of deposit, to pay the transaction cost and make an earlier deposit into the illiquid account. Since such a household can effectively save at the rate of return on the illiquid asset, it chooses to consume less and save more than it would have in the absence of the income windfall. This example illustrates how explicitly modeling W-HtM behavior through transaction costs may alter the MPC even for N-HtM households. The MPCs for both the W-HtM and P-HtM households in the SIM-2 economy are substantial. They are slightly larger for the W-HtM than the P-HtM, particularly for households with a high level of income. As explained in Section 2, since the W-HtM have higher lifetime incomes than the P-HtM, they have higher target consumption and hence spend more out of an unexpected moderately-sized payment.

In the SIM-1 model, the MPCs for HtM households are almost identical to those for

	<u>Model</u>		
	SIM-2	SIM-1	SP-S
\$500 transfer	0.18	0.04	0.35
<u>Size Asymmetry</u>			
\$50 transfer	0.29	0.05	0.35
\$2000 transfer	0.05	0.03	0.35
<u>Sign Asymmetry</u>			
\$500 tax	0.42	0.14	0.36
<u>Income Targeting</u>			
\$500 transfer, bottom tercile	0.26	0.07	0.50
\$500 transfer, top tercile	0.20	0.03	0.34

Table 8: Quarterly aggregate consumption responses for the U.S. using group composition from the 2010 SCF. SIM-2: two-asset, life-cycle, incomplete-market model. SIM-1: one-asset, life-cycle, incomplete-market model. SP-S: spender-saver model. All taxes and transfers are lump-sum, one-time, and unexpected.

P-HtM households in the SIM-2 model, and the MPCs for N-HtM households are, in general, even smaller than those for N-HtM households in the SIM-2 model. In the SP-S model, by construction, the MPCs for the N-HtM households are the same as in the SIM-1 model and are equal to one for HtM households.

Policy simulations for the U.S. We now show that the three models yield very different predictions for the aggregate MPC out of unexpected, one-time, lump-sum transfers/taxes of different amounts. Table 8 reports the policy-experiments results (i.e., the aggregate quarterly consumption responses) for the U.S. using the SCF data from 2010 to estimate the group shares.

We begin by analyzing a policy experiment where every household receives a \$500 transfer, e.g., a stimulus payment. The aggregate MPC of the SIM-2 economy is 0.18. This value is substantially larger than the MPC of the SIM-1 economy (0.04) because the SIM-1 economy, by treating the W-HtM households as N-HtM, misses a large fraction of the population that have high MPCs. The aggregate MPC is highest for the SP-S economy (0.35) because this model implicitly assumes that all P-HtM and W-HtM households all spend the entire \$500. Our discussion of Table 7 suggests that this assumption is extreme: in the SIM-2 economy, HtM households spend on average only 35%-45% of their payments in the quarter they are received.

Table 8 also shows that the degree of size asymmetry in the aggregate MPC differs remarkably across the three models. In the SIM-2 model, the consumption response to

a \$50 windfall is 0.29 while the response to a \$2,000 windfall is only 0.05. The reason for this large asymmetry is the availability of an illiquid savings instrument subject to a transaction cost. For large enough windfalls, many HtM households in SIM-2 may find it optimal to pay the transaction cost and make a deposit into the illiquid asset. However, for small windfalls, it is never optimal to adjust the illiquid asset: households thus face an intertemporal tradeoff governed by the (low) return on the liquid asset, and thus have a large incentive to consume. This size asymmetry is absent from both the SP-S and SIM-1 models. In the SP-S model it is absent because of the assumed rule-of-thumb behavior: the HtM households in the SP-S model always consume their entire transfer, regardless of its size. In the SIM-1 model there is only a modest decline in the MPC with the size of the payment because households always face the same inter-temporal trade off when making their consumption decisions.

The degree of sign asymmetry also differs across the three models. In the SIM-1 and SIM-2 models, the response to a lump-sum tax of \$500 is substantially larger than the response to a \$500 transfer. Even HtM households, who are at a kink in their budget constraints, desire to save some part of a positive windfall if it is large enough to push them off the kink. Negative income changes, however, cannot be smoothed for households at the constraint, and withdrawing from the illiquid account is too expensive to be optimal—recall that in the calibrated SIM-2 model, the transaction cost is \$1,000. In the SP-S model, the response to positive and negative income shocks are essentially the same, since the HtM households have an MPC of 1 regardless of the sign of the shock.

Table 8 reveals that the models have different implications for the optimal degree of “income targeting” of fiscal stimulus transfers for maximizing the aggregate consumption response. A widely held view is that the aggregate consumption response to a fiscal stimulus policy, per dollar paid out, is strongest when the transfers are targeted to households with the lowest income, that is, stimulus payments should be phased out for middle- and high-income households for maximum effect. This view, which is based on the conjecture that HtM households are income-poor, ignores the W-HtM, a group with significantly higher income, as we showed in Sections 5.2 and 6. In line with this observation, the SIM-2 model generates only a very modest decline (0.26 to 0.20) between the MPC out of a \$500 transfer for households in the lowest income tercile and households in the middle income tercile. The corresponding relative declines across income terciles are much larger under the SIM-1 and SP-S models. In the SIM-1 model, the only high-MPC households are the low-income P-HtM; in the SP-S model,

all HtM households are assumed to have the same MPC, while under the SIM-2 model we saw in Table 7 that, among W-HtM households, MPCs are increasing in income.

Implied cross-country variation in effects of policy We now explore what the three models predict for the aggregate response to a \$500 fiscal stimulus check in each of the eight countries in our sample. To do this, we use our survey data to estimate the fraction of households in each country who fall into each of the 27 bins, and then apply these country-specific group weightings to the model-generated MPCs. To illustrate the differences in model predictions, Figure 11 plots the estimated aggregate MPC under the SIM-2 model against the corresponding MPC under the SIM-1 model (triangles) and the SP-S model (circles).

The figure shows striking differences in the amount of cross-country dispersion in the aggregate MPC predicted by the three models. There is much less dispersion in the SIM-1 model compared to the SIM-2 model because, by treating the W-HtM as N-HtM, the SIM-1 model misses most of the cross-country variation in HtM behavior. In contrast, there is more dispersion in the SP-S model than in the SIM-2 model. This is because, by assigning an MPC of 1 to all the W-HtM households, compared to an MPC of 0.44 in the SIM-2 model, the SP-S model exaggerates existing cross-country heterogeneity in the fraction of HtM households.

These experiments clearly illustrate why it is important to think deeply about the behavior of W-HtM households when considering the design of fiscal policies. With respect to the traditional view based on SIM-1 or SP-S models, we demonstrate three lessons: (i) there is limited scope for stimulating aggregate consumption by increasing the transfer size; (ii) the aggregate consumption response to a lump-sum tax is much stronger, in absolute value, than the response to an equal-size transfer; and (iii) targeting stimulus payments exclusively towards low-income families will miss a substantial fraction of liquidity-constrained households.

9 Concluding remarks

We set out to investigate, theoretically and empirically, the behavior of wealthy hand-to-mouth households — an often overlooked, but highly relevant part of the population — and to reflect on their implications for macroeconomic modeling and fiscal policy design. We conclude by taking stock of what we have learnt.

