

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

[PRELIMINARY AND INCOMPLETE - PLEASE DO NOT QUOTE]

Abstract

The *wealthy hand-to-mouth* are households who hold little or no liquid wealth (e.g. cash, checking, and saving accounts), despite owning sizable amounts of illiquid assets (i.e., assets that carry a transaction cost, such as housing, large durables, or retirement accounts). This portfolio configuration implies that these households have large marginal propensities to consume out of small 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) of hand-to-mouth agents based on 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 sheet, and the persistence of hand-to-mouth status over the life cycle. Finally, we discuss the implications of this group of consumers for macroeconomic modelling and policy analysis.

1 Introduction

The life-cycle permanent income hypothesis (LC-PIH) is an invaluable organizing framework for the analysis of both household survey data, and aggregate time series data, on the joint dynamics of income and consumption. However, 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 (Freidman, 1957; Hall, 1978) and the more recent “buffer-stock” versions (Deaton, 1991; Carroll, 1997). Specifically, while it is common to measure a large sensitivity of consumption with respect to anticipated, transitory changes in income, both at the micro and macro level, according to the theory these types of shocks should be largely 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), and Parker, Souleles, Johnson, and McLelland (2011) concluded that, in both episodes, the consumption response was strong: on average between 20 and 40 percent of the stimulus payments are 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 rebate and find comparable effects. A number of additional studies based on micro data also finds substantial consumption responses to anticipated fluctuations in income. Examples are the reaction of consumption to changes in social security taxes in Parker (1999), the analysis of federal tax refunds in Souleles (1999) and, most recently, the consumption response to fiscal stimulus in Singapore estimated by Agarwal and Quian (2013).

Similar anomalies in the joint behavior of income and consumption, relative to the LC-PIH, have been identified from aggregate time series. As discussed, for example, in a series of papers by Campbell and Mankiw (1989, 1990, 1991) which analyze US and cross-country time series, expected changes in aggregate consumption tend to be associated with expected changes in aggregate income. Moreover, expected changes in consumption are uncorrelated with real interest rates, a result which (as long as the IES is not zero) implies a break-down of the forward looking Euler equation holding with equality. Related studies are Attanasio and Weber (1993), and Ludvigson and Michaelides (2001).

The most natural way to account for these facts is through the existence of a sizable share

¹In these two episodes, the US 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 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.

of *hand-to-mouth* (HtM) consumers. HtM consumers have a high marginal propensity to consume (MPC) out of transitory income changes, which can account for the high correlation between consumption and the transitory component of income growth. Moreover, HtM consumers are not on their Euler equation, hence they generate 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 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 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 (1993), Aiyagari (1994), Rios-Rull (1995), and Krusell and Smith (1997), and the so-called spender-saver models with impatient and patient consumers in the tradition of Campbell and Mankiw (1989). 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.

In this paper, we argue that measurements of HtM behavior inspired by this large class of heterogeneous agents models are misleading because they miss what we call *wealthy hand-to-mouth* households. These are households who hold sizable amounts of illiquid wealth (e.g. housing, retirement accounts), and yet optimally choose to consume all of their disposable income every period. Clearly, such households would not be picked up by standard measurements since they own positive (and for many, 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), Kaplan and Violante (2013), and Huntley and Michelangeli (2014). Through the lens of a 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 own significant amounts of illiquid assets. Just like the P-HtM households, W-HtM households have large MPCs out of anticipated small income fluctuations. But, as we show, 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 W-HtM for macroeconomics modelling and policy analysis.

First, we ask why households with sizable net worth 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 simple model based on Kaplan-Violante (2013). 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 hand-to-mouth 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).² The model also provides useful guidance for our empirical strategy. We outline the strategy in detail in Section 3, 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 sheets compare with the N-HtM and P-HtM. Finally, we investigate the persistence of HtM status over time to understand whether being wealthy 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. We describe these data in Section 4. In Section 5, we study US data for which we have several repeated cross-sections starting between 1989 and 2010, as well as a two-year panel in 2007. In Section 6, we present a comparative analysis of Canada, Australia, UK, Germany, France, Italy and Spain with survey data from around 2010.

The analysis based on US data leads to five main findings. First, between 25 and 45 percent of US households are HtM, with our preferred estimate at 35 percent, and over 2/3 of them are wealthy HtM. Therefore, W-HtM represent the vast majority of this group. Second, while households appear to be most frequently P-HtM at young ages, the age profile of W-HtM is hump-shaped and peaks around age 40. Third, 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. Fourth, W-HtM typically hold substantial amount of illiquid wealth, e.g., the median at age 40 is around \$50,000. Finally, the W-HtM status is quite transient. We estimate that it lasts, on average, around 2.5 years.

²This explanation is reminiscent of calculations by Cochrane (1989), and more recently by Browning and Crossley (2001), who show that in some contexts the utility loss from setting consumption equal to income, instead of fully optimizing, is second order.

Comparing the US to the other countries we study, two interesting findings emerge. First, like in the US, in all of the seven other countries W-HtM households are a far greater share of the population than P-HtM households. The age profiles of the three groups are also remarkably similar across countries. Second, there are large cross-country difference in the overall share of HtM households, ranging for over 40% in the UK (with over 30% W-HtM) to less than 15% in Italy (with under 10% W-HtM). We investigate the sources of these cross-country discrepancies to the extent that our data allow.

In Section 7, we argue that in the cast of characters of macroeconomic models, the W-HtM deserve their own separate status and cannot be assimilated neither to the P-HtM nor to the N-HtM. We provide two simulations from the Kaplan and Violante (2013) model to make this point. In the first simulation, we show that if one ignored the W-HtM and considered all households with a minimum amount of net worth as unconstrained, the aggregate consumption response to a fiscal stimulus such as a lump-sum transfer would be severely miscalculated, and its cross-country variation would be vastly underestimated. Moreover, one would draw the wrong conclusions about which demographic group to target in order to obtain the highest bang-for-the-buck from a fiscal stimulus transfer. In the second simulation, we show that if one combined wealthy and poor HtM in a unique group and ignored the fact that W-HtM hold illiquid wealth which is accessible at a cost – as routinely done by spender-saver models where HtM agents are highly impatient and therefore do not save at all – then, the consumption response to a large negative aggregate shock would be largely overestimated. The reason is that, in such a scenario, W-HtM households would withdraw from their illiquid account to smooth consumption, an opportunity unavailable to the poor HtM.

