The Marginal Propensity to Consume in Heterogeneous Agent Models*

Greg Kaplan† and Giovanni L. Violante‡

July 15, 2022

Abstract

What model features and calibration strategies yield a large average marginal propensity to consume (MPC) in heterogeneous agent models? Through a systematic investigation of models with different preferences, dimensions of ex-ante heterogeneity, income processes and asset structure, we show that the most important factor is the share and type of hand-to-mouth households. One-asset models either feature a trade-off between a high average MPC and a realistic level of aggregate wealth, or generate an excessively polarized wealth distribution that vastly understates the wealth held by households in the middle of the distribution. Two-asset models that include both liquid and illiquid assets can resolve this tension with a large enough gap between liquid and illiquid returns. We discuss how such return differential can be justified from the perspective of theory and data.

Keywords: Borrowing Constraints, Consumption, Hand-to-Mouth, Heterogeneity, Income Risk, Liquidity, Marginal Propensity to Consume, Market Incompleteness, Precautionary Saving, Wealth Distribution.

JEL codes: D15, D31, D52, E21, E62, E71, G51

*We thank Chase Abram and Brian Livingston for outstanding research assistance.

†University of Chicago and NBER: gkaplan@uchicago.edu

‡Princeton University, CEBI, CEPR, IFS, IZA, and NBER: violante@princeton.edu

1 Introduction

The marginal propensity to consume (MPC) is the fraction of a small, unanticipated one-time windfall that a household spends within a given time period. In this article we conduct a systematic investigation of the size and determinants of the aggregate MPC in heterogeneous agent incomplete markets models – models whose key features are uninsurable idiosyncratic income risk, a precautionary savings motive, and n endogenous wealth distribution.

The MPC is a central concept in modern macroeconomics. In versions of these models with nominal rigidities, it determines the size of fiscal multipliers (Kaplan and Violante, 2014; Auclert, Rognlie, and Straub, 2018), the transmission mechanism of monetary policy (Auclert, 2019; Kaplan, Moll, and Violante, 2018), the amplification of aggregate shocks (Bilbiie, 2020; Patterson, 2019; Werning, 2015; Kekre and Lenel, 2021), and portfolio choice between risky and safe assets (Kimball, 1990). For these reasons, the usefulness of these models is closely tied to their ability to reproduce the evidence on MPCs.

As we explain below, there exists a tension between data on the household wealth distribution and the MPCs generated by baseline versions of these models. Our goal in this paper is to resolve some of the disagreement and confusion among economists about the sources of this tension and the alternative solutions that have been proposed.

In most of our analysis, we will focus on the quarterly MPC out of a $500 windfall. This size and frequency aligns with most of the empirical evidence on MPCs, and corresponds roughly to the size of fiscal stimulus payment policies. A large body of empirical evidence indicates that the average quarterly MPC on non-durable goods and services out of transitory income changes of $500-$1,000 is between 15% and 25% (see Jappelli and Pistaferri (2010) and Havranek and Sokolova (2020) for surveys).[1] Moreover, this average masks substantial heterogeneity in MPCs. Many households have MPCs that are close to zero, but some households have MPCs not far from 1, with a great deal of variation in between.[2] Some of this heterogeneity is explained by the distribution of liquid wealth and some by fixed individual characteristics, but most is left unexplained by observable data. There is also evidence that households respond more strongly

[1] Three approaches have been used to measure empirical MPCs: (i) quasi-experimental evidence exploiting, for example, historical episodes of economic impact payments (Johnson, Parker, and Souleles, 2006; Parker, Souleles, Johnson, and McClelland, 2013; Kueng, 2018; Aydin, 2022), government shutdowns (Gelman, Kariv, Shapiro, Silverman, and Tadelis, 2020) and lottery wins (Fagereng, Holm, and Natvik, 2019; Golosov, Graber, Mogstad, and Novgorodsky, 2021); (ii) survey instruments that pose hypothetical questions to households about how they would spend income windfalls (Shapiro and Slemrod, 2003, 2009; Jappelli and Pistaferri, 2014; Fuster, Kaplan, and Zafar, 2021; Parker and Souleles, 2019; Bunn, Le Roux, Reinold, and Surico, 2018); and (iii) semi-structural methods that identify transitory income changes (Blundell, Pistaferri, and Preston, 2008; Commault, 2017; Ganong, Jones, Noel, Greig, Farrell, and Wheat, 2020; Andersen, Jensen, Johannesen, Kreiner, Leth-Petersen, and Sheridan, 2021).

[2] See, e.g., Misra and Surico (2014), Gelman (2021), and Lewis, Melcangi, and Pilossoph (2019).

to smaller windfalls than to larger windfalls, and to negative income shocks than to positive income shocks.

But despite their importance for macroeconomics, until recently this evidence on MPCs has made only a guest appearance in modern macroeconomics, because of the representative agent assumption implicit in most models. With a representative agent, or complete markets, there is a single common MPC that is roughly the same size as the interest rate, so these models are unable to speak to the evidence. In the last two decades, macroeconomic models have evolved. Today, the heterogeneous agent incomplete markets framework is widely accepted to be a more useful model of the household sector than the representative agent complete markets framework. As a result, MPCs have regained a important role as a point of connection between models and data. Relative to representative agent models, heterogeneous agent models can deliver two realistic features of MPCs: (i) the aggregate MPC can be much larger than the interest rate, and (ii) the distribution of MPCs can be very dispersed.

From a theoretical perspective, there are three reasons why this class of models can generate large and heterogeneous MPCs: (i) hand-to-mouth households, (ii) precautionary saving, and (iii) ex-ante heterogeneity. First, these models feature poor hand-to-mouth households (HtM, hereafter) – households with zero wealth and borrowers for whom credit constraints bind. These types of households consume the majority of any extra liquidity they receive. Two-asset versions of the models with liquid and illiquid assets also feature *wealthy* HtM households who hold the bulk of their wealth in illiquid assets. These households' consumption responds to small windfalls similarly to the consumption of poor HtM households (Kaplan and Violante, 2014; Kaplan, Violante, and Weidner, 2014).

Second, uninsurable risk, combined with a consumption smoothing motive from concave utility, creates two forces that increase the MPC. If the utility function displays prudence ($u''' > 0$), households save more to avoid the consequences of future negative income shocks (Carroll and Kimball, 1996). Even in absence of prudence, hosueholds who face an occasionally binding liquidity constraints will engage in precautionary saving to avoid hitting the constraint (Zeldes, 1989; Deaton, 1991; Carroll and Kimball, 2001). These precautionary saving motives are especially strong at low levels of wealth, but disappear as wealth grows sufficiently large. As a result precautionary motives lower the level of consumption relative to the certainty equivalent case, which is linear in wealth, but does so more at low wealth levels. The consumption function is therefore concave in wealth and MPCs can be large for low wealth households, even for those who are not hand-to-mouth. Carroll (2001) and Carroll, Holm, and Kimball (2021) contain lucid expositions of these arguments and exhaustive surveys of the literature.

Third, households can differ due to ex-ante characteristics, as opposed to random circum-

stances. Impatience, high willingness to substitute, low returns on saving or behavioral biases can all be sources of high MPCs (Aguiar, Bils, and Boar, 2020; Carroll, Slacalek, Tokuoka, and White, 2017). These sources of heterogeneity affect MPCs both through their affect on the wealth distribution (and therefore the fraction of households in the the former two circumstances) and through their effect on the consumption function at a given level of wealth.

But how close can this class of models come to matching the evidence on MPCs, when constrained to be consistent with key features of the wealth and earnings distributions? Which features of the models are most important for their quantitative performance? By solving multiple versions of the model under different calibration strategies, we arrive at six main findings:

1. One-asset precautionary saving models calibrated to be consistent with measures of aggregate wealth in US data yield aggregate quarterly MPCs between 3% and 5%. This is an order of magnitude bigger than in corresponding representative agent models, but much lower than their empirical counterparts.

2. Calibrations that target liquid wealth or the share of hand-to-mouth households can generate an average MPC that is as large, or larger, than in the data. However, these versions of the model abstract from more than 98% of wealth in the economy and so are of limited use in general equilibrium applications.

