The Wealthy Hand-to-Mouth *

Greg Kaplan
Princeton University, NBER, and IFS
gkaplan@princeton.edu

Giovanni L. Violante
New York University, CEPR, NBER, and IFS
glv2@nyu.edu

Justin Weidner
Princeton University
jweidner@princeton.edu

First version: January 2014 – This version: March 2014

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). This portfolio configuration implies that these households have a high marginal propensity to consume out of transitory income changes –a key determinant of the macroeconomic effects of fiscal policy. The wealthy hand-to-mouth, therefore, behave in many respects like households with little or no net worth, yet they escape standard definitions and empirical measurements based on the distribution of net worth. 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. Using PSID data, we estimate that the wealthy hand-to-mouth have a strong consumption response to transitory income shocks. Finally, we discuss the implications of this group of consumers for macroeconomic modeling and policy analysis.

Keywords: Consumption, Hand-to-Mouth, Household Balance Sheet, Liquidity.

JEL Classification: D31, D91, E21, H31

*We thank Yu Zhang for outstanding research assistance, and Mark Aguiar, Karen Pence, David Romer, and Justin Wolfers for comments. This research is supported by grant no. 1127632 from the National Science Foundation.

1 Introduction

The life-cycle permanent income hypothesis (LC-PIH) is a valuable organizing framework for the analysis of 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 types of income dynamics should be smoothed.

At the microeconomic level, a large body of evidence finds that consumption overreacts to predictable income growth (see Jappelli and Pistaferri, 2010, for a recent survey). Some of the most convincing studies in this literature are based on quasi-natural experiments which measure the consumption response to the U.S. fiscal stimulus payment episodes of 2001 and 2008.[1] Johnson, Parker, and Souleles (2006), Parker, Souleles, Johnson, and McLelland (2013), and Broda and Parker (2012) concluded that, in both episodes, the consumption response was strong: between 20 and 40 percent of the stimulus payments were spent by households on nondurables in the quarter that they are received. Shapiro and Slemrod (2003a, 2003b, 2009) substantiate these studies with qualitative surveys on how consumers use their rebates and find comparable effects. A number of additional studies based on micro data also find sizable consumption responses to anticipated fluctuations in income. Examples are the reaction of consumption to changes in social security taxes (Parker, 1999), the analysis of federal tax refunds (Souleles, 1999) and, most recently, the consumption response to a fiscal stimulus episode in Singapore (Agarwal and Quian, 2013).

Similar anomalies in the joint behavior of income and consumption, relative to the LC-PIH, have been identified from aggregate data. In a series of papers, Campbell and Mankiw (1989, 1990, 1991) analyze U.S. and cross-country time series and show that expected changes in aggregate consumption tend to be associated with expected changes in aggregate income. Moreover, expected consumption growth is uncorrelated with the real interest rate, a result that –as long as the elasticity of intertemporal

[1] In these two episodes, the U.S. Treasury selected the week of the payment based on the second-to-last digit of the taxpayer's social security number, a digit that is effectively randomly assigned. This randomization allows one to identify the causal effect of the fiscal transfer by simply comparing the spending of households that received the rebate earlier with that of households who received it later.

substitution is not zero– implies a break-down of the forward looking Euler equation holding with equality.[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 who spend all of their available resources in every pay-period, and hence do not carry any wealth across periods. HtM consumers have a high marginal propensity to consume (MPC) out of transitory income changes that 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 is 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 only one asset or feature two assets with different risk profiles, but with the same degree of liquidity. Notable examples are the so called Bewley models with uninsurable idiosyncratic risk and credit constraints, in the tradition of Huggett (1996), Aiyagari (1994), Rios-Rull (1995), and Krusell and Smith (1998), and the so-called spender-saver models in the tradition of Campbell and Mankiw (1989) that feature impatient and patient consumers with complete markets. This latter class of models has been recently revived 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. The models of Krusell and Smith (1998) and Carroll, Slacalek, and Tokuoka (2014a, 2014b) combine the spender-saver insight of heterogeneity in patience with an otherwise standard one-asset incomplete-market 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, and therefore consume all of their disposable income every period. Clearly, such households would not be picked up by standard measurements since they own positive –and often

[2]Two related studies are Attanasio and Weber (1993) and Ludvigson and Michaelides (2001).

substantial– amounts of net worth.

Thus to obtain a comprehensive measurement of HtM behavior using cross-sectional survey data on household portfolios, a far better strategy is to use, as a guiding framework, a model with two assets –one liquid asset and one illiquid asset that yields a higher return, but that 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 (2010), 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) who hold little or no liquid wealth and no illiquid wealth, and the wealthy hand-to-mouth (W-HtM) who 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, as we show in the paper, along many other important dimensions W-HtM households are more similar to non HtM (N-HtM) households. As a result, the W-HtM cannot be fully assimilated into either group, and are therefore best represented as a separate class of households.

The goal of this paper is to investigate W-HtM behavior theoretically and empirically, and to reflect on the implications of this peculiar, but sizable, group of households for macroeconomic modeling and policy analysis.

First, we ask why households with significant wealth would optimally choose to consume 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, a portfolio composition with positive amounts of illiquid wealth and zero liquid wealth is optimal. Such wealthy HtM households are better off bearing the welfare loss from income fluctuations rather than smoothing consumption. The reason is that 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 the calculations of Browning and Crossley (2001) who showed that, in the context of a plausibly parameterized life-cycle buffer stock model, the utility loss from setting consumption equal to income, instead of fully optimizing, is second order. Similar calculations were performed by Cochrane (1989) and Krusell and Smith (1998) 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 deal with a number of measurement issues.

Next, we ask how large the share of wealthy HtM households is in the population, what their demographic characteristics are relative to poor and non HtM consumers, and how their balance sheet compares with that of the N-HtM and P-HtM. Finally, we investigate the persistence of HtM status over time to understand whether being W-HtM is a transient or a persistent state during the life cycle of a household. This empirical analysis is based on cross-sectional survey data on household portfolios for a number of countries, specifically, 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 examined these data in the past, 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 around 2010.

The analysis of U.S. data leads to seven main findings: (i) between 25 and 40 percent of U.S. households are HtM, with our preferred estimate at 1/3 of the population; (ii) 1/3 of HtM households are poor HtM and 2/3 are wealthy HtM. Therefore, the W-HtM represent the vast majority of this group; (iii) while households appear to be most frequently P-HtM at young ages, the age profile of the W-HtM is hump-shaped and peaks around age 40; (iv) the W-HtM typically hold sizable amounts of illiquid wealth: for example, the median at age 40 is around $50,000; (v) W-HtM households look a lot like the unconstrained N-HtM in terms of their age-profile of income, and in terms of the shares of illiquid wealth held in housing and retirement account; (vi) the W-HtM status is quite transient: we estimate that it lasts, on average, 2.5 years; (vii) finally, estimates based on net worth miss at least half of HtM consumers.

Comparing the U.S. economy to the other countries we study, some interesting findings emerge. In all of the eight 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. As in the U.S., it is over 30 percent in Canada, U.K., and Germany, but 20 percent or less of the population in Australia, France, Italy, and Spain. We attribute the low number of HtM households in Australia to the compulsory private retirement saving system which means that even low-income young individuals have some illiquid wealth. For the

Euro-area countries, we observe that holdings of consumer debt are minimal suggesting that the substantial liquid savings we observe, even among the income-poor, may act as a buffer stock that substitutes for the expensive and limited access to credit.

In Section 7 we show that the HtM status of a household has strong predictive power for the consumption response to transitory shocks. We apply the identification strategy of Blundell, Pistaferri and Preston (2008) to panel data on income and consumption from the U.S. 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 response between the two groups.

In Section 8, we argue that in the cast of characters of macroeconomic models, the W-HtM deserve their own separate status. We use our empirical estimates of the fraction of households in each of the three HtM groups, together with simulated MPCs from three alternative models, to show that the W-HtM cannot be assimilated to either the P-HtM or the N-HtM. We highlight three areas where models that feature only two types of HtM households provide misguided intuition about the effects of fiscal stimulus policy: the degree of cross-country dispersion in consumption responses, the degree of non-linearity in consumption responses with respect to the size of the stimulus payment, and the optimal degree of phasing of stimulus payments with income for maximizing the aggregate consumption response. Section 9 summarizes and concludes the paper.

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

We start by analyzing a simple two-period model in order to illustrate the determinants of hand-to-mouth behavior. The model also offers some useful guidance for our measurement exercise.

Consider a household that lives for two periods, $t = 1, 2$. Preferences over consumption in the two periods are given by

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

with no discounting. The household starts period 1 with an initial endowment ω and 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 2.4. 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 CRRA 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 $a > 0$.[3] 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.

2.1 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

[3]The 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.

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

$$m_2 = \max\left\{\frac{y_1 - \Gamma + m_1}{2}, 0\right\}. \tag{1}$$

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 (1) 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.

2.2 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

$$m_2 = \max\left\{\frac{y_1 - \Gamma + m_1 - Ra}{2}, 0\right\}. \tag{2}$$

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 (2) implies that the constraint will bind at $t = 2$, 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 (2) 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

$$a = \max\left\{\frac{R^\sigma - \Gamma}{R + R^\sigma}, 0\right\},\ m_1 = \min\left\{\frac{R + \Gamma}{R + R^\sigma}, 1\right\}. \tag{3}$$

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 (3) 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 $c_2 = R^{\sigma} c_1 > \Gamma c_1$.

It is useful to explain the role of the three model parameters in determining wealthy HtM behavior: the relative return on the illiquid asset R, income growth under the low income path Γ, and the elasticity of intertemporal substitution σ. A high relative return R makes the illiquid asset more attractive and induces the agent to absorb consumption swings across periods in order to achieve a higher overall level of consumption. Steep income growth Γ reduces the role of the illiquid asset as a saving instrument, since the slope of the income profile guarantees high consumption later in life already. Finally, the larger the elasticity of intertemporal substitution σ, the more the household is willing to accept consumption fluctuations across periods, and the more likely it is to become W-HtM rather than P-HtM.

Finally, note that since the model is deterministic, some households choose to invest in the illiquid asset, thus diverting resources from liquid wealth, even though they know with certainty that they will be constrained at a future point in time. Nevertheless, they do find it optimal because the welfare gain from the rise in the overall level of consumption more than compensates for the welfare loss from consumption fluctuations between the two dates. In our two period model, we have abstracted from discounting. In a multi-period model with geometric discounting, all the qualitative conclusions remain intact, as we show in Kaplan and Violante (2014). 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 easier to generate wealthy HtM agents.

2.3 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=1 is higher for a W-HtM compared to a

P-HtM.[4]

This finding suggests that, in a general model with more income heterogeneity (as in Kaplan and Violante, 2014), the average MPC among W-HtM households is larger than the average MPC among P-HtM households. We return to this point in Section 8. Finally, note that all the results in this section carry over to the case of an anticipated transfer, as long as the transfer is small enough that it does not change the HtM status at $t = 1$.

2.4 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 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.