Section 8 concludes the paper with some reflections on the implications of our findings for the design of fiscal policy.

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

In this section, we analyze a simple two-period model in order to illustrate the determinants of hand-to-mouth behavior, and to provide 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. Period utility is CRRA with intertemporal elasticity of substitution $\sigma > 0$. 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 $a > 0$

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 and make their consumption and liquid saving decision in period 1. In the second period, they receive an endowment $\Gamma\omega$ and consume this endowment plus their savings in liquid and illiquid wealth. Note that even abstracting from discounting, period 1 and period 2 are intrinsically different because the illiquid asset is only accessible in period 2.

We make the following normalizations and parametric assumptions. 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 assume that $\Gamma > 1$, so Γ captures the growth in the endowment profile over the life cycle. Finally, we consider two possible values for y , $\{y_L, y_H\}$ where $y_L = 0$ and $y_H > R + \Gamma$ and refer to these two cases as “low-income” and “high-income” paths.

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 $a > 0$. 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 easily be accommodated.

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 by backward induction, starting from the consumption decision at $t = 1$. The problem faced by the household at $t = 1$ is

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

which has the solution

$$m_2 = \max \left\{ \frac{y - \Gamma + m_1}{2}, 0 \right\}. \quad (1)$$

The interior solution for m_2 implies a smooth consumption path, $c_1 = c_2 = (y + \Gamma + m_1)/2$, while the corner solution $m_2 = 0$ yields an increasing consumption path, $c_1 = y + m_1$, $c_2 = \Gamma$. Since the liquid asset is the only available asset, the initial portfolio allocation decision is trivial, hence $m_1 = 1$.

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

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

which has the solution

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

The interior solution for m_2 implies a smooth consumption path $c_1 = c_2 = (y + \Gamma + m_1 + Ra)/2$, while the corner solution yields an increasing consumption path ($c_1 = y + m_1$, $c_2 = \Gamma + Ra$). 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 $m_2 = 0$. 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 $m_2 > 0$.

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

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

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 smooth consumption profile $c_1 = c_2 = (y_H + \Gamma + R)/2$.

Under the low income path, when the constraint binds at $t = 1$ and $m_2 = 0$, the problem becomes

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

which has the solution

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

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

When $1 < R \leq \Gamma^{\frac{1}{\sigma}}$, Equation eq3 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 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 $c_2 = \Gamma$. 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 $c_2 = R^\sigma c_1 > \Gamma c_1$.

This simple example clarifies the role of the three key parameters in determining wealthy HtM behavior: the relative return on the illiquid asset R , income growth 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, the more the household is willing to accept consumption fluctuations across periods, and the more likely it is to become wealthy HtM rather than poor HtM.

Finally, note that since the model is deterministic, some households choose to invest in the illiquid asset (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 op-

timal 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.

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 a small, unexpected transfer τ from the government. What is the household's MPC out of this transfer?

A N-HtM household has a 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 $y = y_L = 0$ and optimally chose $a = 0$:

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

which has the solution

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

For any small payment $0 < \tau \leq \Gamma - 1$, this household remains P-HTM and has an MPC of 1. If, instead $\tau > \Gamma - 1$, the household becomes unconstrained and its MPC drops to $(\tau + \Gamma - 1) / 2\tau$, which approaches $1/2$ as τ becomes large.

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

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

which has a solution

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

This household has a MPC of 1 as long as $\tau \leq \Gamma - 1 + (R + 1)a^*$. This condition is looser than

the condition for a P-HtM to have a MPC of 1. Thus in a general model with more income heterogeneity, these findings suggest that the average MPC among W-HTM households will be larger than the average MPC among P-HTM households.

2.4 Unsecured credit

The above discussion ignored the possibility for households to borrow. In this section we extend the model to allow households to access credit to finance consumption at $t = 1$. We assume that households can borrow up to a fraction ϕ of their future endowment 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 $R^b \geq \Gamma$, so that 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

We saw above that 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 as above, and hence $m_2 > 0$.

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

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

which has the solution

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

Since $R^b < \Gamma$, the household always borrows a positive amount. Moreover, if $R^b < \Gamma(1 - 2\phi)$, then the credit limit is binding. In this case, the household has 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 and hence the household has a smooth consumption path with

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

In light of the discussion of 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 her HTM status. For small transfers, with an interior negative position is unconstrained and has a MPC equal to 1/2.

2.4.2 Solution with illiquid asset

As in the case without the illiquid asset, under the high income path the household is not constrained at $t = 2$, so allowing for borrowing has no effect on its decisions. Under the low-income path where $y = y_L = 0$, the household may want to borrow. Its consumption decision at $t = 1$ is

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

which has the solution

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

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

We now analyze the portfolio decision. Since we are interested in characterizing HtM behavior, we zoom in on the case where $m_2 = 0$. In this case, the portfolio problem is:

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

with solution

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

Using the restriction $R^b < \Gamma$, it can be shown that a is strictly positive if $R > \Gamma(1 - \phi) / (1 + \phi)$ and $a = 0$ if $R < R^b(1 - \phi) / (1 + \phi)$. The first parametric configuration corresponds to a household who chooses optimally 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). The second case corresponds to a P-HtM household who has borrowed up to the limit. Both households will have MPC of one out of a small transfer.

3 Identifying hand-to-mouth households in the data

In the previous section, we laid out a simple model to illustrate the emergence of hand-to-mouth behavior. We showed that there are fundamentally 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 can be large: the unsecured credit limit and zero liquid assets. The latter is always a hard constraint. The former 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, like those we describe in Section 4. Broadly speaking, these survey data contain household-level information on income and on different categories of assets and liabilities in the household balance sheet which can be aggregated into net liquid and illiquid wealth.