3. Extensions of the one-asset model that incorporate either ex-ante heterogeneity or behavioral preferences can generate MPCs as large as in the data, while remaining consistent with the amount of aggregate wealth. However, these versions of the model vastly understate the amount of wealth held by households in the middle of the distribution and feature an excessively polarized wealth distribution. For example, median wealth is 5 to 10 times smaller than in the data. We call this the 'missing middle' problem.

4. Judiciously chosen versions of spender-saver models can generate MPCs as large as in the data, without the missing middle problem. However, these models have other implausible implications, such as matching the size and sign dependence of MPCs observed in data.

5. Two-asset precautionary savings models can reconcile all these tensions, but baseline versions requires a sizable gap in financial returns between illiquid and liquid assets.

6. Extensions of the two-asset model that introduce non-financial benefits from holding illiquid assets can simultaneously match the MPC and wealth data, with smaller differences in financial returns. Examples include direct utility flows, such as the benefits of owner-occupied housing, and indirect benefits such as those that arise in models of temptation

where households value the commitment that comes from holding wealth in illiquid assets.

Outline Section 2 analyzes the canonical one-asset incomplete-markets model. Section 3 extends the one-asset model model in several directions. Section 4 analyzes the two-asset model. Section 5 discusses other concepts of MPCs. Section 6 concludes the paper. The Appendix contains additional details about data, models, calibrations, and simulations.

2 One-Asset Models

In this section, we present findings from various versions of a standard one-asset precautionary savings model. Our one asset models are formulated in discrete time, but we also examine a continuous time version for consistency with the two-asset models in Section 4.

2.1 Baseline One-Asset Model

Environment. The economy is populated by a measure one continuum of households who survive each period with probability $(1-\delta)$. Conditional on surviving, a households' time preference factor is given by $\tilde{\beta}$, implying an effective discount factor of $\beta = \tilde{\beta}(1-\delta) < 1$. Period utility is given by $u(c_t)$, where u is strictly increasing and concave, and c_t denotes consumption expenditures. At each date t, households are endowed with labor income y_t which follows an exogenous stochastic process described below. Income draws are IID across households. Households can save, but not borrow, in a risk-free asset b_t with rate of return $R = 1 + r$. The household problem is:

$$\begin{aligned} \max_{\{c_t\}} \quad & \mathbb{E}_0 \quad \sum_{t=0}^{\infty} \beta^t u(c_t) \\ & \text{subject to} \\ c_t + b_{t+1} \quad &= \quad Rb_t + y_t, \quad b_{t+1} \geq 0, \quad b_0 = 0 \end{aligned} \tag{1}$$

The solution to the household problem yields decision rules for consumption $c(b,y)$ and next period's wealth $b'(b,y)$. These decision rules induce a stationary distribution that we denote by $\mu(b,y)$, with associated marginal distribution $\mu(b)$ over wealth.[3]

[3] If the income process has multiple components, then with a slight abuse of notation y should be intended as the vector of all income states.

Parameterization. Consistent with our focus on quarterly MPCs, our baseline discrete time model has a period of one quarter. We set $\delta = 1/200$ so that the expected adult life span is 50 years. In our baseline parameterization we assume a constant-elasticity utility function $u(c) = c^{1-\gamma}$ with $\gamma = 1$ so that $u(c) = \log c$. We set $\underline{b} = 0$ so that there is no borrowing and we set the interest rate to $r = 0.0025$, or 1% per year. This partial equilibrium approach keeps our exercise especially clear because it allows us to move the interest rate independently of the discount factor and to highlight their respective importance in determining MPCs.[4] We model the process for log income as the sum of two orthogonal components, an AR(1) component and an IID component. We assume that shocks to both components arrive stochastically with a Poisson arrival rate of $1/4$, so that shocks are received on average once a year. We estimate the remaining parameters of the income process by matching moments of the household labor income distribution from the Panel Study of Income Dynamics. See Appendix A.1 for details.

Wealth Distribution. We choose the effective discount factor β so that mean wealth in the stationary distribution of the model is consistent with mean wealth in the United States. We express all values as multiples of mean annual household earnings, which we define as labor income plus social security income for retired households. Our income and wealth statistics come from the 2019 Survey of Consumer Finances, from which we exclude households in the the top 5% of the wealth distribution. See Appendix A.3 for details.[5] In the bottom 95% of the population, mean annual earnings is $67,000 and mean net worth is $275,000 or 4.1 times annual earnings.

Table 1 (Panel B) reports key wealth statistics from the data (Column 1) and in the baseline model (Column 2), with the discount factor chosen to hit this target. The implied annualized value for β is 0.98. Median wealth in the model (1.3) is quite close to its value in the data (1.5) despite not being explicitly targeted. The baseline model generates fewer low-wealth households than in the data. A common definition of hand-to-mouth households in the literature is households whose wealth is less than half their monthly income Kaplan, Violante, and Weidner (2014). According to this definition, only 2% of households in the model are hand-to-mouth compared

[4] In the equilibrium of a closed economy the two would be tightly connected, once a target for the wealth-income ratio is chosen. A number of recent papers in the literature (Auclert, Rognlie, and Straub, 2018; Kaplan, Moll, and Violante, 2018; Wolf, 2020) have demonstrated that one can usefully separate macro questions about the transmission of shocks and the effects of policies into (i) a partial-equilibrium response and (ii) general-equilibrium amplification. The analysis in this paper is purely about the size of the initial household response to an income shock, and about how different model assumptions matter, not about its general equilibrium implications.

[5] The top 5% holds 65% of the total net worth in the economy. We exclude this group from our calibration because it is well understood that the simple precautionary savings models we consider here are not well suited to explain the top tails of the wealth distribution (Benhabib and Bisin, 2018; De Nardi and Fella, 2017). Moreover, the top of the wealth distribution has a negligible impact on our definition of the MPC.

	(1) Data	(2) Baseline	(3) E[a]	(4) Median(a)	(5) E[a]	(6) Median(a)	(7) HtM
Quarterly MPC (%)		4.6	2.7	4.3	14.0	33.7	22.0
Annual MPC (%)		14.6	8.5	13.6	40.8	77.4	58.7
Quarterly MPC of the HtM (%)		28.7	26.3	28.4	33.2	42.1	36.6
Share of HtM (%)	14.2	2.5	1.7	2.3	7.5	37.4	14.2
Annualized discount factor		0.980	0.988	0.981	0.945	0.826	0.909
Panel A: Decomposition							
Gap with Baseline MPC			-1.9	-0.3	9.4	29.1	17.4
Effect of MPC Function			-0.8	-0.1	3.1	11.1	6.0
Effect of Distribution			-1.3	-0.2	4.6	12.5	7.7
Interaction			0.2	0.0	1.7	5.6	3.7
Panel B: Wealth Statistics							
Mean wealth	4.1	4.1	9.4	4.6	0.6	0.1	0.3
Median wealth	1.5	1.3	3.5	1.5	0.2	0.0	0.1
$a \leq$ \$1000	15.1	2.5	1.8	2.4	6.2	22.2	10.5
$a \leq$ \$5000	19.5	11.6	7.5	10.9	28.6	62.3	42.1
$a \leq$ \$10000	24.6	18.5	11.8	17.4	44.3	77.2	60.2
$a \leq$ \$50000	37.8	40.3	26.8	38.2	79.3	96.4	91.1
$a \leq$ \$100000	49.4	51.9	36.0	49.4	89.8	99.0	97.1
Wealth, top 10% share	49.9	46.6	44.6	46.1	52.0	56.5	51.6

Table 1: Baseline one-asset model and sensitivity analysis with respect to which moment of the wealth distribution is targeted to set the discount factor.Column (1) contains some key moments of the wealth distribution computed from the 2019 SCF excluding the top 5 percent. Column (2) is the baseline model where we set the discount factor to match mean net worth in our SCF sample. Columns (3) and (5) match mean net worth including the top 5% and mean liquid wealth, respectively. Columns (4) and (6) match median net worth and median liquid wealth. Column (7) matches the share of HtM