2.4.1 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

[4]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 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.

at $t = 1$:

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

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 $c_1 = c_2 = \left(R^b + \Gamma\right)/2$.

In light of the discussion in Section 2.3 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.

2.4.2 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:

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

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 $c_1 = c_2 = \left(R^b m_1 + \Gamma + Ra\right)/2$.

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

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

Using the 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.

3 Identifying hand-to-mouth households in the data

The stylized model of 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: the unsecured credit limit and zero liquid assets. The 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.

Translating these theoretical definitions into empirical measurement poses a number of practical challenges. In this section, we discuss the key problems that arise when trying to identify poor and wealthy HtM households from cross-sectional survey data on household portfolios, such as the data that we use in our empirical analysis and that we describe in Section 4. Broadly speaking, these data contain household-level information on income and on different categories of assets and liabilities that can be aggregated into net liquid and illiquid wealth. We refer the reader to the next section for the exact definition of liquid and illiquid wealth, country by country.

We now describe our strategy for identifying HtM households in the data separately for households at the zero kink and for households at the credit limit. Let y_{it} denote the income of household i in period t, let m_{it} denote holdings of liquid wealth and let a_{it} denote holdings of illiquid wealth. We always assume that income is paid as liquid wealth.

First, consider a HtM household who is at the zero kink. In the model, this household does not borrow or save into liquid assets: it consumes all its cash-on-hand for the period and carries zero liquid wealth between t and $t+1$. Suppose we have information on the length of the pay-period of the earners in the household, or equivalently the frequency of pay (weekly, bi-weekly, monthly, etc...), and we observe the average balance of liquid wealth over this period.

A conservative criterion to identify HtM agents on the zero kink in the data is to count those households in the survey whose average balances of liquid wealth 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" presumes resources being consumed at a constant rate. Specifically, a household is P-HtM at the zero kink if

$$a_{it} = 0, \quad \text{and} \quad 0 \leq m_{it} \leq \frac{y_{it}}{2} \tag{4}$$

and W-HtM at the zero kink if

$$a_{it} > 0, \quad \text{and} \quad 0 \leq m_{it} \leq \frac{y_{it}}{2}. \tag{5}$$

Panel (a) of Figure 1 depicts a graphical representation of the dynamics of income and average cash-on-hand $\bar{m}_{it}$ over a pay period for a HtM individual who starts and

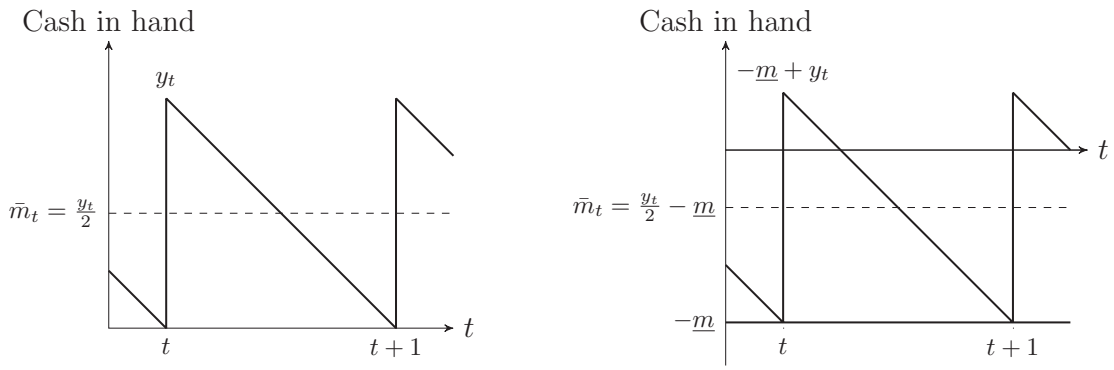

(a) HtM at the zero kink

(b) HtM at the credit limit

Figure 1: Illustration of two cases of HtM behavior

ends the period at the zero kink. Its liquid balances (or cash in hand) peak at y_{it} at the beginning of the pay period, when income is paid into the liquid account, and are depleted constantly until they reach zero at t+1. Average balances over the period are therefore equal to half income.

This estimator 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 average liquid balances above half their earnings. For example, a household who starts the period with positive liquid savings in addition to earnings and ends the period with zero liquid savings is HtM, but its average liquid balance is above half earnings, and so it would not be counted as HtM by this criterion.

Unfortunately, most surveys do not record the frequency of pay for individual earners. In the benchmark measurement exercise, we assume that labor income is received by the household twice a month. In the sensitivity analysis, we report calculations under the assumption that labor income is received weekly or monthly.

In certain surveys, instead of average balances, we observe wealth at a random point in time during the pay-period. Our proposed measurement is still valid, as long as the interview dates are uniformly distributed over the pay-period and are uncorrelated with (y_{it}, m_{it}, a_{it}). In other words, this source of measurement error would tend to underestimate (overestimate) HtM behavior if the interview date is close to the beginning (end) of the pay period, but the error averages to zero.

Next, consider a HtM household at the credit limit $-\underline{m}_{it} < 0$. This is a household who 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

$$a_{it} = 0, \quad m_{it} \leq 0 \quad \text{and} \quad m_{it} \leq \frac{y_{it}}{2} - \underline{m}_{it}, \tag{6}$$

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

$$a_{it} > 0, \quad m_{it} \leq 0 \quad \text{and} \quad m_{it} \leq \frac{y_{it}}{2} - \underline{m}_{it}. \tag{7}$$

Panel (b) 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 individual who starts and ends the period at the credit limit.

When, as in most surveys, individual credit limits are not reported, we assume that the credit limit is a multiple of individual income and we experiment with a plausible range of values. It is easy to see that this criterion is also conservative: a household who 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 (6) and (7).

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

$$0 \leq n_{it} \leq \frac{y_{it}}{2} \quad \text{or,} \quad n_{it} \leq 0 \quad \text{and} \quad n_{it} \leq \frac{y_{it}}{2} - \underline{m}_{it} \tag{8}$$

A 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;). Examples of such expenditures include rent, mortgages or other loan payments, utility bills, school, gym, or club fees, 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 occurs 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.

Finally, whenever the data allow, we also use direct survey questions as alternate estimates of the fraction of HtM households. These questions typically ask (i) whether expenditures over the last month have exceeded income, abstracting from purchases of large durable goods such as housing or cars, and (ii) whether over the past month, the household has saved a positive amount. 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

In this section we describe the cross-sectional survey data that we use to measure HtM behavior. The countries that we 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 order to categorize a household as W-HtM, P-HtM, or N-HtM, we need information on labor income and on the amounts of assets and liabilities held in various categories on the household's balance sheet. We first provide background information on each survey. Next, we describe how the definitions of the balance sheet items differ between surveys. Finally, we present some descriptive statistics on the distribution of asset and liabilities across countries. Appendix A contains information on how to access each of the five data sets.

4.1 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.[5] 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 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.

4.2 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 Cana-

[5]The 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.

dian 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 Statistics Canada (2006) for more information about the 2005 SFS.

4.3 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.

4.4 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 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 Dafin (2009) and Black (2011) for more information on the WAS.

4.5 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 of 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.

	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
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. Age selection for U.K. is 20-79 as age is provided in 5-year bins. Self employment income is not provided in the SFS for Canada.

4.6 Sample selection and data comparability

Each individual survey is tailored to its own country, and as such, the questions asked and the survey definition of particular asset classes will vary. 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 the head of 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.

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, unemployment benefits, worker's compensation, regular private transfers such as child support and alimony, regular public transfers such as 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 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.

[6]The 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.

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]

The main discrepancy in income measurement across surveys is that in Canada income is after taxes, whereas in all other countries the surveys ask for 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 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 holdings of cash. To address this problem, we resort to an imputation procedure. We impute cash holdings for all surveys based on data from the Survey of Consumer Payment Choice (SCPC) administered by the Federal Reserve Bank of Boston (see Foster et al., 2011). We impute cash holdings by taking the ratio of average cash holdings in the SCPC in 2010 to the median value of checking, saving, money market

[7]The 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 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 for its reference period for income which runs from July 1, 2009 to June 30, 2010.

[8]ISAs 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.

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, after the most recent payment, accrue interest.[10] 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, certificate of deposits, 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 using the 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).[11] In the U.K. WAS, we take illiquid wealth to include the value of the main residence, other houses, and land less mortgage and land

[9] Average 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 checking, saving, money market and call accounts and find it to be about 6% in Spain.

[10] As most credit cards in the U.S. feature a one month grace period on purchases which makes them a close substitute for cash in the very short term, we remain conservative and restrict our measure of credit card debt only to debt for those households that do not regularly pay off their balances in full each month. A specific question in the SCF allows us to identify such households.

[11] Superannuation has some features of private retirement accounts, such as 401(k) accounts in the U.S., which we include into illiquid wealth, and some features of public pensions (e.g., the compulsory nature of a minimum contribution) which we exclude from illiquid wealth. As a result, we also offer a sensitivity analysis where we exclude superannuation wealth from illiquid assets.

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 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 up to purchase them, and other non-financial wealth not included in our baseline, such as antiques, artwork, jewels, gold, etc.[12] Note that changing the definition of illiquid wealth only affects the split between poor and wealthy HtM, but not the total number of HtM households.

The reference period for the wealth questions varies across surveys. In the SCF, for most assets it is the interview date, 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 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 date of the interview, and Italy uses December 31, 2010.

4.7 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. 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

[12] In our robustness 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	56944	0.887	113565	0.897	380889	0.984	187157	0.880
Net liquid wealth	1950	0.767	3194	0.769	12551	0.884	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	0.297	0	0.232	0	0.257	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.