Inevitably, the exact definitions of the components of liquid and illiquid wealth vary across surveys. Nonetheless we try to be as consistent as possible. In our baseline definition, liquid wealth includes cash, checking, saving, money markets and call accounts plus directly held mutual funds, stocks, corporate and government bonds net of unsecured debt (mainly revolving balances on credit cards). Illiquid wealth includes housing equity housing net of mortgages and home equity loans, private retirement accounts, life insurance policies, certificate of deposits, and saving bonds. We also examine a narrower definition of liquid wealth which excludes directly held mutual funds, stocks and bonds, and a broader one that includes outstanding debt in home-equity lines of credit. Symmetrically, we explore broader definitions of illiquid wealth that includes the resale value of vehicles net of loans, and the value of businesses for the self-employed. Note that while changing the definition of liquid wealth affects the total fraction of HtM households in the population, changing the definition

of illiquid wealth only affects the split between poor and wealthy HtM. We refer the reader to the next section for a thorough description of these surveys, country by country.

We now describe our strategy for identifying HtM households in the data separately for households at the zero kink and households at the credit limit.

First, consider the case of a HtM household at the zero kink. In the model, such a household does not borrow or save into liquid assets in period t , i.e., it carries zero liquid wealth between t and $t + 1$. The first issue to confront is that, while in the discrete-time model the start of a period coincides with the date when the agent receives her labor income, survey data typically do not contain information about the pay-period of the earners in the household, or even the frequency of pay (weekly, bi-weekly, monthly, etc...). In the benchmark measurement exercise, we'll take an intermediate view that labor income accrues bi-weekly (i.e., every two weeks) to the household, and in the sensitivity analysis we report results for the weekly and monthly paycheck assumption. Let y_{it} be the income of household i in period t .³

Information on liquid and illiquid wealth is typically reported as either the balance at the date of the interview or as the average balance over the preceding period (such as the past week, two weeks, or month). Suppose we observe average holdings of liquid wealth m_{it} and illiquid wealth a_{it} . Clearly, m_{it} is positive for all households (whether HtM or not) because labor income is paid at discrete dates whereas consumption occurs smoothly over the period, i.e., there is always a mismatch in the timing of consumption and earnings *within* a pay period. Thus, a conservative criterion to identify HtM agents in the data is to count those households in the survey whose average balances of liquid wealth are positive (and hence, presumably, they are not borrowing), but are less than half their earnings per pay-period, where “half” presumes resources being consumed at a constant rate. In essence, a household is poor HtM at the zero kink if

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

and wealthy HtM at the zero kink if

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

Note that this estimate is a *lower bound* because, while all non HtM households would always hold average liquid balances above half their earnings, some HtM households may fall in this latter group as well.⁴

³For example, if the survey reports information on household annual income, and the pay-period of reference is 2 weeks, then y_{it} is annual income divided by 26.

⁴For example, a household who starts the period with positive savings in addition to earnings and ends the period with zero liquid savings, is a HtM household in that period, but her average liquid balance would be above half earnings, and she would not be measured as HtM.

If, 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 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 period, but the error averages to zero.

We now consider the measurement of HtM households at the credit limit $\underline{m}_{it} > 0$. If the information on unsecured credit limit is available from the survey, then consistently with the strategy just outlined, we propose to count a household in the survey as poor HtM at the credit limit if

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

and wealthy HtM at the credit limit if

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

When, as in most surveys, individual credit limits are not reported, we assume that the credit limit is a multiple of individual income and experiment with a plausible range of values.

For comparison, 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

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

A recent literature has emphasized the existence of “commitment consumption expenditures,” i.e., expenditures that the household is forced to make every period, unless she incurs in a large transaction cost to modify her previous commitments. Examples are rent, mortgage or other loan payments, utility bills, school and gym fees, alimony, etc. The distinct feature of commitment expenditures is that they are bulk expenditures occurring at a point in time that deplete “discretely” the balance of liquid wealth. Does the presence of such expenditures affect our identification strategy? Let $\bar{c}_{it}$ be the amount of commitment consumption for household i at date t . If $\bar{c}_{it}$ happens at the beginning of the pay period, the criterion to identify a HtM household (say, at the zero kink) should be amended as $m_{it} \leq (y_{it} - \bar{c}_{it})/2$, and if it occurs at the end of the period, the rule should be $m_{it} - \bar{c}_{it} \leq y_{it}/2$. In the first case, our baseline rule overestimates HtM status, and in the second case it underestimates it. If, instead, commitment expenditures occur smoothly over the period or are realized in the middle of the pay period, then the rule should be $m_{it} - \bar{c}_{it}/2 \leq (y_{it} - \bar{c}_{it})/2$ which reduces to our baseline criterion. Throughout the empirical analysis, we verify the robustness of our estimates to the presence of those commitment expenditures that we can measure in our

survey data.

Finally, whenever the survey allows, as an alternate estimate of the fraction of HtM households, we also use survey questions that directly ask whether (i) over the past month, the household has saved a positive amount or whether (ii) its expenditures over the last month have exceeded income, abstracting from purchases of large durable goods such as housing or cars. Counts of HtM households based on these questions provide a useful check of the reliability of our identification strategy based on reported liquid wealth and income. In what follows, we refer to counts of HtM households based on the first type of direct questions as HtM-D1, and to counts based on the second type as HtM-D2.

4 Data on household portfolios

In this section, we describe the datasets that we use to measure HtM behavior. As is clear from the previous section, in order to categorize a household as W-HtM, P-HtM, or N-HtM, we need information on household labor income and on the amounts of each type of assets and liabilities on the household’s balance sheet. Each of the surveys that we use asks questions about the respondent household’s finances, with enough detail to allow us to estimate a household’s HtM status. We first provide a background on each survey, and then describe how the precise definitions of the balance sheet items differ between surveys.

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 since 1983 and collects detailed information on household balance sheets, income, and demographic characteristics for a representative cross-section of US households. Major technical revisions to the survey were implemented in 1989. Since 1992, data have been collected by the National Opinion Research Center at the University of Chicago, who have largely preserved the structure and questions asked from the 1989 survey and maintained a high degree of consistency through the years. As such, we conduct analysis on the 1989 through 2010 surveys. While the survey does 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 later years.

The target population for the survey is all private households residing in the United States 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. 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.