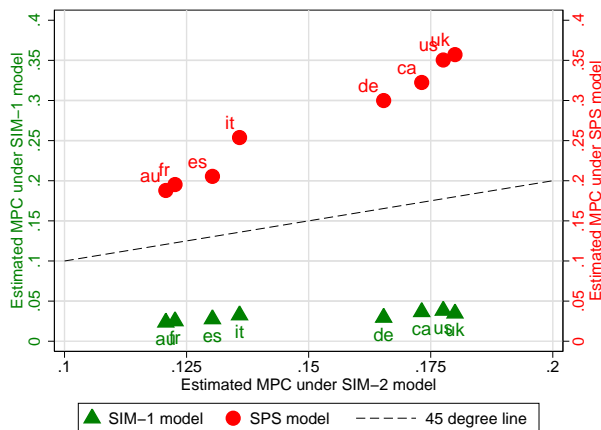


Figure 11: Estimated aggregate consumption response by country under SIM-2 model (x axis), SIM-1 model (triangles, left y axis) and SP-S model (circles, right y axis). The dashed line is the 45 degree line.

Theoretically, we showed that wealthy hand-to-mouth behavior can occur when households face a trade-off between the long-run gain from investing in illiquid assets (i.e., assets that require the payment of a transaction cost for making unplanned deposits or withdrawals), and the short-run cost of having fewer liquid assets available to smooth consumption.

Empirically, we documented that 30 percent of households in the U.S. are hand-to-mouth, and that this fraction has been relatively constant over the past two decades. The share of hand-to-mouth households varies somewhat across the eight countries in our study, from less than 20 percent in Australia and Spain to over 30 percent in the U.K. and Germany. Given our identification strategy, these estimates are likely to be a lower bound. The key finding is that in all countries, the vast majority of hand-to-mouth households, at least two-thirds of them, are wealthy hand-to-mouth, not poor hand-to-mouth.

Who are the wealthy hand-to-mouth? We highlight three features. First, unlike poor hand-to-mouth households, the wealthy hand-to-mouth are not predominantly young households with low income. Rather, the frequency of wealthy hand-to-mouth status has a hump-shaped age profile that peaks in the early forties and an income profile that mirrors strongly that of the non hand-to-mouth. Second, they are not simply poor hand-to-mouth households with very small holdings of illiquid assets. Rather, they hold substantial wealth in housing and retirement accounts, in the same proportions as non hand-to-mouth households. Finally, their hand-to-mouth status is somewhat

more transient than that of the poor hand-to-mouth.

Why does this group of households deserve the attention of economists and policy makers?

Wealthy hand-to-mouth households are important because they have large consumption responses to transitory income shocks —a crucial determinant of the efficacy of many types of fiscal interventions, such as the fiscal stimulus payments that were implemented in the last two recessions. To demonstrate this, we use PSID data to show that the transmission coefficient of transitory income shocks into consumption is significantly larger for wealthy (and poor) hand-to-mouth households than for non hand-to-mouth households.

The wealthy hand-to-mouth thus have consumption responses that, in many ways, are similar to the poor hand-to-mouth, yet have demographic characteristics and portfolio composition that resemble the non hand-to-mouth. This suggests that the three types of hand-to-mouth households each need their own unique place in frameworks that are to be used for analyzing and forecasting the effects of fiscal policy. Macroeconomists need to move beyond one-asset models, such as those in the spirit of Aiyagari (1994), Huggett (1996), and Rios-Rull (1995), since these models assume wealthy hand-to-mouth households are as unconstrained as non hand-to-mouth ones. They also need to move beyond spender-saver models, such as those in the spirit of Campbell and Mankiw (1989), and Eggertsson and Krugman (2012), since these models treat all hand-to-mouth households identically, and thus assume that wealthy hand-to-mouth households are as constrained as poor hand-to-mouth. In particular, by ignoring the fact that wealthy hand-to-mouth can use illiquid assets to buffer large negative shocks, these models exaggerate the financial fragility of this group. In the last section, we ran several fiscal policy experiments to illustrate where misleading inferences would be obtained by using either of these two simpler models of hand-to-mouth behavior.

References

- ABOWD, J. M., AND D. CARD (1989): “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, 57(2), 411–445.
- AIYAGARI, S. R. (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving,” *The Quarterly Journal of Economics*, 109(3), 659–684.
- ALVAREZ, F., L. GUISO, AND F. LIPPI (2012): “Durable Consumption and Asset Management with Transaction and Observation Costs,” *American Economic Review*, 102(5), 2272–2300.
- ANGELETOS, G., D. LAIBSON, A. REPETTO, J. TOBACMAN, AND S. WEINBERG (2001): “The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation,” *Journal of Economic Perspectives*, 15(3), 47–68.
- ATTANASIO, O. P., AND G. WEBER (1993): “Consumption Growth, the Interest Rate and Aggregation,” *The Review of Economic Studies*, 60(3), 631–649.
- BAKER, S. R. (2013): “Debt and the Consumption Response to Household Income Shocks,” *Mimeo*.
- BANKS, J., AND S. TANNER (2002): “Household Portfolios in the United Kingdom,” in *Household Portfolios*, ed. by L. Guiso, M. Haliassos, and T. Jappelli, pp. 219–250. MIT Press.
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): “Consumption Inequality and Partial Insurance,” *American Economic Review*, 98(5), 1887–1921.
- BLUNDELL, R., L. PISTAFERRI, AND I. SAPORTA-EKSTEN (2014): “Consumption Inequality and Family Labor Supply,” Working Paper Series 1656, European Central Bank.
- BRODA, C., AND J. PARKER (2014): “The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption,” *Mimeo*.
- BROWNING, M., AND T. F. CROSSLEY (2001): “The Life-Cycle Model of Consumption and Saving,” *Journal of Economic Perspectives*, 15(3), 3–22.
- CAMPBELL, J. Y., AND N. G. MANKIW (1989): “Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence,” in *NBER Macroeconomics Annual 1989, Volume 4*, NBER Chapters, pp. 185–246. National Bureau of Economic Research, Inc.
- (1990): “Permanent Income, Current Income, and Consumption,” *Journal of Business & Economic Statistics*, 8(3), 265–279.
- (1991): “The Response of Consumption to Income: A Cross-Country Investigation,” *European Economic Review*, 35(4), 723–756.
- CARROLL, C. D. (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, 112(1), 1–55.
- CARROLL, C. D., J. SLACALEK, AND K. TOKUOKA (2014a): “Digestible Microfoundations: Buffer Stock Saving in a Krusell-Smith World,” *Mimeo*.
- (2014b): “The Distribution of Wealth and the Marginal Propensity to Consume,” *Mimeo*.
- CHETTY, R., AND A. SZEIDL (2007): “Consumption Commitments and Risk Preferences,” *The Quarterly Journal of Economics*, 122(2), 831–877.
- CLOYNE, J., AND P. SURICO (2013): “Household Debt and the Dynamic Effects of Income Tax Changes,” Discussion Paper 9649, Centre for Economic Policy Research.