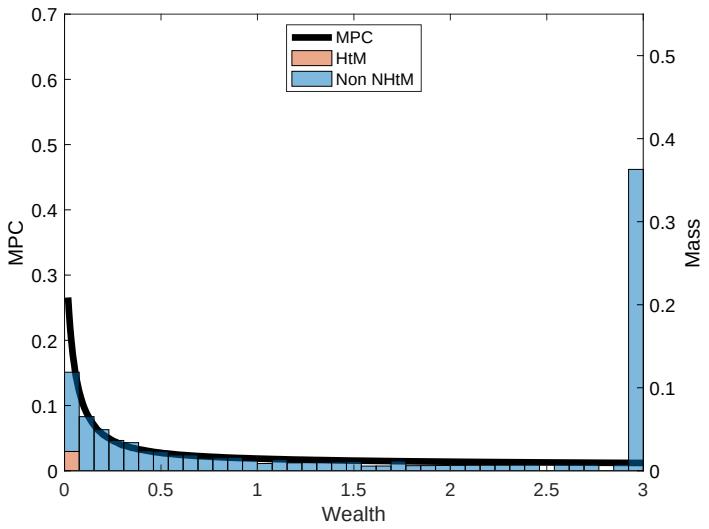


Figure 1: The MPC function (left scale) and the household wealth distribution (right scale) for the baseline one-asset model. Annual labor income normalized to 1.

with 14% in the data, because optimizing households seek to save away from hand-to-mouth regions of the asset space (see line labeled 'Share of HtM').[6]

Marginal Propensities to Consume in the Baseline Model Our main object of interest is the quarterly MPC out of a one-time unanticipated windfall of size x. For a household with state vector (b, y) at the time when the windfall is received, the impact MPC (or MPC at horizon 0) is:

$$\mathfrak{m}_0(x; b, y) = \frac{c(b + x, y) - c(b, y)}{x}. \tag{2}$$

Recall that, without income uncertainty or borrowing constraints, the consumption function is linear in wealth, with constant slope $\mathfrak{m}_0^*$, as given by the formula:[7]

$$c_0 = \mathfrak{m}_0^* \left[Rb_0 + \sum_{t=0}^{\infty} \left(\frac{1}{R} \right)^t y_t \right], \text{ with } \mathfrak{m}_0^* = 1 - R^{-1} [R\beta]^{\frac{1}{\gamma}}. \tag{3}$$

We refer to $\mathfrak{m}_0^*$ as the certainty MPC. It is decreasing in the discount factor and, provided that $\beta R < 1$, is decreasing in the IES $1/\gamma$. With log-utility $(\gamma = 1)$, the MPC is equal to the effective discount rate, $\mathfrak{m}^* = 1 - \beta$, which in our baseline calibration would give a quarterly MPC of 0.5%,

[6] This feature is robust to other definitions. Fewer than 1% have zero or negative wealth, compared with 11% in the data; 3% hold less than $1,000, compared with 15% in the data; and 12% hold less than $5,000, compared with 20% in the data.

[7] See appendix B.3 for the derivation of equation (3).

independently of the size of the windfall x.

When computing the MPC in our models with uninsurable income risk, we set $x = \$500$ in equation (2), which is the approximate size of common fiscal stimulus programs from which MPCs were measured in the literature.[8]

Average MPCs are reported in Table 1. The average quarterly MPC in the baseline precautionary saving model is 4.6%, nearly one order of magnitude larger than the certainty MPC. Figure 1 displays the MPC as a function of wealth for a household with mean income, $\mathfrak{m}_0(\$500; b, \overline{y})$, superimposed over the stationary wealth distribution. Note how the MPC quickly converges to the MPC under certainty, $1-\beta$, as wealth rises. The quarterly MPC of HtM households is very high, nearly 29%, but at a level of wealth of around 0.5 (about \$35,000) the effects of concavity of the consumption function on the MPC have already dissipated: it is only households with very low levels of wealth that contribute to generating an MPC that is substantially above $1-\beta$. This is a theme that will resurface as we explore richer versions of the model.

Decomposition Relative to Certainty Benchmark. To better understand why the precautionary savings model with uncertainty generates a larger average MPC (4.6% vs 0.5%), we propose a decomposition of the difference between the MPCs in the two models. Let $\mathfrak{m}_0^{BC}(b)$ be the MPC function in a model that is identical to the baseline model, except that household income is deterministic and all households receive the average level of income. Despite the absence of income risk, if $\beta R < 1$ the consumption function in this model is concave because of the presence of a borrowing constraint (Helpman, 1981). The average MPC in the baseline model $\overline{\mathfrak{m}}_0$ can then be written as:

$$\overline{\mathfrak{m}}_0 = \underbrace{\mathfrak{m}_0^*}_{\text{Certainty}} + \underbrace{\int_{\mathcal{B}} \left[\mathfrak{m}_0^{BC}(b) - \mathfrak{m}_0^*\right] d\mu(b)}_{\text{Borrowing Constraints}} + \underbrace{\int_{\mathcal{B}\times\mathcal{Y}} \left[\mathfrak{m}_0(b,y) - \mathfrak{m}_0^{BC}(b)\right] d\mu(b,y)}_{\text{Precautionary Savings + Income Risk}}.$$

The first term is the certainty MPC, the second term captures the role of borrowing constraints absent income risk paired with the desire to frontload consumption (i.e. $\beta R < 1$), and the final term captures the additional effect of uninsurable risk and precautionary savings.

This decomposition (see panel A in Table 1) reveals that the bulk of the gap with the certainty MPC is explained by the presence of borrowing constraints together with a declining optimal consumption profile (68%), even absent idiosyncratic uncertainty. The residual (38%) is accounted for by the precautionary saving motive induced by prudence. Thus, perhaps surprisingly, a strong force toward concavity the consumption function emerges even without income

[8] In Section 5 we explore MPCs out of other size windfalls and losses, over other durations.

uncertainty.[9]

2.2 Alternative Calibrations of the Baseline

We conduct an extensive sensitivity analysis of the baseline model. Unless otherwise specified, we always recalibrate the discount factor so that each version of the model has the same ratio of mean wealth to mean earnings of 4.1.

Deviations Between Heterogeneous-Agent Models. To understand why different versions of the baseline model yield different MPCs, we propose a decomposition of the difference in average MPCs across models. The average MPC can be larger in one model than in another either because the consumption function is steeper, or because the distribution of households is more concentrated in steeper regions of the state space. Let μ^* and $\mathfrak{m}_0^*$ be the distribution and average MPC in the baseline model from Section 2.1. We can then write the average MPC in an alternative model as:

$$\begin{aligned}\bar{\mathfrak{m}}_0 \;=\; & \underbrace{\mathfrak{m}_0^*}_{\text{Benchmark MPC}} + \underbrace{\int_{\mathcal{B}\times\mathcal{Y}} \left[\mathfrak{m}_0(b,y) - \mathfrak{m}_0^*\right] d\mu^*(b,y)}_{\text{Consumption Function}} + \underbrace{\int_{\mathcal{B}\times\mathcal{Y}} \mathfrak{m}_0^*(b,y)\left[d\mu(b,y) - d\mu^*(b,y)\right]}_{\text{Distribution}} \\ & + \underbrace{\int_{\mathcal{B}\times\mathcal{Y}} \left[\mathfrak{m}_0(b,y) - \mathfrak{m}_0^*\right]\left[d\mu(b,y) - d\mu^*(b,y)\right]}_{\text{Interaction}}.\end{aligned}$$

The component labeled 'Consumption Function' captures the difference in average MPCs that arises because the consumption functions are different in the two models. The component labeled 'Distribution' captures the difference in average MPCs that arises because the stationary distributions of the two models put different mass in different parts of the state space. The component labeled 'Interaction' arises because of these two effects may reinforce or offset each other.

Target for Wealth-Income Ratio. In Table 1 we report MPCs under alternative aggregate wealth targets. We start by considering a higher target for mean wealth based on the full population without dropping the wealthiest 5% of households, which is \$750,000, corresponding to a ratio of 9.4 to mean earnings (Column 3). The average MPC under this calibration drops to 2.7%, driven mostly by the smaller fraction of low wealth households. With this calibration, median

[9] Of course, in equilibrium it is income uncertainty that pushes the interest rate below the discount rate.

wealth is 3.5, over twice as large as in the data. Next, we target median wealth instead of mean wealth, which gives an average MPC that is almost unchanged from the baseline because median wealth in the baseline model is already close to the value in the data (Column 4).

