NBER WORKING PAPER SERIES

WHO ARE THE HAND-TO-MOUTH?

Mark A. Aguiar
Mark Bils
Corina Boar

Working Paper 26643
http://www.nber.org/papers/w26643

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2020

We thank Nataliya Gimpelson, Stephan Gordeev, and Federico Kochen for excellent research assistance on this project. We also thank Martin Blomhoff Holm, Matthew Knowles, Gianluca Violante, and Ivan Werning for their helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Who Are the Hand-to-Mouth?
Mark A. Aguiar, Mark Bils, and Corina Boar
NBER Working Paper No. 26643
January 2020
JEL No. E21

ABSTRACT

Many households hold little wealth, especially liquid wealth. In precautionary savings models, absent preference heterogeneity, these households should display not only higher marginal propensities to consume (MPCs), but also lower average propensities to consume (APCs) and higher future consumption growth. We see from the PSID that such "hand-to-mouth" households actually display higher APCs and no faster spending growth. They also adjust spending to a greater extent through the number of categories consumed. Consistent with a role for preference heterogeneity, the panel data show that it is the propensity to be hand-to-mouth, not current assets, that predicts high APC, low consumption growth, and other spending differences for the hand-to-mouth. To identify the extent of preference heterogeneity, we consider the model of Kaplan and Violante (2014) with both liquid and illiquid assets, but allow heterogeneity in preferences. To match the data, the vast majority of poor hand-to-mouth must be impatient and have a high intertemporal elasticity of substitution (IES). The richer, but illiquid, hand-to-mouth are disproportionately high IES, though not impatient. Thus a high IES is a key determinant of assets for households typically viewed as hand-to-mouth. The model additionally shows that preferences play a prominent role in differences in MPCs across consumers.

Mark A. Aguiar
Department of Economics
Princeton University
Fisher Hall
Princeton, NJ 08544-1021
and NBER
mark@markaguiar.com

Mark Bils
Department of Economics
University of Rochester
Rochester, NY 14627
and NBER
mark.bils@rochester.edu

Corina Boar
Department of Economics
New York University
19 West 4th Street
New York, NY 10003
and NBER
corina.boar@gmail.com

1 Introduction

This paper explores why some agents hold low levels of wealth relative to income. The question is of interest in its own right, as the answer deepens our understanding of micro consumer behavior. Moreover, it also sheds light on an important component in the design of optimal macro-policy. These so-called "hand-to-mouth" agents are often thought to have a relatively large response to transfers and a muted response to changes in interest rates, and thus feature prominently in the discussion of tax and transfer schemes to boost aggregate output as well as the efficacy of monetary policy.[1] To the extent that macro-policy analysis relies on "getting the micro foundations right," it is crucial to correctly understand why so many consumers find themselves with limited net worth or liquidity relative to income.

The core paradigm of both the micro and macro literatures is the incomplete-markets precautionary saving model in which consumers self-insure against idiosyncratic income risk by saving in a non-contingent asset subject to a borrowing constraint.[2] This model has some key predictions that will guide our empirical work: (i) the consumption policy function is strictly concave in wealth, and hence the marginal propensity to consume (MPC) is strictly decreasing; (ii) the average propensity to consume (APC), which is defined as current consumption over total current income, is increasing in wealth, as low-wealth agents are actively saving away from the constraint; and (iii) expected consumption growth is decreasing in wealth, as low-wealth agents are either constrained or in the process of building up their buffer stock of savings.

These predictions are comparative statics for fixed preferences as well as a given income process and rate of return. If one allows the possibility of preference heterogeneity across households, then these empirical predictions may be obscured, as a household's asset position will reflect its preferences as well history of shocks. Using a simple numerical parameterization, we show that, all else equal (including wealth), MPCs and APCs are decreasing in the discount factor and increasing in the inter-temporal elasticity of substitution (IES). In particular, a relatively impatient consumer at a given level of wealth consumes a larger fraction of income (higher APC) and consumes more of any added income (higher MPC). Similarly, given a discount factor that is less than the inverse gross interest rate, a con-

[1]Many papers have modeled low-asset households as responding more to fiscal policies. Recent examples include Kaplan and Violante (2014), Jappelli and Pistaferri (2014), Farhi and Werning (2017), McKay and Reis (2016), Carroll, Slacalek, Tokuoka and White (2017), Kaplan, Moll and Violante (2018), and Auclert (2019). A number of these papers also make clear that hand-to-mouth agents will display a smaller direct intertemporal substitution response to interest rates, though they may respond more through indirect channels, such as the income and wealth effects, resulting from interest rate changes (e.g., Auclert (2019), Kaplan et al. (2018)).

[2]Just a few of the many papers in this vein are Schechtman and Escudero (1977), Imrohoroglu (1989), Deaton (1991), Carroll (1992), Huggett (1993), and Aiyagari (1994).

sumer with a larger IES has a greater propensity to front-load consumption, and thus also has higher marginal and average propensities to consume. Moreover, all else equal, expected consumption growth is increasing in the discount factor and decreasing in the IES. Therefore, for data composed of consumers with heterogeneous preferences, the relationship between wealth and consumption behavior is muddied. Using the standard model, we argue that measures beyond the household's current asset position, including its measured APC and especially its longer-term propensity to be at low assets, are useful for identifying the role of preference heterogeneity.

Guided by these predictions, we explore empirical consumption behavior using the Panel Study of Income Dynamics (PSID). Our first order of business is to identify in the data those consumers that are likely to behave like the low-wealth agents in the model. Following Zeldes (1989), we define low-net worth households as potentially "hand-to-mouth," denoted $H2M_{NW}$. An alternative subsample of interest are higher net worth individuals that have negligible or negative liquid assets, the group Kaplan, Violante and Weidner (2014) refer to as the "wealthy hand-to-mouth," which we denote $H2M_{LIQ}$. On average, we label about 40 percent of households as hand-to-mouth based on either low net worth or low liquid wealth, with a little over half of these classified as low net-worth.

We employ nine waves of the PSID, following households' income, assets, and spending for a period of up seventeen years. In particular, we can track how a household's hand-to-mouth status evolves over time. That allows us to relate a household's spending behavior separately to its current hand-to-mouth status as well as its propensity to appear hand-to-mouth in other surveys.

Controlling for demographics, we find that the hand-to-mouth, by either measure, do not have higher realized future consumption growth than those with higher net worth or liquidity, despite having higher future income growth and potentially being financially constrained. However, once we include a fixed effect, the low-net-worth display a significantly higher rate of consumption growth going forward, consistent with the basic model.

The importance of including a fixed effect appears throughout the empirical work, speaking to the importance of underlying heterogeneity. To get a sense of why the fixed effect plays a significant role, we measure the fraction of waves in the PSID sample in which each household is $H2M_{NW}$ or $H2M_{LIQ}$. Dropping the fixed effect and adding this frequency control yields estimates that are close to the fixed-effects estimates, suggesting that the heterogeneity that is important for consumption growth projects on the frequency of the hand-to-mouth status. Accordingly, conditional on current $H2M$ status, those that are frequently hand-to-mouth have significantly lower consumption growth on average, while conditional on type (as measured by frequency), current $H2M$ status predicts faster consumption growth.

Viewed through the lens of the standard model, the heterogeneity in consumption growth rates could reflect preference heterogeneity, differences in income risk, or differences in rates of return on financial assets. We find that a large majority of $H2M$ households, especially those $H2M_{NW}$ have considerable liquid debt such as credit card debt. The return on discharging such debts should be quite high. For this reason, we view a low marginal return on savings as the least promising avenue for explaining the low savings rates for low-asset households.

To explore differences in income and consumption volatility, we regress the absolute value of the change in log income and log consumption on lagged hand-to-mouth status. We find that the those households prone to being hand-to-mouth actually exhibit higher income and consumption volatility. Going further, we show more precisely that the frequently $H2M$ exhibit greater transitory income risk. Such risk calls for greater precautionary saving, behavior opposite to what we see in our data.

The standard model's prediction that low-wealth agents have a higher rate of savings is also not present in the data. Specifically, the $H2M_{NW}$ have a higher average propensity to consume out of income. We show this is a persistent trait of those prone to hand-to-mouth status, indicating that the low net-worth households (relative to income) are not temporally in that state due to bad luck, but rather appear to have a low target wealth-to-income level.

Perhaps the most relevant prediction of the standard model for macro policy is that the hand-to-mouth have a higher marginal propensity to consume. We regress the growth in expenditures on the growth in income (and demographic controls) to estimate a marginal propensity to consume out of income (MPCY); the MPCY is the standard MPC plus the impact of current income changes on anticipated future income draws. We estimate a higher MPCY for low-wealth households. Most striking, it is a household's long-run propensity to be hand-to-mouth, not its current status that drives the differences in MPCYs. Household that are frequently hand-to-mouth, either based on net worth or liquid wealth, display higher MPCYs. By contrast, conditioning on those frequencies, currently being hand-to-mouth does not predict a higher MPCY. These findings are consistent with the standard model's prediction that households with either a low discount factor or a high IES will exhibit higher marginal propensities to consume, even if their assets are not below their target level.

The preceding results focused on the aggregate basket of consumption expenditures. The PSID and the Consumer Expenditure Survey (CE) collect expenditure on disaggregated components of expenditure. Using this data, we explore whether hand-to-mouth households allocate their nondurable expenditures differently across categories. We find that, controlling for total expenditures, the hand-to-mouth consume fewer categories. However, in response to a change in expenditure, the hand-to-mouth are much more prone to adjust the number of goods consumed. That is, the hand-to-mouth are much more elastic on the *extensive* *margin* than the non-hand-to-mouth. As with the other spending regularities, it is a household's propensity for being hand-to-mouth, not its current status, that predicts this greater responsiveness in the number of categories.

The fact that the extensive margin is the relevant margin of adjustment for the hand-to-mouth is consistent with a relatively elastic response of total expenditure to movements in inter-temporal prices; that is, a higher IES. We show this formally in a simple example in the appendix, but the intuition is familiar from the role the extensive margin plays in labor supply elasticities (Rogerson, 1988) or risk-preferences in portfolio choice (Grossman and Laroque, 1990, Chetty and Szeidl, 2007). A high IES is also theoretically consistent with the low target net-worth of the hand-to-mouth and the excess volatility of consumption relative to income found in the data.

The empirical facts documented in the paper make a compelling case that preference heterogeneity plays a role in determining hand-to-mouth status. A natural assumption is that those prone to low assets are relatively impatient. However, we show that in the standard one-asset model, a high IES can lead to very similar policy functions as impatience. Moreover, the extensive margin analysis suggests differences that potentially relate to the IES rather than discount rates. To identify these parameters separately, we consider the two-asset framework of Kaplan and Violante (2014) (KV) in which agents allocate wealth between liquid and illiquid assets. As a higher IES implies a willingness to adjust the timing of consumption, more elastic consumers have a lower demand for liquidity all else equal. Using Epstein-Zin preferences, we distinguish this from risk aversion.

We allow for four types of agents, defined by two discount factors (0.90 and 0.95 on an annualized basis) and two magnitudes for the IES. (0.5 and 1.5). We calibrate the respective shares of each type to match the average ratios of net-worth to income and illiquid assets to income, as well as the share of consumers that have low net worth or low liquidity. To match these targets, the model requires three-quarters of the population to have the higher discount factor and a low IES. However, the vast majority of the low-net-worth hand-to-mouth are both relatively impatient and have the higher IES. That this type constitutes over two-thirds of low net-worth individuals is striking, given their calibrated share of the overall population is only one-sixth. Patient but elastic consumers are rare in the overall population, but comprise a quarter of the wealthy but illiquid sub-population. A key takeaway is therefore that the hand-to-mouth (defined by either net worth or liquidity) tend to be relatively elastic with respect to inter-temporal substitution. We also find no role for impatient and inelastic agents; that is, the poor hand-to-mouth almost exclusively reflects the combination of impatience and high inter-temporal elasticity, while a high IES, not a low discount factor (or risk aversion), is most important for generating wealthier households without liquidity.

We find that preference heterogeneity plays a major role in differences in MPCs across consumers for our calibrated model. The poor hand-to-mouth in our two-asset KV model have a high MPC, not only because they have low wealth, but also because they are much more likely to be low-discount-factor/high-IES consumers. The wealthy hand-to-mouth are also relatively high-IES consumers, contributing to a higher MPC. Moreover, their high IES implies that these agents are willing to bear the volatility associated with extreme illiquidity, and therefore have a high MPC for the familiar reason of being hard up against a constraint (or at a kink) in their borrowing.

Finally, we show that the effective IES of hand-to-mouth households can be as high as for other households. While the hand-to-mouth are disproportionately at a borrowing constraint (or kink) muting the intertemporal response, those not literally at a constraint will be especially responsive. This is a compositional effect–being hand-to-mouth largely reflects having a high IES. Thus, except when they are literally at a constraint, these agents will be more sensitive to the intertemporal terms-of-trade implications of monetary policy as well as such policies as cash-for-clunkers or sales-tax holidays. It also implies that delaying insurance, such as until social security eligibility, may not be as costly for the hand-to-mouth.

Our paper intersects with an enormous empirical literature on how spending responds to income. Havranek and Sokolova (2019) reference many of these studies, as they perform a meta analysis of 144 such studies. Our estimates of spending responses to income perhaps most parallel results in Fisher, Johnson, Smeeding and Thompson (2019), who also examine spending responses to income stratifying PSID households by their assets. Our empirical work differs from the many papers in this literature, in that we focus on a much broader set of empirical regularities, including the spending growth, spending volatility, and spending allocation of low-asset households. More importantly, we stress from the panel data that the key predictor for a household's spending is not its current assets, but rather the longer-run positioning of its assets.

Our work also intersects with a large literature identifying preference heterogeneity. Recent examples include Parker (2017), Gelman (2019), Athreya, Mustre-del-Río and Sánchez (2019), and Calvet, Campbell, Gomes and Sodini (2019). Our work is especially complementary to Calvet et al. (2019), who also find support for heterogenity in both the IES and discount factors across a sample of Swedish households.[3]

[3]They draw this conclusion from the heterogeneity in how households reduce savings as the need for precautionary savings decline.

2 H2M in the Canonical IM Model

To motivate our empirical exploration of the hand-to-mouth, we review the canonical savings model in which agents use a non-contingent asset to smooth idiosyncratic income fluctuations. In Section 5, we extend the model to include both liquid and illiquid assets along the lines of Kaplan and Violante (2014); for the current motivational section, we present the standard single-asset environment.

Specifically, suppose agents are infinitely lived and receive a stochastic endowment y_t that follows a first-order Markov chain on support $\{y_1, ..., y_N\}$, with $0 < y_1 < ... < y_N < \infty$. Preferences over consumption streams are given by:

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t u(c_t),$$

where $\beta \in (0, 1)$ and the expectation is conditional on some initial state. We assume u takes the CRRA form:

$$u(c) = \begin{cases} \frac{c^{1-\gamma}}{1-\gamma} \text{ if } \gamma \neq 1 \\ \ln c \text{ if } \gamma = 1. \end{cases}$$

The parameter $\sigma \equiv 1/\gamma$ is the inter-temporal elasticity of substitution (IES). The IES and discount factor will be the focus of the comparative statics below. With power utility, the IES and the coefficient of relative risk aversion are reciprocal pairs. Our language emphasizes the IES, as the extension presented in Section 5 with Epstein-Zin preferences and liquid and illiquid assets indicates that the IES is a key determinant of liquid wealth holdings.

The agent has access to a non-contingent savings vehicle that has gross return $R = 1+r < \beta^{-1}$. The environment is partial equilibrium and we take R to be a primitive of the model. The level of assets is restricted to be above some threshold $\underline{a} \leq 0$. At the start of period t, the agent has resources ("cash on hand") $x_t \equiv Ra_t + y_t$ to allocate to consumption and next period's assets a_{t+1}. The agent's problem can be expressed in recursive form by:

$$V(x, y) = \max_{c, a' \geq \underline{a}} u(c) + \beta \mathbb{E} \left[V(Ra' + y', y') \mid y \right]$$
$$\text{subject to } a' + c \leq x.$$

Let $\mathcal{C}(x, y)$ denote the associated optimal consumption function. Letting $c_t = \mathcal{C}(x_t, y_t)$, consumption satisfies the usual Euler Equation:

$$\mathbb{E}_t \left[\beta R \left(\frac{c_{t+1}}{c_t} \right)^{-\frac{1}{\sigma}} \right] \leq 1, \tag{1}$$

with equality when $a_{t+1} > \underline{a}$.