- COCHRANE, J. H. (1989): “The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives,” *American Economic Review*, 79(3), 319–337.
- DEATON, A. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59(5), 1221–1248.
- EGGERTSSON, G. B., AND P. KRUGMAN (2012): “Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach,” *The Quarterly Journal of Economics*, 127(3), 1469–1513.
- EYMANN, A., AND A. BÖRSCH-SUPAN (2002): “Household Portfolios in Germany,” in *Household Portfolios*, ed. by L. Guiso, M. Haliassos, and T. Jappelli, pp. 291–340. MIT Press.
- FOSTER, K., S. SCHUH, AND H. ZHANG (2013): “The 2010 Survey of Consumer Payment Choice,” Research Data Reports 13-2, Federal Reserve Bank of Boston.
- FRIEDMAN, M. (1957): *A Theory of the Consumption Function*. Oxford and IBH Publishing Company.
- GALÍ, J., J. D. LÓPEZ-SALIDO, AND J. VALLÉS (2007): “Understanding the Effects of Government Spending on Consumption,” *Journal of the European Economic Association*, 5(1), 227–270.
- GUIO, L., M. HALIASSOS, AND T. JAPPELLI (2002): *Household Portfolios*. MIT Press.
- HALL, R. E. (1978): “Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence,” *Journal of Political Economy*, 86(6), 971–987.
- HUGGETT, M. (1996): “Wealth Distribution in Life-Cycle Economies,” *Journal of Monetary Economics*, 38(3), 469–494.
- HUNTLEY, J., AND V. MICHELANGELI (2014): “Can Tax Rebates Stimulate Spending in a Life-Cycle Model?,” *American Economic Journal: Macroeconomics*, 6(1), 162–189.
- JAPPELLI, T., AND L. PISTAFERRI (2010): “The Consumption Response to Income Changes,” Working Paper 15739, National Bureau of Economic Research.
- JOHNSON, D., J. PARKER, AND N. SOULELES (2006): “Household Expenditure and the Income Tax Rebates of 2001,” *The American Economic Review*, 96(5), 1589–1610.
- JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2013): “Household Leveraging and Deleveraging,” Working Paper 18941, National Bureau of Economic Research.
- KAPLAN, G., AND G. VIOLANTE (2010): “How Much Consumption Insurance Beyond Self-Insurance?,” *American Economic Journals: Macroeconomics*, 2(4), 53–87.
- (2014a): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, forthcoming.
- (2014b): “A Tale of Two Stimulus Payments: 2001 vs 2008,” *American Economic Review (Papers and Proceedings)*, 104(5).
- KRUSELL, P., AND A. SMITH (1996): “Rules of Thumb in Macroeconomic Equilibrium: A Quantitative Analysis,” *Journal of Economic Dynamics and Control*, 20(4), 527–558.
- (1997): “Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns,” *Macroeconomic Dynamics*, 1(02), 387–422.
- (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5), 867–896.

- LAIBSON, D., A. REPETTO, AND J. TOBACMAN (2003): “A Debt Puzzle,” in *Knowledge, Information, and Expectations in Modern Economics: In Honor of Edmund S. Phelps*, ed. by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, pp. 228–266. Princeton University Press.
- LUDVIGSON, S. C., AND A. MICHAELIDES (2001): “Does Buffer-Stock Saving Explain the Smoothness and Excess Sensitivity of Consumption?,” *The American Economic Review*, 91(3), 631–647.
- LUSARDI, A., D. SCHNEIDER, AND P. TUFANO (2011): “Financially Fragile Households: Evidence and Implications,” *Brookings Papers on Economic Activity*, 42(1), 83–150.
- MACURDY, T. E. (1982): “The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis,” *Journal of Econometrics*, 18(1), 83–114.
- MISRA, K., AND P. SURICO (2013): “Consumption, Income Changes and Heterogeneity: Evidence from Two Fiscal Stimulus Programmes,” Discussion Paper 9530, Centre for Economic Policy Research.
- OECD (2013): “Gross Pension Replacement Rates,” in *Pensions at a Glance 2013: OECD and G20 Indicators*. OECD Publishing.
- PARKER, J. A. (1999): “The Reaction of Household Consumption to Predictable Changes in Social Security Taxes,” *The American Economic Review*, 89(4), 959–973.
- PARKER, J. A., N. S. SOULELES, D. S. JOHNSON, AND R. MCCLELLAND (2013): “Consumer Spending and the Economic Stimulus Payments of 2008,” *American Economic Review*, 103(6), 2530–53.
- PENCE, K. (2011): “Comment on ‘The Financially Fragile Households: Evidence and Implications’,” *Brookings Papers on Economic Activity*, 42(1), 141–150.
- RÍOS-RULL, J.V. (1995): “Models with Heterogeneous Agents,” in *Frontiers of Business Cycles Research*, ed. by T. F. Cooley, pp. 98–125. Princeton University Press.
- SHAPIRO, M., AND J. SLEMRUD (2003a): “Consumer Response to Tax Rebates,” *American Economic Review*, 93(1), 381–396.
- (2003b): “Did the 2001 Tax Rebate Stimulate Spending? Evidence from Taxpayer Surveys,” *Tax Policy and the Economy*, 17, 83–109.
- (2009): “Did the 2008 Tax Rebates Stimulate Spending?,” *American Economic Review*, 99(2), 374–379.
- SHORE, S. H., AND T. SINAI (2010): “Commitment, Risk, and Consumption: Do Birds of a Feather Have Bigger Nests?,” *The Review of Economics and Statistics*, 92(2), 408–424.
- SOULELES, N. (1999): “The Response of Household Consumption to Income Tax Refunds,” *American Economic Review*, 89(4), 947–958.
- TELYUKOVA, I. A. (2013): “Household Need for Liquidity and the Credit Card Debt Puzzle,” *Review of Economic Studies*, 80(3), 1148–1177.
- VANDONE, D. (2009): *Consumer Credit in Europe: Risks and Opportunities of a Dynamic Industry*. Springer Verlag Berlin Heidelberg.
- ZHANG, C. Y. (2014): “Monthly Budgeting Heuristics: Evidence from ‘Extra’ Paychecks,” *Mimeo*.

Online Appendix

A A simple model of wealthy hand-to-mouth behavior

In this Appendix, we provide a more detailed analysis of the model in Section (2) of the paper.

Consider a household that lives for periods, $t = 0, 1, 2$, but consumes only in the last two periods of life. Preferences over consumption at $t = 1, 2$ are given by

$$v_0 = u(c_1) + u(c_2)$$

with no discounting. In period 0, the household has an initial endowment ω and makes a portfolio allocation decision. Two assets are available as saving instruments. First, there is an illiquid asset a that pays off a gross return R before the consumption decision in period 2, but cannot be accessed at the time of the consumption decision in period 1. Second, there is a liquid asset m that can be accessed before the consumption decision in both periods, but pays a return $1 < R$. For now, we do not allow the agent to borrow, i.e. take negative a position in the liquid asset, but we relax this assumption in Section A.4. After the initial portfolio allocation decision, households receive income y_1 and make their consumption and liquid saving decision in period 1. In the second, and last, period, they receive income y_2 and consume this endowment plus their savings in liquid and illiquid wealth. The only two decisions to characterize are therefore the initial portfolio allocation decision, and the consumption/saving decision at $t = 1$.

We make the following normalizations and parametric assumptions. Period utility u is CES with intertemporal elasticity of substitution $\sigma > 0$. We set the initial endowment ω to 1, so the initial portfolio allocation (m_1, a) has the interpretation of shares of wealth invested in liquid and illiquid wealth. We set $y_2 = \Gamma > 1$ and we allow two possible values for y_1 , $\{y_L, y_H\}$ where $y_L = 0$ and $y_H > R + \Gamma$. We refer to these two cases as “low-income” and “high-income” paths. The low income path is increasing and the high income path is decreasing.