4.2 Euro area: HFCS

Our data for France, Germany, Italy and Spain come from the Household Finance and Consumption Survey (HFCS). The HFCS is a joint project 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 across 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 will typically go into more detail on some aspect of the core questionnaire that

the member country wishes to explore. For example, Spain asks series on non-core questions designed to examine methods by which households pay their bills. We use data from this module on cash holdings in Spain to impute cash holdings in France, Germany, and Italy.

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, using individual data on taxable wealth to create the wealthy sample. The wealthy sample was divided into four strata and sampled proportionally according to the relative size of the strata. Germany uses regional level taxable income, and oversampled small municipalities and, in larger municipalities, street sections with average income over a threshold. Spain defined eight wealth strata based on individual taxable wealth, and were oversampled progressively at higher rates. Italy did not oversample in any way.

4.3 United Kingdom: WAS

Our data for the United Kingdom come from the Wealth and Assets Survey (WAS). The WAS is a longitudinal survey that was implemented by the Office of National Statistics (ONS) who interviewed households across the UK. The survey was intended to measure the economic well-being of households in the UK, 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 in July 2008 to June 2010. In our analysis, we use data from Wave Two.

The questionnaire was divided into two parts. The first part was the household questionnaire which was completed by one person in the household designated to be the household reference person and collected household level information on household demographics, as well as information about household assets and liabilities. The second part of the questionnaire was an individual questionnaire administered to each adult aged 16 or over in the household and asked in depth questions about economic status, education, employment, benefits, and

individual financial assets, just to name a few.

4.4 Canada: SFS

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

The surveyed households are a representative sample of all private households in Canadian provinces. The SFS also uses a dual frame sample design. The main sample was a sample selected from the Labour Force Survey sampling frame. In order to over-sample high income households, the second sample was 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 were asked questions regarding income, demographics, education, and employment. Persons aged 25 and older were asked about employer pension plans and benefits. Persons aged 45 and over were asked about retirement information. Questions regarding household assets and liabilities were asked to the household member deemed most knowledgeable on the subject.

4.5 Australia: HILDA

Our data for Australia come from the Household, Income and Labour Dynamics in Australia (HILDA) Survey. HILDA is a broad social and economic longitudinal survey, with particular attention paid to family and household formation, income and work. The Survey is managed by the Melbourne Institute of Applied Economic and Social Research at the University of Melbourne. Wave One of the survey was implemented in 2001, and households in the survey have been interviewed annually since.

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, we look at data from Wave Two (2002), Six (2006), and Ten (2010) which 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 Wave Two, Six, and Ten. The Person Questionnaires are asked to all members of the household age 15 years and older, and collects information on family background, education, employment, and income among other things.

4.6 Data comparability and sample selection

As mentioned in Section 3, we distinguish between household assets and liabilities based on their relative liquidity. Each individual survey is tailored to its 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 defining income, liquid, and illiquid wealth across surveys.

In the SCF, we consider liquid assets to be cash, checking, saving, money market and call accounts plus directly held mutual funds, stocks, corporate and government bonds. The SCF does not ask questions about cash holdings, and as such we impute cash holdings based on data from the Survey of Consumer Payment Choice administered by the Federal Reserve Bank of Boston (see Foster et al., 2011). For the HFCS, liquid assets are considered cash, sight (also called current, draft, or checking) accounts, mutual fund holdings, shares in publicly traded companies, and corporate or government bond holdings. Questions on cash holdings are not in the core questionnaire model and only information from Spain is available. For France, Germany, and Italy, we impute cash holdings from the Spanish data based on age, gender, marital status, number of children, annual income, and employment status. In the U.K. WAS, liquid assets include Individual Savings Accounts, savings accounts, and holdings of shares and gilts. 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).

We consider liquid debt in the SCF to be the balance on all credit cards (after the most recent payment) which accrue interest. As most credit cards feature a one month “grace period” on purchases which makes them a substitute for cash in the (very) short term, we restrict our measure of credit card debt to households that do not regularly pay off their balances in full each month. 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 accrue interest. In the WAS, liquid debt is credit card debt, plus any

balances on store cards, hire purchases, and mail orders. Liquid debt in the SFS is credit card and installment debt. Liquid debt in the HILDA is credit card debt.

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 (401(k)s, IRAs, thrift accounts, and future pensions), life insurance policies, certificate of deposits, and saving bonds. The definition 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, life insurance policies, certificate of deposits, and saving bonds. In the U.K. WAS, we take illiquid wealth to include pensions, insurance products, National Savings products, and the value of the main residence, other houses, and land less mortgage and land debt. Illiquid wealth in the Canadian SFS is the value of the principal residence and other real estate 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 properties plus life insurance policies and superannuation (government-supported, compulsory private retirement funds).⁵

In choosing our definition of income, we make an attempt to include all labor earnings plus government transfers (which 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, unemployment benefits, worker’s compensation, regular private transfers such as child support and alimony, regular public transfers such as food stamps and SSI, and regular income from other sources excluding investment income. Income in the HFCS is gross income from wages and salaries, unemployment benefits, regular private transfers such as child support and alimony, and regular public transfers. For the U.K. WAS, we define income as net employee earnings, plus any public benefits such as the Jobseeker’s Allowance and Maternity Allowance. Income in the Canadian SFS is after-tax total income. In HILDA, income is wages and salaries, regular private transfers such as child support and alimony, and public benefits such as the Australian Government Parenting Payment.

The reference period for income and wealth questions also 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

⁵Superannuation has some features of private retirement accounts (such as 401K in the U.S.) which we always 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 altogether.