We define the *marginal propensity to consume* (MPC) as $\partial \mathcal{C}/\partial x$. As is well known (see Carroll and Kimball, 1996), in this environment $\mathcal{C}$ is a strictly increasing and concave function of x. Hence, the MPC is well-defined almost everywhere and decreasing in the level of assets.

A voluminous literature uses this fact to proxy MPC with some measure of wealth (or liquid wealth). A complication is that the consumption function is defined for a particular utility function, borrowing constraint, and process for income. It is not generally the case that the MPC is monotonic in assets when comparing across people with distinct preferences at a point in time, or, if the process for income changes, across time for the same person.

More generally, consider an individual with relatively low assets compared to the population average. Such a low asset position may be a consequence of a string of low income realizations. Alternatively, the individual may have preferences that imply a low target level of assets. A potentially useful concept to sort between these forces is the *average propensity to consume* (APC) out of income:

$$APC(x, y) \equiv \frac{\mathcal{C}(x, y)}{ra + y},$$

which is well defined whenever $ra + y > 0$.[4]

To understand the mapping between APC and resources, we can rewrite the flow budget constraint as:

$$\frac{a' - a}{ra + y} = 1 - APC,$$

where a is start-of-period assets and a' is next period's asset position. Hence, assets are accumulating when $APC < 1$. That is, whether $APC \gtrless 1$ is a signal that the agent's asset position is above or below their "target".

If asset dynamics are monotone (that is, $a' - a$ is decreasing in a given y) and $r \approx 0$, then APC is increasing in a for a given level of income. That is, APC is monotonically related to x conditional on income, and hence to MPC. In many datasets that include consumption

[4]In the simulations below, we assume $r\underline{a} + y > 0$, so the APC is always well defined. In the empirical work, we do encounter negative denominators, but not zero. However, only very few households have a negative APC and we exclude them from the sample.

expenditures, we have arguably more reliable measures of income than wealth; so APC, being normalized by flow income, is a useful empirical proxy for x and MPC.

Given $\beta R < 1$, the stationary process for y_t, and the specification of preferences, there exists a unique ergodic distribution for (a_t, y_t), denoted $\lambda(a, y)$.[5] Integrating the budget constraint over the distribution λ, we define the *target wealth* a^* by:

$$a^* = \frac{1}{r}\left(\int \mathcal{C}(Ra + y, y) d\lambda - 1\right).$$

That is, a^* is the mean of the ergodic distribution for assets, integrating the ergodic distribution over both assets and income. Similarly, we can let $x^* \equiv Ra^* + \bar{y}$ denote the ergodic mean of x_t, where $\bar{y}$ is the ergodic mean of the endowment process.

To present the key predictions of the model used in the empirical analysis, we solve the model numerically. Specifically, following Krueger, Mitman and Perri (2016), we postulate that income evolves according to:

$$\begin{aligned} \ln y_t &= z_t + \varepsilon_t \\ z_t &= \rho z_{t-1} + \eta_t, \end{aligned}$$

where $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$, $\eta_t \sim N\left(0, \sigma_\eta^2\right)$ and we set $\rho = 0.97$, $\sigma_\eta = 0.20$ and $\sigma_\varepsilon = 0.23$.

We set $R = 1.04$ and consider two alternative values for $\beta \in \{0.90, 0.95\}$. We also let the IES take on two possible values $\sigma \in \{0.5, 1.5\}$.[6] Note that for a given individual, preferences are fixed, but, given the two-dimensional heterogeneity, there are four possible preference configurations. We load the heterogeneity on preference parameters, but the alternative discount factors can also proxy for (permanent) differences in financial returns. That is, if an individual has access to a high-return savings vehicle that another individual with the same preferences lacks, then βR will differ across the two in the same fashion as a difference in discount factors. A separate issue is if returns vary by scale, say, due to a fixed cost of access. In this case, the level of assets will also reflect variation in returns, mitigating the negative impact of wealth on expected consumption growth described below. We discuss this point in more detail in the empirical section.

In Figure 1 Panel (a), we plot $C(., \bar{y})$ against x, where $\bar{y}$ is the ergodic mean of the income process. The fact that y is held constant implies that a is varying as we move along the horizontal axis. The four lines represent the four alternative preference parameterizations. The large dot represents the ergodic mean x^* for each set of preferences.

[5] See Açikgöz (2016).

[6] Havrnek (2015) conducts a meta-study of IES estimates in the literature, reporting a median of 0.5. Kaplan and Violante (2014), citing the asset-pricing literature, choose an IES of 1.5.

Figure 1: The Consumption Policy Function

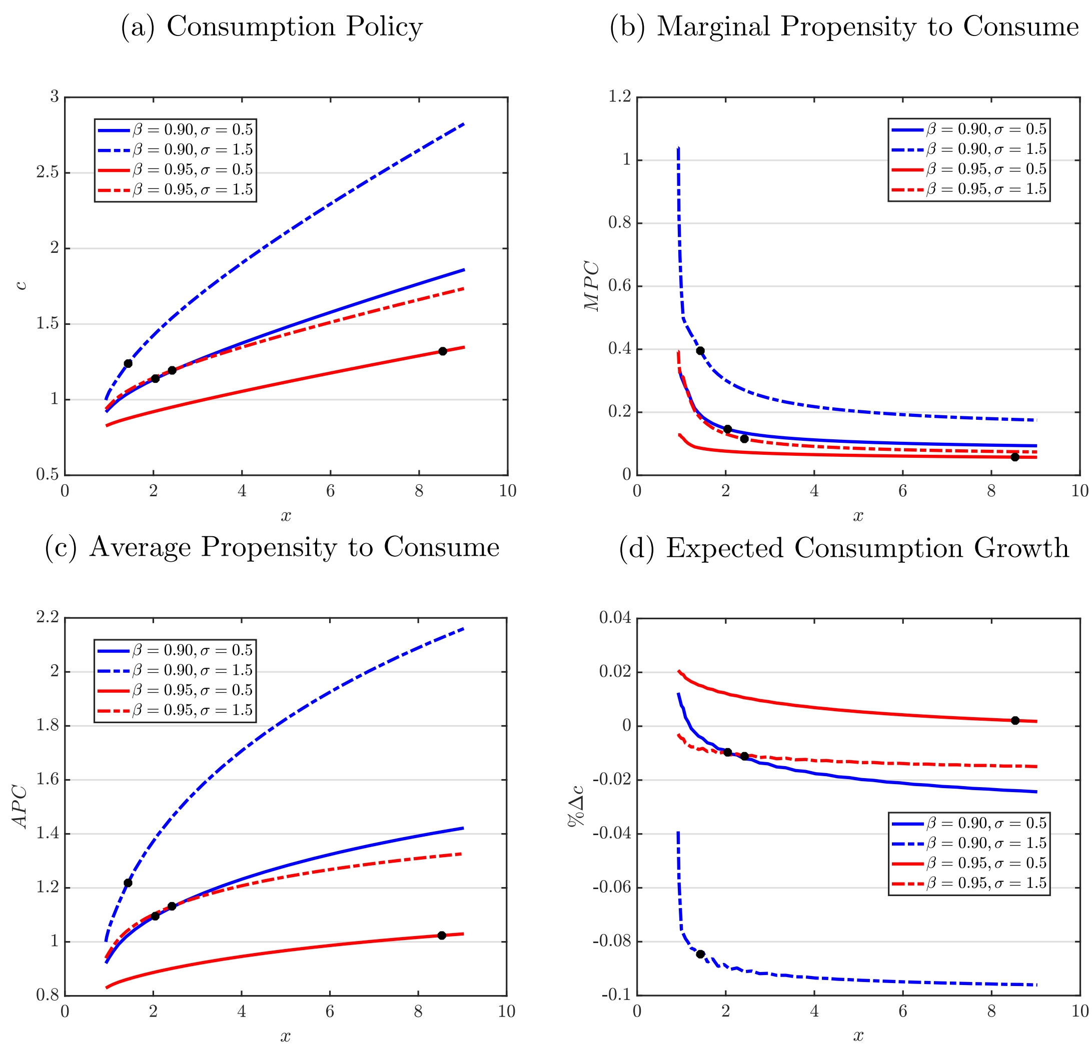

Note: Panel (a) depicts $\mathcal{C}(\cdot, y)$ as a function of x, with y set to its mean value in all panels. Panel (b) depicts $\partial\mathcal{C}(x, y)/\partial x$ as a function of x. Panel (c) depicts $\mathcal{C}(x, y)/(y + ra)$ as a function of x. Panel (d) depicts $\mathbb{E}\ln\left(\frac{\mathcal{C}(x', y')}{\mathcal{C}(x,y)}\right)$ as a function of x, where the expectation is over y' conditional on $y = \overline{y}$ with $x' = Ra' + y'$ and $a' = x - \mathcal{C}(x, a)$.

The consumption functions are strictly increasing and concave. For a given level of resources, consumption is decreasing in β and increasing in the IES. The reason that consumption is increasing in the IES is that $\beta R < 1$ in all specifications. Hence, there is an incentive to front-load consumption, and this incentive increases with the inter-temporal elasticity of substitution.

Note that two alternative preference specifications yield very similar consumption functions and associated target assets. Specifically, the case of $(\beta = 0.9, \sigma = 0.5)$ and $(\beta = 0.95, \sigma = 1.5)$ track each other closely. This highlights that relative impatience and a high elasticity both push towards a low target level of assets.

This analysis suggests (at least) four possibilities for why an individual has low levels of assets. One is the traditional model of uniform preferences, in which a low-wealth agent, due to a sequence of low income draws, is temporarily below their target wealth. The second, also familiar from the literature, is that the agent is relatively impatient. The third, which is perhaps less familiar, is that the agent has a high elasticity of inter-temporal substitution (assuming that $\beta R < 1$). A fourth possibility is differences in the income process, particularly a high anticipated growth rate or low volatility, which reduces the demand for precautionary savings. The latter three generate low wealth positions due to low target wealth, while the first is due to "bad luck." Identifying these four forces in the data will be one focus of the empirical work.

In Panel (b), we plot $MPC(., \bar{y})$ against x. As already discussed, for each preference specification, the MPC is strictly decreasing in x, reflecting the concavity of $\mathcal{C}$. Looking across preferences at a given level of resources, the MPC is decreasing in β and increasing in the IES.

For an intuition for why MPC is so sensitive to preferences, consider an increase of x in period t holding constant y. To map out the response of consumption, consider the unconstrained Euler Equation $\mathbb{E}_t \left\{ (\beta R)^k (c_{t+k}/c_t)^{-\gamma} \right\} = 1$ for $k \geq 1$. A small increase in x given y can be accommodated by increasing c_{t+k} proportionally at all $k \geq 0$. While this new allocation satisfies the Euler Equation wherever it holds, it is not a true optimal response as it ignores the fact that in some states the Euler Equation may not be satisfied with equality. Nevertheless, it is a useful thought exercise to provide some intuition. As the increase is proportional in every period, the initial level change in consumption (the discrete version of the MPC) is increasing in the initial level of consumption. For a given x, the more front loaded is consumption, the greater the MPC. Hence, for a given x, the MPC is increasing in the IES and decreasing in β, as depicted in the figure.

Interestingly, for low β preferences, the IES has a fairly big impact on the MPC for asset levels near the target. For the relatively patient specifications, the impact on MPC is primarily intermediated through the large difference in target wealth. This reflects that as $\beta R \to 1$, there is less of an incentive to inter-temporally substitute consumption, and the IES becomes less relevant in determining steady state assets.

Panel (c) of Figure 1 depicts the APC as a function of x. For fixed preferences, this relationship is monotonic.[7] However, for a given x, there is substantial variation across preference specifications. Recall that an APC of one implies stationary assets given an income level; hence, the APCs at the respective mean assets are clustered near one (it is not exact as APC is a nonlinear function of x, and also because x^* is integrated over the ergodic distribution of y as well as x). We will use the fact that comparing individuals across wealth positions for a given APC identifies variation due to preferences.

Panel (d) plots the expected consumption growth as a function of x, still for $y = \overline{y}$.[8] Expected log consumption is computed by integrating over the policy function for each possible draw of next period's earnings, with transition probabilities conditional on mean y today and next period assets governed by the policy function associated with $(x, \overline{y})$. For intuition, if consumption growth is approximately log-normally distributed, we can log-linearize the Euler Equation (1) as:

$$\mathbb{E}_t \Delta \ln c_{t+1} \gtrapprox \sigma \ln(\beta R) + \frac{1}{2\sigma} Var_t(\Delta \ln c_{t+1}). \tag{2}$$

This suggests that a consumer will have relatively large expected consumption growth if (i) they are constrained; (ii) they are relatively patient (high βR); (iii) they have a relatively low IES (low σ, assuming $\beta R < 1$); and/or (iv) they have a relatively large demand for precautionary savings (a large conditional variance of consumption growth scaled by risk aversion $1/\sigma$).

Taking stock, the standard model predicts that, for a given set of preferences, low-wealth households should exhibit higher MPCs, higher expected rates of consumption growth, and lower APCs. But, these predictions are sensitive to preference heterogeneity. The usual approach is to assume ex ante identical individuals. In that environment, low wealth is due to the idiosyncratic history of income realizations. Moreover, there is a tight mapping between resources on hand and MPC. Allowing for preference heterogeneity muddies the picture. In the remainder of the paper, we present a series of empirical exercises to elucidate what determines who has low wealth.

[7] The relationship is monotonic for the mean income level, which is the case depicted. However, for the lowest income states we do find that the APC eventually declines in a as a becomes very large. This reflects the presence of a in the denominator of the APC.

[8] By definition, the average consumption growth integrated over the ergodic distribution of x and y is zero. Integrated over income states, the average expected consumption growth at x^* is much closer to zero.

3 Data

Our empirical work is primarily conducted on the PSID, employing its biennial surveys from 1999 to 2015. 1999 was the onset of the PSID measuring wealth in each survey. It also initiated the PSID including spending more broadly than on food and housing. The data appendix discusses our variable constructions and sample restrictions in detail. Here we highlight the key variables for our analysis of earnings, income, wealth, and expenditures.

In the next section we identify households as hand-to-mouth, following Zeldes (1989) and Kaplan et al. (2014), by assets relative to a measure of earnings. Our earnings measure equals labor income, net of payroll taxes, plus government transfers received. We also consider a broader measure of after-tax income, for instance for calculating a household's APC. After-tax income is the sum of earnings and transfer income, net profits from business or farm, and net income from assets, minus the family's federal and state income tax liabilities calculated by TAXSIM. For homeowners we include 6 percent of the home value as implicit rent, while subtracting associated property taxes, mortgage interest, and home insurance. Our division of assets by liquidity largely follows Kaplan et al. (2014). For liquid net worth we sum checking and savings balances, money market funds, certificates of deposit, treasury bills, and stocks outside of pension funds, while subtracting debts in the forms of credit and store cards, student loans, medical or legal debt, and debt owed to family. Illiquid assets reflect home and other real estate equity, IRA/pension holding, non-government bonds, insurance equity, and the net value of any business, farm, or vehicles. All nominal variables are converted to 2009 CPI-deflated dollars.

Measured expenditures include shelter, utilities, food, gasoline, health insurance and medical expenses, education, child care, public transportation, and vehicles spending for purchases, repairs, insurance and parking. Spending on shelter equals rent payments for renters; for homeowners we set it to 6 percent of respondent's valuation of the home. Summing categories, expenditures average 58.3 percent of after-tax income for our sample. Unless stated otherwise, our results reflect spending on all these categories. But our results are robust if we exclude spending on the durable categories of vehicle purchases and repair.

Our sample reflects only the PSID's nationally representative core sample, including its "split-off" families and PSID sample extensions to better represent dynasties of recent immigrants. (All results reflect PSID longitudinal family weights.) We restrict our sample to households with heads ages 25 to 64 and for which we can measure hand-to-mouth status from wealth and earnings for at least three surveys. We exclude households with less than $2,000 in annual earnings plus transfers, after-tax income, or expenditures. The data appendix details the impact of these sample restrictions.

4 Empirical Results

4.1 Identifying the Hand-to-Mouth

As depicted in Figure 1, consumption in the standard model is highly non-linear at low levels of resources, with the MPC quickly declining in wealth and then "flattening out" at higher asset positions. Various measures have been introduced to identify which individuals in the data are likely to be on their highly non-linear regions of their consumption functions. The early and influential paper by Zeldes (1989) distinguished individuals by net worth. Specifically, Zeldes considered an individual "constrained" if their net worth was less than two months of labor earnings. Following Zeldes, we therefore define an agent as hand-to-mouth based on net worth (denoted $H2M_{NW}$) if their net worth is less than two months labor earnings.