In Columns 5 and 6 we report results from what is known in the literature as a *liquid wealth calibration*. Rather than measuring wealth in the SCF as total net worth, we only include liquid wealth, which we define as bank accounts and directly held stocks and bond net of credit card debt (see Appendix A.3 for details). The logic behind this approach is that liquid wealth is a better measure of those funds that households can readily access to smooth consumption against unexpected income fluctuations. The ratio of mean liquid wealth to mean income is only 0.56. Choosing the discount factor to match this alternative target for mean wealth, the quarterly MPC is 14% (Column 5), which is much higher than in the baseline and closer to empirical estimates. When we target median liquid wealth ($3,100 or 0.05 times as average earnings), the average quarterly MPC rises to 33% (Column 6). A related strategy is to choose the discount factor to match directly the fraction of hand-to-mouth households in the data. When we target 14% of hand-to-mouth households, the average quarterly MPC is 22%, in between the mean and median liquid wealth calibrations (Column 7).

In all these three calibrations, the discount factor is much smaller than in the baseline. The lower discount factor raises both the number of households with low wealth around the concave region of the consumption function as well as the certainty MPC, $(1-\beta)$, i.e. the slope of the consumption function for high wealth levels. The MPC decomposition in panel A confirms this intuition: the steeper consumption function and the shift of the distribution both play similar roles in accounting for the higher MPC.

Despite the apparent success of liquid wealth calibrations at generating high MPCs, it is important to note that these calibrations are difficult to integrate into modern dynamic macroeconomic models, because they necessitate abstracting from essentially the entire stock of aggregate assets owned by the household sector. This limits the usefulness of these calibrations in general equilibrium models with capital (either land, housing or productive capital). For example, the calibration that matches mean liquid wealth in the data (Column 5) effectively abstracts from 85% of the wealth in our baseline sample (which excludes the top 5%), or 98% of total wealth. The calibrations in Columns 6 and 7 miss an even larger share of total wealth in the economy.

Interest rate Table 2 shows that lowering the interest rate from 1% p.a. to 0% p.a has a negligible effect on the MPC. Raising the interest rate to 5% p.a. increases the average quarterly MPC by around half a percentage point. A higher interest rates leads to a calibration with a lower discount factor, which raises the certainty MPC and makes the consumption function steeper.

	(1) Baseline	(2) r = 0%	(3) r = 5%	(4) RRA=0.5	(5) RRA=6
Quarterly MPC (%)	4.6	4.4	5.2	5.3	3.0
Annual MPC (%)	14.6	13.9	16.8	16.4	10.6
Quarterly MPC of the HtM (%)	28.7	28.5	29.5	29.0	26.6
Share of HtM (%)	2.5	2.4	2.5	3.3	0.6
Annualized discount factor	0.980	0.990	0.944	0.986	0.840
Panel A: Decomposition					
Gap with Baseline MPC		-0.2	0.6	0.7	-1.6
Effect of MPC Function		-0.2	0.7	0.1	0.5
Effect of Distribution		-0.0	-0.1	0.6	-2.1
Interaction		-0.0	0.0	0.0	0.1
Panel B: Wealth Statistics					
Mean wealth	4.1	4.1	4.1	4.1	4.1
Median wealth	1.3	1.4	1.4	1.1	2.3
$a \leq$ \$1000	2.5	2.5	2.4	3.4	0.6
$a \leq$ \$5000	11.6	11.6	11.1	13.9	3.1
$a \leq$ \$10000	18.5	18.4	17.8	21.4	6.7
$a \leq$ \$50000	40.3	40.1	39.5	43.3	24.8
$a \leq$ \$100000	51.9	51.6	51.2	54.2	38.8
Wealth, top 10% share	46.6	46.2	45.9	48.6	36.3

Table 2: Baseline one-asset model and sensitivity analysis with respect to the interest rate r (annualized values) and the relative risk aversion (RRA) coefficient γ in the utility function. Column (1): baseline model with $r = 1\%$ and coefficient of relative risk aversion equal to 1. In columns (2) and (3) we solve the model with different values of the interest rate. In columns (4) and (5) we solve the model with different values of the coefficient of relative risk aversion.

This is confirmed by the decomposition which shows that the entirety of the difference in MPC relative to the baseline is due to the consumption function.

Curvature in Utility Table 2 also shows that changing the curvature parameter of the CRRA utility function γ away from $\gamma = 1$ has only a small effect on the average MPC. Higher risk aversion and lower intertemporal elasticity of substitution (higher γ) strengthens the precautionary savings motive and concavifies the consumption function. Because of this higher desire for saving, a lower discount factor is required to generate the same amount of aggregate wealth. These forces increase the MPC at all wealth levels. There is, however, an offsetting force: the stronger precautionary motive means that there are fewer households close to the borrowing constraint in the stationary distribution, for a given aggregate amount of wealth. Column 5 shows that for a coefficient of relative risk aversion of $\gamma = 6$, this latter effect dominates and the average quarterly MPC falls from 4.6% to 3.0%. The decomposition shows that the 1.6 percentage point lower MPC is comprised of a 0.5 percentage point increase from the lower discount factor, offset by a 2.1 percentage point fall from the smaller fraction of low wealth households. For example, relative to the baseline where $\gamma = 1$, median wealth is almost twice as large in the $\gamma = 6$ economy. With CRRA utility, this analysis does not separate the effects of changing risk aversion from those of changing the IES. In Section 3 we generalize preferences and perform this analysis.

Model Frequency. In Table E.1 in the Appendix we also report results for the baseline calibration with a model period of one year rather than one quarter, and for a continuous time version of the model. The average annual MPC in the annual model is 14.3%, essentially the same as the average annual MPC in the quarterly model, which is 14.6%. Note that this value is lower than what one would obtain by cumulating the quarterly MPC over four quarters, using the formula $\mathfrak{m}_4 = (1 + \mathfrak{m}_0)^4 - 1$, which is often erroneously used in the literature. Applying this formula would yield an annual MPC of 19.7%, or 35% higher than the correct annual MPC. Simply multiplying the quarterly MPC by 4 would give an annual MPC of 18.4%, or 26% higher than the correct MPC.

The quarterly MPC in the continuous time version of the model is 3.0%, slightly lower than our baseline discrete time model.[10]

[10]Discrete and continuous time models necessitate different specifications of the income process. We have adopted analogous processes calibrated to match the same targets, so as to enable comparisons between the two. Table A.2 contains parameter estimates for the continuous time counterpart of the income processes. Wealth and MPC statistics are very similar which ensures that, when we move to the continuous time two-asset model in Section 4, the differences are not being driven by the switch to continuous time.

Income Process. We also considered several alternative income processes to the one in our baseline model, including: (i) assuming that income shocks arrive on average once each quarter, rather than once each year; (ii) estimating the arrival rates of the shocks alongside the other parameters by targeting the kurtosis of income growth rates at different lags, in addition to the variance of income growth rates; (iii) estimating an annual income process and converting the parameter estimates into quarterly values using the approach in Krueger, Mitman, and Perri (2016); (iv) eliminating transitory shocks. We also considered alternative processes in the annual version of the model including: (i) a version with higher persistence in the AR(1) component; and (ii) a version with individual-specific fixed effects. Details for all these statistical representations of income dynamics can be found in Table A.2 and a summary of the results can be found in Table E.2, in the Appendix.

For most of these alternative specifications of the income process, the MPCs are extremely close to the baseline model. The only versions that differ significantly from the baseline are when we assume no transitory shocks and when we assume a high persistence in the AR1 component. Without transitory shocks, households are willing to hold less precautionary wealth because they face less easily insurable income risk. As a result, this calibration has a share of hand-to-mouth consumer of 18% compared to 2% in the baseline and a larger average MPC. In the second specification, the AR1 component is very close to a random walk. As a result, the estimation loads all the short-term volatility onto a higher transitory variance. Both forces contribute toward a lower MPC: because shocks are near-permanent a higher discount factor is needed to match aggregate wealth, and because of the sizable transitory variance the distribution of wealth has less mass near the credit limit.