Our characterization of hand-to-mouth behavior concerns the asset position at the time of the $t = 1$ consumption decision. We define a household as not hand-to-mouth (N-HtM) if, after consuming at $t = 1$, it holds a positive amount of liquid assets, i.e. $m_2 > 0$ and $a \geq 0$. We define a household as poor hand-to-mouth (P-HtM) if, after

consuming at $t = 1$, it does not hold any liquid or illiquid assets, i.e. $m_2 = 0$ and $a = 0$. We define a household as wealthy hand-to-mouth (W-HtM) if, after consuming at $t = 1$, it holds a positive amount of illiquid assets but no liquid assets, i.e. $m_2 = 0$ and $a > 0$.³¹ Therefore, the $t = 1$ consumption/saving decision determines whether an agent is HtM, and the initial portfolio allocation determines whether a HtM agent is poor or wealthy HtM.

A.1 Solution without illiquid asset

We begin by analyzing a special case where there is no illiquid asset. In this case we refer to the liquid asset as *net worth*. We solve the model backwards, starting from the consumption decision at $t = 1$. The problem faced by the household at $t = 1$ is

$$\begin{aligned} v_1(m_1) &= \max_{c_1, m_2} u(c_1) + u(m_2 + \Gamma) \\ &\quad s.t. \\ c_1 + m_2 &= y_1 + m_1 \\ m_2 &\geq 0 \end{aligned}$$

which has the solution

$$m_2 = \max \left\{ \frac{y_1 - \Gamma + m_1}{2}, 0 \right\}. \quad (A1)$$

The interior solution for m_2 implies a perfectly smooth consumption path, $c_1 = c_2 = (y_1 + \Gamma + m_1)/2$ because there is no discounting and the interest rate on the liquid asset (the only saving vehicle available at $t = 1$) is 1. The corner solution $m_2 = 0$ yields an increasing consumption path, $c_1 = y_1 + m_1$, $c_2 = \Gamma$. Since the liquid asset is the only available asset, the initial portfolio allocation decision is trivial, and $m_1 = 1$.

Thus there two cases, depending on the income path. Under the low income path with $y_L = 0 < \Gamma - 1$, equation (A1) reveals that the constraint binds at $t = 1$ and the household is P-HtM with an increasing consumption profile. Under the high income path with $y_H > R + \Gamma > \Gamma - 1$, the constraint is not binding and the household is N-HtM with a smooth consumption profile.

³¹The final case, $m_2 > 0$ and $a = 0$, which is another form of N-HtM behavior, is never optimal given the assumptions above, but could be easily accommodated.

A.2 Solution with illiquid asset

We now turn to the general two-asset model. At $t = 1$ the consumption decision is

$$\begin{aligned} v_1(m_1, a) &= \max_{c_1, m_2} u(c_1) + u(m_2 + Ra + \Gamma) \\ &\quad s.t. \\ c_1 + m_2 &= y_1 + m_1 \\ m_2 &\geq 0 \end{aligned}$$

which has the solution

$$m_2 = \max \left\{ \frac{y_1 - \Gamma + m_1 - Ra}{2}, 0 \right\}. \quad (\text{A2})$$

The interior solution for m_2 implies a smooth consumption path $c_1 = c_2 = (y_1 + \Gamma + m_1 + Ra)/2$, while the corner solution yields the consumption pair $(c_1 = y_1 + m_1, c_2 = \Gamma + Ra)$. Note that under the low income path $y_L = 0 < \Gamma - 1 \leq \Gamma - m_1 + Ra$ for any feasible pair (a, m_1) . Therefore, equation (A2) implies that the constraint will bind at $t = 1$, regardless of the initial portfolio allocation, and $m_2 = 0$. In this case, the household is therefore HtM. Instead, under the high income path, $y_H > R + \Gamma \geq \Gamma - m_1 + Ra$ for any pair (a, m_1) . Hence equation (A2) implies that the constraint will not bind at $t = 2$, regardless of the initial portfolio allocation, and $m_2 > 0$. In this case, the household is N-HtM.

Next, consider the initial portfolio allocation decision. Under the high income path, when the constraint is not binding, the problem is

$$\begin{aligned} v_0 &= \max_{a, m_1} u \left(\frac{y_1 + \Gamma + m_1 + Ra}{2} \right) \\ &\quad s.t. \\ 1 &= a + m_1 \end{aligned}$$

It is immediate to see that the objective function is steeper in a than in m_1 because of the higher rate of return on the illiquid asset. Hence the household invests all of its initial endowment in the illiquid asset and we have a corner solution with $a = 1$. In this case, the household is N-HtM with a perfectly smooth consumption profile $c_1 = c_2 = (y_H + \Gamma + R)/2$.

Under the low income path ($y_1 = y_L = 0$) the constraint binds at $t = 1$ and $m_2 = 0$.

The problem becomes

$$\begin{aligned} v_0 &= \max_{a, m_1} u(m_1) + u(Ra + \Gamma) \\ &\quad s.t. \\ 1 &= a + m_1 \end{aligned}$$

which has the solution

$$a = \max \left\{ \frac{R^\sigma - \Gamma}{R + R^\sigma}, 0 \right\}, \quad m_1 = \min \left\{ \frac{R + \Gamma}{R + R^\sigma}, 1 \right\}. \quad (\text{A3})$$

Note that the portfolio allocation decision will always imply $m_1 > 0$ since the household needs liquidity at $t = 1$ for consumption. Thus, it only remains to determine when $a = 0$ and when $a > 0$.

If $1 < R \leq \Gamma^{\frac{1}{\sigma}}$, equation (A3) implies that $a = 0$ and the household is P-HtM. In this case the return on the illiquid asset is not large enough for the household to tolerate the large jump in consumption between $t = 1$ and $t = 2$ that would occur if it were to save some of the initial endowment in illiquid wealth. Hence $c_1 = 1$ and $c_2 = \Gamma$, and therefore $c_2 = \Gamma c_1$. When $R > \Gamma^{\frac{1}{\sigma}}$, we instead have an interior solution for the portfolio allocation, and the agent is W-HtM with consumption $c_1 = (R + \Gamma) / (R + R^\sigma)$ and $c_2 = R^\sigma c_1 > \Gamma c_1$.

A.3 Implications for the MPC out of an unexpected income transfer

Suppose that after the initial portfolio allocation decision, but before the consumption decision at $t = 1$, the household receives an unexpected transfer τ from the government. What is the household's MPC out of this transfer? A N-HtM household has an MPC of exactly $1/2$, since it smooths the payment equally across the two periods. Next, consider the problem of a household who, in absence of the transfer would be P-HtM, i.e., it faces $y_1 = y_L = 0$ and optimally chose the portfolio allocation ($m_1 = 1, a = 0$):

$$\begin{aligned} v_1(1, 0) &= \max_{c_1, m_2} u(c_1) + u(m_2 + \Gamma) \\ &\quad s.t. \\ c_1 + m_2 &= \tau + 1 \\ m_2 &\geq 0 \end{aligned}$$

which has the solution

$$m_2 = \max \left\{ \frac{\tau - \Gamma + 1}{2}, 0 \right\}.$$

For any small payment $0 < \tau < \Gamma - 1$, this household remains P-HtM and has an MPC of 1. Its consumption path is: $c_1 + 1 + \tau$, $c_2 = \Gamma$. If, instead, $\tau \geq \Gamma - 1$, the household becomes unconstrained, consumption equals $(\tau + \Gamma + 1)/2\tau$ in both periods and its MPC out of the transfer drops to $(\tau + \Gamma - 1)/2\tau$ which approaches $1/2$ as τ increases.