Kaplan et al. (2014) (henceforth KVW) pursue an alternative measure of constraints. These authors focus on liquidity rather than wealth. In particular, they divide assets into net illiquid wealth versus net liquid wealth. They define an individual as constrained if liquid wealth is close to a borrowing limit or if liquid wealth is close to zero. Specifically, they define the borrowing limit of an individual to be 18.5% of annual earnings. An individual is considered constrained if they have a negative liquid net worth position within one week of earnings of the borrowing limit, or if they have positive liquid wealth but the position is less than one week of earnings. The latter criteria is designed to identify individuals at a "kink" in the budget set near zero liquid assets due to the difference between borrowing and saving interest rates. Note that KVW's definition focuses only on liquid net worth, and is designed to include agents with high levels of total net worth (the "wealthy hand-to-mouth"). We therefore assign an individual to be wealthy hand-to-mouth (denoted $H2M_{LIQ}$) if they are not $H2M_{NW}$, but have liquid net worth that satisfies the KVW criteria.

In our PSID sample, an average of 40.2% of the households (pooling across waves) are hand-to-mouth, with 22.7% denoted $H2M_{NW}$ and 17.5% denoted $H2M_{LIQ}$. That is, 17.5% of the sample is liquidity constrained according to the KVW definition, but have sufficient total net worth to be considered unconstrained by the Zeldes measure. As Table 1 shows, 16.6% of the sample is both net worth and liquidity constrained, which we assign to the low-net-worth $H2M_{NW}$ category; that is, 74% of the $H2M_{NW}$ group would also satisfy the KVW liquidity-constrained definition.[9] We also constructed these shares for the seven waves of the

[9] For comparison, Zeldes classified 29% of his (earlier) PSID sample as hand-to-mouth using his net-worth definition. KVW classify roughly 31% of their Survey of Consumer Finance sample as liquidity constrained, compared to 34.1% in our PSID sample (spread over both measures). KVW classify 20% percent as "wealthy hand-to-mouth," compared to 17.5% of our PSID sample.

Survey of Consumer Finance (SCF) from 1998 to 2016. The respective household shares for not hand-to-mouth, $H2M_{NW}$, and $H2M_{LIQ}$ are 62.5%, 25.0%, and 12.5%, similar to our counts from the PSID.

Table 1: Hand-to-Mouth Groups

	Not $H2M$	$H2M_{NW}$	$H2M_{LIQ}$
Shares	59.7%	22.7%	17.5%

	By LIQ (KVW)	
By NW (Zeldes)	Not $H2M$	$H2M$
Not $H2M$	59.7%	17.5%
$H2M$	6.1%	16.6%

4.2 Characteristics of the H2M

Table 2 provides some summary statistics for the hand-to-mouth. Specifically, we compute statistics based on whether a household is designated as one of our H2M measures in a given year. This implies that the same household may be represented in multiple columns, albeit in different waves of the survey.

First compare the hand-to-mouth based on net worth ($H2M_{NW}$) in the second column to those not hand-to-mouth in the first column. These hand-to-mouth are younger on average by 7 years, their earnings and incomes are only half as much, and of course their wealth, both liquid and illiquid is much lower. Turning to the wealthy hand-to-mouth ($H2M_{LIQ}$), we see they are not really so wealthy. In particular their median net worth is less than 30 percent of that of the households classified as not hand-to-mouth by either measure. In all other variables they are intermediate to the groups in columns 1 and 2. They more closely resemble those not hand-to-mouth in age, but better resemble the poor hand-to-mouth in terms of earnings and income.

One possible source of differences in consumption growth rates is heterogeneity in expected rates of return. In particular, if low-asset households face a lower marginal return on savings, this could push them towards lower expected consumption growth. The variable High Liquid Debt in Table 2 reports the fraction of households, by group, that have balances on credit cards, store credit, student loans, medical or legal bills, or loans from family that sum to at least one month of household earnings. A large share of $H2M$ households exhibit such high debts. More exactly, two-thirds of $H2M_{NW}$ and over half of $H2M_{LIQ}$ households do, compared to only a fourth for those not hand-to-mouth. The bulk of such debts, especially credit card debt provide a high return to such households.[10] For this reason, we do not explicitly pursue a low marginal return as an explanation for lower expected consumption growth for low-asset households. But we note that the heterogeneity in preference discount factors we entertain below could alternatively be interpreted as combined heterogeneity in discount factors and market rates of return.

Table 2: Summary Characteristics of the Hand-to-Mouth

	Not $H2M$	$H2M_{NW}$	$H2M_{LIQ}$
Age	46.6	39.8	44.5
Income	96,660	46,781	64,058
Earnings	88,908	44,664	55,695
Liq Wealth (median)	13,918	- 7,719	- 2,731
Net Worth (median)	175,823	- 2,498	50,389
High Liquid Debt	24.7%	66.1%	54.4%
Sample Shares	59.7%	22.7%	17.5%

Note: All figures in 2009 dollars. Sample is PSID 1999-2015, with H2M status observed at least three times. High Liquid Debt equals one for households with credit card, store credit, student loans, medical or legal bills, or loans from family that sum to a month's or more of earnings, zero otherwise.

Table 3 reports transition probabilities into and out of the H2M categories. As the PSID waves are two years apart, we report two-year and four-year transition rates. The large diagonal elements reflect significant persistence in H2M status. This persistence is particularly striking given that reporting errors for household income and wealth no doubt create spurious transitions. In particular Panel A shows that those classified as H2M by net worth in a given year have a 65% chance of remaining in that status and a 16% chance of transitioning to the liquidity based measure of $H2M$ status; this leaves only a 19% chance that after two years an individual has become "unconstrained" by either definition. From Panel B we see that, even four years out, $H2M_{NW}$ households exhibit a 58% probability of retaining that status, and only a 24% probability of being not hand-to-mouth by either

[10] PSID surveys for 2011 and later allow us to separately identify credit card debt. Households categorized as hand-to-mouth are twice as likely to have credit card debt of a month's earnings or more. The incidence is 29.2% and 30.6% respectively for $H2M_{NW}$ and $H2M_{LIQ}$, compared to 14.5% for households not hand-to-mouth by net worth or liquid wealth.

measure. By contrast, for those households unconstrained at t by either measure, these probabilities are respectively 6% and 81%. This persistence is consistent with findings in Athreya et al. (2019) that less than 10 percent of households, those with repeat episodes of delinquent loans, account for half of all experiences of such financial distress. For the wealthier $H2M_{LIQ}$, there is a 42% chance of transiting out of H2M status altogether at $t+2$ and a 48% chance of doing so by $t+4$. While higher than the corresponding probability for the poor $H2M_{NW}$, these are still quite low given that, at the median, these households have net worth exceeding $\$50,000$ (as reported in Table 2).

Table 3: Transition Probabilities of H2M Status

	Not $H2M(t)$	$H2M_{NW}(t)$	$H2M_{LIQ}(t)$
		Panel A	
Not $H2M(t+2)$	.823	.194	.424
$H2M_{NW}(t+2)$	.059	.648	.176
$H2M_{LIQ}(t+2)$	.119	.158	.400
		Panel B	
Not $H2M(t+4)$	.808	.236	.480
$H2M_{NW}(t+4)$	.064	.584	.181
$H2M_{LIQ}(t+4)$	.128	.180	.339

Note: Sample is PSID 1999-2015, with H2M status observed at least three times.

4.3 Income and Consumption Growth

Standard consumption smoothing arguments suggest that an individual who expects their income to be higher in the future should draw down assets or increase debt today. Hence, low-wealth individuals may anticipate relative faster income growth compared to their higher wealth counterparts. Similarly, in the model of Section 2, an individual that is constrained will have an Euler equation that holds with an inequality. All else equal, this suggests a relatively higher anticipated growth in consumption. As shown in Figure 1 Panel (d), even when the borrowing constraint is not binding and the Euler equation holds with equality, a low-wealth individual has relatively high growth in consumption going forward, as the agent rebuilds its target level of assets.

With this logic in mind, we turn to the empirical relationship between $H2M$ status and the growth of income and consumption. We regress the (annualized) log growth in income (earnings plus financial income) and the log growth in consumption (durables plus nondurables) on dummies indicating hand-to-mouth status (one each for our two $H2M$ measures), as well as demographic controls described below. Specifically, growth rates are the log difference between year t and the subsequent wave in year $t + 2$, where we divide the log difference by two to compute the annualized growth rate. Hand-to-mouth status is defined as of year t. Hence, the specification asks whether on average the hand-to-mouth have relatively higher income or consumption growth over the subsequent two years.

In our baseline specification, we do not include individual fixed effects. Throughout the paper, our baseline demographic controls for specifications in which the dependent variable is in levels consist of a cubic in age, three race dummies, a marital status dummy, and five family size dummies. In the specifications that include fixed effects, we drop the race dummies. If the dependent variable is a change or growth rate, the demographic controls are a quadratic in age, two dummies for the change in marital status, and the change in the number of family members. All regressions include dummies for the year of the survey.

The results are reported in Table 4. The first three columns report estimates for income growth and the final three columns report consumption growth. In each panel, Column (1) does not include fixed effects, while Column (2) does. We postpone discussing Column (3)'s specification.

For income growth, the specification without fixed effects suggests that both the $H2M_{NW}$ and the $H2M_{LIQ}$ each have about a one percentage point higher growth rate than the non-H2M. Adding the fixed effect increases the differences by a factor of two to three.

Turning to consumption growth, absent fixed effects there is little difference between the $H2M_{NW}$ and those with higher net worth, while the wealthier hand-to-mouth ($H2M_{LIQ}$) show about a one percentage point lower consumption growth than the other groups. But these effects are significantly different controlling for household fixed effects. Controlling for fixed effects, low-net-worth households have a significantly higher rate of consumption growth; in particular, the coefficient on $H2M_{NW}$ is 2.5 log points. Including a fixed effect controls for (permanent) differences between households, including differences in preferences or rates of return. Including fixed effects, the point estimate on $H2M_{LIQ}$ is essentially zero, suggesting that the wealthier hand-to-mouth do not have different consumption growth going forward than wealthy households with more liquid wealth. One interpretation of the latter fact is that the wealthy hand-to-mouth do not need to reduce consumption in order to build up a buffer stock of precautionary saving, as they have the option of converting illiquid to liquid wealth in response to large shocks to income.

The respective third columns in each panel of Table (4) provide a sense of why including a fixed effect has a large impact on the baseline estimates. We drop the fixed effect and add controls for how frequently a household is denoted hand-to-mouth in the sample. This captures whether a household tends to have low net worth or liquidity on average. Those with a relatively frequent rate of hand-to-mouth status are interpreted as having a low target for net worth or liquidity; however, given the relatively short time sample, some of the average undoubtedly reflects realized shocks as well as ex ante targets. Including the frequency of hand-to-mouth status brings the baseline coefficients in line with the fixed-effects estimates. Moreover, conditional on current hand-to-mouth status, an individual that frequently finds itself designated as $H2M_{NW}$ or $H2M_{LIQ}$ tends to have significantly *lower* income and consumption growth. The relatively low rate of consumption growth is consistent with $H2M$ individuals having a lower discount factor (relative to return on savings) or a higher IES. Interestingly, households that tend (on average) to have low net worth or liquidity do not have a steeper income profile, as suggested by consumption smoothing logic; in fact, the opposite is true.

Table 4: Consumption and Income Growth for the Hand-to-Mouth

	Income Growth			Consumption Growth		
	(1)	(2)	(3)	(1)	(2)	(3)
$H2M_{NW}$	.012	.033	.038	.003	.025	.025
	(.004)	(.007)	(.008)	(.004)	(.006)	(.007)
$H2M_{LIQ}$	.010	.025	.028	-.009	.003	.002
	(.004)	(.006)	(.006)	(.004)	(.005)	(.005)
Fraction time $H2M_{NW}$			-.045			-.038
			(.009)			(.009)
Fraction time $H2M_{LIQ}$			-.043			-.025
			(.008)			(.007)
Fixed Effects	No	Yes	No	No	Yes	No

Note: Sample size is 19,105. Growth rates are annualized. Not H2M group is omitted in all regressions. Sample is PSID 1999-2015, with H2M status observed at least three times. Regressions include controls for changes in family size or marital status, age, and year. Standard errors are clustered at household level.

From Table 2, households that are frequently hand-to-mouth are on average younger and have lower earnings. In Appendix A2.1 we examine rates of spending growth dividing the sample between households with heads ages 25 to 39 versus age 40 to 64. (Note, the regressions reported in the text tables also include age controls.) Households frequently hand-to-mouth continue to display notably lower spending growth; but for $H2M_{LIQ}$ households this effect is more pronounced for younger households. In the appendix we also examine spending growth after dividing households with respect to long-term earnings.[11] Of course, there is no presumption that differences in preferences are orthogonal to earnings.[12] But it is useful to examine high and low-earnings households separately because it partially controls for: (a) the possibility of scale effects in savings returns, and (b) a differential importance of government savings (e.g., Social Security) in discouraging savings. Separating by long-term income, we continue to see lower rates of spending growth for households that are frequently hand-to-mouth; but this effect is smaller and not statistically significant for the higher-income households that are hand-are-mouth only due to low liquid assets.

An additional control that is implied by the standard model is the volatility of income growth. That is, from (2), if H2M households face less risk, all else equal, their consumption growth rates will be lower. However, it is not the case that hand-to-mouth agents face less risk, a point we document in the next subsection.

4.4 Income and Consumption Volatility

In the model of Section 2, low-wealth individuals are subject to higher anticipated consumption volatility given the absence of a buffer stock of savings. An alternative not considered in Section 2 is that low-wealth individuals have a lower volatility of income and therefore do not desire a large amount of precautionary savings. Finally, if the hand-to-mouth are up against a hard constraint in contiguous periods, their consumption will necessarily track income; that is, if a household is literally "hand to mouth," then $c = y + ra$. This suggests that an interesting moment is the volatility of the difference in consumption growth from income growth as a function of hand-to-mouth status.

To explore these relationships in our PSID sample, we regress the absolute value of income growth $|\Delta \ln (y + ra)|$, consumption growth $|\Delta \ln c|$, and the absolute value of the difference $|\Delta \ln c - \Delta \ln (y + ra)|$ on lagged hand-to-mouth status. Specifically, if the growth rate is calculated between waves t and $t + 2$ of the survey, the hand-to-mouth indicator is computed using the period t data. This specification is designed to capture whether a hand-to-mouth household today faces greater or lesser uncertainty about the future than a

[11] Long-term earnings is defined a household's average natural log of earnings after removing a cubic function of the head's age and year dummies.

[12] In fact, Dynan, Skinner and Zeldes (2004) argue that an important reason that lower-income households save less may reflect a lower demand for precautionary savings. This aligns with our conclusions, and those of Calvet et al. (2019) that the behavior of hand-to-mouth households is consistent with a high IES.

non-H2M household. All regressions also include a quadratic in age and dummies for change in family size, change in marital status, and year.

The results are reported in Table 5. The first column reports that both the poor and wealthy hand-to-mouth face greater uncertainty regarding future income growth. This difference is reduced substantially once we include individual fixed effects (Column 2). Similarly, hand-to-mouth households face greater consumption uncertainty over the subsequent two years, but this difference becomes negligible with the inclusion of fixed effects (Columns 3-4). The final two columns concern volatility over the consumption-to-income ratio. The hand-to-mouth exhibit significantly more volatility in this ratio. Including fixed effects again reduces the coefficients by more than half. For quantitative context, the mean of the dependent variables are 0.139, 0.146, and 0.190 for $|\Delta \ln (y + ra)|$, $|\Delta \ln c|$, and $|\Delta \ln c - \Delta \ln (y + ra)|$, respectively. Thus a coefficient on the $H2M$ dummies of 2.5% represents roughly an increase of 13-18% of the mean level.

Table 5: The Volatility of Income and Consumption Part I

	$\lvert\Delta \ln (y + ra)\rvert$		$\lvert\Delta \ln c\rvert$		$\lvert\Delta \ln c - \Delta \ln (y + ra)\rvert$	
	No FE	FE	No FE	FE	No FE	FE
$H2M_{NW}$	.021	.005	.027	.010	.025	.011
	(.004)	(.004)	(.004)	(.004)	(.004)	(.005)
$H2M_{LIQ}$	.025	.013	.010	-.002	.020	.009
	(.004)	(.004)	(.003)	(.003)	(.005)	(.004)

Note: Sample size is 19,351. Not H2M group is omitted in all regressions. Sample is PSID 1999-2015, with H2M status observed at least three times. Regressions include controls for changes in family size or marital status, age, and year. Standard errors are clustered at household level.

In Table 6, we replace the fixed effects with the fraction of time spent as hand-to-mouth. We see that households that are frequently in hand-to-mouth status tend to have more volatility in consumption, income, and the ratio of consumption-to-income. Conditional on the fraction of time spent as hand-to-mouth, current hand-to-mouth status has no predictive power regarding subsequent volatility.