Survival and Bequests. We also examined sensitivity to different assumptions about the survival rate, how the assets of the deceased are distributed and how the assets of new-born households are determined. As long as the discount factor is always recalibrated to match the same amount of aggregate wealth, none of these assumptions matters. See Table E.3 in the Appendix.

3 Extensions of the One-Asset Model

3.1 Ex-Ante Heterogeneity

In this section, we extend the one-asset model to allow for various forms of ex-ante heterogeneity.

	(1) Baseline	(2) Het β	(3) Het β	(4) Het β	(5) Het β	(6) Het r	(7) Het r	(8) S-S
Set of β		0.93, . . . , 1.01	0.85, . . . , 1.01	0.85, . . . , 1.01	0.85, . . . , 1.01			
Set of effective β		0.91, . . . , 0.99	0.83, . . . , 0.99	0.83, . . . , 0.99	0.83, . . . , 0.99			
Switch probability β		0.00	0.00	0.02	0.10			
Set of r						-1, 1, 3	-3,1,5	
Quarterly MPC (%)	4.6	11.7	18.6	14.3	6.7	7.5	11.6	16.8
Annual MPC (%)	14.6	33.4	47.8	38.1	19.6	22.8	33.1	26.1
Quarterly MPC of the HtM (%)	28.7	33.8	38.5	37.0	33.0	31.0	33.8	85.4
Share of HtM (%)	2.5	6.7	14.2	9.9	4.1	4.0	6.8	14.8
Annualized discount factor	0.980	0.953	0.915	0.933	0.963	0.970	0.954	0.392
Panel A: Decomposition								
Gap with Baseline MPC		7.1	14.0	9.7	2.1	2.9	7.0	12.2
Effect of MPC Function		2.4	4.3	3.4	1.2	1.2	2.7	1.9
Effect of Distribution		3.2	6.0	4.1	0.8	1.4	3.1	3.3
Interaction		1.4	3.7	2.1	0.1	0.3	1.2	7.0
Panel B: Wealth Statistics								
Mean wealth	4.1	4.1	4.1	4.1	4.1	4.1	4.1	4.1
Median wealth	1.3	0.3	0.2	0.3	1.1	0.6	0.3	1.0
$a \leq$ \$1000	2.5	5.5	10.0	7.6	3.8	3.7	5.5	14.6
$a \leq$ \$5000	11.6	24.1	35.9	28.2	14.9	16.8	23.7	22.8
$a \leq$ \$10000	18.5	36.5	49.7	39.9	22.2	26.6	35.8	28.7
$a \leq$ \$50000	40.3	64.4	74.5	64.1	43.5	52.6	62.5	46.6
$a \leq$ \$100000	51.9	73.8	81.0	72.2	54.6	63.6	71.5	56.2
Wealth, top 10% share	46.6	79.2	87.9	75.7	51.4	64.7	74.0	50.9

Table 3: One-asset model with heterogeneity in the effective discount factor β (annualized values) and in the the rate of return r (annualized values). Column (1): baseline with common discount factor and interest rates. Columns (2) and (3): models with fixed ex-ante heterogeneity in discount factors across households. Columns (4) and (5) solve models where the individual discount factor switches randomly. Columns (6) and (7): models with fixed ex-ante heterogeneity in interest rates across households. Column (8): spender-saver model.

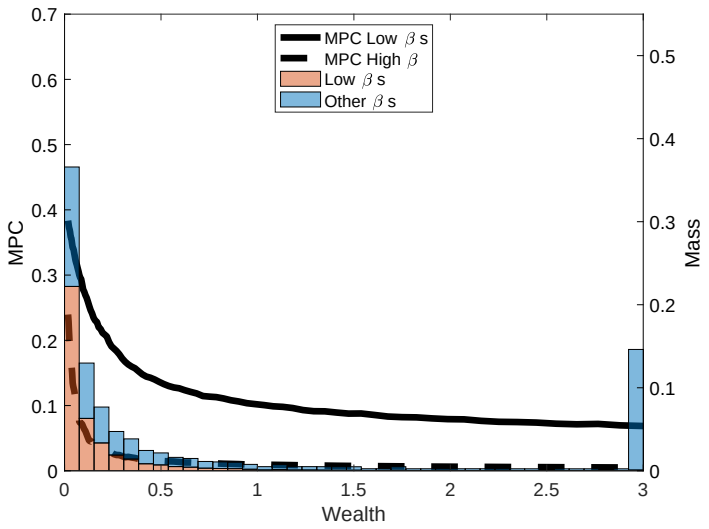


Figure 2: The MPC function (left scale) and the household wealth distribution (right scale) for the model with large heterogeneity in discount factors across households. Annual labor income normalized to 1.

3.1.1 Heterogeneity in Discount Factors

We start by allowing for heterogeneity in households' discount factors β. We consider a a discretized uniform distribution for β (annualized) with 5 equally spaced grid points between $[\beta - 2\Delta, \beta + 2\Delta]$. We choose the the mid-point β to match average wealth of 4.1 as in our previous calibrations. Table 3 reports results for versions with a moderate amount of heterogeneity ($\Delta = 0.01$, Column 2) and with a large amount of heterogeneity ($\Delta = 0.02$, Column 3). For both calibrations, the median effective discount factor is substantially lower than in the baseline economy. The reason is that in these economies there is a subset of very patient households who have a strong intertemporal savings motive, and it is these households who hold the bulk of aggregate wealth. This can be seen from the much higher share of wealth held by the top 10% households (88% and 79%), compared to the model without discount factor heterogeneity (47%). The model can therefore match the target for mean wealth while still allowing for a large fraction of households to be impatient. Consequently, the models with discount factor heterogeneity match well the very bottom of the wealth distribution. For example in the calibration with $\Delta = 0.02$, 14% of households are hand-to-mouth, as in the data.

The average quarterly MPC in the models with discount factor heterogeneity are much higher than in the baseline model, and with enough heterogeneity, these models can approach the target empirical values. With $\Delta = 0.02$, the average MPC is nearly 19%, four times as large as the baseline model. While both the shape of the MPC function and the stationary wealth distribution

contribute to the higher MPC, the majority of the effect comes from the larger fraction of low wealth households (see the decomposition in panel A). Figure 2 displays the MPC functions for high and low β households overlaid on the stationary wealth distribution for the two groups of households. The figure illustrates how the impatient households are not only amassed near to the borrowing constraint, but have higher MPCs at all levels of wealth.

We also examined a version of the model in which households switch randomly between different discount factors. We assumed that a household draws a new value of β with probability p, independent of its current value.[11] The models with stochastic discount factors generate a lower average MPC compared with a model with the same stationary distribution of β, but with fixed heterogeneity. The reason is that in the model with stochastic discount factors there is a weaker correlation between wealth and discount factors. In the stochastic β model, some low-wealth households who were previously impatient then become patient and quickly accumulate wealth. The wealth distribution is therefore less concentrated at the bottom, which is evident in the lower share of HtM households and higher value of median wealth.

Although the models with discount factor heterogeneity feature a large average MPC while generating the key features of the two tails of the wealth distribution, these models typically fail in reproducing the wealth distribution everywhere in between. For example in the model with $\Delta = 0.02$, despite matching the fraction of wealth-poor households (14% hand-to-mouth, 10% with wealth less than \$1,000) the model's wealth distribution is excessively squeezed toward the bottom and has far too many households with wealth just above this threshold: median wealth is ten times smaller than in the data and 75% of households have wealth below \$50,000, compared with only 38% in the data. This problem, which we label the 'missing middle' is a recurring issue that arises in many of the versions of the model with heterogeneity that we examine below. Comparing the wealth distributions in Figures 1 and 2 clearly illustrates this shortcoming.