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, 2002). The median value of housing equity relative to median annual income is especially remarkable in Italy and Spain, where this ratio 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 PAYG 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. 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.[13] 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 occupational pension plan of their employer, rather than opt in (see Banks and Tanner, 2002, for more details). Finally, the proportion of households with life insurance in their portfolio is a lot 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 United States, 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

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

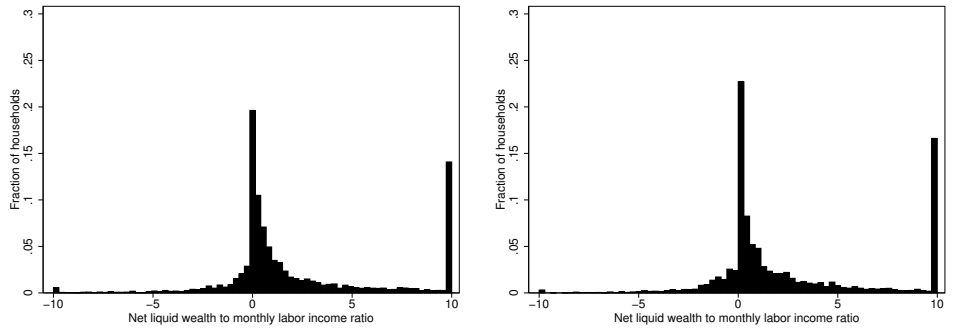

(a) United States

(b) Canada

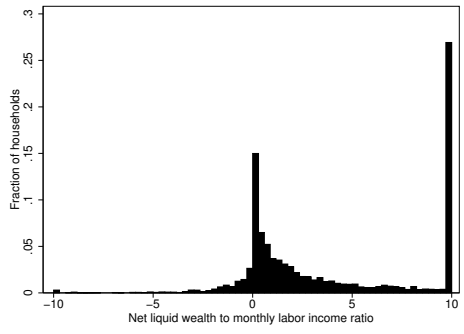

(c) Australia

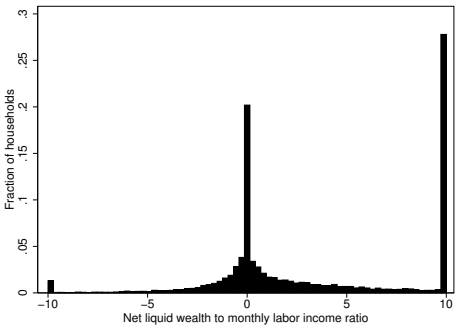

(d) United Kingdom

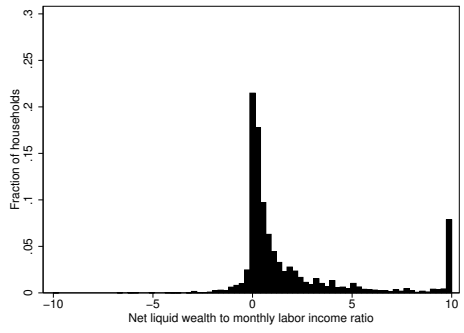

(e) Germany

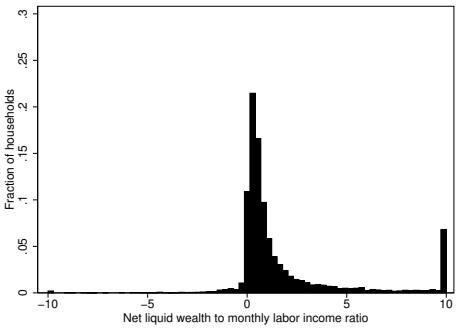

(f) France

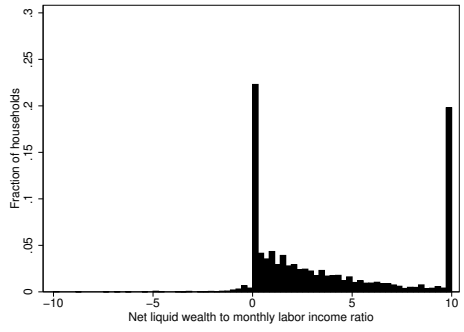

(g) Italy

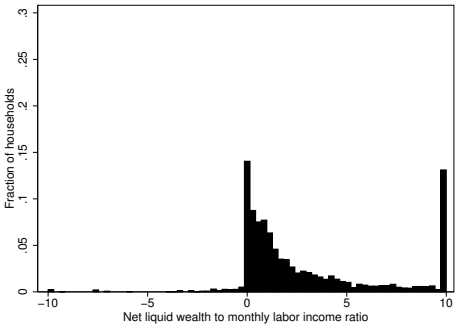

(h) Spain

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

detail. Lastly, we exploit the longitudinal dimension of the 2007-2009 waves of the SCF to study the persistence of HtM status over the life cycle.

5.1 The share of HtM households

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. Recall that our benchmark definition sets the pay frequency to two weeks, and uses equations (4) to (8) with the credit limit $\underline{m}_{it}$ set to one month of income.

Our estimates indicate that, on average, 31% of U.S. households are HtM over this period. Of these, roughly 1/3 are poor HtM and 2/3 are wealthy HtM. This is the first main result of our paper: the vast majority of hand-to-mouth households are not poor, but rather own illiquid assets.[14]

Looking at changes over time across the two decades our data cover, the fraction of HtM households remains fairly stable and the split between poor and wealthy does not change significantly. Panel (b) plots the share of W-HtM households who each own housing, retirement wealth, or both. About a half of W-HtM have both, about a third have positive housing and no retirement wealth, and a sixth have positive retirement wealth and no housing.

The first line of Table 3 also reports the share of U.S. households that are HtM when using net worth as an index of wealth. We find that less than 14% of households are HtM in terms of net worth, and thus looking at the wealth distribution through the eyes of net worth misses over half of the HtM households in the United States.

5.1.1 Robustness

Figure 4 and Table 3 summarize our sensitivity analysis.

Panel (a) of Figure 4 provides an estimate of HtM behavior using a combination of sequential questions in the SCF aimed at assessing whether "over the past year, [household] spending exceeded, or was about the same as, income, and such expenditures did included purchases of a home or automobile or spending for any investments."[15] Based

[14] A small number of HtM households (0.5 percent of the population over the whole sample period) have negative illiquid wealth because they have negative equity on their house. Currently, they are included among the P-HtM.

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

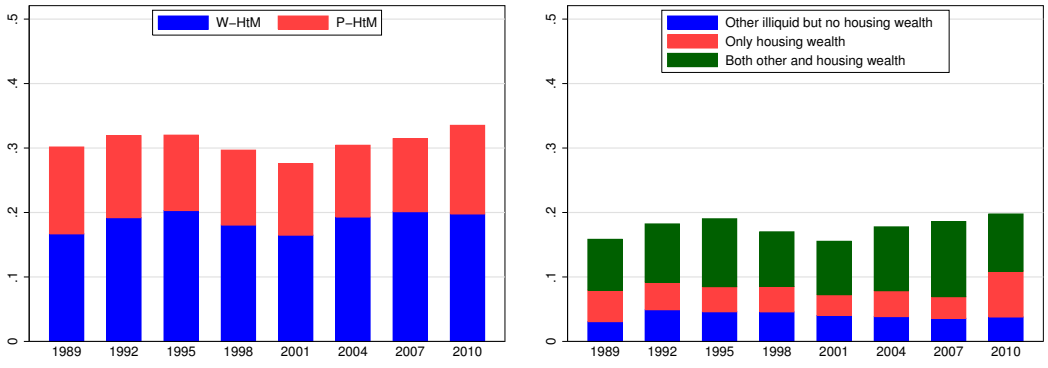

(a) Share of total, wealthy, and poor HtM (b) W-HtM by portfolio composition

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

on this definition, the share of HtM households is around 40-45 percent. W-HtM households account for 3/4 of the total, and fluctuations in this measure over time follow very closely those in the baseline definition of Figure 4(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." Just under 25 percent of households are HTM according to this definition.

It is very reassuring that the baseline count of HtM households sits in between the counts based on these two direct questions. Our baseline calculations refer to the current HtM status for the households. The first set of direct questions asks about the past year, so if there were periods when the household spent more than its income, even if now it is no longer HtM, it would answer that question positively. Conversely, the second set of direct questions asks about the usual HtM status, and therefore if the current HtM status is perceived as transitory, the household would answer negatively.

In panel (b) of Figure 4, 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 point, with all the drop being accounted for by a lower number of W-HtM households. Panel (c) 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 W-HtM account for most of the difference with the baseline. Symmetrically, the sixth 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.

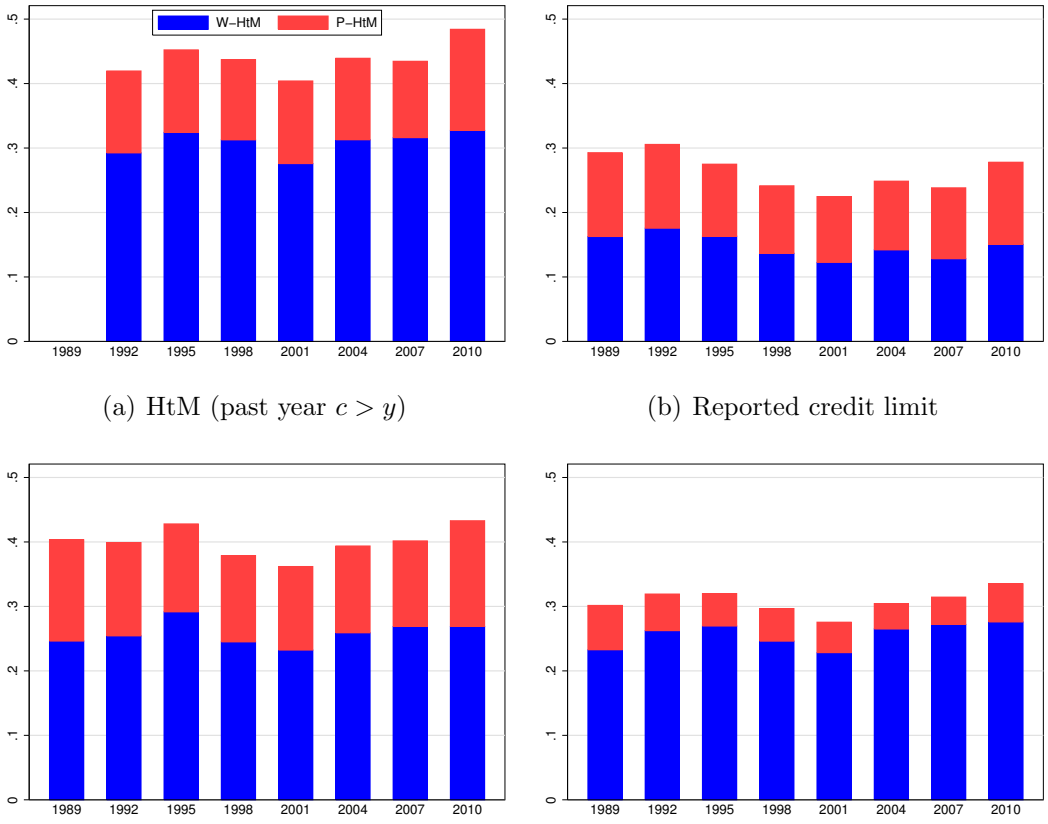

(a) HtM (past year $c > y$)

(b) Reported credit limit

(c) Pay-period of 1 month

(d) Vehicles in illiquid wealth

Figure 4: Time series of fraction of HtM households in the U.S., alternate definitions.

Panel (d) in Figure 3 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 shows that using a higher illiquid wealth threshold in the definition of W-HtM ($1,000 instead of $1) has no impact on our findings. Including business equity, or directly held stocks and bonds, or other valuables (artwork, antiques, jewels, etc.) among illiquid assets has small effects relative to the baseline.

Table 3 also contains other sensitivity analyses. Changing the definition of liquid debt by including used up HELOCs –while simultaneously increasing the credit limit by the total available line of credit– increases the fraction of HtM households by 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.[16] As expected, under this income measure, the total fraction of HtM households declines, but quantitatively this effect is very small.

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 percent.