One caveat with the volatility measures is that they also capture error in the measures of consumption and income. Such error is undoubtedly significant (see, for example, Bound, Brown, Duncan and Rodgers, 1994, Aguiar and Bils, 2015, and Carroll, Crossley and Sabelhaus, 2015). However, it is less clear that the magnitude of the error varies with hand-to-mouth status, which is relevant for the exercises performed above. To explore this, we posit that measurement error is *iid* over different waves of the survey. If the persistence of

Table 6: The Volatility of Income and Consumption Part II

	$\lvert\Delta \ln (y + ra)\rvert$	$\lvert\Delta \ln c\rvert$	$\lvert\Delta \ln c - \Delta \ln (y + ra)\rvert$
$H2M_{NW}$	.003	.009	.008
	(.006)	(.005)	(.006)
$H2M_{LIQ}$	.009	-.002	.006
	(.005)	(.004)	(.006)
Frac. time	.030	.031	.029
$H2M_{NW}$	(.009)	(.007)	(.009)
Frac. time	.040	.027	.036
$H2M_{LIQ}$	(.010)	(.008)	(.012)

Note: Sample size is 19,351. Not H2M group is omitted in all regressions. Sample is PSID 1999-2015, with H2M status observed at least three times. Regressions include controls for changes in family size or marital status, age, and year. Standard errors are clustered at household level.

the true variable is similar across hand-to-mouth status, then the observed autocorrelation will be lower for the group with the greater mis-measurement.

Table 7 reports the estimated auto-regressive coefficient for the growth of income and consumption for each hand-to-mouth status. Specifically, we compute the correlation of growth between years $t - 4$ and $t - 2$ and the growth between $t - 2$ and t for each group defined by period $t - 2$ hand-to-mouth status. Looking across the rows, there is little evidence that the hand-to-mouth have a significantly lower autocorrelation for either income or consumption. Under the assumption that the true process is the same for all groups, this suggests that classical measurement error (in logs) is not more or less severe for the hand-to-mouth households.

4.5 More on Income Processes

A natural hypothesis is that some agents face less idiosyncratic income risk and in response hold less precautionary savings. Similarly, some could face a steeper life-cycle profile of earnings, and therefore desire to borrow rather than save. Both points suggest differences in income processes as a potential explanation for the persistent hand-to-mouth. However, the preceding analysis indicates that the frequently hand-to-mouth have lower income growth and more volatility. Given the plausibility of the income-difference hypothesis, a closer study is warranted. Hence, in this subsection we explore in greater detail the relative income processes for the frequently hand-to-mouth.

Table 7: Autocorrelation of Income and Spending Growth

	Not $H2M$	$H2M_{NW}$	$H2M_{LIQ}$
$\rho(\Delta \ln y^d_t, \Delta \ln y^d_{t-1})$	-.336	-.251	-.390
	(.022)	(.023)	(.031)
$\rho(\Delta \ln c_t, \Delta \ln c_{t-1})$	-.372	-.380	-.348
	(.012)	(.021)	(.025)

Note: Sample is PSID for 1999-2015, with seeing H2M status at least 3 times. Regressions include usual controls. Standard errors are clustered at household level. For brevity of exposition, we denote $y^d \equiv y + ra$.

To this end, we divide the PSID sample in half based on the frequency of hand-to-mouth status over the sample. The median household finds itself in either H2M status 2 times. We therefore group the 55.6 percent of the sample that are $H2M$ 2 times or less in the "Below Median" category, and the remaining 44.4 percent in the "Above Median."[13]

For each subsample we estimate a standard specification of the process for household earnings. Specifically, we estimate the following specification:

$$\ln y_{ijt} = \beta_t + \mathbf{X_{ijt}} \boldsymbol{\gamma} + \tilde{y}_{ijt},$$

where y_{ijt} is total after-tax labor earnings for household i in wave t with a head of age j, β_t is a year fixed effect, $\boldsymbol{X}$ is a vector of demographic controls, and $\tilde{y}_{ijt}$ is the residual income that we analyze below in a second stage. The vector $\boldsymbol{X}$ includes a cubic in age (normalized to zero at 25), as well as dummy controls for educational attainment (for less than high school, high school, some college, and college degree), the number of earners $(0, 1$ or $2)$, family size $(1, 2, 3, 4,$ or $5+)$, and three race categories. We estimate the specification for each subsample separately. The results are reported in Table 8. The coefficients are similar across the two groups. The frequently hand-to-mouth have a slightly shallower age profile, in line with Table 4, but the difference is small relative to the standard errors.

We now study the residual income process $\tilde{y}_{ijt}$. Specifically, we estimate the standard specification:

$$\tilde{y}_{ij} = \alpha_i + \epsilon_{ij} + \nu_{ij}$$
$$\epsilon_{ij} = \rho \epsilon_{ij-1} + \eta_{ij},$$

[13]We have also performed the analysis splitting only on the frequency of just $H2M_{NW}$ status and found similar patterns.

Table 8: First-Stage Income Process Parameters

	Rarely $H2M$	Frequently $H2M$
Age-25	.057	.046
	(.005)	(.005)
$(\text{Age-25})^2$	-.002	-.002
	(<.001)	(<.001)
$(\text{Age-25})^3$	(<.001)	(<.001)
	(<.001)	(<.001)
Observations	14,438	11,515

Note: Not reported are controls for education, household status, number of earners, family size and time.

where α_i is a mean-zero household fixed effect with cross-sectional variance σ_α^2; ϵ is a persistent AR(1) process with parameter ρ; and ν and η are mean-zero *iid* shocks with variances σ_ν^2 and σ_η^2, respectively. All shocks are assumed to be orthogonal. We estimate the parameter vector $\{\rho, \sigma_k\}$, where $k = \alpha, \nu, \eta$ by method of moments. We follow Tonetti (2011) and use the (available) covariances between agents of ages j and j' as moments.[14] The results are reported in Table 9, along with bootstrapped 95% confidence intervals.

We see that households frequently hand-to-mouth have less volatile innovations to the AR(1) coefficient, but with slightly more persistence. Together the estimates indicate an unconditional variance of ϵ is 0.26 for the rarely $H2M$ compared to 0.20 for the frequently $H2M$. Thus the frequently $H2M$ have slightly less long-run uncertainty. On the other hand, the *iid* component of earnings is over five times as volatile for the frequently $H2M$.

To put these differences in perspective, we simulated the standard model under each of the residual income processes, that estimated for households infrequently hand-to-mouth and that estimated for those frequently so, separately for each of our four preference specifications. This reveals whether frequent hand-to-mouth status reflects the differential income process for these households. To focus on the differences in uncertainty, we normalized both income processes to have an unconditional mean of one and set the borrowing constraint at zero for each exercise. The fixed effect is set equal to zero. In Table 10, we report the frequency of $H2M_{NW}$ status (computed as $a' < y/6$, as in the data) and average wealth divided by average income.

[14] See Tonetti (2011) for details. We are grateful to Chris Tonetti for sharing the Matlab code to compute the estimates.

Table 9: Residual Income Process Parameters

	Rarely $H2M$	Frequently $H2M$
ρ	.909	.948
	[.484 .958]	[.781 .961]
σ_η^2	.045	.020
	[.017 .141]	[.016 .067]
σ_ν^2	.022	.131
	[.000 .127]	[.077 .144]
σ_α^2	.161	.063
	[.082 .214]	[.035 .101]

Note: The table reports point estimates of the parameters of the income process and the 95% confidence interval in brackets.

Table 10: Simulated Moments for Alternative Income Specifications

	Frequency H2M		Wealth/Earnings	
	(1)	(2)	(1)	(2)
$\beta = 0.95, \sigma = 0.5$	.016	.008	9.647	8.137
$\beta = 0.95, \sigma = 1.5$	.260	.126	1.885	1.447
$\beta = 0.90, \sigma = 0.5$	.271	.135	.148	.925
$\beta = 0.90, \sigma = 1.5$	.955	.616	.058	.207

Note: The columns labelled (1) use the residual income process parameters estimated for the Rarely Hand-to-Mouth reported in Table 9, setting $\sigma_\alpha^2 = 0$; the columns labelled (2) use the income parameters estimated for the Frequently Hand-to-Mouth. $H2M$ status in the model is whether the agent has assets a less than two months of annual earnings, $y/6$. Wealth/Earnings is the ergodic mean of a, where the mean of income has been normalized to one in each simulation.

In all simulations, the model's predictions for frequency of $H2M$ status is opposite the data. That is, the empirical income process for the frequently hand-to-mouth predicts less frequent $H2M$ status in the model. There is no clear pattern with respect to average wealth. Recall that the frequently hand-to-mouth income process has significantly more transitory risk. This generates a large incentive to save near the borrowing constraint and quickly exit the low-asset region. This explains the infrequent hand-to-mouth status. Far from the constraint, however, the fact that the empirical hand-to-mouth income process has slightly less persistent risk plays a role, making the prediction for the long-run mean of wealth depend on the preference specification.

The evidence presented in this and the preceding subsection indicates that differences in income processes is not a compelling explanation for why some agents are frequently hand-to-mouth. The frequently $H2M$ have greater transitory income risk that calls for greater precautionary saving, all else equal. This predicted pattern of behavior is the opposite of what we observe in our data.

4.6 The Average Propensity to Consume

For a given set of preferences, and assuming monotonic dynamics, the average propensity to consume is a useful guide to the distance to target wealth. In particular, agents below their target wealth should have a low APC. In Table 11, we report the results of regressing APC at time t in the PSID on lagged $t-2$ measures of hand-to-mouth status. The measures and specifications are the same as in Table 4.

In Column (1), we find that the average propensity to consume is *higher* for hand-to-mouth than for non-$H2M$ households. That is, the hand-to-mouth are not actively building up their buffer stock of liquid savings relative to their wealthier or more liquid counterparts. This is counter to the model of Section 2 if preferences across agents are identical. Controlling for fixed effects, reported in Column (2), essentially eliminates any difference between the groups. Column (3) documents that households that frequently find themselves as hand-to-mouth (and controlling for current $H2M$ status) have a relatively high APC. This is consistent with the hypothesis that the average hand-to-mouth individual has a significantly lower target level of assets or liquidity (relative to income).[15]

[15] In Appendix A2.1 we examine APCs by hand-to-mouth status dividing households by age and by long-term earnings. The results from Table 11 hold across age groups. Separating by long-term income, we continue to see much higher APCs for households that are frequently hand-to-mouth, though the magnitudes are not as large as in Table 11 pooling all households.

Table 11: The Average Propensity to Consume

	(1)	(2)	(3)
$H2M_{NW}$	.065	-.007	-.022
	(.010)	(.010)	(.015)
$H2M_{LIQ}$	.077	.005	.005
	(.012)	(.009)	(.011)
Fraction			.157
time $H2M_{NW}$			(.022)
Fraction			.191
time $H2M_{LIQ}$			(.025)
Fixed effects	No	Yes	No

Note: Sample is PSID for 1999-2015, with seeing H2M status at least 3 times. Regressions include usual controls. Standard errors are clustered at household level.

4.7 The Marginal Propensity to Consume out of Income

Conditional on preferences, the model of Section 2 clearly predicts the MPC is declining in cash on hand. The empirical evidence on the relationship is mixed.[16] As noted in Section 2, with preference heterogeneity, there is no clear predicted relationship in pooled data.

Consider estimating the response of consumption to a change in income. Suppose we observe $\Delta c_t \equiv c_t - c_{t-1}$. From the consumption function, we have $\Delta c_t = \mathcal{C}(x_t, y_t) - \mathcal{C}(x_{t-1}, y_{t-1})$. Approximating the consumption function around (x_{t-1}, y_{t-1}), we have

$$\begin{aligned} \mathcal{C}(x_t, y_t) - \mathcal{C}(x_{t-1}, y_{t-1}) &\approx \mathcal{C}_x \Delta x_t + \mathcal{C}_y \Delta y_t \\ &= MPC * R\Delta a_t + MPCY \Delta y_t, \end{aligned} \tag{3}$$

where C_x and C_y are evaluated at (x_{t-1}, y_{t-1}), the second line uses $x_t = Ra_t + y_t$, and we define $MPCY \equiv MPC + \mathcal{C}_y$ as the marginal propensity to consume out of a change in income. Recall that holding x constant, $\mathcal{C}_y$ captures consumption responses due to anticipated future income following today's y realization. If y is iid, $\mathcal{C}_y = 0$, while if y is persistent and the individual is not constrained, then $\mathcal{C}_y > 0$. Hence, we expect $MPCY \geq MPC$ to hold in the data.

[16]See Havranek and Sokolova (2019) for a recent meta-study.

To scale the responses, divide (3) through by $y_{t-1} + ra_{t-1}$ to obtain:

$$\frac{\Delta c_t}{y_{t-1} + ra_{t-1}} \approx MPC \frac{R\Delta a_t}{y_{t-1} + ra_{t-1}} + MPCY \frac{\Delta y_t}{y_{t-1} + ra_{t-1}}, \tag{4}$$

where MPC and MPCY are evaluated at (x_{t-1}, y_{t-1}). Recalling that a_t denotes start of period t assets, we have $\Delta a_t = ra_{t-1} + y_{t-1} - c_{t-1}$. Hence, Δa_t is in the individual's $t-1$ information set. If the growth rate of income is *iid* (that is, log income is a random walk), then the first term on the right of (4) is orthogonal to $\Delta y_t/(y_{t-1} + ra_{t-1})$.

Under the random walk assumption, a regression of consumption growth on income growth, within a sample of individuals of similar (x_{t-1}, y_{t-1}) and consumption functions, provides an estimate of MPCY. In the model of Section 2, this is monotonically decreasing in assets a for a given set of preferences.

To explore this prediction empirically, we begin by simply regressing the (normalized) change in consumption on the (normalized) change in income separately for each of our three groups: the non-$H2M$, $H2M_{NW}$, and $H2M_{LIQ}$. All regressions also include year dummies and controls for age and the changes in family size and marital status. In the regressions, the normalization factor is the average of total income in wave t and the previous wave of the PSID. We first estimate the MPCY using OLS. One concern is that the change in earnings (which is a right-hand side variable) is measured with error. To address this, we also run a separate specification in which we instrument for changes in income using lagged employment status, changes in weeks worked and in annual hours for both the head and the spouse. This will help correct for errors in variables, but it must also be kept in mind that this source of variation may have different implications for the persistence of the income change, thus changing the true coefficient.

The results using OLS are reported in Table 12. The point estimates in Column (1) suggest a considerably higher MPCY for low net-worth individuals compared to those not hand-to-mouth, but only modestly higher for the low-liquidity. The estimated MPCYs are actually fairly low for all households, as anticipated by the empirical literature (e.g., Blundell, Pistaferri and Preston, 2008, Straub, 2019, Fisher et al., 2019). Column (2) includes our measures for frequency of hand-to-mouth status. With these controls, current hand-to-mouth status no longer predicts MPCY; rather, conditional on current status, those that are frequently hand-to-mouth tend to have higher MPCYs. The final column implies that those with a higher APC tend to have a lower MPCY. In our benchmark model, a high APC is associated with a level of assets above one's target, which should imply a lower MPC.

The IV results reported in Table 13 tell a similar story to the OLS, but the baseline magnitudes are higher. This is consistent with attenuation bias in the OLS specification.

Nevertheless, the point estimates are quite small; specifically, a one percent increase in income implies only an 0.08% increase in spending relative to income for the non-$H2M$. Moreover, if income is persistent, then the MPCY estimated in these regressions should be an upper bound on the MPC out of a transfer.

Table 12: Marginal Propensity to Consume out of Income (OLS)

Dependent variable is $\Delta c/(y + ra)$

	(1)	(2)	(3)
$\Delta y/(y + ra)$	.039	.019	.069
	(.008)	(.009)	(.014)
$\Delta y/(y + ra) \times H2M_{NW}$	.046	-.017	-.032
	(.019)	(.025)	(.025)
$\Delta y/(y + ra) \times H2M_{LIQ}$	.012	-.030	-.034
	(.020)	(.022)	(.022)
$\Delta y/(y + ra)$		.119	.156
$\times$Fraction time $H2M_{NW}$		(.039)	(.035)
$\Delta y/(y + ra)$		.093	.119
$\times$Fraction time $H2M_{LIQ}$		(.041)	(.039)
$\Delta y/(y + ra)\times$ APC			-.076
			(.018)

Note: H2M's are lagged (t-2) values, APC is Tornqvist of current and lagged APCs, income is Tornqvist of current and lagged income. Regressions include controls for age, changes in family size and marital status, and year. HtM status observed at least three times.