Finally, consider the problem of a household who, in absence of the transfer would be W-HtM, i.e., it faces $y_1 = y_L = 0$ and optimally chose $a = a^* = (R^\sigma - \Gamma) / (R + R^\sigma) > 0$:

$$\begin{aligned} v_1(1 - a^*, a^*) &= \max_{c_1, m_2} u(c_1) + u(m_2 + Ra^* + \Gamma) \\ &\quad s.t. \\ c_1 + m_2 &= \tau + (1 - a^*) \\ m_2 &\geq 0 \end{aligned}$$

The solution to this problem is:

$$m_2 = \max \left\{ \frac{\tau - \Gamma + (1 - a^*) - Ra^*}{2}, 0 \right\}.$$

This household has a MPC of 1 as long as $\tau \leq \Gamma - 1 + (R + 1)a^*$. This condition is weaker than the condition for a P-HtM to have a MPC of 1 because the income (and consumption) ratio between $t = 1$ and $t = 2$ is higher for a W-HtM compared to a P-HtM.³²

A.4 Unsecured credit

We now extend the model and allow households to access credit to finance consumption at $t = 1$. We assume that households can borrow up to a fraction $\phi \leq 1$ of their future income Γ and that the interest rate on borrowing is $R^b > 1$. Hence the credit limit is $\underline{m} = \phi\Gamma/R_b$. To make the exercise interesting, we impose the additional restriction that $R^b < \Gamma$, which ensures that a household with the low income path will always borrow a positive amount. Indeed, the no-borrowing case studied above can be interpreted as

³²Put differently, the shadow value of an additional unit of income at $t = 1$ is higher for the W-HtM than for the P-HtM. If we let λ be the shadow value of a unit of income in period 1, for a P-HtM agent we have $\lambda = u'(1 + \tau) - u'(\Gamma)$ and for a W-HtM agent we have $\lambda = u'(\tau + (1 - a)) - u'(Ra + \Gamma)$, which is larger.

a model where borrowing is allowed but $R^b \geq \Gamma$, and credit is so expensive that no household ever uses it. Since the role of the intertemporal elasticity of substitution is well understood from the previous analysis, we impose $\sigma = 1$ (i.e., logarithmic utility) to simplify the exposition.

A.4.1 Solution without illiquid asset

Under the high income path, the household is not constrained and chooses to save some of its high income into the liquid asset at $t = 1$. Since the borrowing constraint is not binding, the solution with borrowing is unchanged and $m_2 > 0$.

Under the low income path, the problem is more interesting. In this case, $m_2 \leq 0$ and at $t = 1$:

$$\begin{aligned} v_1 &= \max_{c_1, m_2} \log(c_1) + \log(R^b m_2 + \Gamma) \\ &\quad s.t. \\ c_1 + m_2 &= 1 \\ m_2 &\geq -\frac{\phi\Gamma}{R_b} \end{aligned}$$

which has the solution

$$m_2 = \max \left\{ -\frac{\Gamma - R^b}{2R^b}, -\frac{\phi\Gamma}{R_b} \right\}.$$

Since $R^b < \Gamma$, the household always borrows a positive amount. Moreover, if $R^b < \Gamma(1 - 2\phi)$, then the credit limit is binding. The household is forced to choose an increasing consumption path, $c_1 = 1 + \phi\Gamma/R_b$, $c_2 = \Gamma(1 - \phi)$. If, instead, $\Gamma > R^b \geq \Gamma(1 - 2\phi)$, the solution for m_2 is negative and interior: by borrowing, it can perfectly smooth consumption at the level $c_1 = c_2 = (R^b + \Gamma)/2$.

In light of the discussion in Section A.3 about MPCs, only the household at the credit limit has a MPC equal to 1, and only if the transfer is small enough not to change its HtM status. For small transfers, a household with an interior negative position is unconstrained and has a MPC equal to 1/2.

A.4.2 Solution with illiquid asset

Once again, under the high income path the household is not constrained at $t = 1$, so allowing for borrowing has no effect on its decisions. Under the low-income path where $y_1 = y_L = 0$, the household may want to borrow at $t = 1$. Its consumption decision at

$t = 1$ is:

$$\begin{aligned} v_1(m_1, a) &= \max_{c_1, m_2} \log(c_1) + \log(R^b m_2 + Ra + \Gamma) \\ &\quad s.t. \\ c_1 + m_2 &= m_1 \\ m_2 &\geq -\frac{\phi\Gamma}{R_b} \end{aligned}$$

which has the solution

$$m_2 = \max \left\{ \frac{R^b m_1 - \Gamma - Ra}{2R^b}, -\frac{\phi\Gamma}{R_b} \right\}$$

If $R^b < \Gamma$, then for every feasible portfolio allocation (m_1, a) , the first argument of the max operator in the above equation is negative, and hence $m_2 < 0$. The credit limit is binding when $R^b < [\Gamma(1 - 2\phi) + Ra]/m_1$, i.e., when borrowing is sufficiently cheap. In this case, consumption is given by $c_1 = m_1 + \phi\Gamma/R_b$ and $c_2 = \Gamma(1 - \phi) + Ra$. When borrowing is sufficiently expensive, i.e. $R^b \geq [\Gamma(1 - 2\phi) + Ra]/m_1$, the solution for m_2 is interior in the negative range and consumption is fully smoothed with $c_1 = c_2 = (R^b m_1 + \Gamma + Ra)/2$.

We now analyze the portfolio decision at $t = 0$. Since we are interested in characterizing HtM behavior, we focus on the case where the borrowing constraint binds, i.e. $m_2 = -\phi\Gamma/R_b$. In this case, the portfolio problem is:

$$\begin{aligned} v_0 &= \max_{a, m_1} u(m_1 + \phi\Gamma/R_b) + u(Ra + \Gamma(1 - \phi)) \\ &\quad s.t. \\ 1 &= a + m_1 \end{aligned}$$

with solution

$$a = \max \left\{ \frac{R + [(1 + R/R^b)\phi - 1]\Gamma}{2R}, 0 \right\}.$$

Using the assumed restriction $R^b < \Gamma$, it can be shown that $a = 0$ if $R < R^b(1 - \phi)/(1 + \phi)$ and a is strictly positive if $R > \Gamma(1 - \phi)/(1 + \phi)$. The former parameter configuration corresponds to a P-HtM household who has borrowed up to the credit limit. The latter parameter configuration corresponds to a household who chooses to save into the illiquid asset and then borrows up to its credit limit. This is a W-HtM household with negative liquid wealth (at the credit limit). Both households will have an MPC of 1 out of a small transfer.

B Measurement bias

Without loss of generality, let the pay-period be the unit interval $[0, 1]$, where 0 is the beginning and 1 the end. Denote a generic point within the pay period as t . Let y_0 be income paid at the beginning of the pay-period (by definition), s_0 be the liquid saving transferred from the previous pay-period, and s_1 be the end of the pay-period liquid saving. Initial and final balances of liquid wealth are therefore:

$$\begin{aligned}m_0 &= y_0 + s_0 \\m_1 &= s_1.\end{aligned}$$

We always assume that, during the pay-period, consumption expenditures c_t occur at a constant rate. Recall that this is optimal if, as we have assumed so far, there is no discounting and the return on the liquid asset is one. Then, at every $t \in [0, 1]$, we have that $c_t = y_0 + s_0 - s_1$. Moreover, at a point t the balances of liquid wealth are

$$m_t = y_0 + s_0 - c_t \cdot t = (y_0 + s_0)(1 - t) + s_1 t.$$

Note that, since t is uniformly distributed over the unit interval,

$$m_t \sim U[s_1, y_0 + s_0], \tag{B1}$$

and its mean (average balances of liquid wealth) is

$$\bar{m} = (y_0 + s_0 + s_1) / 2. \tag{B2}$$

We start by analyzing the measurement of HtM households when what we observe is average balances of liquid wealth.