3.1.2 Heterogeneity in Rates of Return

In columns (6) and (7) of Table 3, we report results for an economy with fixed heterogeneity in r, uniformly distributed over $\{-1\%, 1\%, 3\%\}$ p.a. (Column 6), and over $\{-3\%, 1\%, 5\%\}$ p.a. (Column 7). Heterogeneity in rates of return generates similar results to heterogeneity in discount factors. For example in the calibration in Column 7, 7% of households are hand-to-mouth, the top 10% share is 74% and the aggregate quarterly MPC reaches nearly 12%. These economies, however, also feature the missing middle problem, with median wealth levels much lower than

[11] We set $\Delta = 0.02$ and consider annual switching rates of 0.02, so that the expected duration of a discounting regime has the same expected duration of a lifetime (Column 4); and of 0.1., so that the expected duration is equal to a decade (Column 5).

in the data.

Heterogeneity in Risk Aversion and Elasticity of Intertemporal Substitution We now allow for heterogeneity in the curvature parameter γ in the CRRA utility function. The distribution of γ is a discretized uniform with 5 geometrically spaced grid points in the interval $[\frac{1}{\bar{\gamma}}, \bar{\gamma}]$. In Table E.4 in the Appendix, we report results for a moderate amount of heterogeneity ($\bar{\gamma} = e^2 = 7.4$, Column 3) and a large amount of heterogeneity ($\bar{\gamma} = e^3 = 20.1$, Column 4). In both cases we set the mid-point to $\gamma = e^0 = 1$ as in the baseline model. With enough heterogeneity in γ, these calibrations yield very large average MPCs, because of the high shares of households in the tails of the wealth distribution. For example the model in Column 3, which has values of γ ranging from 0.14 to 7.4, gives an average quarterly MPC of 17%. Most of the wealth is held by the high risk aversion / low IES households, and the top 10% wealth share is around 67%. The large MPC is driven by the low risk aversion / high IES households, and 20% of households are hand-to-mouth. The decomposition of panel A reveals that it is indeed this different distribution of households relative to the baseline model, rather than the shape of the MPC function, that is the most important factor in generating a larger MPC. But like the previous calibrations with heterogeneity, this economy suffers from the same missing middle problem: a disproportionate share of households is wealth-poor and median wealth is much lower than in the data.

To distinguish the roles of risk aversion and intertemporal substitution in generating the high degree of wealth concentration and large MPC in this last calibration, we generalize our preferences to those of Epstein and Zin (1991):

$$U_t = \left\{ (1-\beta)\, c_t^{1-\theta} + \beta \left(\mathbb{E}_t \left[U_{t+1}^{\frac{1-\theta}{1-\gamma}} \right] \right)^{1-\gamma} \right\}^{\frac{1}{1-\theta}} \tag{4}$$

where γ is the coefficient of relative risk aversion and $1/\theta$ the IES. In Table E.5 in the Appendix we report results for various values of γ and θ. Without preference heterogeneity, varying either of these two parameters has only a very small impact on the average MPC once the discount factor is recalibrated to match the same average wealth target.

Allowing for heterogeneity in risk aversion, holding the IES constant, has barely any effect on the average MPC. However, with heterogeneity in the IES, the model is able to generate large average MPCs. These results are reported in Table 4 and imply that the findings with CRRA preferences described above are being driven by the heterogeneity in the IES rather than in risk aversion. For example, with a geometric uniform distribution of θ ranging from $e^{-3} = 0.05$ to $e^3 = 20$, the average quarterly MPC is above 20%. Households with a high IES are willing to

	(1) Baseline	(2) Het IES	(3) Het IES	(4) Het IES	(5) Temptation	(6) Temptation	(7) Het Temptation
Set of IES		$e^{-1},\ldots,e^{1}$	$e^{-2},\ldots,e^{2}$	$e^{-3},\ldots,e^{3}$			
Set of Temptation					0.01	0.05	0, 0.05, 0.1
Quarterly MPC (%)	4.6	6.6	10.9	20.8	6.4	19.1	21.4
Annual MPC (%)	14.6	20.1	29.9	39.2	19.0	46.2	47.7
Quarterly MPC of the HtM (%)	28.7	30.5	34.4	66.2	31.4	39.3	45.1
Share of HtM (%)	2.5	4.0	11.1	20.9	4.0	20.4	29.3
Annualized discount factor	0.980	0.977	0.975	0.974	0.982	0.990	0.988
Panel A: Decomposition							
Gap with Baseline MPC		2.0	6.3	16.2	1.8	14.5	16.8
Effect of MPC Function		0.9	2.1	3.1	0.4	2.7	3.9
Effect of Distribution		1.0	3.3	6.0	1.2	7.4	8.3
Interaction		0.2	0.9	7.1	0.2	4.4	4.5
Panel B: Wealth Statistics							
Mean wealth	4.1	4.1	4.1	4.1	4.1	4.1	4.1
Median wealth	1.3	0.9	0.5	0.3	1.0	0.1	0.1
$a \leq$ \$1000	2.5	3.7	8.0	19.0	3.9	14.9	22.2
$a \leq$ \$5000	11.6	15.4	24.3	33.3	16.7	42.5	42.8
$a \leq$ \$10000	18.5	23.8	33.7	41.7	25.0	52.6	50.8
$a \leq$ \$50000	40.3	47.3	55.9	60.9	46.2	69.4	66.3
$a \leq$ \$100000	51.9	58.1	64.7	68.3	55.9	73.6	71.5
Wealth, top 10% share	46.6	55.7	62.8	67.9	49.8	70.6	70.1

Table 4: One-asset model with heterogeneous Epstein-Zin preferences and one-asset model with Gul-Pesendorfer preferences. Column (1): baseline model with common IES across households. Columns (2)-(4): model with ex-ante heterogeneity in IES. Columns (5)–(6): model with Gul-Pesendorfer preferences, and common temptation parameter. Column (7): model with Gul-Pesendorfer preferences, and heterogeneous temptation parameter across households.

absorb consumption fluctuations, and so hold only small buffer stocks of wealth. However, this calibration suffers from a missing middle problem similar to other versions of the model with ex-ante heterogeneity.

3.2 Behavioral Preferences

In this section we explore two alternative behavioral models for preferences that have been proposed in the literature as ways to generate higher MPCs: temptation and present bias. Our main message is that both models are effectively models of endogenous discount rates, with the discount rate depending on the level of assets. As a result, these models have they same strengths and weaknesses in terms of average MPCs and the wealth distribution as the model with ex-ante discount factor heterogeneity in Section 3.1.

3.2.1 Temptation and Self-control

Gul and Pesendorfer (2001) proposed a model of temptation and self-control in which consumers are tempted to consume according to a preference specification that overweights current consumption, but can exert some degree of self-control. These preferences have recently gained popularity in the quantitative macro literature (Krusell, Kuruşçu, and Smith Jr, 2010, 2002; Attanasio, Kovacs, and Moran, 2020; Nakajima, 2017; Pavoni and Yazici, 2017). We consider the limiting formulation in which consumers are tempted to consume according to preferences that place no weight on future consumption. The degree of temptation is governed by an additional parameter $\varphi \in [0, \infty]$. It is straightforward to show that for $\varphi > 0$, the temptation and self-control model generates a modified Euler equation in which the effective discount factor is decreasing in the average propensity to consume, which has the effect of making poor households act as if they are more myopic than wealthier households. For example, with log utility, the effective discount factor becomes

$$\beta \mathbb{E}\left[1 - \frac{\varphi}{1+\varphi}\left(\frac{c'}{x'}\right)\right]$$

where x is cash on hand. We describe the setup formally in Appendix C.1.

The last three columns of Table 4 report results for values of the temptation parameter $\varphi \in \{0.01, 0.05\}$. The calibrations with a high degree of temptation can generate a large MPC without any ex-ante heterogeneity, while still matching the target for mean wealth. For example, with $\varphi = 0.05$ the average quarterly MPC is 19%. The decomposition reveals that the shape of the consumption function plays only a minor role and that the large MPC is due to a different

shape of the wealth distribution compared to the baseline model: wealth-poor households have a lower effective discount rate than high wealth households, which leads to a large mass near the borrowing constraint. In this calibration, median wealth is around \$10,000, i.e. less than one-tenth than in the data. Allowing for heterogeneity in φ, centered around a median value of 0.05 has very little effect on either the average MPC or the wealth distribution. The one-asset model with temptation thus also suffers from the missing middle problem.