4.8 Consumption Baskets: The Extensive Margin of Expenditure

The framework of Section 2 suggests that an individual may have low wealth because of poor luck or because of preference heterogeneity. Moreover, the latter can also be due to differences in time preference versus differences in the IES. In this subsection, we look at detailed consumption behavior of low-wealth individuals to shed some light on this question. One complication is that if agents are truly constrained in terms of borrowing, then it is difficult to obtain an estimate of time preference or the IES. For this reason, we shall pursue an indirect route.

Table 13: Marginal Propensity to Consume out of Income (IV)

Dependent variable is $\Delta c/(y + ra)$

	(1)	(2)	(3)
$\Delta y/(y + ra)$	.082	.045	.102
	(.020)	(.022)	(.042)
$\Delta y/(y + ra) \times H2M_{NW}$	.038	-.076	-.059
	(.039)	(.062)	(.061)
$\Delta y/(y + ra) \times H2M_{LIQ}$	.066	-.023	-.027
	(.047)	(.052)	(.052)
$\Delta y/(y + ra)$		.188	.184
$\times$Fraction time $H2M_{NW}$		(.087)	(.087)
$\Delta y/(y + ra)$		.211	.190
$\times$Fraction time $H2M_{LIQ}$		(.101)	(.100)
$\Delta y/(y + ra)\times$ APC			-.076
			(.060)

Note: H2M's are lagged (t-2) values, APC is Tornqvist of current and lagged APCs, income is Tornqvist of current and lagged after-tax income. Instruments are employment status at (t-2) and changes in weeks worked and annual hours for each of head and for spouse. Regressions include controls for age, changes in family size and marital status, and year. HtM status observed at least three times.

This route focuses on the "extensive" margin for various consumption categories; that is, whether a broad category is consumed at all.[17] Those close to adjustment on the extensive margin may exhibit a highly elastic response of total expenditure to changes in inter-temporal prices. The fact that adjustment on the extensive margin may alter the price elasticity is a familiar concept in economics. This idea has been applied to macro-labor markets by Rogerson (1988) and to portfolio choice by Grossman and Laroque (1990). Chetty and Szeidl (2007) make a related argument in the context of risk preference in the presence of consumption commitments. In the appendix, we provide a simple example of how the results presented in this subsection may shed light on the relative IES of those prone to adjusting expenditure on the extensive margin.

Consider nondurable expenditure at two adjacent dates, c_{t-1} and c_t. Suppose expenditure

[17]See Michelacci, Paciello and Pozzi (2019) for an analysis of the extensive margin of expenditure in the Kilts-Nielsen Consumer Panel.

Table 14: Nondurable Expenditure Categories

Category	Share	Positive	Add	Drop
Food at home	0.346	1.00	0.00	0.00
Food away	0.154	0.95	0.03	0.03
Gasoline	0.138	0.93	0.03	0.03
Car insurance	0.106	0.93	0.03	0.02
Health insurance	0.076	0.73	0.09	0.09
Education	0.055	0.30	0.11	0.13
Doctors	0.034	0.80	0.09	0.09
Prescription drugs	0.022	0.80	0.09	0.09
Childcare	0.019	0.12	0.04	0.05
Hospital	0.015	0.28	0.15	0.16
Other transport	0.015	0.30	0.17	0.18
Other utilities	0.009	0.17	0.08	0.11
Bus & train	0.005	0.08	0.03	0.03
Parking	0.003	0.09	0.05	0.04
Taxi	0.002	0.05	0.03	0.02

in period t is divided among N_t discrete goods: $c_t = \sum_{n=1}^{N_t} x_{n,t}$, where $x_{n,t}$ is the amount devoted to good n in period t. A similar decomposition can be done for period $t-1$.

In this spirit, we divide non-durable expenditure into the categories listed in Table 14.[18] For each category, we list the share of total non-durable expenditure as well as the average fraction of individuals that spend a strictly positive amount on that category in a given two-year time frame. The final two columns are the average probability of addition or deletion of that category, respectively. We can have a finer decomposition of expenditure into distinct categories using the Consumer Expenditure Survey (CE). In the appendix we report the CE counterpart of Table 14.

Our first exercise explores whether the hand-to-mouth consume a different number of distinct categories conditional on total expenditure. That is, do the hand-to-mouth allocate a given level of expenditure differently between the number of goods versus average spending per good? To explore this, we regress the log number of categories with positive expenditure on log total expenditure as well as our $H2M$ dummies, adding year dummies and demographic controls. Note that under time-separability, the number of goods is a static decision conditional on total expenditure, and is therefore independent of time preference and borrowing constraints.

The results are reported in Table 15. Columns one and two use the PSID and our

[18]We exclude basic utilities like water, heat, and electricity as these may be included in rental contracts.

benchmark measures of $H2M$ status. The CE does not have as detailed wealth data as the PSID, so we cannot construct identical measures of hand-to-mouth status in the CE. Instead, we construct a measure that uses only liquid assets, which is denoted in the table as $H2M_{KVW}$, and is defined as having liquid assets within one week of earnings of the borrowing limit or a positive position less than one week of earnings.[19] The final column estimates the impact of $H2M_{KVW}$ on the log number of categories in our CE sample. For comparison, we construct the same measure in the PSID sample and report the CE regression implemented in the PSID sample in the third column.

Table 15: Number of Categories Consumed

Dependent variable is $\ln N_t$

	PSID	PSID	PSID	CE
$\ln c$	.204	.201	.202	.456
	(.006)	(.006)	(.006)	(.002)
$H2M_{NW}$	-.051	-.028		
	(.008)	(.006)		
$H2M_{LIQ}$	-.036	-.007		
	(.006)	(.005)		
Fraction time $H2M_{NW}$		-.043		
		(.014)		
Fraction time $H2M_{LIQ}$		-.080		
		(.017)		
$H2M_{KVW}$			-.043	-.115
			(.005)	(.002)
Observ.	26,178	26,178	26,178	192,299

Note: Categories restricted to nondurables and services. Households on average spend on 8.9 of 17 categories in PSID data, on 12.1 of 27 categories in the CE. Have to see H2M status at least three times in PSID.

The estimates show a clear pattern that low-wealth and low-liquidity households consume fewer categories of goods conditional on a given level of total expenditure. This pattern is established for both measures of $H2M$ status, as well as for $H2M_{KVW}$ status in the PSID and in the CE. The effect is also economically significant. For instance, comparing the

[19] This is exactly the hand-to-mouth measure in Kaplan et al. (2014).

coefficients for total expenditure and for being hand-to-mouth based on net worth, we see that the $H2M_{NW}$ coefficient is opposite in sign and one-fourth the magnitude of that for $\ln c$. That implies that being the status of $H2M_{NW}$ predicts the same impact on number of categories as a 25% reduction in total spending.

The second column of the table includes the frequency that we see a household to be hand-to-mouth in the PSID. Households that are frequently hand-to-mouth, by either measure, are the ones who consume fewer categories of goods, conditional on total expenditure. Once we include a household's fraction of time as $H2M$, the current hand-to-mouth status is much less relevant, especially being hand-to-mouth based only on liquidity, which becomes statistically insignificant. Thus, as with the results above detailing consumption growth, consumption volatility, and APCs, it is a household's propensity to be hand-to-mouth, not its current hand-to-mouth status, that is key to predicting household spending behavior.

Our next set of results concern the elasticity of the number of goods to a change in total expenditure. To explore this, we regress $\Delta \ln N_t$ on $\Delta \ln c_t$, interacting the growth of total expenditure with our $H2M$ indicators. The regressions also include the $H2M$ dummies as well as the usual demographic controls. This specification tests whether *at the margin* of changing total nondurable expenditure, the hand-to-mouth allocate additional expenditure differently than the non-hand-to-mouth along the extensive versus intensive margins.

The results are reported in Table 16. The benchmark results reported in Column (1) indicate that the category elasticity is higher for the low-net-worth hand-to-mouth.[20] The second column of Table 16 includes the frequency with which a household is hand-to-mouth. We see that those who are prone to being hand-to-mouth, by either measure, are much more likely to adjust on the extensive margin in response to a change in total expenditure. Moreover, once we include the fraction of time as $H2M$, the current hand-to-mouth status is no longer relevant.

The results of Tables 15 and 16 are therefore consistent with heterogeneity across consumers in the relevance of the extensive margin of consumption, and this heterogeneity is correlated with the target level of assets. As noted above, the elasticity of total expenditure to a relative price (or interest rate) change may be sensitive to whether the adjustment is primarily occurring along the extensive or intensive margins. The empirical results therefore suggest that differences in the effective IES are a plausible candidate for explaining why some households are prone to hand-to-mouth status. In the next section, we will allow for such differences in calibrating preference heterogeneity.

[20] In Appendix A2.2 we explore this further by decomposing the growth in nondurable consumption between $t-1$ and t into an intensive change, reflecting the change in spending on goods consumed in both periods, and an extensive margin reflecting the adding and dropping goods. There we see that the extensive margin is about twenty percent more important for $H2M_{NW}$ households.

Table 16: Regression of $\Delta \ln N_t$ on $\Delta \ln c_t$

Dependent variable is $\Delta \ln N_t$

	(1)	(2)
$\Delta \ln c$	.132	.116
	(.006)	(.007)
$\Delta \ln c \times H2M_{NW}$	.036	-.019
	(.012)	(.019)
$\Delta \ln c \times H2M_{LIQ}$	.013	-.022
	(.013)	(.016)
$\Delta \ln c \times$ Fraction time $H2M_{NW}$		.089
		(.025)
$\Delta \ln c \times$ Fraction time $H2M_{LIQ}$		.074
		(.027)

Note: H2M status at least three times. HtM status and controls included.

5 Calibrating Preference Heterogeneity

The facts documented above make a compelling case that preference heterogeneity plays a role in determining hand-to-mouth status. In this section, we allow for preference heterogeneity in a model of savings, calibrating the shares of each "type" of agent. We can then address how preference heterogeneity affects the wealth distribution, focusing on why some agents have low wealth or liquidity as well as key related outcomes, such as the distribution of effective MPCs and IESs. The latter are crucial for the proper understanding and design of fiscal and monetary policies.

As noted in Section 2, it is difficult to distinguish impatience (β) from a high IES (σ) when the interest rate is less than the discount rate. Moreover, the extensive margin analysis suggests differences that potentially relate to the IES, but not discount factor; so it is important to include this potential source of heterogeneity.

One approach would be to introduce interest rate "shocks" to the standard model of Section 2 and try to match the responses to the corresponding empirical moments. As is well known, such an exercise is difficult empirically, particularly when consumers are not necessarily on their Euler equation. Even if they are, it is not obvious which interest rate is the empirically relevant rate for consumption-savings decisions of the hand-to-mouth.

An alternative is to move to a two-asset model, using portfolio choices to distinguish the IES from the discount factor. A consumer with a high IES is sensitive to interest rate differences across assets. Moreover, a high-IES consumer has a lower need for liquidity. That is, a high-IES agent cares less about *when* in time consumption takes place, and therefore is less concerned if assets are tied up in illiquid investments. It is also useful to separate a high IES from a low coefficient of relative risk aversion, as achieved by Epstein-Zin preferences. The model of Kaplan and Violante (2014) (henceforth KV) uses Epstein-Zin preferences and includes a meaningful liquidity decision; so we therefore employ their model.

5.1 Model Environment

In this section, we briefly recap the key elements of the KV model, which we take directly from their paper and augment by allowing for preference heterogeneity.[21]

Agents live for 58 years, or $J = 232$ quarters, of which $J^w = 152$ are spent earning income and $J^r = 80$ in retirement. Consumer i has Epstein-Zin preferences given recursively at age j by:

$$V_{ij} = \left[(c_{ij}^{\phi} s_{ij}^{1-\phi})^{1-1/\sigma_i} + \beta_i \left\{ \mathbb{E} V_{i,j+1}^{1-\gamma} \right\}^{\frac{1-1/\sigma_i}{1-\gamma}} \right]^{\frac{1}{1-1/\sigma_i}},$$

where c is non-durable consumption; s is the service flow from durables (to be described below); γ is the coefficient of relative risk aversion, which we set to 4 as in KV; and σ is the inter-temporal elasticity of substitution. Note that we allow the time-preference parameter β and the IES to vary by individual. Following KV, we set $\phi = 0.85$, which is based on the ratio of housing expenditures to total consumption in the US National Income and Product Accounts.

Earnings are given by the exogenous process:

$$\ln y_{ij} = \chi_j + \alpha_i + z_{ij},$$

where χ_j is a deterministic function of age, j, and α_i is an individual fixed effect. The idiosyncratic risk is represented by z_{ij}, which is a random walk. KV estimate a fourth-order polynomial for χ_j using the PSID. The variance of the fixed effect is set to 0.18 and the mean-zero quarterly innovation variance for the random walk z_{ij} is set equal to 0.003 to match the age profile of the cross-sectional variance in earnings. The income process implies that mean earnings equal $53,000 dollars.

An individual invests their wealth in two assets. The liquid asset, m, has an annual

[21]We thank the authors for sharing their code.

after-tax return given by:

$$r_m(m) = \begin{cases} 5.77\% \text{ if } m \in [\underline{m}_j, 0); \\ -1.41\% \text{ if } m \geq 0, \end{cases}$$

where $\underline{m}_j$ is an age- and income-specific borrowing limit.[22] The illiquid asset $a \geq 0$ has a higher return, but consumers must pay a fixed transaction cost κ to alter their stock of a. Again following KV, we set the after-tax return of the illiquid asset to 2.21% and κ to $\$1,000$.[23]

In addition to the financial return on illiquid assets, there is a service flow from the stock of illiquid assets. Specifically, $s_{ij} = \zeta a_{ij} + h_{ij}$, where $h_{ij} \geq -\zeta a_{ij}$ represents housing services obtained from a rental market. KV set $\zeta = 1\%$, quarterly. The units of rental housing h are normalized such that the relative price of c to h is one.

Consumption and rental housing are taxed at a rate of 7.2%. Earnings and assets are taxed, with the tax rate a function of earnings and the consumer's portfolio: $\mathcal{T}(y_{ij}, m_{ij}, a_{ij})$. Retirees receive social-security benefits given by $p(\chi_{J^w}, \alpha_i, z_{i,J^w})$. The benefits are taxed using the same function $\mathcal{T}$ as workers face, with p replacing y as the first argument. We refer the reader to KV for the exact functional forms and parameterization of $\mathcal{T}$ and p.

5.2 Calibrating the Hand-to-Mouth

We solve the consumer's problem for four sets of preference parameters. Specifically, as in Section 2 we let $\sigma_i \in \{0.5, 1.5\}$ and (annualized) $\beta_i \in \{0.90, 0.95\}$. In all cases, we set risk aversion to $\gamma = 4$, as in KV. Recall that the illiquid asset yields 2.21% plus a service flow, and hence consumers are impatient relative to the financial return. The relatively patient, inelastic specification ($\beta = 0.95, \sigma = 0.5$) is a standard representative-agent parameterization in macroeconomics. KV set all agents to $\beta = 0.941, \sigma = 1.5$, and hence our patient *elastic* specification is roughly comparable to the KV calibration. The two types with $\beta = 0.90$ are included to allow for extreme impatience as the driving force behind hand-to-mouth behavior.

Given the partial equilibrium nature of the model, we can solve and simulate each type of agent separately. To obtain an ergodic distribution, we solve and simulate each preference type's life-cycle 50,000 times.

[22] Specifically, for $j \leq J^w$, $\underline{m}_j = 0.74 y_j$. For $j > J^w$, $\underline{m}_j = 0$.

[23] KV calibrate the fixed cost to match the fact that one-third of households in the Survey of Consumer Finance are hand-to-mouth. We keep their value of κ, but will calibrate preference parameters to match the shares of hand-to-mouth observed in our PSID sample.

Table 17 reports key moments for each preference configuration. In particular, we compute the probability of our two hand-to-mouth status measures, as well as the ratio of total net worth to income and illiquid wealth to income. The first row are the comparable moments from our PSID sample. The last row is our calibration with heterogeneous types, which we discuss below.

Table 17: Type-Specific Moments

Type	Prob of $H2M_{NW}$	Prob of $H2M_{LIQ}$	NW/Inc $\frac{a+m}{y}$	Illiq NW/Inc $\frac{a}{y}$
Data (PSID)	22.7%	17.5%	4.04	3.27
$\beta = 0.95, \sigma = 0.5$	8.1%	18.9%	3.52	3.43
$\beta = 0.95, \sigma = 1.5$	7.3%	43.6%	5.30	5.21
$\beta = 0.90, \sigma = 0.5$	70.6%	9.4%	0.60	0.59
$\beta = 0.90, \sigma = 1.5$	96.5%	0.9%	-0.04	0.03
Calibration	22.0%	18.8%	3.21	3.15

The simulated moments indicate that relative patience (that is, $\beta = 0.95$) is important for generating realistic average wealth-to-income ratios. As one would expect, the specifications with $\beta = 0.90$ generate very low levels of wealth. Accordingly, high-β agents are relatively unlikely to be hand-to-mouth at any point in time.