B.1 Measurement through average balances

We first analyze the measurement of HtM households at the zero kink, then that of households at the credit limit.

B.1.1 Kink at zero

We abstract from debt for now, so $s_0 \geq 0$ and $s_1 \geq 0$. In the model, a household is HtM when $s_1 = 0$. According to our identification strategy of Section 3, a household is HtM when $\bar{m} \leq y_0/2$. Given (B2), this condition that corresponds to

$$s_1 + s_0 \leq 0.$$

Thus, our estimator is a lower bound because, for this condition to be true it must be that $s_0 = s_1 = 0$, whereas the true definition just requires $s_1 = 0$. Put differently,

$$\Pr[\bar{m} \leq y_0/2 | s_1 = 0] = \Pr[s_0 = 0 | s_1 = 0] < 1$$

and

$$\Pr[\bar{m} \leq y_0/2 | s_1 > 0] = \Pr[s_1 + s_0 = 0 | s_1 > 0] = 0$$

so, we always underestimate the number of true HtM and never mistaken a non HtM agent for a HtM.

B.1.2 Kink at the credit limit

Let the credit limit be $\underline{m} > 0$. In the model, a household is HtM when $s_1 = -\underline{m}$. According to our identification strategy, a household is HtM when $\bar{m} \leq y_0/2 - \underline{m}$, a condition that corresponds to

$$(s_1 + s_0)/2 \leq -\underline{m},$$

which, once again means our estimator is a lower bound because, for this condition to be true it must be that $s_0 = s_1 = -\underline{m}$, whereas the true definition just requires $s_1 = -\underline{m}$. Put differently,

$$\Pr[\bar{m} \leq y_0/2 - \underline{m} | s_1 = -\underline{m}] = \Pr[s_0 = -\underline{m} | s_1 = -\underline{m}] < 1$$

and

$$\Pr[\bar{m} \leq y_0/2 - \underline{m} | s_1 > -\underline{m}] = \Pr[(s_1 + s_0)/2 = -\underline{m} | s_1 > -\underline{m}] = 0$$

so, once again, we always underestimate the number of true HtM and never mistaken a non HtM agent for a HtM.

B.2 Measurement through random balances

Suppose now we observe balances m_t at a random point in time t over the period $[0, 1]$. Recall that in this case m_t is uniformly distributed as defined by (B1).

B.2.1 Kink at zero

We identify as HtM a household who has $m_t \leq y_0/2$, an event occurring with probability

$$\Pr [m_t - y_0/2 \leq 0] = \frac{y_0/2 - s_1}{y_0 + s_0 - s_1}.$$

If $s_1 = 0$, then the household is truly HtM, but we would catch him only with probability $y_0/[2(y_0 + s_0)]$, so we are missing some truly HtM households. In other words

$$\Pr [m_t - y_0/2 \leq 0 | s_1 = 0] = \frac{y_0/2}{y_0 + s_0} < 1.$$

If the household is not HtM, i.e., $s_1 > 0$, then we would mistaken it for a HtM household only if $s_1 < y_0/2$ because

$$\Pr [m_t - y_0/2 \leq 0 | s_1 > 0] = \frac{y_0/2 - s_1}{y_0 + s_0 - s_1}.$$

To sum up, we may mistaken a non HtM household for a HtM only if its end-of period savings are below a half of pay-period income.

B.2.2 Kink at the credit limit

In this case, we call HtM anyone with $m_t \leq y_0/2 - \underline{m}$. The probability we identify it as HtM is

$$\Pr [m_t - y_0/2 + \underline{m} \leq 0] = \frac{y_0/2 - \underline{m} - s_1}{y_0 + s_0 - s_1}.$$

If $s_1 = -\underline{m}$, and it is truly a HtM household, then

$$\Pr [m - y_0/2 + \underline{m} \leq 0 | s_1 = -\underline{m}] = \frac{y_0/2}{y_0 + s_0 + \underline{m}} < 1,$$

since $s_0 \geq -\underline{m}$, which means that we miss some truly HtM households.

If the household is not HtM, i.e., $s_1 > -\underline{m}$, then we would call him mistakenly HtM only if $\underline{m} + s_1 < y_0/2$ because

$$\Pr [m_t - y_0/2 \leq 0 | s_1 > 0] = \frac{y_0/2 - \underline{m} - s_1}{y_0 + s_0 - s_1}.$$

To sum up, we may mistaken a non HtM household for a HtM only if its end-of period balances of liquid wealth are less than half-income away from the credit limit.

C Survey data on household portfolios

The countries included in our study are the U.S., Canada, Australia, the U.K., and the four largest economies in the Euro area: Germany, France, Italy, and Spain. In this Appendix, we provide background information on each survey.

C.1 United States: SCF

Our data for the United States come from the Survey of Consumer Finances (SCF). The SCF is sponsored by the Board of Governors of the Federal Reserve System in cooperation with the Statistics of Income Division of the Internal Revenue Service (IRS). The survey has been conducted every three years and collects detailed information on household balance sheets, income, and demographic characteristics for a representative cross-section of U.S. households. We conduct analysis on the 1989 to 2010 surveys.³³ While the surveys do not normally follow households over time, there is a panel component to the 2007 survey where a subset of households were contacted and re-surveyed in 2009. See Bricker et. al (2011) for more information on the 2007-2009 panel of the SCF.

The target population for the survey is all private households residing in the U.S. at the time of data collection. The SCF uses a dual frame sample design. Households in the first frame are intended to provide representative coverage of various characteristics of households in the United States. Households in the second frame are drawn from statistical records derived from tax information provided by the IRS and are intended to disproportionately select relatively wealthy households. This oversampling design allows the SCF to more accurately measure the distribution and composition of wealth

³³The survey started in 1983, but major technical revisions to the survey were implemented in 1989 and the structure and questions have largely been preserved since then. Since 1992, data have been collected by the National Opinion Research Center at the University of Chicago.

for the population as a whole, given the extreme right skewness in the distribution of holdings for many asset classes.

The main interviewee is the household head. The head is defined as the core individual in single households, the male in mixed-sex couples, and the older individual in same-sex couples. In the case of couples, either member can be interviewed and the data are rearranged after to define the household head in this way. Summary information is then collected about all other household members. Labor market, pension, and demographic data on the spouse or partner of the respondent are also collected. See Kennickell (2005) for more information of the sample design of the SCF.

C.2 Canada: SFS

Our data for Canada come from the Survey of Financial Security (SFS). The SFS is a cross-sectional survey implemented by Statistics Canada in 1999 and 2005, and is intended to provide a comprehensive picture of net worth of Canadian households. In our analysis, we use data from 2005. The survey asks questions on the value of all major financial and non-financial assets and liabilities.

The surveyed households are a representative sample of all private households in Canadian provinces. Like the SCF, the SFS uses a dual frame sample design. The main sample is a sample selected from the Labour Force Survey sampling frame. In order to over-sample high income households, the second sample is drawn from geographic areas in which there are a large proportion of family units with total income over a certain threshold.