To summarize, we estimate that, depending on the exact definition, between 20 and 40 percent of U.S. households are HtM, but most importantly around 2/3 of them are wealthy HtM, i.e. they own positive illiquid wealth in spite of their low, or even negative, holdings of liquid wealth. In contrast, estimates of HtM status in terms of net worth vary between 6 and 17 percent, and therefore miss more than half of the HtM households in the United States. Overall, these findings are consistent with related, but different, measurements by Lusardi, Schneider and Tufano (2012), who document that nearly one half of U.S. households would probably be unable to come up with $2,000 in 30 days.

5.2 The demographics of HtM groups

We now turn to the demographic characteristics of the three groups of HtM households. For this analysis, we pool the 1989 to 2010 waves of the SCF in order to maximize sample size, and we use the baseline definition from Figure 3.

Figure 5 plots the share of the population that is W-HtM and P-HtM by age.[17] Not surprisingly, 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 lifecycle. The share in the residual group of N-HtM individuals

[16]The 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 from gross income federal income taxes. We assumed each household files their actual marital status and claims all their children as dependent. As an upper bound, we have also computed the case where they all file as single without dependents.

[17]These 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 somewhat noisy.

	P-HtM	W-HtM	N-HtM	HtM	HtM-NW
Baseline	0.121	0.188	0.691	0.309	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	—
Reported credit limit	0.115	0.147	0.738	0.262	0.127
1 year income credit limit	0.104	0.120	0.776	0.224	0.109
Weekly pay period	0.106	0.147	0.748	0.252	0.118
Monthly pay period	0.142	0.258	0.600	0.400	0.164
Vehicles as illiquid assets	0.052	0.257	0.691	0.309	0.061
Higher illiquid wealth cutoff	0.121	0.188	0.691	0.309	0.137
Businesses as illiquid assets	0.114	0.189	0.697	0.303	0.129
Direct as illiquid assets	0.120	0.212	0.668	0.332	0.137
Other valuables as illiquid assets	0.117	0.192	0.691	0.309	0.132
Includes usually paid off cc debt	0.121	0.192	0.688	0.312	0.137
HELOCs as liquid debt	0.120	0.178	0.702	0.298	0.135
Usual income	0.120	0.194	0.687	0.313	0.136
Disposable income - Reported	0.121	0.184	0.695	0.305	0.137
Disposable income - Single	0.120	0.183	0.698	0.302	0.135
Comm. cons. - beg. of period	0.101	0.162	0.737	0.263	0.115
Comm. cons. - end of period	0.149	0.269	0.581	0.419	0.173

Table 3: Robustness results for fraction HtM in each category in the SCF pooled 1989-2010. Higher illiquid wealth cutoff requires households to have above $1,000 in illiquid assets to be considered W-HtM. 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.

increases steadily from 50 percent at age 22 to 80 percent in retirement.

The first three panels of Figure 6 report some demographic characteristics of the three HtM groups by age. N-HtM households have on average one year of education more than the W-HtM who, in turn, have one more year of education than the P-HtM. In terms of marital status, H-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

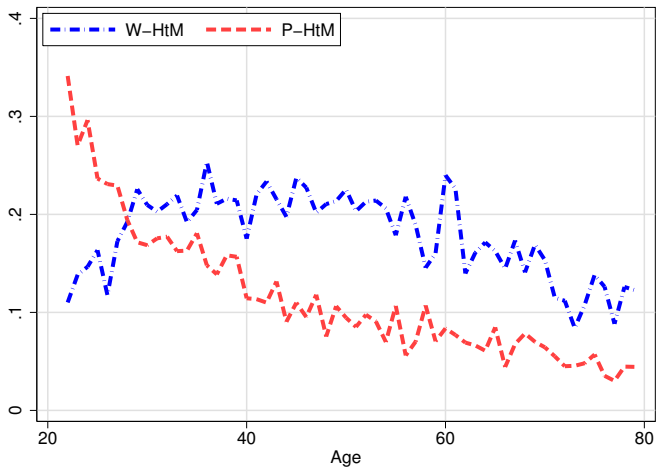

Figure 5: Age profile of fraction of HtM households in the U.S., pooled 1989-2010.

households.

Figure 6(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 who earn on average $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.

5.3 The portfolio composition of HtM groups

Figure 7 digs deeper into the composition of the balance sheets of the three groups of HtM households. Panel (a) shows that median liquid wealth holdings are zero for the P-HtM and, perhaps unexpectedly, slightly negative for the W-HtM. N-HtM households have substantial holdings of liquid wealth, peaking at around $15,000 before retirement. Panel (b) reveals that the 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 7 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

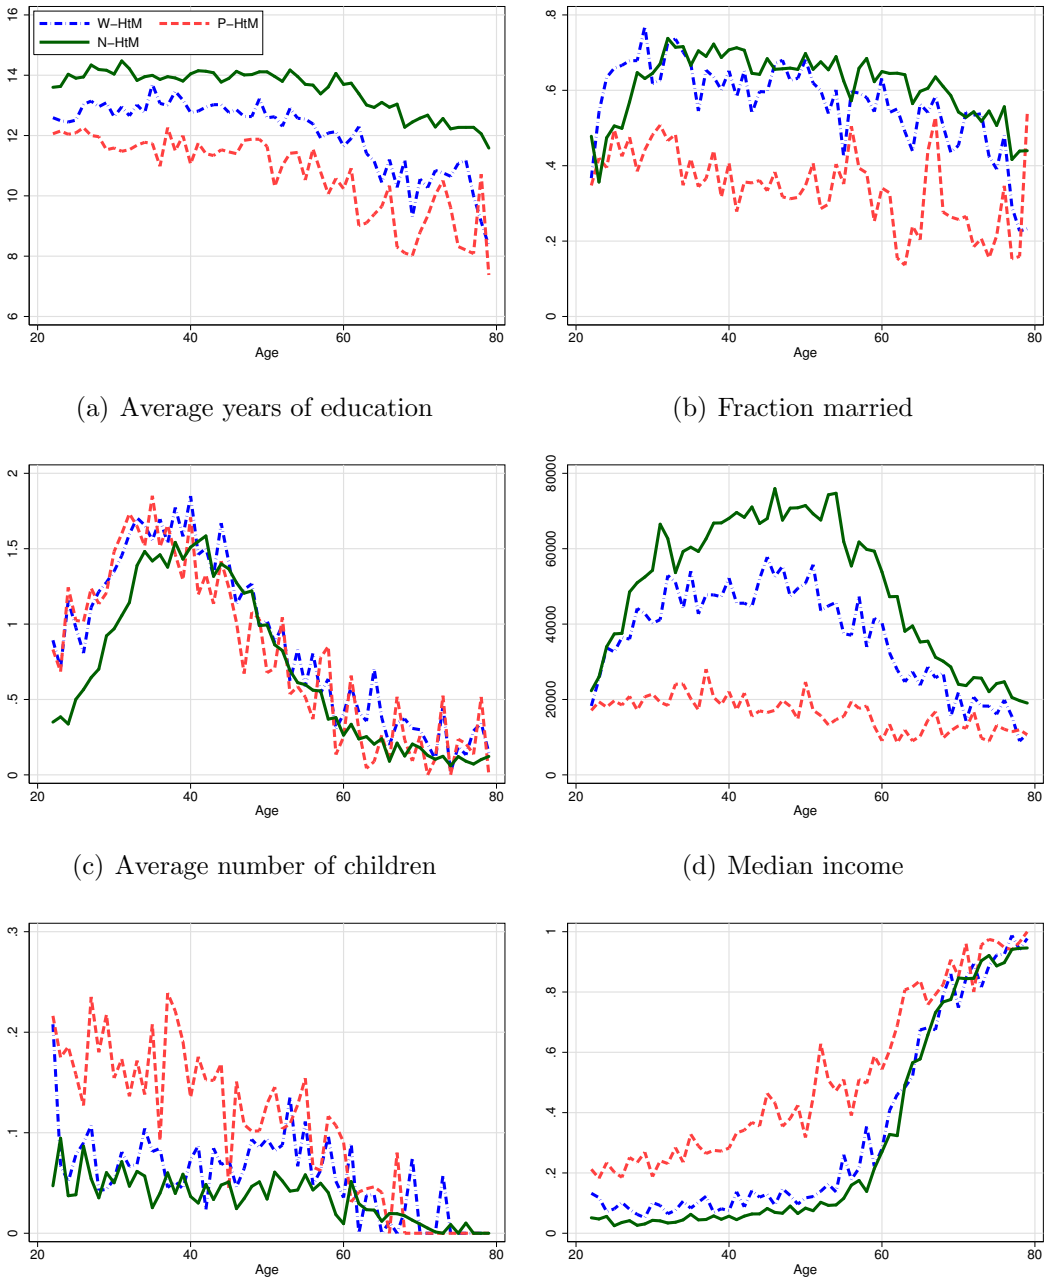

(a) Average years of education

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

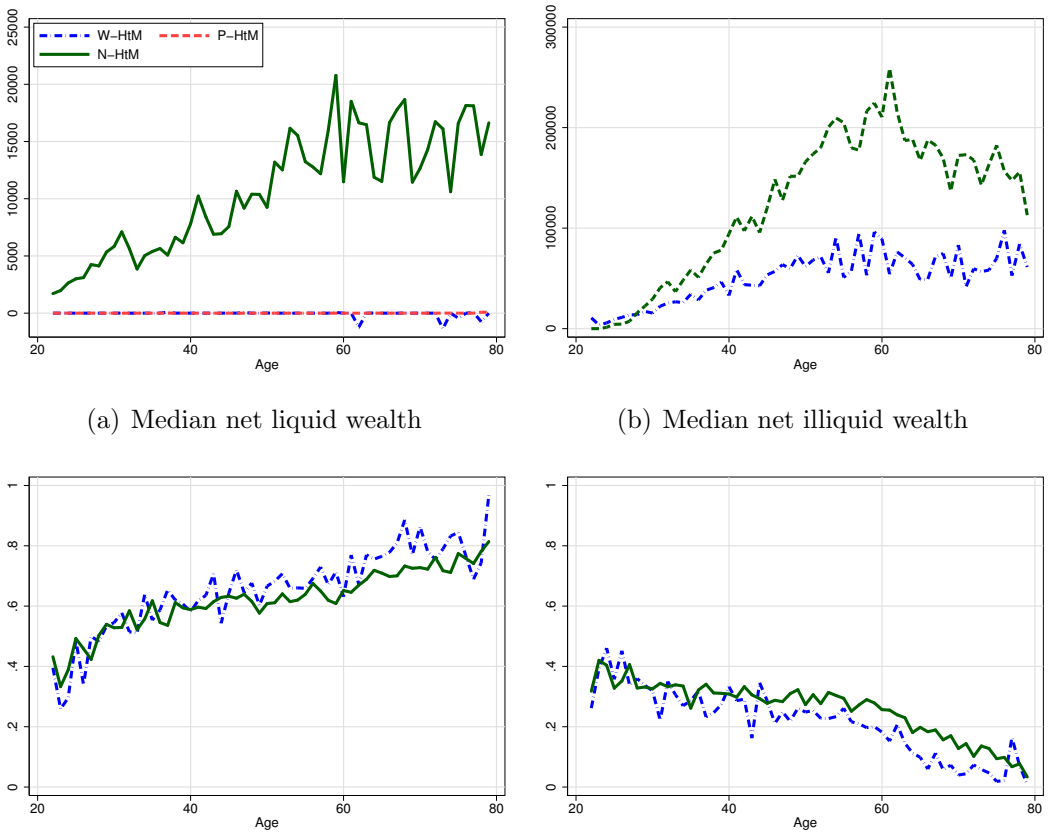

(a) Median net liquid wealth

(b) Median net illiquid wealth

(c) Mean frac. of ill. wealth in housing

(d) Mean frac. of ill. wealth in ret. account

Figure 7: 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.

on top of each other, indicating that the portfolio allocation of these two groups is nearly identical. The key difference is that N-HtM are income and wealth richer, and that they hold a lot more liquid wealth relative to their income.