For all types, most wealth is held in the illiquid asset. Holding β constant, a higher σ generates a small decrease in the fraction of wealth held in the liquid asset. Thus, a high σ plays an important role in generating illiquid hand-to-mouth agents. For example, the $\beta = 0.95, \sigma = 1.5$ agents are more than twice as likely to find themselves in the $H2M_{LIQ}$ state relative to $\beta = 0.95, \sigma = 0.5$ agents. Hence, relative patience combined with a high IES generates frequent "wealthy hand-to-mouth status." Recall that the $\beta = 0.95, \sigma = 1.5$ agents are close to the preference specification used in KV, consistent with their finding that a relatively high share of consumers are wealthy hand-to-mouth.

We now simulate the model with heterogeneous types. To do so, we pool the four types assuming type i has relative share π_i and then compute aggregate simulated moments. We select the relative shares to match the empirical moments in the top row of Table 17. There are four moments and three share parameters (one is redundant from the adding up to one constraint), and we select shares to minimize the equally weighted squared percent deviation between the pooled simulated moments and empirical moments.

The bottom row of Table 17 reports the simulated moments, while the first column of

Table 18 reports the shares for each type chosen that map to those moments. In terms of fit, the calibration does well in matching the probability of $H2M_{NW}$ status (22% in the model versus 22.7% in the data). It slightly over-represents (by 1.3 percentage points) the wealthy hand-to-mouth. The ratio of total wealth to income is lower in the simulation (3.21 versus 4.04), but the ratio of illiquid net worth to income is quite close (3.15 versus 3.27). Note that preference heterogeneity is crucial for matching these moments. Standard preferences, patient and not so elastic, do well in matching aggregate wealth, but generate few poor hand-to-mouth. Interestingly, simply reducing patience, while maintaining the lower IES, is rejected by the moment matching exercise. Finally, solely introducing a high IES, maintaining patience, generates too many wealthy hand-to-mouth, but few poor hand-to-mouth.

In terms of relative shares, Table 18 indicates that 73 percent of the agents are relatively patient and inelastic (that is, $\beta = 0.95, \sigma = 0.5$). From the preceding table, patience is important for generating a realistic wealth-to-income ratio. The calibration indicates that 11.3% of the population is the $\beta = 0.95, \sigma = 1.5$ type. As these agents spend nearly half their time in $H2M_{LIQ}$ status, they cannot be the majority of the population given the relatively small share of wealthy hand-to-mouth in the PSID sample (17.5%).

The remainder of the population is impatient but elastic agents ($\beta = 0.90, \sigma = 1.5$). Interestingly, the moment matching sets the share of impatient-inelastic agents to zero. Recall that in the one asset model, it was difficult to tell apart impatience and elasticity. However, in the two asset model, the role of liquidity clearly indicates that the impatient agents are also relatively elastic.

Table 18: Calibrated Preference Heterogeneity

	Calibrated Share	Share of Not $H2M$	Share of $H2M_{NW}$	Share of $H2M_{LIQ}$
$\beta = 0.95, \sigma = 0.5$	73.0%	90.0%	26.8%	73.4%
$\beta = 0.90, \sigma = 0.5$	0.0%	0.0%	0.2%	0.0%
$\beta = 0.95, \sigma = 1.5$	11.3%	9.4%	3.9%	25.9%
$\beta = 0.90, \sigma = 1.5$	15.7%	0.7%	69.2%	0.7%

The remaining columns of Table 18 report the preference composition of each hand-to-mouth state in the pooled cross-section. For example, 90% of the non-hand-to-mouth agents are the “standard” patient-inelastic type. These types hold much of the wealth in

Table 19: Aggregate MPC and IES

	Weighted Average	Not $H2M$	$H2M_{NW}$	$H2M_{LIQ}$
MPCY	0.64	0.57	0.89	0.55
MPC out of \$500	0.07	0.00	0.22	0.11
IES	0.54	0.60	0.68	0.21

the economy. Relative to their underlying 73% share of the total, they are infrequently poor hand-to-mouth, but only slightly under-represented in the wealthy hand-to-mouth sub-sample. Conversely, the patient-elastic types have a share of $H2M_{LIQ}$ that is more than double their population share. The vast majority (73.4%) of the poor hand-to-mouth are the impatient-elastic types, which is more than four times their 15.7% share of the general population.

The bottom line of the calibration is that the hand-to-mouth have a disproportionately high relative IES, with the poor and wealthy hand-to-mouth separated by their respective discount factors. The poor hand-to-mouth are predominantly low β-high σ agents, representing more than two-thirds of $H2M_{NW}$ but only 15.7% of the population. The majority of the agents are "standard" types, with relatively high discount factors and low IES; these agents are important for wealth accumulation, but relatively less important for hand-to-mouth status. That said, given that the latter group is such a large share of the population and hold much of the wealth, they constitute a large share of the wealthy hand-to-mouth.

5.3 Implications of Preference Heterogeneity

Using the calibrated model we can address how preference heterogeneity affects objects key to macroeconomic policy; specifically, the cross-sectional distribution of marginal propensities to consume as well as sensitivities of consumption to movements in interest rates.

The first row of Table 19 reports the marginal propensity to consume out of income (MPCY). Specifically, as in the empirical exercise behind Table 12, we regress the log change in consumption on the log change in income. As income in the model is a random walk, an innovation in income has a relatively large effect on permanent income. Hence, MPCYs in the model are quite large. The coefficient using the full simulated sample is 0.64; that is, consumption growth is on average 64 percent of the contemporaneous income growth. This number is lower for the non-hand-to-mouth. For the sub-sample of the poor hand-to-mouth ($H2M_{NW}$), the MPCY is nearly ninety percent. Despite the lack of liquidity, the $H2M_{LIQ}$ has an MPCY similar in magnitude to the non-H2M.

The second row computes the marginal propensity to consume out of a \$500 transfer. This is a pure transfer, and not offset by increased taxation. Using the policy functions, we compute the difference in consumption between having a liquid asset position of m versus m plus 500, dividing this difference by \$500. This is equivalent to randomly assigning a fraction of the simulated population an additional \$500 and then regressing the change in consumption on the amount received (either zero or \$500). The coefficient is therefore the percentage consumed by the recipients.

The average MPC is quite small, only 7%. For the non-hand-to-mouth, the MPC is zero. This reflects that the non-H2M are relative wealthy and liquid, and therefore close to the boundary at which they re-adjust their portfolios. For 14 percent of the non-H2M, the transfer induces a reallocation toward illiquid assets. Given that a fixed cost is paid for this adjustment, the agent simultaneously cuts back on consumption to increase the amount deposited. The MPC for the hand-to-mouth is larger. In particular, the $H2M_{NW}$ and $H2M_{LIQ}$ have MPCs of 0.22 and 0.11, respectively. Johnson, Parker and Souleles (2006) find that non-durable consumption increased by 37 cents on a dollar in response to the income tax rebates of 2001, while Parker, Souleles, Johnson and McClelland (2013) found an increase of non-durable spending of 12.8 cents on a dollar in response to the economic stimulus payments of 2008. In neither case was this response monotonic in liquid assets.

Table 20: Breakdown of Aggregate MPC

	Population Share	MPC out of \$500	MPC if not $H2M$	MPC if $H2M_{NW}$	MPC if $H2M_{LIQ}$
$\beta = 0.95, \sigma = 0.5$	73.0%	.015	.004	.070	.038
$\beta = 0.95, \sigma = 1.5$	11.3%	.132	-.018	.179	.292
$\beta = 0.90, \sigma = 1.5$	15.7%	.268	.020	.275	.354

Table 20 shows how preference heterogeneity plays a role in the MPCs. Each row is a particular type and the first column repeats the type's population share. The remaining columns report the average MPC, the MPC if not $H2M$, and the MPC conditional on being $H2M_{NW}$ and $H2M_{LIQ}$, respectively. The first-row concerns the most common type, that is, patient but inelastic consumers. Their MPC is small on average (0.015), essentially zero if not hand-to-mouth, and between four and seven percent if hand-to-mouth. Thus, the standard macro preferences generate very small MPCs, even conditional on low net-worth or low liquidity. Recall from Section 2 that conditional on assets, MPCs are higher if consumers are less patient or more elastic. This is borne out in the remaining rows of Table 20. The second row is the patient, elastic type. This type has an average MPC of 0.132. Recall that this type is similar to that of the original KV paper. Their benchmark rebate coefficient is 0.15, slightly higher than that reported in Table 20. The difference largely stems from the fact that the original KV agents are less patient ($\beta = 0.94$ vs. $\beta = 0.95$).[24] The final row of Table 20 reports the MPCs for the impatient, elastic types. These MPCs are quite high on average, primarily reflecting that they are prone to be in a hand-to-mouth status and have a high MPC conditional on being hand-to-mouth. Thus, the results indicate that not only asset positions, but the interaction of low assets or liquidity combined with high elasticities and/or low discount factors especially matters for MPCs.

Another perspective on how preference heterogeneity impacts MPCs is through a variance decomposition exercise. In the simulated population, the variance of MPCs across individuals is 0.036. We can decompose this number into the amount contributed by within-type variation and that explained by type differences. Letting $i \in I$ index individuals, $j \in J$ type, and s_j denoting the population share of type j, we have

$$Var(MPC_i) = \sum_j s_j Var(MPC_i|j) + \sum_j s_j \left(\mathbb{E}[MPC_i|j] - \mathbb{E}[MPC]\right)^2 .$$

The first term on the right is the weighted sum of within-type variances. The second term captures how the conditional mean MPC varies by type. For the simulation, the average within-type variance is 0.028, or 76% of the total. The remaining 24% is accounted for by differences in means by type, which totals 0.009. Hence, nearly a quarter of the total variation is accounted for by preferences heterogeneity. Even for within-type heterogeneity, the nature of preferences plays an important role. In particular, the within-type variance for the patient/inelastic agents is only 0.014, while it is 0.049 and 0.086 for the impatient and patient elastic types, respectively. That is, a high inter-temporal elasticity is crucial to generating a very dispersed within-type cross-sectional distribution of MPCs.

The final row of Table 19 reports the average IES. Because the consumer's problem is solved for a fixed set of interest rates, we do not have a policy function with respect to interest rate changes. This would require solving the model with a specific stochastic process for interest rates. We therefore consider the thought experiment of a one-time unanticipated small change in the one-period liquid interest rate (for both borrowing and saving). For consumers at an interior solution to their Euler Equation, this response is governed by σ.

[24] A second difference is the KV rebate coefficient is for a slightly different experiment. Specifically, they randomly assign a rebate to half the population in quarter t and to the remaining half in period $t + 1$. Thus, everyone is eventually treated, and the control group in the first quarter increases consumption in anticipation of the transfer. This lowers the rebate coefficient. Performing our experiment using the KV calibration generates an unanticipated transfer MPC of 0.16 out of $500.

Those with zero liquid assets, which face a discontinuity in the interest rate schedule, and those at the age-specific borrowing constraint, are assumed to be completely insensitive to small changes in interest rates and are assigned an IES of zero.[25] We abstract from the fact that in a richer model with interest rate shocks, changes in interest rates may also induce portfolio rebalancing, altering the response of consumption to changes in the interest rate.

The result of this experiment for the total population is an implied IES of 0.54, which is close to the σ of the majority of the population and in line with the standard one-agent macroeconomic calibration. The effective IES of the poor H2M is relatively high, at 0.68. This reflects that the majority of this group has $\sigma = 1.5$, combined with the fact that only 36.8 percent of the $H2M_{NW}$ are at the constraint. Thus, the poor hand-to-mouth may be relatively responsive to changes in the inter-temporal terms of trade, despite the presence of borrowing constraints.

The $H2M_{LIQ}$ tend to collect at zero liquid assets, due to the discontinuity in the interest rate schedule, or at the borrowing constraint; in particular, 63.7 percent of the $H2M_{LIQ}$ are at zero or the lower bound. Hence, despite the fact that more than a quarter of this sub-group has $\sigma = 1.5$, the average effective IES is only 0.21.

The heterogeneity suggests that the hand-to-mouth are not relatively insensitive to interest rate changes; if anything, they are more responsive. This abstracts from capital gains or losses due to interest rate policy shifts (as in Auclert (2019)). Nevertheless, it suggests that hand-to-mouth agents may be effective targets of policies that affect the inter-temporal terms of trade such as the Car Allowance Rebate System (aka cash-for-clunkers) of 2009, or sales-tax holidays. The result is also relevant for the timing of insurance. For example, the progressivity of social security benefits provides some insurance against adverse lifetime income shocks, but are not paid out until the eligibility age. Lack of precautionary savings suggests the delay in insurance may be costly for the hand-to-mouth; however, a relatively high inter-temporal elasticity of substitution suggests there is a countervailing force.

The calibration also suggests that MPCs are modest; indeed, the aggregate MPC is negligible and that of the wealthy unconstrained consumers is negative. However, the results do reinforce the conventional wisdom that low-wealth and low-liquidity agents have a relatively high MPC, both because of the usual concavity of the consumption function as well as the relatively high σ of these agents.

[25]This implies that those at the knife edge case of an interior optimum at zero liquidity or the constraint are strictly constrained.

5.4 Revisiting Consumption Behavior: Model versus Data

In Table 4 we documented that absent controls for type (either fixed effects or frequency of $H2M$ status), the hand-to-mouth do not have a higher growth rate of consumption on average. This runs counter to the standard model without preference heterogeneity discussed in Section 2. However, once we control for fixed effects or the frequency of $H2M$ status, we find in the data that the $H2M_{NW}$ have stronger growth on average. In particular, those prone to hand-to-mouth status tend to have a lower growth rate of consumption on average, making this an important control variable in determining the impact of current $H2M$ status.

Table 21: Consumption Growth: Data vs. Model

	Data			Model		
	(1)	(2)	(3)	(1)	(2)	(3)
$H2M_{NW}$	.003 (.004)	.025 (.006)	.025 (.007)	-.007	.004	.005
$H2M_{LIQ}$	-.009 (.004)	.003 (.005)	.002 (.005)	-.007	-.007	-.007
Fraction time $H2M_{NW}$			-.038 (.009)			-.017
Fraction time $H2M_{LIQ}$			-.025 (.007)			.004
Fixed effects	No	Yes	No	No	Yes	No

Note: The columns labeled "Data" reproduce the empirical estimates in Table 4. The columns labeled "Model" report the corresponding estimates based on model generated data.

In Table 21, we replicate the empirical exercise as well as perform the same set of regressions in our model simulation. The first column for the Model results indicates that current $H2M$ status (for both measures) predicts a weakly lower consumption growth rate. The second column includes fixed effects and increases the coefficient on $H2M_{NW}$ by 1.1 percentage points. The coefficient on $H2M_{LIQ}$, however, remains unchanged. Column (3) drops the fixed effect but adds the fraction of time spent in hand-to-mouth status as an additional control. We see that the agents that are frequently $H2M_{NW}$ have a lower growth rate on average, consistent with the data. Including this control generates a similar outcome on $H2M_{NW}$ as adding the fixed effect.

Interestingly, $H2M_{LIQ}$ is not sensitive to the inclusion of a fixed effect or the addition of the frequency control. Moreover, those that are frequently $H2M_{LIQ}$ tend to have a higher average consumption growth rate, contrary to the data. This reflects in part that in the model the $H2M_{LIQ}$ are richer on average than their empirical counterparts. In particular, they are relatively patient compared to the average discount factor in the simulated population. In the data, the liquidity constrained are much closer to the poor hand-to-mouth (both in terms of assets and behavior) than they are similar to the non-H2M. In the model, this is reversed; the $H2M_{LIQ}$ tend to look like the non-H2M.

Table 22: Average Propensity to Consume: Data vs. Model

	Data			Model		
	(1)	(2)	(3)	(1)	(2)	(3)
$H2M_{NW}$	.065 (.010)	-.007 (.010)	-.022 (.015)	.138	.016	.015
$H2M_{LIQ}$	.077 (.012)	.005 (.009)	.005 (.011)	-.006	.003	.004
Fraction time $H2M_{NW}$			.157 (.022)			.160
Fraction time $H2M_{LIQ}$			.191 (.025)			-.032
Fixed effects	No	Yes	No	No	Yes	No

Note: The columns labeled "Data" reproduce the empirical estimates in Table 11. The columns labeled "Model" report the corresponding estimates based on model generated data.