	US		CA		AU		UK	
	Median	Frac. Pos.	Median	Frac. Pos.	Median	Frac. Pos.	Median	Frac. Pos.
Earnings plus benefits (age 22-59)	46000	0.972	50091	1.000	58615	0.958	22200	0.973
Net Worth	60703	0.887	114477	0.876	293462	0.984	187146	0.880
Net liquid wealth	2064	0.768	2569	0.712	9089	0.875	2000	0.630
Cash, checking, saving, MM	2787	0.925	2792	0.864	6477	0.978	2500	0.767
Directly held stocks	0	0.148	0	0.109	0	0.352	0	0.161
Directly held bonds	0	0.015	0	0.106	0	0.015	0	0.155
Revolving credit card debt	0	0.299	0	0.412	0	0.296	0	0.404
Net illiquid wealth	55000	0.764	103308	0.752	269231	0.939	175543	0.843
Housing net of mortgages	31000	0.636	65894	0.648	192308	0.714	81999	0.678
Retirement accounts	1500	0.526	893	0.518	46154	0.860	58327	0.765
Life insurance	0	0.191	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.
Earnings plus benefits (age 22-59)	29801	0.963	26149	0.981	20457	0.900	26662	0.936
Net Worth	45396	0.976	103397	0.976	153257	0.996	204071	0.985
Net liquid wealth	1352	0.922	1468	0.960	4683	0.976	3311	0.970
Cash, checking, saving, MM	1246	1.000	1283	1.000	4018	1.000	2913	1.000
Directly held stocks	0	0.110	0	0.153	0	0.048	0	0.108
Directly held bonds	0	0.054	0	0.015	0	0.150	0	0.014
Revolving credit card debt	0	0.229	0	0.077	0	0.053	0	0.083
Net illiquid wealth	39207	0.880	98254	0.923	140423	0.807	195598	0.885
Housing net of mortgages	0	0.487	81601	0.614	135022	0.719	185074	0.849
Retirement accounts	0	0.240	0	0.039	0	0.089	0	0.035
Life insurance	5	0.496	0	0.381	0	0.199	0	0.252

Table 1: All dollar figures converted to 2010 US dollars using PPP conversion factor (GDP) to market exchange rate ratio from the World Bank. Data for Canada is first adjusted to 2005 USD and the to 2010 USD using the US CPI-RS.

the “usual” amounts for 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.

For the SCF, the reference period for wealth questions is the interview date. For some assets such as checking accounts, if the respondent was unsure, the interview could prompt for an average balance over the month. France, Germany, and Spain use the date of the interview as the reference period for wealth variables in the HFCS. Italy uses December 31, 2010. The WAS asks for current balances or values of assets and liabilities. The SFS asks for information on assets and debts for “a time as close as possible to the date of the interview”, and the HILDA asks for current balances.

In all surveys, we restrict analysis to households in which the head of household is between 22 and 79 years of age.⁶ We also drop a household from the analysis if the household has negative income. In Section 5 where we look at the United States, we pool all years for the SCF, yielding about 167 thousand observations after sample selection. In Section 6 where we document international evidence, we use the 2010 SCF in which there are 30 thousand.

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

The HFCS has about 68 thousand households for France, 16 thousand for Germany, 35 thousand for Italy, and 28 thousand for Spain. Wave Two of the WAS has about 18 thousand households. The 2005 SFS has about 5 thousand households. The HILDA has about 6 thousand households for the Wave Ten (2010).

Table 1 reports, for each country in the sample, some basic statistics about household income, liquid and illiquid wealth holdings, and portfolio composition.

[TO BE COMPLETED]

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 sensitivity of our estimates to alternative assumptions about the frequency of pay, liquidity of different asset classes, and availability of credit, as described in Section 2. We then analyze the key demographic characteristics of N-HtM and HtM households, and we examine their portfolio composition in more detail. Lastly, we exploit the longitudinal dimension of the 2007-2009 waves of the SCF to study the persistence of HtM status.

5.1 The share of HtM households

Figure 1 plots the fraction of wealthy and poor HtM households in the US population over the period 1989-2010, according to the benchmark definition which sets the pay frequency to two weeks, and uses equations (4) – (8) with $\underline{m}_{it}$ being the self-reported credit limit in the SCF. Our estimates indicate that around 35% of US households are HtM. 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 is not poor, but owns illiquid assets.

Looking at changes over time, across the two decades we study there are some notable fluctuations in the fraction of HtM households, from 30% in 2001 to 37% around the 1991 and the 2008 recessions. The measure of P-HTM and W-HTM tends to comove over time, albeit weakly.

In panels a) and b) of Figure 2, we verify the robustness of our estimates with respect to the tightness of the credit limit. Panel a) imposes a tight individual-level credit limit of one month of income and panel b) a laxer individualized limit of one year of income. By comparing the two panels, we draw three major conclusions. First, with the tight credit limit, the share of HtM households is slightly above our baseline estimates, or 37%. Under the loose




Figure 1: Time-series of fraction of HtM households in United States.

limit, instead, it drops to 25%. Second, this drop is entirely accounted for by W-HTM, while the measure of poor HtM agents is not much affected by the value of the credit limit used in the threshold rules and always fluctuates around 10%. Third, the fluctuations over time in the share of HtM households is robust to the tightness of the credit limit.

Panels c) and d) of Figure 2 provides an estimate of HtM behavior based on two direct questions in the SCF. The first question (X3015, X3016) asks households “Which of the following statements comes closest to describing your saving habits?” And we label HtM a household who responds that “Don’t save - usually spend more than (or as much as) income.” Panel d) shows that just below 25 percent of households are HTM according to this definition. Interestingly, the split between P-HTM and W-HTM is similar to that in Figure 1.

The second method identifies HtM households by using a combination of sequential questions (X7510, X7509, X7508) 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.”⁷ Based on this definition (see panel d), the share of HtM households is around 40-45 percent. W-HtM households account for 3/4 of the total and fluctuations of this measure over time follows very closely the baseline definition of Figure 1.

⁷These questions were included in the survey starting from 1992.




Figure 2: Fraction of HtM households in United States, alternate definitions.

To sum up, we estimate that, depending on the exact definition, between 25 and 45 percent of US households are HtM, but most importantly around 2/3 of them are wealthy HtM, i.e. own positive illiquid wealth in spite of their low holdings, often even negative, of liquid wealth.