5.4 Persistence of HTM status

How persistent is the HtM status of a household? The 2007-2009 panel of the SCF allows us to address this question. Table 4 reports the 2-year transition matrix across the three HtM statuses for U.S. households. P-HtM households have a 52% chance of still being P-HtM two years later, but also a 37% chance of becoming N-HtM households within the same period. The W-HtM status appears to be the most transitory: the probability of remaining W-HtM two years later is approximately 17%. To put it

07 → 09	P	W	N
P	0.520	0.114	0.366
W	0.058	0.169	0.773
N	0.028	0.052	0.92

Table 4: Transition matrix for the 2007-2009 panel of the SCF. Fraction of households classified as the row status in 2007 and the column status in 2009.

differently, the expected length of W-HtM status is around 29 months. Finally, the N-HtM state is the least transient of the three, and it is almost an absorbing state.

6 Cross-country evidence

The previous section showed that around 30% of households in the United States are HtM, with 1/3 of which are P-HtM and 2/3 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 –Canada, Australia, and the U.K.– and the four most populated European countries, Germany, France, Italy, and Spain. While data is available for more than one point in time for most of these countries –including panel data for Australia and the U.K.– 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 8(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 European 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.

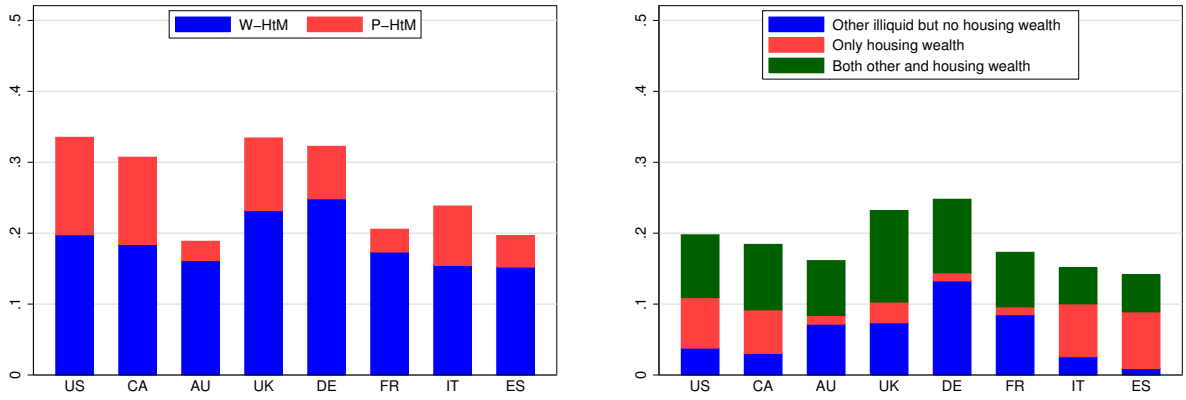

(a) Share of total, wealthy, and poor HtM

(b) W-HtM by portfolio composition

Figure 8: Fraction of HtM households across countries

Even for the Euro area countries, the fraction of W-HtM among the HtM households exceeds 2/3. For all eight countries, Figure 8(a) shows there are 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 8(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.

What explains the fact that Australia and the Euro area countries, on the other, have a smaller fraction of HtM households than in the U.S.? Figure 8(a), shows that this difference is largely accounted for by differences in the fraction of P-HtM households. Table 5 reveals that for Australia this discrepancy can be traced to the very high share of the population that owns private retirement wealth. Over 86% of N-HtM households in Australia have assets in a retirement account compared with 62% in the U.S. As explained in Section 4, the high ownership rate 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, closing more than half of the gap with the United States.

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	-0.631	-0.486	-0.399	-1.332	-0.012	-0.092	-0.002	-0.228
Frac. neg. liquid wealth	0.167	0.180	0.122	0.289	0.051	0.072	0.007	0.102
Frac. neg. illiquid wealth	0.140	0.114	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	0.003	0.000	0.021	0.000	0.059	0.122	0.000	0.011
Mean liquid wealth / income	-1.172	-0.893	-0.783	-1.902	-0.019	-0.206	-0.067	-0.266
Frac. neg. liquid wealth	0.362	0.369	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	0.839	0.733	0.377	0.312	0.000	0.389	1.000	1.000
Mean housing / illiquid wealth	0.610	0.587	0.406	0.398	0.369	0.471	0.968	0.985
Frac. pos. housing equity	0.782	0.800	0.547	0.667	0.464	0.505	0.834	0.933
Median retire / illiquid wealth	0.034	0.253	0.563	0.659	0.000	0.066	0.044	0.023
Mean retire / illiquid wealth	0.316	0.403	0.583	0.583	0.093	0.175	0.390	0.036
Frac. pos. retirement account	0.533	0.655	0.924	0.862	0.271	0.020	0.041	0.021
N-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Median liquid wealth / income	1.737	2.324	3.404	4.690	1.133	0.878	3.701	2.129
Mean liquid wealth / income	15.144	12.440	34.667	38.504	6.144	4.730	8.918	6.977
Frac. neg. liquid wealth	0.102	0.086	0.061	0.149	0.081	0.042	0.016	0.017
Frac. neg. illiquid wealth	0.025	0.035	0.004	0.004	0.006	0.007	0.001	0.003
Median housing / illiquid wealth	0.651	0.774	0.709	0.432	0.534	0.839	1.000	1.000
Mean housing / illiquid wealth	0.593	0.647	0.598	0.459	0.446	0.604	0.965	0.948
Frac. pos. housing equity	0.714	0.722	0.771	0.785	0.533	0.654	0.770	0.877
Median retire / illiquid wealth	0.237	0.204	0.273	0.545	0.000	0.029	0.036	0.021
Mean retire / illiquid wealth	0.329	0.343	0.387	0.524	0.052	0.081	0.174	0.083
Frac. pos. retirement account	0.620	0.570	0.878	0.849	0.259	0.045	0.107	0.042

Table 5: 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.

In the Euro area countries, the lower fraction of P-HtM implies that households hold more liquid wealth relative to their income. Table 5 shows that this difference can be in part attributed to differences in liquid debt, a fact that we also highlighted in the discussion of Figure 2. 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. Presumably, lower access to unsecured credit in Europe means that there are more incentives for households to hold 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.

6.1 Robustness

Table 6 contains an extensive sensitivity analysis on our definitions of P-HtM and W-HtM households that parallels that displayed in Table 3.

Questions on whether household spending exceeded income in the past year are present in all surveys. Like in 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 consistent 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.[18]

Shortening the pay-period to a week and extending it to a month, from the bi-weekly baseline, has a very 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, adding other non-financial assets (valuables, collectibles, jewels, etc.) in the definition of illiquid wealth definition shifts 12 and 5 percent of households from poor to wealthy HtM, respectively. Including HELOCs among liquid

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

P-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Baseline	0.138	0.123	0.027	0.103	0.074	0.032	0.083	0.044
In past year, c > y	0.157	0.181	0.020	0.093	0.090	—	0.156	0.091
1 year income credit limit	0.118	0.102	0.024	0.078	0.070	0.030	0.083	0.040
Weekly pay period	0.117	0.105	0.022	0.098	0.058	0.021	0.080	0.036
Monthly pay period	0.164	0.150	0.033	0.111	0.086	0.048	0.091	0.061
Vehicles as illiquid assets	0.059	0.083	0.012	0.065	0.052	0.002	0.028	0.024
Businesses as illiquid assets	0.131	0.117	0.027	0.102	0.071	0.031	0.076	0.043
Direct as illiquid assets	0.137	0.121	0.027	0.102	0.074	0.032	0.083	0.045
Other valuables as illiquid assets	0.133	0.008	0.025	0.099	0.071	—	0.034	0.044
HELOCs as liquid debt	0.134	0.127	—	0.103	0.074	0.032	0.083	0.044
Disposable income	0.136	—	—	0.103	—	—	0.080	—
Comm. cons. - beg. of period	0.114	—	—	—	0.066	0.025	0.076	0.036
Comm. cons. - end of period	0.174	—	—	—	0.092	0.050	0.090	0.064
W-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Baseline	0.198	0.184	0.161	0.232	0.248	0.173	0.155	0.152
In past year, c > y	0.327	0.409	0.189	0.251	0.392	—	0.474	0.596
1 year income credit limit	0.131	0.117	0.115	0.135	0.229	0.157	0.147	0.141
Weekly pay period	0.151	0.146	0.115	0.211	0.161	0.087	0.142	0.119
Monthly pay period	0.269	0.250	0.229	0.276	0.370	0.354	0.188	0.220
Vehicles as illiquid assets	0.276	0.224	0.177	0.269	0.270	0.204	0.211	0.173
Businesses as illiquid assets	0.201	0.190	0.163	0.232	0.251	0.173	0.158	0.154
Direct as illiquid assets	0.214	0.219	0.193	0.246	0.303	0.198	0.165	0.162
Other valuables as illiquid assets	0.203	0.300	0.164	0.235	0.252	—	0.204	0.153
HELOCs as liquid debt	0.188	0.113	—	0.154	0.238	0.166	0.147	0.140
Disposable income	0.195	—	—	0.237	—	—	0.149	—
Comm. cons. - beg. of period	0.169	—	—	—	0.219	0.127	0.148	0.138
Comm. cons. - end of period	0.280	—	—	—	0.344	0.336	0.173	0.199

Table 6: Robustness results for fraction P-HtM and W-HtM in each category. Higher illiquid wealth cutoff requires households to have above 1,000 local currency units in illiquid assets to be considered W-HtM. 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. 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.

debt has no effect, except in Canada, where the share of HtM increases by 8 percentage points.

Our baseline measure of income is 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.[19] In these three countries, the effect is minor.

6.2 Demographic characteristics and persistence of HtM behavior across countries

Age profiles of the fraction of poor and wealthy HtM households in each country are shown in Figure 9. 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 is flat. There are some marked differences in the age profiles of the W-HtM which are 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 is decreasing 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 why in Spain, in spite of the large share of homeowners among the W-HtM, their share falls 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 35-45 and 45-55 year olds faced over the past two decades.

Figure 10, which shows age-income profiles for each country by HtM status, 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 three age paths are all very similar.

[19]For 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.