3.2.2 Present Bias

Starting with Laibson (1997), the other commonly adopted departure from the consumption-saving problem with standard preferences is to assume some form of present bias, such as quasi-hyperbolic discounting. We adopt the continuous time formulation as in Laibson, Maxted, and Moll (2021), known as *instantaneous gratification*. Unlike the model of temptation and self-control, these preferences are not time consistent and so one needs to make an assumption about whether consumers are naive, meaning that they are unaware that their future selves will have different preferences, or are sophisticated, meaning that they are aware of the time inconsistency and play a game against their future selves. Following Laibson, Maxted, and Moll (2021), and to give the model its best shot at generating a large average MPC, we assume that households are naive. We describe the setup formally in Appendix C.3.

The instantaneous gratification model features one additional parameter $\zeta \leq 1$, which measures the extent of present bias. In Table E.6 in the Appendix, we report results for values of $\zeta \in \{0.9, 0.8, 0.7\}$, alongside the baseline continuous time one-asset model (corresponding to the special case $\zeta = 1$). The results suggest that once the instantaneous gratification model is recalibrated to match the same level of aggregate wealth, present bias has a negligible effect on the average MPC. In fact with $\zeta = 0.7$, the average MPC is *lower* than in the model with exponential discounting. The MPC decomposition reveals that there are two offsetting effects. On the one hand, the instantaneous gratification models has a stationary distribution with more very low wealth households (e.g., 19% of households with less than \$1,000, compared with 2.5% in the baseline model). On the other hand, in order to match the same level of aggregate wealth, the calibrated value of the effective discount factor is much larger in the model with present bias (0.997 p.a. vs 0.985 p.a.). This means that everywhere above the borrowing constraint the MPC is lower in the model with present bias. Figure 3 illustrates these differences. The decomposition shows that these two effects roughly cancel out.

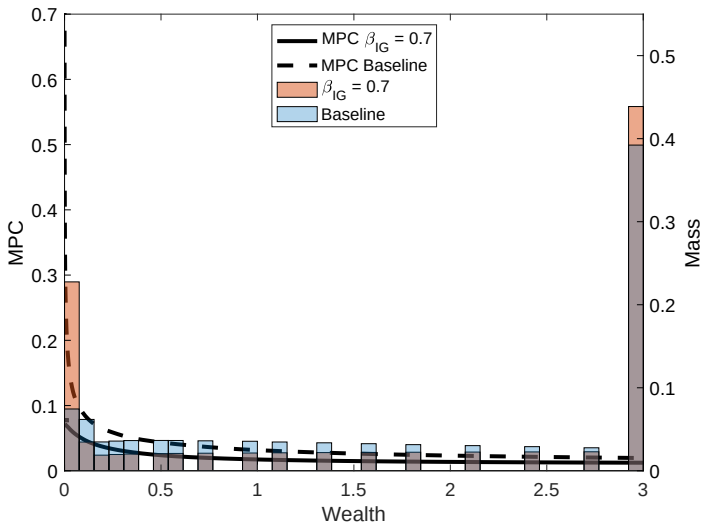


Figure 3: The MPC function (left scale) and the household wealth distribution (right scale) for the model with preferences featuring present bias. Annual labor income normalized to 1.

3.3 Taking Stock

When the one-asset model is calibrated to match the amount of wealth held by the poorest 95% of US households, it generates an average quarterly MPC of around 4%. This value is an order of magnitude larger than the MPC of the representative agent model, but still much smaller than empirical estimates.

Without preference heterogeneity or behavioral preferences, the baseline model can only generate an average quarterly MPC of around 20% when it is parameterized to match an amount of aggregate wealth that is 10 to 20 times smaller than its empirical counterpart for the US, for example in calibrations based on liquid wealth, rather than total wealth. However, these calibrations necessarily ignore almost the entire aggregate capital stock, and so are of limited use for general equilibrium applications.

When the one-asset model is extended to allow for ex-ante heterogeneity in discount factors, rates of return, or elasticity of intertemporal substitution across households, it can generate quarterly MPCs of 20% or higher in some calibrations. In all these versions of the model, the reason for the larger MPC is that the presence of different household types makes it possible for the stationary distributions to contain a substantial number of hand-to-mouth households, while still being consistent with an average level of wealth as high as in the data. Yet, all these models have far too many households who are not quite poor enough to be hand-to-mouth, but still have very little wealth – the 'missing middle'. For example, the fraction of households with less than $50,000 is typically much higher than in the data, and the wealth of the median household is 5

to 10 times smaller than in the data.

When preferences are generalized to allow for temptation, the model behaves like one where discount factors are decreasing in wealth. This feature allows it to achieve generate a high average MPCs, but suffers from the same missing middle problem that plagues versions with ex-ante heterogeneity. Versions of the model with present bias introduces offsetting forces when calibrated to match aggregate net worth and the average MPC remains small.

3.4 Spender-Saver Models

The challenge for one-asset models is therefore to simultaneously generate a sufficiently large amount of wealth in the aggregate as well as a large enough fraction of hand-to-mouth households, while still generating enough households in the middle of the distribution. There is one form of ex-ante heterogeneity in the one-asset model that succeeds in this respect. Inspired by the spender-saver model of Campbell and Mankiw (1989), these are versions of the model with discount factor heterogeneity in which there are two groups of households: one with a very low discount factor (the spenders) and one with a high discount factor (the savers). Judiciously chosen calibrations of these models can match all of these feature of the data.

To illustrate this, the last column in Table 3 reports results from a calibration in which 15% of households have a discount factor of 0.4, and the remaining 85% of households have a discount factor that is chosen so that the ratio of average wealth to average income is 4.1 as in the data. In this spender-saver model, 14.8% of households are hand-to-mouth and the average quarterly MPC is 17%. Moreover, median wealth is 0.96, which is not too far from its empirical counterpart of 1.54, and around 47% of households have less than $50,000 of wealth, compared with 39% in the data. The missing middle problem is therefore nowhere near as extreme as in other versions of the one-asset model with ex-ante heterogeneity.

Nonetheless, despite its success at generating a large average MPC and simultaneously matching the median and mean of the wealth distribution, this extreme form of heterogeneity has a major drawback that limits its usefulness for counterfactual policy analysis. For example, fiscal transfers targeted towards the impatient group are extremely powerful at triggering and immediate response of aggregate spending. Because of this feature, the spender-saver model displays an intertemporal MPC function with a steep peak coinciding with the time of the disbursement and little subsequent propagation. We return to these comparisons in Section 5.

4 Two-Asset Models

In this section, we extend the precautionary savings model to include two assets: a liquid asset with low return and an illiquid asset with higher return but subject to adjustment frictions. We demonstrate that the two-asset model can go a long way toward resolving the intrinsic tension in the one-asset model even in the absence of ex-ante heterogeneity.

In two-asset heterogeneous agent models, households can separate their different savings motives into distinct assets. Precautionary and smoothing motives against small, regular income fluctuations induce households to accumulate a buffer of liquid assets, but since the return on this asset is low, a strong negative intertemporal motive ($\beta R << 1$) prevents households from accumulating large amounts of liquid wealth. The bulk of household savings is done in the high-return illiquid asset, motivated both by a positive intertemporal motive and by precautionary and smoothing motives against large, infrequent income fluctuations. At endogenous intervals, households move funds between their liquid and illiquid accounts as desired. For instance, in response to a large negative income shock, a household who has exhausted their buffer of liquid wealth can pay a fee or exert effort to withdraw funds from their illiquid account. Similarly, a household who has experienced a long stream of positive income growth and has accumulated excess liquid wealth can transfer funds to their illiquid account, which pays a higher return and is therefore a better vehicle for long-run saving.[12]

Since most households end up holding the majority of their wealth in illiquid assets, which are of limited use for short-term consumption smoothing, they expose themselves to potentially larger consumption fluctuations than if they saved only in liquid assets. However, the welfare cost of these fluctuations is second-order, relative to the first-order gain from earning a higher return on illiquid savings. Because of this trade-off, the model generates *wealthy* hand-to-mouth households who have positive, and sometimes substantial, holdings of illiquid wealth, but only small amounts of liquid wealth. Wealthy hand-to-mouth households co-exist alongside the traditional *poor* hand-to-mouth households, who hold very little wealth, either liquid or illiquid. Adopting the same definition of a hand-to-mouth household as in previous sections (less than half of monthly income), 27% of households were wealthy hand-to-mouth in the 2019 SCF, in addition to the 14% of poor hand-to-mouth households. Therefore, in total 41% of US households are hand-to-mouth. Since wealthy hand-to-mouth households also have a high MPC out of small one-time windfalls, it is the presence of this additional group of hand-to-mouth households that enables the two-asset model to generate a substantial average MPC while remaining

[12]The first scenario corresponds, for example, to a household who suffers a job loss and extracts some equity from its house or withdraws from its retirement account to smooth expenditures. The second scenario corresponds to a household that saves enough to make a down payment on a house.

consistent with the distributions of liquid and illiquid wealth in the data.