5.1.1 Robustness

Table 2 offers a robustness analysis on the fraction of HtM households in the population, and on their split between poor and wealthy. Across all the variations in Table 2, the total share of HtM households is always within the 30-40 percent range, with only few exceptions on both sides of the range. According to the direct question about saving habits (HtM-D1) described in Section 3, only 24% of households report that they do not “usually save” and are classified as HtM according to this criterion. However, answers to the other direct question (HtM-D2), arguably better articulated, leads to one of the highest estimates of hHtM families, close to 44%. Even when we impose a very loose credit limit (6 months of income), the estimated

	W-HtM	P-HtM	N-HtM	HtM	HtM-NW
Baseline	0.243	0.131	0.626	0.374	0.152
HtM-D1	0.156	0.086	0.758	0.242	—
HtM-D2	0.309	0.127	0.565	0.435	—
Weekly	0.207	0.118	0.675	0.325	0.136
Monthly	0.300	0.148	0.552	0.448	0.174
Comm. Cons. - Beg.	0.220	0.113	0.667	0.333	0.132
Comm. Cons. - End	0.295	0.150	0.555	0.445	0.175
6 months of credit	0.125	0.103	0.772	0.228	0.110
Reported credit limit	0.214	0.127	0.659	0.341	0.145
Higher illiquid wealth cutoff	0.231	0.143	0.626	0.374	0.152
Cars as illiquid wealth	0.322	0.052	0.626	0.374	0.064
Businesses as illiquid wealth	0.246	0.128	0.626	0.374	0.147
Stocks as illiquid assets	0.276	0.130	0.594	0.406	0.152
HELOCs as liquid debt	0.265	0.131	0.604	0.396	0.152

Table 2: Robustness results for fraction HtM in each category in the SCF pooled 1989-2010. Higher illiquid cutoff refers to using $a_{it} > 1000$ to be considered W-HtM.

fraction of HtM households is close to 30%.

The split between W-HtM and P-Htm is also quite robust, with the former group systematically accounting for 2/3 of HtM households. A notable exception is when we add vehicles (net of car loans) to illiquid wealth. Over 90% of families owns at least one car (see Table 1 in Section 4), which turns 10% of households previously classified as poor HtM into wealthy HtM.

The last column of the table reports measures of HtM in terms of net worth. The key finding is that this indicator, which is centered around 15%, consistently misses 60% of HtM households in the economy.

5.2 The demographics of HtM groups

We now turn to the demographic characteristics of the three group of households, poor HtM, wealthy HtM and non HtM. For this analysis, we pool all the years together to maximize sample size, and use the baseline definition of Figure 1. Panel a) of Figure 3 plots the size of these three groups by age. Not surprisingly, the bulk of P-HtM behavior is observed at the early stages of the life-cycle. The fraction of P-HtM drops sharply between age 20 and age 30, and keeps falling steadily over the lifecycle until reaching roughly 5 percent in retirement. The fraction of W-HtM households is, instead, markedly hump shaped in age: it peaks at around age 40 when 25 percent of US households are HtM, and remains above 10 percent throughout the lifecycle. The share in the residual group of N-HtM individuals increases




Figure 3: Age profile of the demographics of the HtM

steadily from 45 percent at age 25 to 70 percent in retirement.

Panels b) and c) plot the education level and the marital status of the three 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, non and wealthy HtM are indistinguishable, whereas the figure shows that poor HtM are 30 percent less likely to be married.

Panel d) plots the age profile of median income across the three groups. Poor HtM households are income-poor, with median income around \$20,000 (in \$2010) during working years, while N-HtM are high-income households earning \$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.




Figure 4: Age profile of the portfolio composition of the HtM

5.3 The portfolio composition of HtM groups

Figure 4 digs deeper into the composition of the balance sheet of these three group of households, poor HtM, wealthy HtM and non HtM. Panel a) shows that median liquid wealth holdings are zero for P-HtM and, perhaps unexpectedly, slightly negative for W-HtM. N-Htm households have, instead, substantial holdings of liquid wealth peaking at around 20,000 before retirement. Panel b) reveals that the W-HtM group holds significant amounts of illiquid wealth: for example, median holdings at age 40 are already \$50,000. Therefore, wealthy HtM are not just poor HtM with small amounts of savings in less liquid assets.

The next two panels of Figure 4 articulate this observation further. Panels c) and d) plot the age profile of the average fraction of illiquid wealth held in housing and retirement accounts by wealthy HtM and non HtM households. The conclusion is striking: the lines are on top of each other, indicating that the portfolio allocation of these two groups is nearly identical. 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.

07 $\rightarrow$ 09	P	W	N
P	0.526	0.116	0.358
W	0.065	0.182	0.753
N	0.026	0.048	0.926

Table 3: 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.

5.4 Persistence of HtM status

A natural question to ask is: how persistent is the HtM status long the life cycle of an individual? The 2007-2009 panel of the SCF allows us to address this question. Table 3 reports the 2-year transition matrix across the three HtM statuses for US households. Households in the poor HtM status have a 53% chance of still being P-HtM two years later, but also a 36% chance of becoming non HtM households within the same period. The wealthy HtM status appears to be the most transitory: the probability of remaining W-HtM two years later is less than 20%. Put it differently, the expected length of W-HtM status is around 30 months. Finally, the non HtM state is the least transient of the three, and it almost an absorbing state.

6 Cross-country Evidence

The previous section showed that around 35% of households in the United States are HtM, with 1/3 of these P-HtM and 2/3 of these W-HtM. In this section we use household portfolio data from seven other industrialized countries to assess whether the prevalence of W-HtM households is a feature of the wealth distribution that is common across countries, and if so, whether the characteristics of W-HtM in terms of demographics and balance sheets are similar to those in the United States. As discussed in Section 4, we focus our attention on three other Anglo-Saxon countries - Canada, Australia and the United Kingdom - and the four largest European countries - France, Germany, 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 UK), 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 United Kingdom 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 5 shows the fraction of P-HtM and W-HtM households in each country. For all 8 countries, there are more W-HtM households than P-HtM. Thus a wide-spread feature of international household portfolios is that a complete characterization of the fraction of the




Figure 5: Fraction of HtM households across countries

population that is likely to exhibit HtM behavior requires going beyond those with just low net worth.

Although the existence of HtM is common across countries, there are large differences in the fraction of these households and the breakdown between P-HtM and W-HtM. A number of features of this cross-country distribution are apparent from Figure 5. First, there is a striking similarity in the prevalence and breakdown of HtM household in the United States, Canada and the United Kingdom. These three countries have the biggest fraction of both P-HtM and W-HtM households. Second, Australia is somewhat of an outlier among the Anglo-Saxon countries since although over 20% of the population are W-HtM, there are almost no P-HtM households. Third, the European countries have smaller fractions of both types of HtM households than the US, Canada and UK. Table 4, which contains an extensive robustness analysis on our definitions of HtM households, shows that these general findings are extremely robust.