All individuals older than 15 years of age in the household are asked questions regarding income, demographics, education, and employment. Questions regarding household assets and liabilities are asked to the household member deemed most knowledgeable on the subject. See Pensions and Wealth Surveys Section (2006) for more information about the 2005 SFS.

C.3 Australia: HILDA

Our data for Australia come from the Household, Income and Labour Dynamics in Australia (HILDA) Survey. The Survey is managed by the Melbourne Institute of Applied Economic and Social Research at the University of Melbourne. HILDA is a broad social and economic longitudinal survey, with particular attention paid to family and household formation, income and work. Wave One of the survey was implemented

in 2001, and households in the survey have since been interviewed annually.

The original sample for the HILDA survey was a large national probability sample of Australian households occupying private dwellings. All members of the households providing at least one interview in Wave 1 form the basis of the panel to be pursued in each subsequent wave. The sample has been gradually extended to include any new household members resulting from changes in the composition of the original household.

In addition to regular questions about economic and subjective well-being, the survey features special modules covering specific topics. In particular, Waves Two (2002), Six (2006), and Ten (2010) contain data from the wealth module that examines the composition of household's balance sheets.

Data for our analysis come from the Household Form and the Person Questionnaire. The Household Form records basic information about the composition of the household. The Household Questionnaire is administered primarily to one member of the household, and covers child-care, housing, household spending, and the wealth modules in Waves Two, Six, and Ten. The Person Questionnaires are asked to all members of the household aged 15 years and older, and collects information on family background, education, employment, and income among other things. See Watson and Wooden (2002) for more information on the HILDA.

C.4 United Kingdom: WAS

Our data for the United Kingdom come from the Wealth and Assets Survey (WAS). The WAS is a longitudinal survey that is conducted by the Office of National Statistics (ONS). The survey is intended to measure the economic well-being of households in the U.K., by documenting the level of household savings and debt, lifecycle accumulation of wealth, and participation in pension schemes.

For the first wave, the survey aimed to sample all persons living in private households in Great Britain. The WAS also uses a dual frame design, using the first frame to meet precision targets, and the second frame to over-sample the top wealth decile. The sample for the first frame was drawn from the Royal Mail's database of all addresses in the UK. Households where at least one member was likely to have total financial wealth above a certain threshold were flagged by Her Majesty's Revenue and Customs. Flagged households were sampled in such a way that they had two and a half times higher probability of being sampled than non-flagged households. Wave One was conducted from July 2006 to June 2008, and attempts were made to contact respondents for a follow-up interview two years later for Wave Two. About two-thirds of cooperating

households from Wave One completed the Wave Two interview from July 2008 to June 2010. In our analysis, we use data from Wave Two.

The questionnaire is divided into two parts. The first part is the household questionnaire which is completed by one person in the household designated to be the household reference person, and collects household-level information on household demographics, as well as information about household assets and liabilities. The second part of the questionnaire is an individual questionnaire administered to each adult aged 16 or over in the household, and asks in-depth questions about economic status, education, employment, benefits, and individual financial assets. See Daffin (2009) and Black (2011) for more information on the WAS.

C.5 Euro area: HFCS

Our data for Germany, France, Italy and Spain come from the Household Finance and Consumption Survey (HFCS). The HFCS is a joint project administered by all of the central banks of the Eurosystem and three National Statistical Institutes. The survey provides detailed information on balance sheets, demographics, and other economic variables for households in Euro area countries. Fieldwork in the various countries was conducted between November 2008 and August 2011.

The HFCS is conducted and financed by each participating institution. For some member countries, a previous wealth survey had already existed, and for others, an entirely new survey had to be set up. The HFCS represents an effort towards gradual harmonization of the content of the surveys across the member countries. The survey will be conducted in each country every two to three years.

The core questionnaire, asked in every country, is composed of three parts. The first comprises questions regarding the household as a whole and contains questions regarding household assets and liabilities, transfers, and consumption-saving decisions. This part is answered by one member of the household deemed to be the main respondent. The second part of the questionnaire is asked to all members of the household and collects basic demographic information. The final part of the questionnaire is given only to members of the household over 16 years of age and covers information regarding employment, pension entitlements, and labor-market income.

There are also a set of standardized, non-core extension modules that the member countries are allowed to include at their discretion in addition to the core questionnaire. These non-core questions typically go into more detail on some aspect of the core questionnaire that the member country wishes to explore. For example, Spain asks

questions that are designed to examine methods by which households pay their bills.

The target population for the survey is all private households and their current members residing in the national territory at the time of data collection. The sampling design, however, is chosen by each participating country. France uses a dual frame design, exploiting individual data on taxable wealth to create the wealthy sample. The wealthy sample is divided into four strata and sampled proportionally according to the relative size of the strata. Germany uses regional level taxable income, and oversamples small municipalities and, in larger municipalities, street sections with average income over a threshold. Spain defines eight wealth strata, based on individual taxable wealth, that are oversampled progressively at higher rates. Italy did not oversample in any way. See Eurosystem Household Finance and Consumption Network (2013a) and (2013b) for more information on the HFCS.

C.6 Data access

The datasets can be accessed at the following websites:

SCF: <http://www.federalreserve.gov/econresdata/scf/scfindex.htm>

SFS: <http://www23.statcan.gc.ca/imdb/p2SV.pl?Function=getSurvey&SurvId=1706&InstaId=8244>

HILDA: <http://www.melbourneinstitute.com/hilda/>

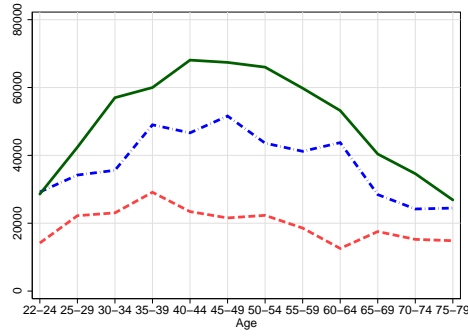
WAS: <http://discover.ukdataservice.ac.uk/catalogue/?sn=7215&type=Data%20catalogue>

HFCS: http://www.ecb.europa.eu/home/html/researcher_hfcn.en.html

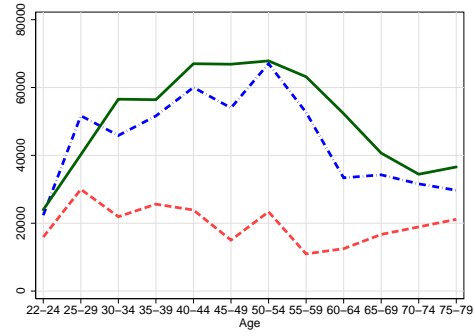
The SCF, SFS, and WAS datasets are open to public use. Applications for access to the HILDA and HFCS datasets are available at their respective websites.