In Table 22, we revisit how the average propensity to consume varies with $H2M$ status. The first three columns of the table reproduce the empirical results reported in Table 11. The final three columns are the counterparts using the model's simulated population. As was the case in the previous table, the model's $H2M_{NW}$ track the behavior of their empirical counterparts. In particular, they have a relatively high APC with no controls for preference heterogeneity. However, including a fixed effect or the frequency of hand-to-mouth status lowers this difference substantially. This reflects that those prone to $H2M_{NW}$ status tend to have a high APC on average, even when not currently hand-to-mouth.

The model again fails to match the empirical behavior of the liquidity-constrained hand-to-mouth. In the model, those that are frequently $H2M_{LIQ}$ tend to have an APC that is similar to the non-hand-to-mouth (if not lower). Conversely, in the data those prone to be $H2M_{LIQ}$ have a much higher average APC than the non-H2M.

Finally, in Table 23 we revisit the marginal propensity to consume out of income (MPCY). In the model, MPCYs tend to be larger than in the data. This is not surprising, given that income is a random walk in the model and is measured without error. In the data, those that are frequently $H2M_{NW}$ tend to have a larger MPCY, a pattern confirmed in the model. Moreover, controlling for frequency of $H2M_{NW}$ raises the estimated MPCY of the current $H2M_{NW}$. Again, however, the $H2M_{LIQ}$ results generate a conflict between model and data.

Taking stock, the two-asset KV quantitative model, re-calibrated only to allow a limited set of plausible patience and inter-temporal elasticity parameters, does surprisingly well in matching the rich empirical patterns documented above. The quantitative model is particularly suited to match the behavior of the frequently $H2M_{NW}$, the sub-sample that, when assets are low, empirically shows the lowest expected consumption growth as well as highest average and marginal propensities to consume. The model allows us to pinpoint a willingness to substitute inter-temporally as a crucial characteristic behind the observed patterns of the hand-to-mouth in the data.

Table 23: Marginal Propensity to Consume out of Income: Data (IV) vs. Model

	Data		Model	
	(1)	(2)	(1)	(2)
$\Delta y/(y+ra)$	.082 (.020)	.045 (.022)	.333	.344
$\Delta y/(y+ra) \times H2M_{NW}$	.038 (.039)	-.076 (.062)	.229	.159
$\Delta y/(y+ra) \times H2M_{LIQ}$	.066 (.047)	-.023 (.052)	.037	.042
$\Delta y/(y+ra)$ $\times$ Fraction time $H2M_{NW}$		.188 (.087)		.078
$\Delta y/(y+ra)$ $\times$ Fraction time $H2M_{LIQ}$		.211 (.101)		-.074

Note: The columns labeled "Data" reproduce the empirical estimates in Table 13. The columns labeled "Model" report the corresponding estimates based on model generated data.

6 Conclusion

A workhorse model of savings in the literature has consumers self-insuring their income risk via a non-contingent asset subject to a borrowing constraint. This model has a number of predictions for the spending of low-asset households; namely, they should exhibit a lower average propensity to consume (APC) and higher expected consumption growth, as well as a higher marginal propensity to consume (MPC). But these predictions are masked or muddled in the data if differences in household asset holdings also reflect heterogeneity in preferences. We show that, if either a low discount factor or a high IES is what drives households to hold few assets, then these hand-to-mouth households can display higher APCs and lower spending growth. Moreover, such households have a relatively high MPC, even if they are at their target level of assets.

We see in the data that hand-to-mouth households actually display higher APCs and no faster spending growth than households with more assets, consistent with an important role for preference heterogeneity. These statements apply to households we label as hand-to-mouth based on their low net worth, or those we label as hand-to-mouth based only on their lack of liquid assets (the wealthier hand-to-mouth). In addition, low-asset households show more volatility in their spending, even conditioning on income, and adjust their spending to a greater extent by varying the number of categories consumed. The latter finding cannot be explained by heterogeneity in discount factors, but is consistent with low-asset households exhibiting a higher IES *because* they have a more active extensive margin to vary consumption.

Strikingly, all the "puzzles" we see in spending by hand-to-mouth households project on a household's long-run average propensity to be hand-to-mouth, not on its current asset position. That is, it is households that tend to hold few assets, in terms of liquid or illiquid, that display higher APCs, lower spending growth, more volatile spending, and more volatility in terms of categories of spending. Furthermore, it is households that are often hand-to-mouth, especially by net worth, that display the largest responses of spending to changes in their earnings. We view these findings as consistent with important, relatively stable, differences in preferences for hand-to-mouth versus other households.

To identify the source and consequences of this preference heterogeneity, we consider the two asset setting in Kaplan and Violante (2014) where agents allocate wealth between liquid and illiquid assets, but allow for heterogeneity in both households' discount factors and IES. A high-IES household, all else equal, will hold less liquidity to avoid short-term fluctuations in their consumption, thereby allowing us to distinguish a high IES from impatience. To match empirical targets, we find that nearly three quarters of the model population must exhibit the higher discount factor of 0.95 as well as the lower IES of 0.5. But to explain the significant share of households in the data with low net worth or low liquid wealth, the model requires that the balance of the population have the high IES of 1.5, with a subset also having a low discount factor. Most notably, the vast majority of the poor hand-to-mouth are both impatient and have a high IES. The wealthy hand-to-mouth, while not impatient, disproportionately have a high IES. Thus, we conclude that a high IES is at least as important as impatience in determining those households that are typically treated empirically as hand-to-mouth.

We find, from our two-asset model, that preference heterogeneity plays an important role in differences in MPCs across consumers. Finally, our model shows that the poor hand-to-mouth exhibit a higher effective IES than the balance of the population. This reflects the large share of poor hand-to-mouth with inter-temporally elastic preferences. Thus poorer households may be more responsive to inter-temporal terms of trade, despite their proximity to borrowing constraints.

References

Açikgöz, Ömer T., "On the Existence and Uniqueness of Stationary Equilibrium in Bewley Economies with Production," *working paper*, 2016.

Aguiar, Mark and Mark Bils, "Has Consumption Inequality Mirrored Income Inequality?," *American Economic Review*, September 2015, *105* (9), 2725–2756.

Aiyagari, S. R., "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, August 1994, *109* (3), 659–684.

Athreya, Kartik, José Mustre-del-Río, and Juan M. Sánchez, "The Persistence of Financial Distress," *The Review of Financial Studies*, October 2019, *32* (10), 38513883.

Auclert, Adrien, "Monetary Policy and the Redistribution Channel," *American Economic Review*, June 2019, *109* (6), 2333–2367.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, December 2008, *98* (5), 1887–1921.

Bound, John, Charles Brown, Greg J. Duncan, and Willard L. Rodgers, "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data," *Journal of Labor Economics*, 1994, *12* (3), 345–368.

Calvet, Laurent E., John Y. Campbell, Francisco J. Gomes, and Paolo Sodini, "The Cross-Section of Household Preferences," 2019.

Carroll, Christopher D., "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 1992, *1992* (2), 61.

_ **and Miles S. Kimball**, "On the Concavity of the Consumption Function," *Econometrica*, 1996, *64* (4), 981–992.

_ **, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White**, "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, 2017.

_ **, Thomas F. Crossley, and John Sabelhaus**, *Improving the Measurement of Consumer Expenditures*, University of Chicago Press, 2015.

Chetty, Raj and Adam Szeidl, "Consumption Commitments and Risk Preferences*," *The Quarterly Journal of Economics*, 05 2007, *122* (2), 831–877.

Deaton, A., "Saving and Liquidity Constraints," *Econometrica*, 1991, *59* (4), 1221–1248.

Dynan, Karen E., Jonathan Skinner, and Steven P. Zeldes, "Do the Rich Save More?," *Journal of Political Economy*, 2004, *112* (2), 397–444.

Farhi, Emmanuel and Ivan Werning, "Fiscal Multipliers: Liquidity Traps and Currency Unions," *Handbook of Macroeconomics*, 2017, *2*, 2417–2492.

Fisher, Jonathan, David Johnson, Timothy Smeeding, and Jeffrey P. Thompson, "Estimating the Marginal Propensity to Consume Using the Distributions of Income, Consumption, and Wealth," 2019.

Foster, Kevin, Scott Schuh, and Hanbing Zhang, "The 2010 Survey of Consumer Payment Choice," *Federal Reserve Bank of Boston Research Review*, 2013, *13.2.*

Gelman, Michael, "What Drives Heterogeneity in the Marginal Propensity to Consume? Temporary Shocks vs Persistent Characteristics," 2019.

Grossman, Sanford J. and Guy Laroque, "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica*, 1990, *58* (1), 25–51.

Havranek, Tomas and Anna Sokolova, "Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say "probably not"," *Review of Economic Dynamics*, May 2019.

Havrnek, Tom, "Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting," *Journal of the European Economic Association*, 2015, *13* (6), 1180–1204.

Huggett, Mark, "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of Economic Dynamics and Control*, September 1993, *17* (5-6), 953–969.

Imrohoroglu, Ayse, "Cost of Business Cycles with Indivisibilities and Liquidity Constraints," *Journal of Political Economy*, 1989, *97* (6), 1364–1383.

Jappelli, Tullio and Luigi Pistaferri, "Fiscal Policy and MPC Heterogeneity," *American Economic Journal: Macroeconomics*, October 2014, *6* (4), 107–136.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, December 2006, *96* (5), 1589–1610.

Kaplan, Greg and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 2014, *82* (4), 1199–1239.

_ , Benjamin Moll, and Giovanni L. Violante, "Monetary Policy According to HANK," *American Economic Review*, March 2018, *108* (3), 697–743.

_ , Giovanni L. Violante, and Justin Weidner, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 2014, *45* (1 (Spring), 77–153.

Krueger, D., K. Mitman, and F. Perri, "Macroeconomics and Household Heterogeneity," in "Handbook of Macroeconomics," Elsevier, 2016, pp. 843–921.

McKay, Alisdair and Ricardo Reis, "The Role of Automatic Stabilizers in the U.S. Business Cycle," *Econometrica*, 2016, *84* (1), 141–194.

Michelacci, Claudio, Luigi Paciello, and Andrea Pozzi, "The Extensive Margin of Aggregate Consumption Demand," 2019.

Parker, Jonathan A., "Why Don't Households Smooth Consumption? Evidence from a $25 Million Experiment," *American Economic Journal: Macroeconomics*, October 2017, *9* (4), 153–183.

_ , Nicholas S. Souleles, David S. Johnson, and Robert McClelland, "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, October 2013, *103* (6), 2530–53.

Rogerson, Richard, "Indivisible labor, lotteries and equilibrium," *Journal of Monetary Economics*, January 1988, *21* (1), 3–16.

Schechtman, Jack and Vera L.S. Escudero, "Some results on an income fluctuation problem," *Journal of Economic Theory*, 1977, *16* (2), 151 – 166.

Straub, Ludwig, "Consumption, Savings, and the Distribution of Permanent Income," 2019.

Tonetti, Christopher, "Notes on Estimating Earnings Processes," *working paper*, 2011.

Zeldes, Stephen P., "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 1989, *97* (2), 305–346.

Appendices for "Who Are the Hand-to-Mouth"

A1 Data Appendix

A1.1 Description of PSID sample

Our primary data source is the Panel Study of Income Dynamics (PSID) biannual surveys from 1999 to 2015. The advantage of the PSID for our purposes is that it provides measures of income, assets, and expenditures. Income measures were a focus of the PSID from it onset in 1968 (hence its name). The PSID introduced a module to measure assets and liabilities in 1984 that reappeared every five years. Beginning 1999 the PSID includes a wealth module in every survey. That is one reason we begin our sample in 1999. The second reason is that the PSID first began surveying households on a number of expenditure categories, beyond food and housing, with the 1999 survey.

We focus here on our variable constructions for the key variables of earnings, income, wealth, and expenditures. We then detail the sample restrictions we employ.

The analysis separately considers earnings income and a broader measure of after-tax income that includes net income from assets, including owner-occupied housing. We measure earnings by wage and salary income, net of payroll taxes, plus the head of households's labor component of income from any unincorporated business, and one half of family farm income. We add to these earnings any receipts of government transfer payments from AFDC, supplemental security income, other welfare payments, veteran's pensions, unemployment benefits, worker's compensation, or social security benefits. To construct after-tax income, we first sum taxable income (earnings, net profits from business or farm, and income from assets), transfer income, and social security income for the husband and wife as well as other family members. From this we subtract the family's federal and state income tax liabilities as measured by the TAXSIM program. For homeowners we then add 6 percent of the respondent's assessed value of their home to account for the implicit rent on their home, while subtracting payments for property taxes, mortgage interest, and home insurance.

We define a household net worth as the sum of its liquid and illiquid assets net of debts. We treat liquid net worth as the sum of balances in checking or savings accounts, money market funds, certificates of deposit, holdings of treasury bills and other government savings bonds, the value of stocks outside of pension funds, minus the value of all debts. The values for checking or savings accounts, money market funds, certificates of deposit, treasuries and other government bonds are multiplied by 1.055, to reflect cash holdings that are not reported in the survey. See Foster, Schuh and Zhang (2013) for justification. Illiquid wealth is the sum of a household's home equity, equity in other real estate, holdings of IRAs and other pensions, the value of bonds (not including treasury or other government bonds), insurance holdings, the value of any business or farm net of debts, and the value of any vehicles (including motor homes), boats, and trailers net of debt owed. These distinctions for liquid versus illiquid assets largely follow Kaplan et al. (2014), while fitting within the grouping of assets within the PSID questionnaire. Our stratification of households into not hand-to-mouth, hand-to-mouth by net worth, and hand-to-mouth by liquid wealth are based on these measures of assets relative to our broad measure of earnings, as discussed at the beginning of Section 4.

Our base measure for expenditures includes spending categories for shelter, utilities (by type), food for consumption at home, food for consumption away from home, gasoline, health insurance, medical expenses (separately for doctors, hospitals, and prescription drugs), education, child care, purchases or lease of vehicles, vehicle repair, vehicle insurance, parking, and public transportation (by type). Spending on shelter reflects rent payments for renters; for homeowners we set it to 6 percent of respondent's valuation of the home. Our measure of nondurable and services spending excludes spending on vehicles or their repair from the base measure. Beginning with the 2005 survey, the PSID added categories for home repairs, home furnishings, clothing, vacations, recreation, and telecommunications. We consider robustness to this broader measure in judging households spending relative to income (APC). Our base measure of expenditures relative to after-tax income averages 58.3 percent. For the broader measures of expenditures available beginning with the 2005 survey, this average is 73.2 percent.

In addition to controlling for year and age effects, our regression analysis includes controls for marital status (single or married/cohabiting), race (three values), and family size. Family size takes five distinct values, with 5 representing family sizes of 5 and above. For regressions in growth form, e.g., growth rate of expenditures over two years, the controls reflect year and age dummies, and a set of dummies for the conceivable changes in marital status and for whether family size increased, stayed the same, or decreased.

Respondents report earnings, income, and taxes for the previous calendar year, whereas they report assets and liabilities as of the interview. Expenditures are reported for differing time frames. Among categories available from the 1999 survey, education spending is for the prior calendar year, health spending (including health insurance) for the previous two calendar years, and vehicle spending for since the survey two years prior. Other categories are in terms of the household's usual (typically monthly) expenditures. We treat these variables as aligned with respect to the previous calendar year, with assets viewed as end of period.

We deflate nominal variables by the corresponding CPI measured in 2009 dollars.

Our sample reflects only the PSID's nationally representative core sample (i.e., we use the Survey Research Center sample, excluding the Survey of Economic Opportunity.). This sample includes "split-off" families from the orignal sample as well as the PSID sample extensions to better represent the families of immigrants and recent immigrants. Throughout the analysis we employ the PSID longitudinal family weights, which are designed to correct for non-random sample attrition as well as failures to draw an entirely random sample.

We restrict our sample to households with heads ages 25 to 64. We exclude households with less than $2,000 (2009 $'s) in any of annual earnings (including transfer receipts), after-tax income, or annual expenditures. We also exclude households with extreme responses on expenditures in which food purchases for consumption at home are zero, or spending on housing and food (home and away) is less than 5% or greater than 90% of total expenditures. Finally, we include only households whose hand-to-mouth status, from their earnings and asset information, can be measured for at least three surveys. Table A1 displays the impact of these restrictions sequentially for our resulting sample, both in terms of households and number of observations. It also shows the sample impact of examining two-year growth rates, such as income or expenditure growth.