6.1 Explaining cross-country differences

What explains these differences across countries? Part of the differences can be understood by examining portfolio characteristics by HtM status in each country, which are shown in Table 5. For each country we report the fraction of N-HtM, P-HtM and W-HtM who have positive

N-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Baseline	0.602	0.596	0.772	0.571	0.741	0.853	0.882	0.858
Weekly	0.660	0.639	0.815	0.585	0.829	0.924	0.928	0.901
Monthly	0.524	0.528	0.722	0.546	0.624	0.698	0.842	0.783
Comm. Cons. - Beg.	0.648	—	—	—	0.789	0.898	0.892	0.882
Comm. Cons. - End	0.527	—	—	—	0.635	0.699	0.840	0.795
6 months of credit	0.743	0.801	0.865	0.771	0.805	0.884	0.897	0.880
Max credit	0.636	—	—	—	—	—	—	—
Higher illiquid wealth cutoff	0.602	0.596	0.772	0.571	0.741	0.853	0.882	0.858
Cars as illiquid wealth	0.602	0.596	0.772	0.571	0.741	0.853	0.882	0.858
Businesses as illiquid wealth	0.602	0.596	0.772	0.571	0.741	0.853	0.882	0.858
Stocks as illiquid assets	0.582	0.548	0.731	0.555	0.693	0.834	0.874	0.850
HELOCs as liquid debt	0.576	0.516	—	—	0.740	0.852	0.880	0.857
P-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Baseline	0.152	0.151	0.018	0.127	0.053	0.023	0.044	0.033
Weekly	0.134	0.140	0.017	0.124	0.036	0.011	0.020	0.025
Monthly	0.175	0.171	0.021	0.130	0.073	0.039	0.056	0.049
Comm. Cons. - Beg.	0.130	—	—	—	0.040	0.014	0.037	0.025
Comm. Cons. - End	0.176	—	—	—	0.085	0.045	0.065	0.053
6 months of credit	0.120	0.093	0.016	0.082	0.041	0.020	0.043	0.027
Max credit	0.147	—	—	—	—	—	—	—
Higher illiquid wealth cutoff	0.162	0.156	0.025	0.130	0.083	0.062	0.047	0.037
Cars as illiquid wealth	0.062	0.088	0.008	0.077	0.036	0.023	0.011	0.015
Businesses as illiquid wealth	0.146	0.144	0.018	0.126	0.053	0.022	0.043	0.033
Stocks as illiquid assets	0.151	0.150	0.019	0.126	0.054	0.023	0.044	0.033
HELOCs as liquid debt	0.149	0.153	—	—	0.053	0.023	0.044	0.033
W-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Baseline	0.246	0.253	0.210	0.302	0.206	0.123	0.075	0.109
Weekly	0.206	0.221	0.169	0.291	0.135	0.065	0.052	0.074
Monthly	0.301	0.300	0.257	0.324	0.303	0.264	0.102	0.169
Comm. Cons. - Beg.	0.221	—	—	—	0.171	0.088	0.071	0.093
Comm. Cons. - End	0.297	—	—	—	0.280	0.256	0.096	0.152
6 months of credit	0.137	0.105	0.119	0.147	0.154	0.096	0.060	0.093
Max credit	0.217	—	—	—	—	—	—	—
Higher illiquid wealth cutoff	0.236	0.248	0.203	0.299	0.176	0.085	0.072	0.104
Cars as illiquid wealth	0.336	0.316	0.220	0.352	0.223	0.123	0.108	0.127
Businesses as illiquid wealth	0.251	0.260	0.210	0.303	0.206	0.124	0.076	0.109
Stocks as illiquid assets	0.268	0.302	0.250	0.319	0.253	0.143	0.082	0.117
HELOCs as liquid debt	0.276	0.331	—	—	0.207	0.125	0.076	0.110

Table 4: Robustness results for fraction HtM in each category. Higher illiquid cutoff refers to using $a_{it} > 1000$ to be considered W-HtM.

holdings of retirement assets, positive housing equity, and negative liquid wealth. We also report the mean and median shares of retirement assets and housing equity in households' illiquid wealth portfolios, and the mean and median liquid wealth to monthly income ratio.

First, consider the low fraction of poor HtM households in Australia, especially when compared with the other Anglo-Saxon countries. Table 5 reveals that this can be traced to the very high fraction of the population that own retirement account in Australia. Around 85% of N-HtM households in Australia have assets in a retirement account compared with 63% in the US, and 96% of W-HtM households in Australia have assets in a retirement account compared with 57% in the US. The high rate of retirement account ownership in Australia is in part due to the superannuation regulations, which require all employers to contribute to tax-deferred retirement accounts on behalf of their employees. When we exclude superannu-