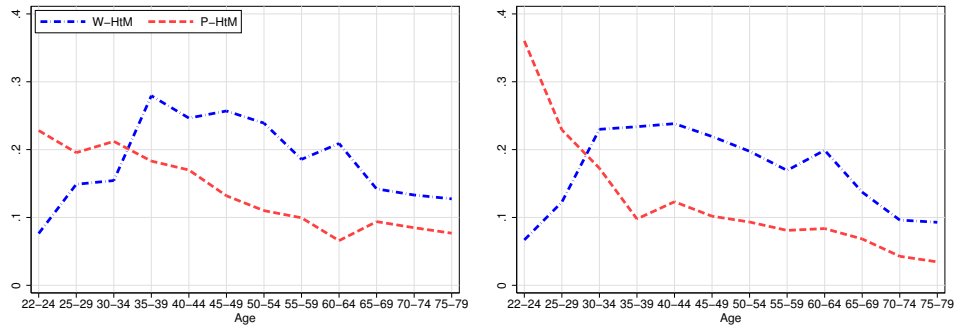

(a) United States

(b) Canada

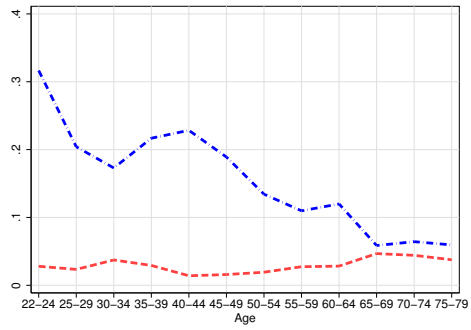

(c) Australia

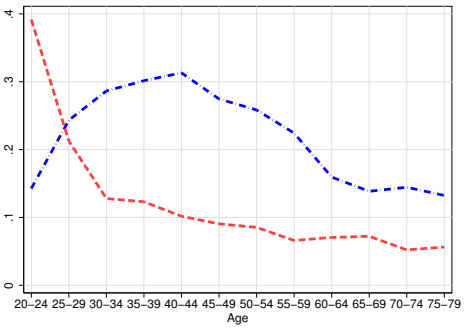

(d) United Kingdom

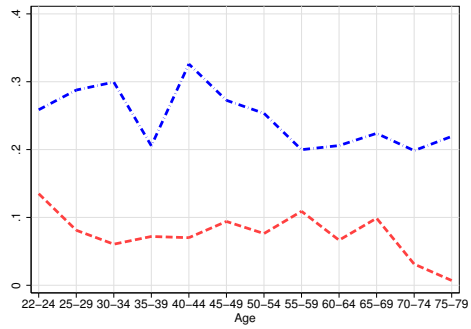

(e) Germany

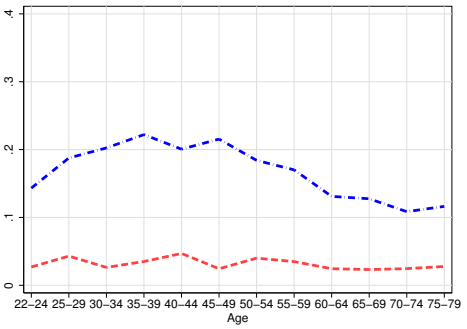

(f) France

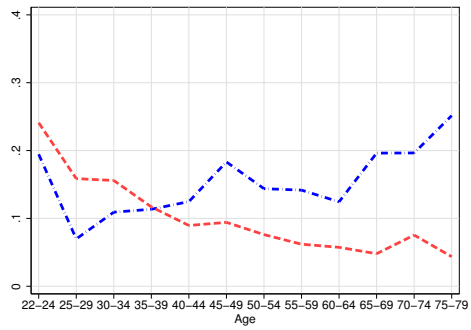

(g) Italy

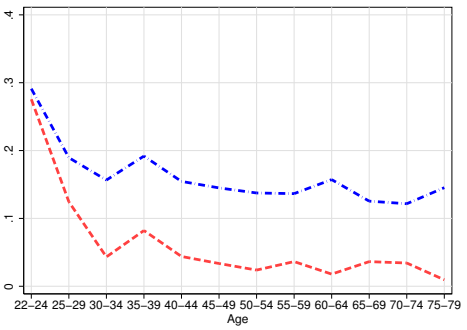

(h) Spain

Figure 9: Age profile of fraction of HtM households by country.

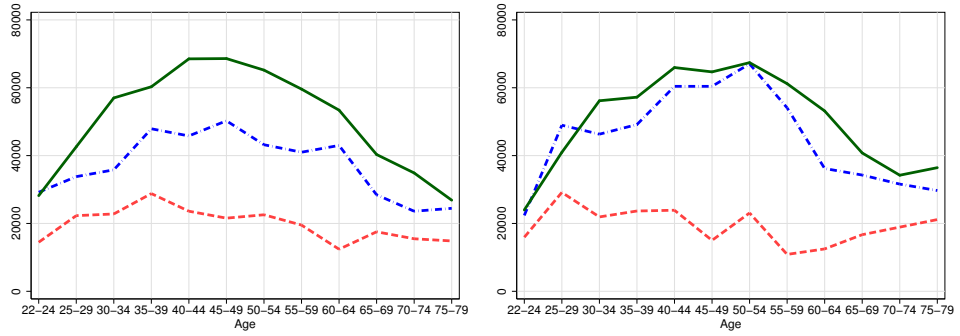

(a) United States

(b) Canada

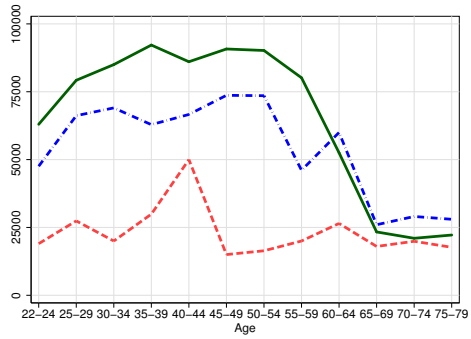

(c) Australia

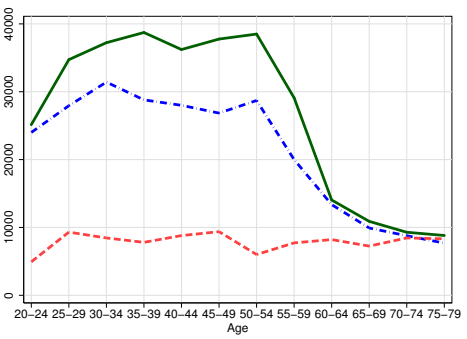

(d) United Kingdom

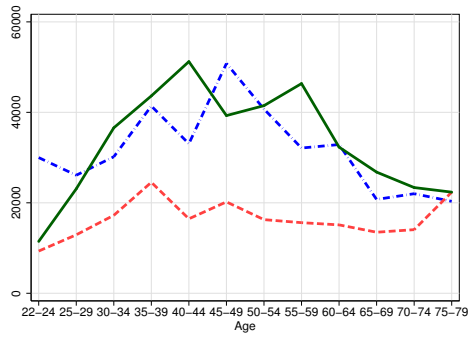

(e) Germany

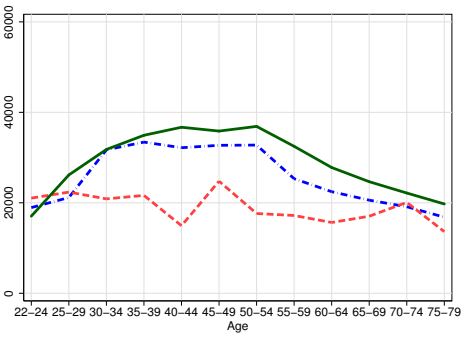

(f) France

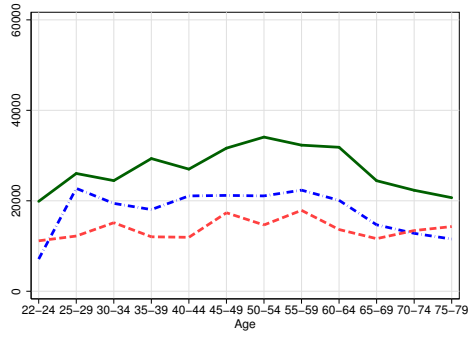

(g) Italy

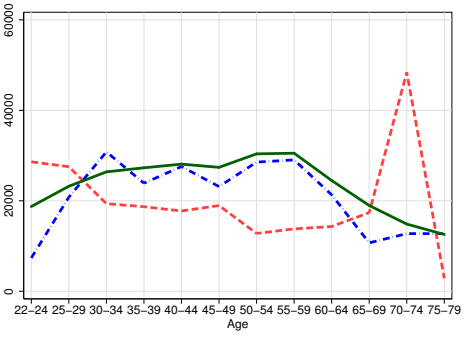

(h) Spain

Figure 10: Age profile of median income by HtM status by country.

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

In the previous sections we documented a sizable presence of wealthy HtM households across a number of countries. In this section we show evidence that, as predicted by the theory presented in Section 2, these households have a large MPC with respect to transitory income shocks. From the Panel Study of Income Dynamics (PSID), we 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 and sample selection. 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 all this 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 starting in 1997. In 1999, the survey augmented the consumption information available to researchers, which now covers over 70% 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.[20] This addition makes the PSID the only large scale representative U.S. panel to include income, consumption, and assets data.

From the PSID Core Sample, we drop households with missing information on race, education, or state of residence, and those whose income grows more than 500%, falls by more than 80%, 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

[20] Until 1999, the Wealth Files supplemented the annual survey every five years. Starting from 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.

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 (2013). 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.[21]. 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. In our PSID sample, 25 percent of households are W-HtM, in line with the SCF estimates. The share of the P-HtM is 21 percent, and hence is somewhat larger than 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 components and transitory I.I.D. component. Hence, income growth is given by

$$\Delta y_{it} = \eta_{it} + \Delta \varepsilon_{it}, \tag{9}$$

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

[21]The main discrepancies with the SCF definitions are that (i) we do not attempt a cash imputation and (ii) CDs and saving bonds are in liquid, instead of illiquid, wealth.

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

$$\widehat{MPC}_t = \frac{cov(\Delta c_{it}, \Delta y_{i,t+1})}{cov(\Delta y_{i,t}, \Delta y_{i,t+1})}. \tag{10}$$

Under the assumption that consumption growth in period t is uncorrelated with the transitory and permanent shocks at $t+1$, i.e.:

$$cov\left(\Delta c_{it}, \eta_{i,t+1}\right) = cov\left(\Delta c_{it}, \varepsilon_{i,t+1}\right) = 0, \tag{11}$$

this estimator uncovers the marginal propensity to consume out of a transitory shock, i.e.:

$$MPC_t = \frac{cov\left(\Delta c_{it}, \varepsilon_{it}\right)}{var(\varepsilon_{it})} \tag{12}$$

Condition (11) means that the household has no foresight, or no advanced information, about future shocks. Under this condition, (10) is a consistent estimator of the pass-through coefficient of transitory shocks into consumption in (12). The estimator is implemented by an IV regression of Δc_{it} on Δy_{it}, instrumented by $\Delta y_{i,t+1}$. Kaplan and Violante (2010) show that the presence of tight borrowing constraints does not bias the estimate of the transmission coefficient for transitory shocks. This is 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 non HtM households.