Table A1: Impact of Sample Restrictions

Restriction	Households	Observations
Ages 25 to 64	7,572	35,766
Earnings & Inc.≥$2000	7,476	34,641
Expenditures≥$2000	7,472	34,565
No odd spending	7,341	33,325
H2M status 3+ times	4,907	27,134
WRT 2-year changes	4,767	21,569

Note: PSID data, 1999 to 2015 survey waves.

A1.2 Description of the Consumer Expenditure Survey Sample

We employ a sample derived from the Consumer Expenditure Surveys (CE) to augment our evidence from the PSID that low-asset houshold concentrate their spending on a narrower set of categories. We make use of the CE data on income, assets, and expenditure from survey years 1996 to 2016.

The CE surveys households on their expenditures for up to four consecutive quarters. We only include households that were surveyed in the fourth interview. In this interview they are asked about asset and debt holdings and about their income over the previous 12 months. The CE does not provide information on households pensions, and is more limited than the PSID in collecting information on some other forms of illiquid assets. Therefore, we stratify households by assets only with respect to their liquid assets, using the definition in Kaplan et al. (2014) for households that are hand-to-mouth. The asset information for liquid assets parallels that in the PSID. It includes the household's balances in checking or savings accounts, money market funds, certificates of deposit, holdings of treasury bills and other government savings bonds, the value of stocks outside of pension funds. Debts include credit card and store credit debt, student loans, and medical or personal loans.

We express a household's assets relative to its annual income. That income sums household earnings, farm and business income, retirement payments including from social security, government transfers, and alimony receipts. Our income measure is before taxes. (The corresponding CE variable for after-tax income was eliminated after 2015.)

Each household's quarterly expenditures are divided into many distinct categories. To examine the share of categories purchased, we exclude durable categories as well as utilities (e.g., water, gas, electricity, trash collection) that we view as tied to the choice of housing. The resulting 27 nondurable categories are listed in Table A2. The table reports the fraction of households that spend on the category in a quarter, as well as the average fraction adding or dropping the category in a quarter.

A2 Additional Empirical Results

A2.1 Results Splitting by Age and by Long-term Income

Table A3 reports spending growth and APC for the H2M, by age group. The first age group includes households with heads ages 25 to 39. The second age group includes households with heads ages 40 to 64.

Table A4 reports spending growth and APC for the H2M, by long-term earnings. Long-term earnings are defined as a household's average natural log of earnings after removing a cubic function of the head's age and year dummies.

Table A2: Nondurable Expenditure Categories, CE

Category	Share	Positive	Add	Drop
Food at home	0.289	1.00	0.00	0.00
Motor fuel	0.113	0.92	0.01	0.01
Food away from home	0.097	0.86	0.06	0.06
Telephone services	0.072	0.96	0.02	0.02
Health insurance	0.067	0.61	0.04	0.03
Motor vehicle insurance	0.052	0.61	0.12	0.12
Professional services	0.035	0.52	0.14	0.13
Video and audio – services	0.034	0.74	0.04	0.04
Tuition, other school fees, and childcare	0.034	0.19	0.06	0.06
Tobacco and smoking products	0.026	0.28	0.04	0.04
Public transportation	0.020	0.24	0.11	0.11
Tenants' and household insurance	0.016	0.33	0.08	0.08
Personal care services	0.015	0.68	0.10	0.10
Club memberships, fees for sports, lessons or instructions	0.014	0.33	0.10	0.10
Prescription drugs	0.014	0.44	0.12	0.11
Information processing other than telephone services	0.013	0.52	0.07	0.06
Alcoholic beverages at home	0.012	0.40	0.09	0.09
Admission tickets	0.012	0.45	0.14	0.14
Motor vehicle fees	0.012	0.49	0.19	0.18
Lodging away from home – hotel	0.011	0.19	0.12	0.12
Household operations – nondurables	0.010	0.24	0.09	0.09
Laundry and other apparel services	0.009	0.40	0.10	0.11
Financial and legal services	0.009	0.29	0.12	0.13
Alcoholic beverages away from home	0.008	0.32	0.10	0.10
Hospital and related services	0.005	0.06	0.05	0.05
Elderly care, funeral and dating services	0.002	0.02	0.02	0.02
Medical supplies	0.001	0.02	0.01	0.01

A2.2 More on the Extensive Category Margin

The results of Tables 15 and 16 indicate that the hand-to-mouth consume fewer distinct categories, but move into and out of categories more elastically. To explore this further, we decompose the growth in nondurable consumption between $t-1$ and t into three components: the change in spending on goods consumed in both periods (the "intensive" margin); the addition of new goods; and the dropping of old goods. In particular, suppose individual i consumes N_k categories of goods in period $k = t-1, t$. Let I denote the set of categories consumed in both $t-1$ and t. Let A denote the set of categories added in period t, and D the set of categories dropped between $t-1$ and t. Hence, $N_{t-1} = I \cup D$ and $N_t = I \cup A$.

Table A3: Consumption Growth and APC for H2M, by Age

	Consumption Growth				APC			
	Ages 25 to 39		Ages 40-64		Ages 25 to 39		Ages 40-64	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(9)
$H2M_{NW}$	.003	.020	.005	.032	.049	-.056	.079	.012
	(.006)	(.009)	(.006)	(.011)	(.013)	(.023)	(.015)	(.019)
$H2M_{LIQ}$	-.007	.009	-.011	-.001	.066	-.012	.088	.015
	(.006)	(.008)	(.005)	(.007)	(.016)	(.018)	(.015)	(.014)
Fraction time $H2M_{NW}$		-.032		-.046		.200		.115
		(.011)		(.013)		(.032)		(.031)
Fraction time $H2M_{LIQ}$		-.038		-.016		.191		.196
		(.013)		(.009)		(.033)		(.032)

Note: Sample is PSID for 1999-2015, with seeing H2M status at least 3 times. Regressions include usual controls. Standard errors are clustered at household level.

Table A4: Consumption Growth and APC for H2M, by Long-Term Income

	Consumption Growth				APC			
	Bottom third		Top two-thirds		Bottom third		Top two-thirds	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(9)
$H2M_{NW}$	.016	.039	.004	.018	.013	-.024	.032	-.007
	(.008)	(.014)	(.005)	(.008)	(.022)	(.033)	(.009)	(.012)
$H2M_{LIQ}$	.005	.021	-.012	-.006	.085	.034	.027	-.009
	(.009)	(.012)	(.004)	(.005)	(.026)	(.026)	(.009)	(.011)
Fraction time $H2M_{NW}$		-.043		-.027		.073		.072
		(.016)		(.010)		(.047)		(.018)
Fraction time $H2M_{LIQ}$		-.037		-.012		.137		.104
		(.016)		(.008)		(.047)		(.023)

Note: Sample is PSID for 1999-2015, with seeing H2M status at least 3 times. Regressions include usual controls. Standard errors are clustered at household level.

Let $x_{n,k}$ denote expenditure on good n in period $k = t-1, t$, always expressed in period t prices. We decompose growth in expenditure between $t-1$ and t as follows:

$$\begin{aligned}\frac{c_t - c_{t-1}}{0.5(c_t + c_{t-1})} &= \frac{\sum_{n\in N_t} x_{n,t} - \sum_{n\in N_{t-1}} x_{n,t-1}}{0.5(c_t + c_{t-1})} \\ &= \underbrace{\frac{\sum_{n\in I}(x_{n,t} - x_{n,t-1})}{0.5(c_t + c_{t-1})}}_{\text{Intensive}} + \underbrace{\frac{\sum_{n\in A} x_{n,t}}{0.5(c_t + c_{t-1})}}_{\text{Add}} + \underbrace{\frac{-\sum_{n\in D} x_{n,t-1}}{0.5(c_t + c_{t-1})}}_{\text{Drop}}.\end{aligned}$$

To obtain the contribution of the sub-components, we individually regress the three measures on the right-hand side of this decomposition on the total growth rate of non-durable expenditure defined on the left-hand side. Mechanically, the coefficients from the three regressions will add up to one. We run this decomposition for the pooled group of individuals in the PSID, as well as the non-hand-to-mouth and the hand-to-mouth separately.

Table A5 reports the results. The estimates indicate that the H2M households are relatively prone to adding and dropping goods as they adjust expenditure while those with higher wealth tend to operate more on the intensive margin. (The p-value of the tests that the elasticities are the same across the two groups are all well below one percent.)

Table A5: Decomposition of Spending Growth by H2M status

Status	Not $H2M$	$H2M_{NW}$	$H2M_{LIQ}$
Intensive	0.736	0.677	0.734
Add	0.116	0.138	0.114
Drop	0.148	0.186	0.152

Note: Regressions include controls for age, year, and changes in family size or marital status. HtM status observed at least three times.

A3 The Extensive Margin and the IES: A Simple Two-Good Example

In this appendix we provide a simple model that links the extensive margin analysis of Section 4.8 and the inter-temporal elasticity of substitution. The example is designed to deliver a transparent and plausible explanation about why certain consumers are prone to be highly elastic at the margin in terms of total expenditure's response to relative price (including interest rate) movements.

Suppose there are two goods, c_1 and c_2 and utility is given by $u(c_1, c_2 - \underline{c}y)$, where y is average income and $\underline{c}$ is a parameter that captures a minimum consumption level as a fraction of income. To make things simple, suppose both goods trade at price 1 and period income is $y = 1$. We shall contrast two individuals with differing $\underline{c}$.

To make things concrete, let the indirect utility function over expenditure be given by:

$$v(x) = \max_{c_1, c_2} c_1^{\rho} + \mathbb{1}_{[c_2 > \underline{c}y]}(c_2 - \underline{c}y)^{\rho}$$
$$\text{subject to } c_1 + c_2 \leq x,$$

where $\mathbb{1}_x$ is an indicator function that takes value if x is true and zero otherwise, and ρ is a parameter. The key static decision is to consume both goods versus only good 1. Figure A1 plots $v(x)$ for $\underline{c} = 1$ (left panel) and $\underline{c} = 0.25$ (right panel). The decision between one versus two goods is to choose the max of the two alternatives (where the two-good option is depicted by the dashed line). Of course, the switch from one good to two occurs at much lower expenditures levels for low $\underline{c}$. That is, conditional on spending, the agent with low $\underline{c}$ is likely to consume fewer categories, just like the $H2M$ in the data. The important point is whether the point at which the agent switches is far or close to the typical level of expenditure. Keep in mind that the decision in the static problem is invariant to monotonic transformations. Hence, one can make the convex kink as dramatic or negligible as one wishes without altering the decision of whether to consume the second good.

Figure A1: $v(x)$ as a function of x

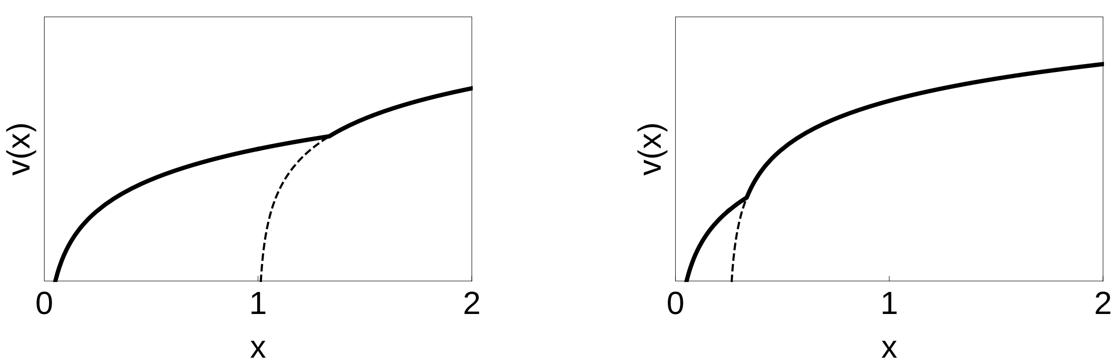

(a) $\underline{c} = 1$ (b) $\underline{c} = 0.25$

To see how this affects the inter-temporal problem, suppose the individual has the following time-separable utility over two periods, $t = 1, 2$:

$$V(x_1) + \beta V(x_2), \tag{5}$$

where

$$V(x) = \frac{v(x)^{\frac{1-\gamma}{\rho}}}{1-\gamma}.$$

Given a deterministic income process, y_t, and an interest rate R, the consumer's problem is to maximize (5) subject to $x_1 + R^{-1}x_2 = y_1 + R^{-1}y_2$. Setting $\beta = R^{-1} = 0.98$ and $y_1 = y_2 = y = 1$, Figure A2 plots the value of the objective as we vary x_1 and letting $x_2(x_1) = (2+r)y - Rx_1$. (We also set $\rho = 1/3$, $\gamma = 1.01$.)

Figure A2: $V(x_1) + \beta V(x_2(x_1))$ as a function of x_1

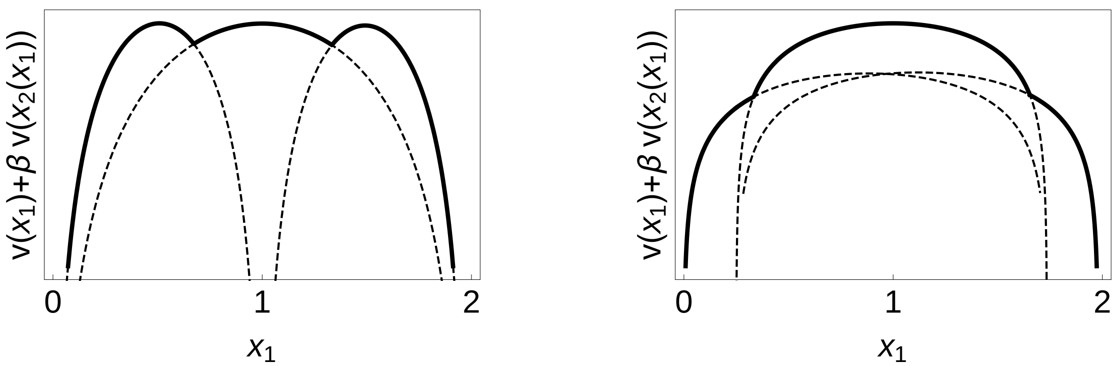

(a) $\underline{c} = 1$ (b) $\underline{c} = 0.25$

From left to right in Panel (a) of Figure A2, the dashed lines denote the value from (i) consuming one good in period 1 and two goods in period 2; (ii) consuming one good in both periods; and (iii) consuming two goods in period one and one good in period two. As drawn, the individual is indifferent over the three choices. The important point is that small movements in inter-temporal prices may lead to large shifts in first-period expenditure. In the right panel, the dashed lines denote the value from (i) consuming one good in period 1 and two goods in period 2; (ii) consuming two goods in both periods; and (iii) consuming two goods in period one and one good in period two. Here, consuming both goods in both periods is clearly optimal, and doing so is robust to small movements in the interest rate. In other words, the agent with low $\underline{c}$ is more likely to adjust at the extensive margin, similarly to the $H2M$ in the data.

Figure A3 plots the indifference curves for the two period problem. That is, it depicts points (x_1, x_2) for constant $V(x_1) + \beta V(x_2) = \bar{V}$ for various values of $\bar{V}$. It also includes the budget line. Again, the non-convex portion is relevant for the left-hand $\underline{c}$, but not the right. Small changes in R (the slope of the budget line), can have a big effect on the inter-temporal spending of the agent with low $\underline{c}$, who willingly shifts spending over time via adding or dropping the second good. Note that this is akin to this agent having a higher IES.

Figure A3: Indifference Curves

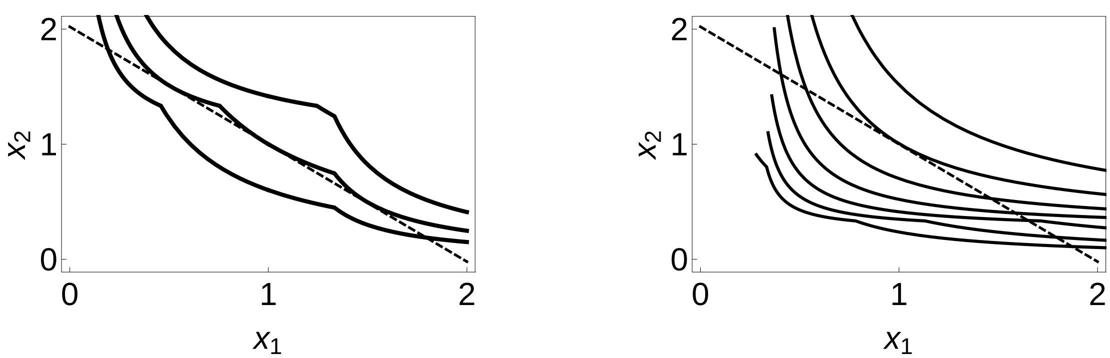

(a) $\underline{c} = 1$ (b) $\underline{c} = 0.25$