P-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Median liquid wealth / income	0.000	0.000	0.013	0.000	0.009	0.077	0.000	0.000
Mean liquid wealth / income	-0.676	-0.680	-0.399	-1.332	-0.012	-0.092	-0.002	-0.228
Frac. neg. liquid wealth	0.185	0.274	0.122	0.289	0.051	0.072	0.007	0.102
Frac. neg. illiquid wealth	0.140	0.117	0.043	0.037	0.039	0.071	0.001	0.164
W-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Median liquid wealth / income	0.001	0.000	0.023	0.000	0.059	0.122	0.000	0.011
Mean liquid wealth / income	-1.225	-1.167	-0.775	-1.902	-0.019	-0.206	-0.067	-0.266
Frac. neg. liquid wealth	0.382	0.467	0.305	0.464	0.082	0.105	0.054	0.090
Frac. neg. illiquid wealth	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Median housing / illiquid wealth	0.818	0.747	0.408	0.312	0.000	0.389	1.000	1.000
Mean housing / illiquid wealth	0.608	0.592	0.414	0.398	0.369	0.471	0.968	0.985
Frac. pos. housing equity	0.786	0.800	0.555	0.667	0.464	0.505	0.834	0.933
Median retire / illiquid wealth	0.045	0.242	0.533	0.659	0.000	0.066	0.044	0.023
Mean retire / illiquid wealth	0.318	0.398	0.576	0.583	0.093	0.175	0.390	0.036
Frac. pos. retirement account	0.545	0.645	0.925	0.862	0.271	0.020	0.041	0.021
N-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Median liquid wealth / income	1.707	2.077	3.410	4.690	1.133	0.878	3.701	2.129
Mean liquid wealth / income	15.177	12.233	34.785	38.504	6.144	4.730	8.918	6.977
Frac. neg. liquid wealth	0.117	0.138	0.065	0.149	0.081	0.042	0.016	0.017
Frac. neg. illiquid wealth	0.025	0.034	0.004	0.004	0.006	0.007	0.001	0.003
Median housing / illiquid wealth	0.650	0.768	0.708	0.432	0.534	0.839	1.000	1.000
Mean housing / illiquid wealth	0.593	0.645	0.597	0.459	0.446	0.604	0.965	0.948
Frac. pos. housing equity	0.713	0.721	0.770	0.785	0.533	0.654	0.770	0.877
Median retire / illiquid wealth	0.237	0.209	0.274	0.545	0.000	0.029	0.036	0.021
Mean retire / illiquid wealth	0.329	0.345	0.387	0.524	0.052	0.081	0.174	0.083
Frac. pos. retirement account	0.618	0.572	0.878	0.849	0.259	0.045	0.107	0.042

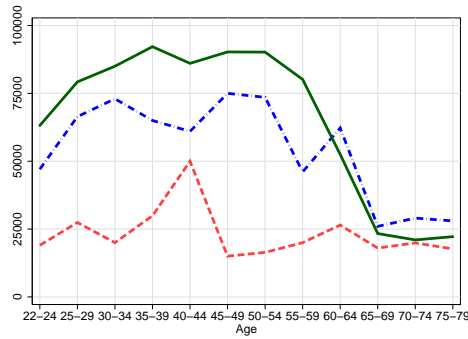
Table D1: Portfolio characteristics by country and HtM status. To reduce the sensitivity outliers, means are computed after trimming the overall top and bottom 0.1 percent of that statistic's distribution.



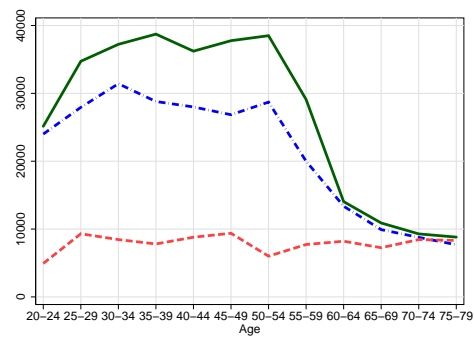
(a) United States



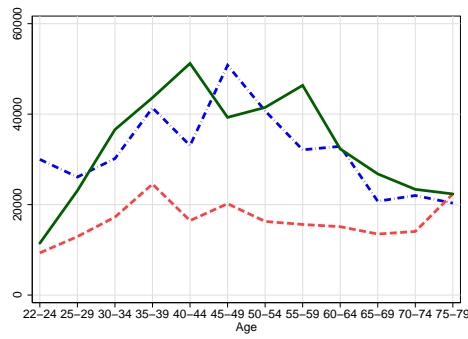
(b) Canada



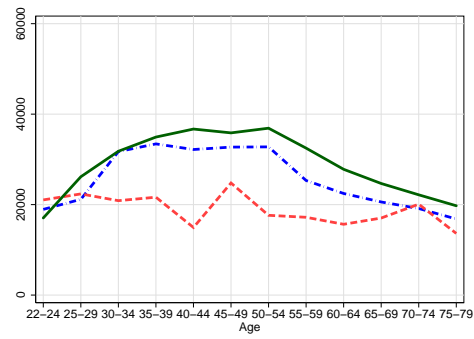
(c) Australia



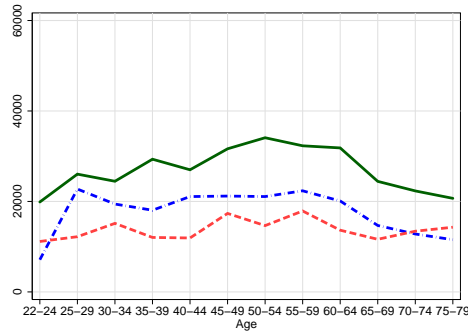
(d) United Kingdom



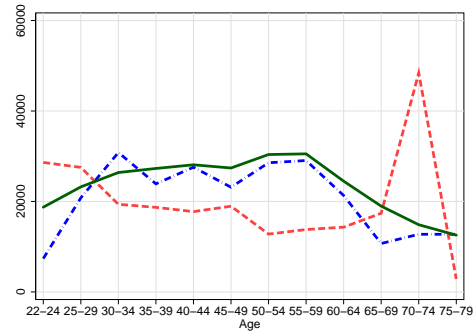
(e) Germany



(f) France



(g) Italy



(h) Spain

Figure D1: Age profile of median income by HtM status by country.

References

- BLACK, O. (2011): “Wealth in Great Britain, Main Results from the Wealth and Assets Survey 2008/10,” Discussion paper, Office for National Statistics.
- BRICKER, J., B. BUCKS, A. KENNICKELL, T. MACH, AND K. MOORE (2011): “Surveying the Aftermath of the Storm: Changes in Family Finances from 2007 to 2009,” Finance and Economics Discussion Series 17, Board of Governors of the Federal Reserve Bank.
- DAFFIN, C. (2009): “Wealth in Great Britain, Main Results from the Wealth and Assets Survey 2006/08,” Discussion paper, Office for National Statistics.
- EUROSYSTEM HOUSEHOLD FINANCE AND CONSUMPTION NETWORK (2013a): “The Eurosystem Finance and Consumption Survey: Methodological Report for the First Wave,” Statistics Paper Series 1, European Central Bank.
- (2013b): “The Eurosystem Finance and Consumption Survey: Results from the First Wave,” Statistics Paper Series 2, European Central Bank.
- KENNICKELL, A. B. (2005): “The Good Shepherd: Sample Design and Control for Wealth Measurement in the Survey of Consumer Finances,” Discussion paper, Board of Governors of the Federal Reserve Bank.
- PENSIONS AND WEALTH SURVEYS SECTION (2006): “The Wealth of Canadians: An Overview of the Results of the Survey of Financial Security, 2005,” Discussion paper, Statistics Canada.
- WATSON, N., AND M. WOODEN (2002): “The Household, Income and Labour Dynamics in Australia (HILDA) Survey: Wave 1 Survey Methodology,” HILDA Project Technical Paper Series 1, Melbourne Institute of Applied Economic and Social Research, University of Melbourne.