Results. Table 7 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 dissipates entirely 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 non 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 many W-HtM households with high

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

Table 7: 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 moth instead of two weeks.

MPCs. Based on this latter household classification, one would conclude that there is no evidence of a differential response of consumption to income shocks based on HtM status. A classification based on liquid and illiquid wealth, instead, finds economically significant differences.

The remaining rows in Table 7 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.

7.1 Additional evidence

Although none of the existing empirical investigations of the consumption reaction to income shocks has explicitly tried to separate wealthy and poor HtM, as we did here, some of them offer, indirectly, evidence on wealthy HtM behavior.

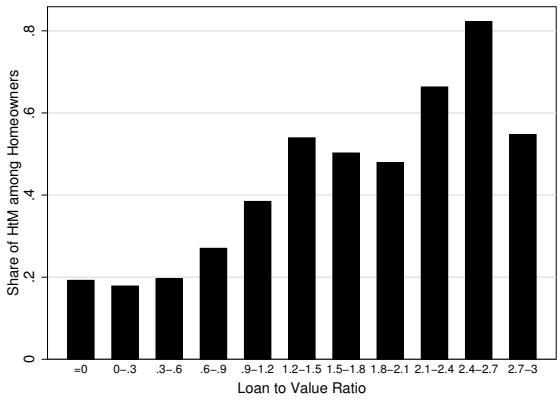

Figure 11: Share of W-HtM households among homeowners by loan-to-value ratio. U.S. SCF 1989-2010.

In some of these studies, the sample is split between homeowners with large mortgages and those with small leverage ratios. With respect to hand-to-mouth behavior, it is intuitive to conjecture that the former group is more likely to be wealthy HtM because the regular mortgage payments absorb a significant fraction of disposable income. Figure 11 shows that, in the U.S., this is indeed the case: among homeowners, the fraction of W-HtM, according to our definition, grows steadily with the loan to value ratio.

A number of recent papers find evidence of consumption responses to income shocks whose strength increases in the degree of indebtedness of the homeowner. Misra and Surico (2013) expand on the research of Johnson, Parker and Souleles (2006) and Parker, Souleles, Johnson and McLelland (2014) on the U.S. fiscal stimulus payment episodes of 2001 and 2008. They conclude that, for both stimulus programmes, 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. Their key finding is that homeowners with high mortgage debt 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 highly-indebted households with illiquid assets are more sensitive to income fluctuations, and attributes this result largely to borrowing and liquidity constraints.

A different sort of evidence on wealthy HtM behavior comes from the Danish fiscal

stimulus of 2009. At the end of 2009, the Danish government allowed households to withdraw from their own pension funds –accounts that would otherwise be illiquid until age 65– and transform them into liquid resources available for consumption. Kreiner, Lassen, and Pedersen (2012) measured household consumption expenditures before and after the policy as well as their marginal interest rate on loans and credit cards. They conclude that households with a high borrowing cost consumed the largest share of their withdrawals.

8 Implications of the wealthy hand-to-mouth 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 in the cast of hand-to-mouth characters, W-HtM households require their own unique place. Viewing the household portfolio data through the lens of a model with only two types of households -the non HtM and the HtM– leads to a distorted view of the effects of fiscal stimulus policies on household consumption, compared with the effects one obtains when viewing the data though the lens of a model with all three types of households.

Our approach is to use simulated marginal propensities to consume (MPCs) from three alternative structural life-cycle models, in conjunction with the empirical estimates of the shares of HtM households from Sections 5 and 6. We use calibrated versions of each model to predict average MPCs for households with each type of HtM status. We then use the estimated fractions of the population in each of these groups to compare the predicted overall average MPCs for each country, under each of the three models.

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 which 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. We 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, Souleles, Johnson, and McLelland (2013), respectively. Here, it suffices to say that the version of the model we use here does not allow borrowing and has a transaction cost of $1,000.

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. When using this model to compute aggregate MPCs, we impute the model implied MPCs for N-HtM households to the W-HtM households in the data. This is because, when the data is viewed through the lens of this model, net worth is the relevant measure of wealth, and by this measure the wealthy and non HtM are equivalent.

The third model, which we label SPS, is a spender-saver model in the spirit of Campbell and Mankiw (1989) and, more recently, Eggertson and Krugman (2012), and Justiniano, Primiceri and Tambalotti (2013). In the SPS model, one group of 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 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 SPS model, the W-HtM and the P-HtM households are considered to be the spenders, while the N-HtM households are considered to be the savers. When we compute the aggregate MPC for the SPS economy we impute the MPC for the N-HtM households in the SIM-1 model to the N-HtM households in the data, and we impute an MPC of one to the W-HtM and P-HtM households in the data.

These three alternative economies can thus be thought of as follows. SIM-2 is an economy in which all 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, and the only households with a high MPC are the P-HtM. SPS is a "liquid wealth economy" where both the W-HtM and the P-HtM are treated identically as HtM households, who have an MPC that is always equal to 1.

For each of the three models, we first simulate a cohort of households. For each simulated household we compute their consumption response to a one-time unexpected

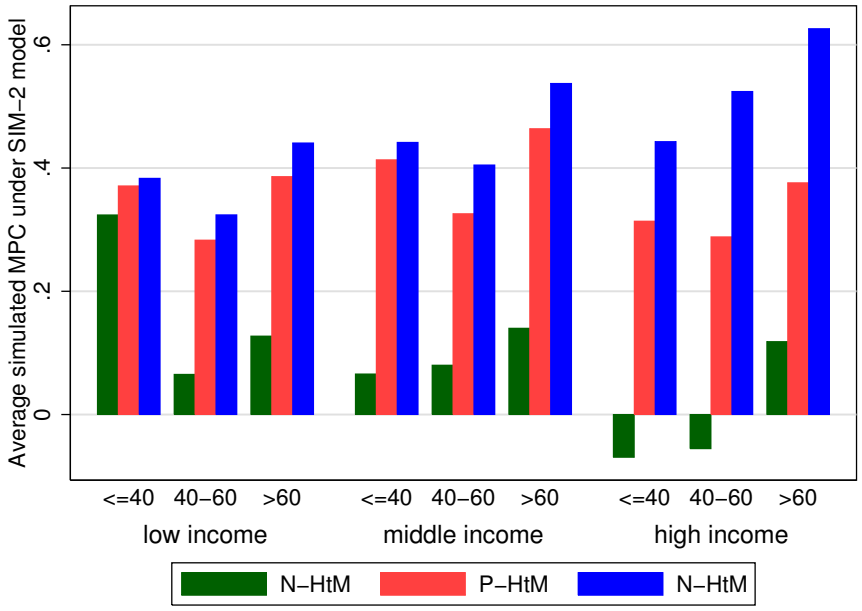

Figure 12: Average simulated MPC under SIM2-model by age group, income tercile and HtM status.

cash windfall or cash loss of different amounts ($50, $500, $2,000). We then divide the simulated population into 27 bins, based on three income terciles, three age groups (22-40, 40-60, 60+) and the three HtM groups. For each of these bins we compute the average consumption response from the model.

Figure 12 shows the average MPCs out of a $500 windfall in the SIM-2 economy for each of the 27 bins. The MPCs are close to zero for N-HtM households, except for households who are both income poor and very young. 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. The reason is that, as explained in Section 2,

	Model		
	SIM-2	SIM-1	SPS
\$500 windfall	0.18	0.04	0.35
Size asymmetry			
\$50 windfall	0.29	0.05	0.35
\$2000 windfall	0.05	0.03	0.35
Sign asymmetry			
\$500 loss	0.42	0.14	0.36
Income targeting			
\$500 windfall, bottom tercile	0.26	0.07	0.50
\$500 windfall, top tercile	0.20	0.03	0.34

Table 8: Estimated aggregate consumption responses for US using demographics and HtM composition from 2010 SCF.

since the W-HtM have higher lifetime incomes than the P-HtM households, they have higher desired consumption and hence spend more out of a moderately sized payment.[22]

The simulated MPCs in the other two models are similar to those in Figure 7.1. In the SIM-1 model, the MPCs for P-HtM households are almost identical to those for P-HtM households in the SIM-2 model, and the MPCs for N-HtM households –and hence for the W-HtM as well– are only slightly larger than those for N-HtM households in the SIM-2 model (and are never negative). For the SPS model, the MPCs for the N-HtM households are the same as in the SIM-1 model and are equal to one for both groups of HtM households.

Implied cross-country variation in MPCs We compute an aggregate MPC for each of the eight countries in our data. To do this, we 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. The implied aggregate quarterly consumption responses for the U.S. using the SCF data from 2010 are shown in Table 8. For a \$500 windfall, the aggregate MPC of the SIM-2 economy is 0.18. This is substantially larger than for the MPC of the SIM-1 economy (0.04)

[22]In the model there are no simulated P-HtM households who are in the top tercile of the income distribution and are over 60 years old. This is almost true in the data too. For example, in the U.S. in 2010, less than 0.1% of households fall in this category. For these households we impute the MPC as the average between the the MPC for P-HtM aged 40-60 in the top income tercile, and the MPC for P-HtM aged 60+ in the middle income tercile.

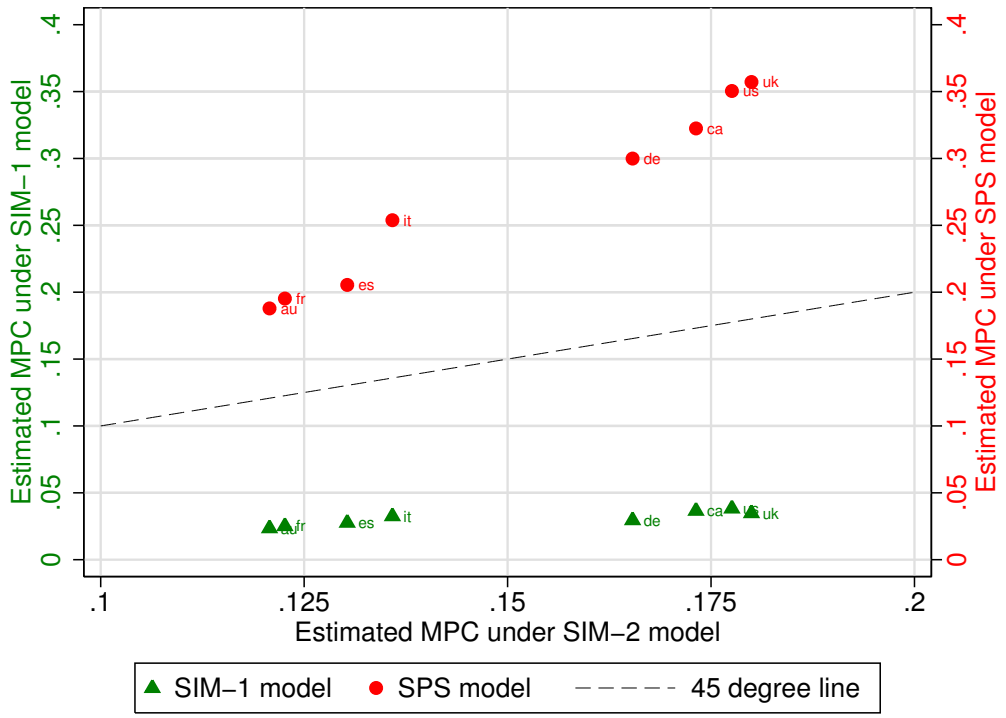

Figure 13: Estimated aggregate consumption response by country under SIM-2 model (x axis), SIM-1model (triangles, left y axis) and SPS model (circles, right y axis).

because the SIM-1 economy treats the W-HtM households as N-HtM and so misses a large part of the population who have a high MPC. The MPC is highest for the SPS economy (0.35), because the SPS model implicitly assumes that the P-HtM and W-HtM households all spend the entire $500, while we saw in Figure 12 that this is an exaggeration: in the SIM-2 economy, they only spend 30%-40% of the payment on average in the quarter they receive it.