N-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Median liq. wealth / income	2.337	2.896	3.600	8.527	1.189	1.022	3.306	2.116
Mean liq. wealth / income	56.934	15.836	73.908	56.916	12.972	91.485	12.880	10.916
Frac. neg. liquid wealth	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Median housing / illiq. wealth	0.657	0.769	0.724	0.452	0.583	0.848	0.996	1.000
Mean housing / illiq. wealth	0.651	0.648	0.601	0.479	0.471	0.619	0.868	0.915
Frac. pos. housing equity	0.731	0.744	0.767	0.833	0.537	0.654	0.746	0.870
Median retire / illiq. wealth	0.229	0.207	0.258	0.529	0.000	0.000	0.000	0.000
Mean retire / illiq. wealth	0.290	0.341	0.383	0.505	0.042	0.003	0.022	0.005
Frac. pos. retirement account	0.625	0.583	0.851	0.877	0.236	0.042	0.093	0.036
P-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Median liq. wealth / income	0.000	0.000	0.007	0.000	0.098	0.130	0.135	0.068
Mean liq. wealth / income	-0.668	-1.053	-0.900	-1.726	-0.063	-10.679	-0.079	-0.272
Frac. neg. liquid wealth	0.256	0.415	0.309	0.434	0.253	0.226	0.051	0.192
Median housing / illiq. wealth	1.000	1.000	1.319	1.000	1.091	1.212	5.333	1.000
Mean housing / illiq. wealth	1.603	1.411	1.511	8.096	1.376	1.217	5.333	1.022
Frac. pos. housing equity	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Median retire / illiq. wealth	0.000	0.000	-0.319	0.000	0.000	0.000	0.000	0.000
Mean retire / illiq. wealth	-0.575	-0.407	-0.511	-7.095	-0.201	0.000	0.000	0.000
Frac. pos. retirement account	0.068	0.028	0.090	0.011	0.027	0.000	0.000	0.000
W-HtM								
	US	CA	AU	UK	DE	FR	IT	ES
Median liq. wealth / income	0.000	-0.270	0.000	-0.314	0.085	0.123	0.094	0.078
Mean liq. wealth / income	-1.494	-0.937	-1.974	-2.347	-0.773	-0.989	-0.303	-0.950
Frac. neg. liquid wealth	0.502	0.607	0.495	0.606	0.312	0.281	0.285	0.209
Median housing / illiq. wealth	0.799	0.727	0.467	0.301	0.000	0.000	0.999	1.000
Mean housing / illiq. wealth	0.586	0.576	0.411	0.374	0.328	0.377	0.794	0.863
Frac. pos. housing equity	0.794	0.809	0.583	0.670	0.433	0.449	0.816	0.940
Median retire / illiq. wealth	0.073	0.255	0.508	0.674	0.000	0.000	0.000	0.000
Mean retire / illiq. wealth	0.342	0.416	0.575	0.606	0.102	0.008	0.028	0.001
Frac. pos. retirement account	0.565	0.657	0.957	0.869	0.309	0.025	0.099	0.032

Table 5: Portfolio characteristics by country and HtM status.

ation accounts as a component of wealth, we find that the fraction of W-HtM and P-HtM in Australia are 12.5, and 10.3 respectively.

Second, the much lower of fraction of both W-HtM and P-HtM households in the European countries compared with the US, Canada and the UK implies that households in these European countries hold more liquid wealth relative to their income. Table 5 shows that this difference can be in part traced to differences in lower holding of liquid debt. The fraction of W-HtM households with negative liquid wealth in the European countries ranges from 21% in Spain to 31% in Germany, compared with 50% in the US and Australia and 61% in the UK and Canada. Presumably tighter limits on unsecured credit in Europe mean that there is more incentive for households in these countries to hold liquid wealth for transaction and precautionary reasons. In Italy, which has the highest fraction of N-HtM households of these eight countries, the fraction of P-HtM households with negative liquid wealth is only 5%,

which is by far the lowest in our data. This suggests unsecured credit is less easily substituted for liquid wealth in Italy, and to a lesser extent in Germany France and Spain, than in the Anglo-Saxon countries.

Some of the differences in portfolio holdings across countries are evident in the age profiles of the fraction of households of each HtM type. These are shown in Figure 6. The US, Canada and the UK all have a very similar breakdown of the HtM groups by age: the fraction of P-HtM households declines monotonically with age, while the fraction of W-HtM households is humped shape with a peak in the early- to mid-forties. The age profile for Australia differs in that many more young people are W-HtM, and many less are P-HtM, compared with the other Anglo-Saxon countries. This is directly related to the high fraction of people holding retirement accounts in Australia, even for very young households. The European countries have very similar age profiles of HtM status to the US, Canada, with roughly parallel shifts downward due to the higher fraction of N-HtM households in these countries. The hump shaped age profile for the W-HtM is also less pronounced for the European countries.


[TO BE COMPLETED]

7 Implications for macroeconomic modelling and policy analysis


[TO BE COMPLETED]

8 Concluding remarks


[TO BE COMPLETED]




(a) US




(b) CA




(c) AU




(d) UK




(e) DE



(f) FR



(g) IT



(h) ES

Figure 6: Age profile of fraction HtM by country

References

- 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.
- BLACK, O. (2011): “Wealth in Great Britain, Main Results from the Wealth and Assets Survey 2008/10,” *Office for National Statistics*.
- BRICKER, J., B. BUCKS, A. KENNICKELL, T. MACH, AND K. MOORE (2011): “Surveying the aftermath of the storm: changes in family finances from 2007 to 2009,” Discussion paper, Board of Governors of the Federal Reserve System (US).
- 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.
- CAMPBELL, J. Y., AND N. G. MANKIW (1990): “Permanent income, current income, and consumption,” *Journal of Business & Economic Statistics*, 8(3), 265–279.
- CAMPBELL, J. Y., AND N. G. MANKIW (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.
- CHETTY, R., AND A. SZEIDL (2007): “Consumption Commitments and Risk Preferences,” *The Quarterly Journal of Economics*, 122(2), 831–877.
- 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*.
- DAVID LAIBSON AND ANDREA REPETTO AND JEREMY TOBACMAN (2003): “A Debt Puzzle,” in *Knowledge, Information, and Expectations in Modern Economics: In Honor of Edmund S. Phelps*, ed. by Philippe Aghion and Roman Frydman and Joseph Stiglitz and Michael Woodford, pp. 228–266. Princeton University Press.
- 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.
- 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 (2011): “Can Tax Rebates Stimulate Spending in a Life-Cycle Model?,” *Mimeo*.
- 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 (2013a): “Household Leveraging and Deleveraging,” NBER Working Papers 18941, National Bureau of Economic Research, Inc.
- (2013b): “Household Leveraging and Deleveraging,” *NBER Working Papers*, 18941.
- KAPLAN, G., AND G. VIOLANTE (2012): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Mimeo*.
- KENNICKELL, A. B. (2005): “The Good Shepherd: Sample Design and Control for Wealth Measurement in the Survey of Consumer Finances,” *Mimeo*.
- (2009): “Ponds and Streams: Wealth and Income in the US, 1989 to 2007,” *Mimeo*.
- KRUSELL, P., AND A. SMITH, JR. (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5), 867–896.
- 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.

- 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.
- 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*.
- 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.