The differences in the level of the MPCs predicted by the three models are driven by the differences in the fraction of households to whom the models assign high MPCs. Since there is substantial variation in the fraction of HtM households across countries, this implies that the three models also yield very different predictions for the degree of cross-country dispersion in the MPC to a $500 fiscal stimulus check. To illustrate these differences, Figure 13 plots the estimated aggregate MPC under the SIM-2 model against the corresponding MPC under the SIM-1 model (triangles) and the SPS model (circles).

The figures show striking differences in the amount cross-country dispersion in the

MPC across the three models. There is much less dispersion under the SIM-1 model than under the SIM-2 model. Cross-country dispersion arises from differences in the fraction of HtM households. 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 under the SPS model than under the SIM-2 model. This is because, by assigning an MPC of 1 to all the W-HtM households, compared with around 0.3-0.4 in the SIM-2 model, the SPS model exhibits too much of a difference between the MPC of HtM and N-HtM households than does the SIM-2 model. Cross-country heterogeneity in the fraction of HtM households is thus more strongly reflected in the MPC under the SPS model than under the SIM-2 model.

Other implications We now highlight three additional areas where the three classes of models yield different predictions because of the different types of consumption behavior that they assign to the P-HtM, W-HtM, and N-HtM.

First, Table 8 shows that the degree of size asymmetry in the aggregate MPC differs remarkably across the three models. Under 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 small windfalls, households in SIM-2 face an inter-temporal tradeoff that is governed by the (low) return on the liquid asset and thus have a large incentive to consume. However, for large enough windfalls, many HtM households may find it optimal to pay the transaction cost and make a deposit into the illiquid asset.

This size asymmetry is absent from both the SIM-1 and SPS models. In the SPS model it is absent because of the assumed rule-of-thumb behavior: the HtM households in the SPS model always consume their entire windfall, regardless of the size. In the SIM-1 model there is some decline in the MPC with the size of the payment, but it is much more modest than in the SIM-2 model because households always face the same inter-temporal trade off when making their consumption decisions. These experiments clearly illustrate why it is important to think deeply about the response of W-HtM households when considering the design of stimulus policies. Viewing the data through the lens of either the SIM-1 or SPS model would lead one to conclude that there is far more scope for stimulating consumption by increasing the size of stimulus payments, than is predicted by the SIM-2 model.

Second, the degree of sign asymmetry differs across the three models. Under the SIM-1 and SIM-2 models, the response to a \$500 negative income shock is substantially larger

than the response to a $500 positive income shock. The reasons is that even HtM households, who by definition are at a kink in their budget constraints, desire to save some part of a positive windfall if the windfall is large enough to push them off the kink. Negative income shocks, 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. Under the SPS model, the response to a positive and negative shock are essentially the same, since the MPC is driven by the HtM households who have an MPC of 1 to both positive and negative shocks.

Third, Table 8 reveals different implications of "income targeting" across the three models. A widely held view is that the aggregate consumption response to a fiscal stimulus policy, per dollar paid out, is larger when the payments are made to households with lower income, i.e. stimulus payments should be phased out for middle- and high-income households for maximum effect. This view is based on the observation that P-HtM households have very low income, as illustrated in Figure 6d. However, Figure 6d also showed that W-HtM households have much higher income than P-HtM households, yet we saw in Section 7, as well as in Figure 12, that the W-HtM are the households who are most likely to spend their stimulus payments. For this reason, the SIM-2 model generates only a very modest decline (0.26 to 0.20) between the average MPC out of a $500 windfall 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 SPS models. In the case of the SIM-1 model this occurs because the only high MPC households are the P-HtM; in the case of the SPS model this occurs because all HtM households are assumed to have the same MPC, while under the SIM-2 model we saw in Figure 12 that among W-HtM households, MPCs are increasing in income. We thus conclude that explicitly considering the consumption response of W-HtM households may lead one to prefer stimulus policies with substantially less phasing out at middle-income levels.

9 Concluding remarks

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

Theoretically, we showed that W-HtM behavior can arise 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. Because of a wedge between the interest rates at which households can borrow and the return that households earn on liquid savings, access to credit does not qualitatively change the conclusion that W-HtM typically have high propensities to consume out of small transitory income shocks.

Empirically, we found that around 30% to 40% of households in the US are HtM, and that this fraction has been relatively constant over the past two decades. We showed that the majority of HtM households, around two-thirds, are W-HtM, not P-HtM. Although the total fraction of HtM households varies somewhat across the eight countries that we have studied, from less than 20% in Australia and Spain to above 30% in the US and the UK, in all countries the W-HtM constitute the vast majority of HtM households.

Who are the W-HtM? We highlight two key findings. First, unlike P-HtM households, the W-HtM are not predominantly young households with low income. Rather they have a humped-shaped age profile that peaks in the early forties, and an income profile that mirrors strongly the income of the N-HtM. Second, the W-HtM are not simply P-HtM with very small holdings of illiquid assets. Rather they hold substantial wealth in housing and retirement accounts in the same proportions as N-HtM households.

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

W-HtM households are important because they tend to have a large consumption response to transitory income shocks, which is a key determinant of the efficacy of many types of fiscal interventions, such as the fiscal stimulus payments that were used in the last two recent recessions. To demonstrate this, in Section 7 we utilized the identification strategy of Blundell, Pistaferri and Preston (2008), to show that the consumption response to transitory income shocks is significantly larger for W-HtM and P-HtM households than for N-HtM households. But splitting households according to whether they were HtM based on their holdings of total assets yields identical consumption responses for HtM and non-HtM households.

The W-HtM thus have consumption responses that are similar to the P-HtM, yet have demographic and financial characteristics that resemble the N-HtM. This suggests that

the three types of HtM households each need their own unique place in frameworks that are to be used for understanding and forecasting the effects of fiscal policy. Macroeconomists need to move beyond one-asset models, i.e. those in the spirit of Aiyagari (1994), Huggett (1995) and Rios-Rull (1996), since these models assume W-HtM households act the same way as N-HtM households. Spender-saver models., i.e those in the spirit of Campbell and Mankiw (1989) and Eggertsson and Krugman (2012), are not a panacea, since these models assume W-HtM households act identically to P-HtM households. Moreover, spender-saver models do not allow for the fact that the W-HtM households hold illiquid wealth to affect their behavior.

In Section 8, we used an alternative framework from Kaplan and Violante (2014), which explicitly models W-HtM behavior through the presence of both liquid and illiquid assets, to illustrate two examples where misleading inferences would be obtained by using either of the two simpler models of HtM households. First we showed that the drop in the consumption response to a fiscal stimulus payment as the size of the payment increases, is much steeper in the model that allows for W-HtM behavior than in the models that do not. Second, we showed that the model that allows for W-HtM behavior implies that to maximize the aggregate consumption response to fiscal stimulus payments, the payments should feature more moderate phasing out with household income.

References

Abowd, J. M., and D. Card (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57(2), 411–445.

Agarwal, S., and W. Qian (2013): "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore," *Working paper*.

Aiyagari, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 109(3), 659–84.

Alvarez, F., L. Guiso, and F. Lippi (2010): "Durable Consumption and Asset Management with Transaction and Observation Costs," *NBER Working Paper*, 15835.

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.

Black, O. (2011): "Wealth in Great Britain, Main Results from the Wealth and Assets Survey 2008/10," *Office for National Statistics*.

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887–1921.

Borsch-Supan, A., and A. Eymann (2002): "Household Portfolios in Germany," in *Household Portfolios*, ed. by L. Guiso, M. Haliassos, and T. Jappelli, pp. 291–340. MIT Press.

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," Discussion paper, Board of Governors of the Federal Reserve System (US).

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," *CEPR Discussion Papers*, (9649).

Cochrane, J. H. (1989): "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives," *American Economic Review*, 79(3), 319–37.

Daffin, C. (2009): "Wealth in Great Britain, Main Results from the Wealth and Assets Survey 2006/08," *Office for National Statistics*.

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.

Eurosystem Household Finance and Consumption Network (2013a): "The Eurosystem Finance and Consumption Survey: methodological report for the first wave," *ECB Statistics Paper Series*, 1.

——— (2013b): "The Eurosystem Finance and Consumption Survey: results from the first wave," *ECB Statistics Paper Series*, 2.

Foster, K., E. Meijer, S. Schuh, and M. Zabek (2011): "The 2009 Survey of Consumer Payment Choice," *Federal Reserve Bank of Boston Public Policy Discussion Paper*, 11-1.

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.

Guiso, 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), 971–87.

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," Discussion paper, 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," *NBER Working Papers*, 18941.

Kaplan, G., and G. Violante (2010): "How Much Consumption Insurance Beyond Self-Insurance?," *American Economic Journals: Macroeconomics*, 2, 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*.

Kennickell, A. B. (2005): "The Good Shepherd: Sample Design and Control for Wealth Measurement in the Survey of Consumer Finances," *Mimeo*.

Kreiner, C. T., D. D. Lassen, and S. Leth-Petersen (2012): "Consumer Responses to Fiscal Stimulus Policy and Householdsâ Cost of Liquidity," *CEPR Discussion Papers*, (9161).

Krusell, P., and A. Smith, Jr. (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," *CEPR Discussion Papers*, (9530).

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 (2011): "Consumer spending and the economic stimulus payments of 2008," Discussion paper, National Bureau of Economic Research.

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. Slemrod (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," *The American Economic Review*, 89(4), 947–958.

Statistics Canada (2006): "The Wealth of Canadians: An Overview of the Results of the Survey of Financial Security, 2005," *Pension and Wealth Research Paper Series*.

Vandone, D. (2009): *Consumer Credit in Europe: Risks and Opportunities of a Dynamic Industry*. Springer Verlag Berlin Heidelberg.

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.

Appendix

A Data sources

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.