4.1 Baseline Two Asset Model

Environment We write the two-asset model in continuous time. The economy is populated by a continuum of households who discount the future at the effective rate $\rho = \tilde{\rho} + \delta < 1$ where $\tilde{\rho}$ is the rate of time preference and δ is the death rate. Flow utility is given by $u(c_t)$, where u is strictly increasing and concave, and c_t denotes consumption expenditures. At each date t, households are endowed with log labor income $y_t \in Y$, which follows an exogenous stochastic process described in Appendix B.3. Households can save, but not borrow, in two assets: (i) a liquid asset b with return r^b; and (ii) an illiquid asset a with return $r^a > r^b$. At rate χ, households receive an opportunity to rebalance their financial portfolio which they can take by paying a fixed transaction cost κ. Between rebalance dates, the illiquid asset accumulates in the background and the household solves a standard consumption-savings problem out of their liquid assets. The HJB equation for the household problem is:

$$\begin{aligned}
\rho v(a,b,y) &= \max_c u(c) + v_b(a,b,y)\dot{b} + v_a(a,b,y)\dot{a} + \mathcal{A}v(a,b,y) \\
&\qquad +\chi\left[v^*(a,b,y) - v(a,b,y)\right] \\
&\quad \text{subject to} \\
\dot{b} &= r^b b + y - c, \quad b \geq 0 \\
\dot{a} &= r^a a, \quad a \geq 0
\end{aligned}$$

where $\mathcal{A}$ is the infinitesimal generator of the stochastic process for income. The last term in the HJB equation describes the gain from rebalancing the asset portfolio. The value function after the rebalancing opportunity arises, $v^*(a,b,y)$, is defined by

$$v^*(a,b,y) = \max\left\{\omega(a+b,y), v(a,b,y)\right\}$$

Upon receipt of an adjustment opportunity, a household will choose to rebalance their portfolio if the gains from rebalancing exceed the transaction cost κ. In this case, the household can choose any feasible combination of liquid and illiquid assets (b', a') and its value is given by

$$\begin{aligned}
\omega(a+b,y) &= \max_{a',b'} v(a',b',y) \\
&\quad \text{subject to} \\
a'+b' &\leq a+b-\kappa, \quad a', b' \geq 0
\end{aligned}$$

The solution to this problem yields decision rules for consumption $c(a,b,y)$, portfolio rebalancing $p(a,b,y) \in \{0,1\}$, and for the optimal rebalanced portfolio $a'(a,b,y)$ and $b'(a,b,y)$, as well as the stationary distribution of households $\mu(a,b,y)$.

Parameterization Relative to the continuous time one-asset model, there are three additional parameters: (i) the arrival rate for rebalancing opportunities χ; (ii) the illiquid asset return r^a; and (iii) the transaction cost κ. In our baseline model we set $\chi = 3$ so that households get an opportunity to rebalance on average once per month (the model period is one quarter). We also set the (annualized) liquid return $r^b = -2\%$ (roughly corresponding to a zero nominal return). We then choose the the discount rate ρ, the illiquid rate r^a and the transaction cost κ to match three targets: (i) a ratio of mean total net worth to mean earnings of 4.1, as in the one-asset model; (ii) a total share hand-to-mouth households (both wealth and poor) of 41%; and (iii) a share of poor hand-to-mouth households of 14%.

Table 5 reports the calibrated parameters and wealth statistics in our baseline two asset model. The model matches well all three targets. To achieve this, the calibration requires a sizable gap between the liquid and illiquid returns. We return to this feature of the calibration in the next section. The calibrated transaction cost is around \$1,500. On average, 9% of households rebalance their portfolios each quarter. Although the model generates less aggregate liquid wealth than in the data (0.23 vs 0.56), median liquid wealth matches almost exactly its empirical value, around 5% of average annual labor income. The empirical liquid wealth distribution is extremely right-skewed and there is no force in the model that can deliver this feature of the data.[13]

Marginal Propensities to Consume in the Two Asset Model The average quarterly MPC in the baseline two-asset model model is 16.1% (Table 5, Column 2), around 5 times larger than the corresponding average MPC in the continuous time one-asset model (Column 1). A simple back-of-the-envelope calculation helps in rationalizing this magnitude. Multiplying the average MPC for poor hand-to-mouth households (24%) by their share (13%), the average MPC for wealthy hand-to-mouth households (30%) by their share (27%) and the average MPC for non hand-to-mouth households (7%) by their share (60%), and adding up these three components, we obtain an average MPC close to 16%. Figure 4 (left-panel) illustrates the average MPCs as a function of liquid wealth holdings.

To shed some light on the sources of the higher average MPC in the two-asset model com-

[13]This feature, however, is not crucial in determining the aggregate MPC because once households have enough liquid assets, their MPC converges to the certainty level (approximately equal to the discount rate).

	(1) Baseline 1-asset	(2) Baseline 2-asset	(3) $r_b = -3\%$	(4) $r_b = 0\%$	(5) $r_a = 5\%$	(6) $r_a = 7.5\%$
Rebalance arrival rate		3.00	3.00	3.00	3.00	3.00
Quarterly MPC (%)	3.0	16.1	17.7	12.3	14.0	16.5
Annual MPC (%)	11.5	41.2	43.7	33.9	37.3	40.9
Quarterly PHtM MPC(%)	33.7	24.3	25.4	27.8	25.8	25.3
Quarterly WHtM MPC (%)		29.8	31.4	22.2	25.6	27.5
Mean MPC at Mean Wealth (%)	1.7	7.0	7.3	5.9	5.8	7.1
Prob. HtM status at year t and year t+1	0.5	0.9	0.9	0.8	0.8	0.9
Panel A: Calibrated Variables						
Annualized discount factor	0.985	0.937	0.937	0.937	0.948	0.926
Annualized liquid return	1.0	-2.0	-3.0	0.0	-2.0	-2.0
Annualized illiquid return		6.2	6.2	6.2	5.0	7.5
Rebalance cost ($)		1456.8	1456.8	1456.8	1456.8	1456.8
Panel B: Targeted Statistics						
Mean total wealth	4.1	4.1	4.1	4.1	4.1	4.1
Share hand-to-mouth		39.5	43.1	28.7	34.6	41.6
Share poor hand-to-mouth	2.0	13.2	15.2	8.5	10.2	15.7
Panel C: Decomposition						
Gap with Baseline MPC			1.6	-3.9	-2.1	0.4
Effect of MPC function			0.5	-3.0	-1.3	-0.8
Distributional Effect			1.0	-2.7	-1.3	0.2
Interaction			0.1	1.8	0.5	0.9
Panel D: Wealth Statistics						
Mean liquid wealth		0.2	0.2	0.3	0.2	0.2
Median total wealth	1.6	1.1	1.1	1.1	1.1	0.9
Median liquid wealth		0.1	0.0	0.1	0.1	0.1
$b \leq$ \$1000		33.0	36.6	23.7	28.7	35.2
$b \leq$ \$5000		54.7	57.0	47.7	52.5	55.3
$b \leq$ \$10000		66.1	67.8	60.7	65.0	66.4
$w \leq$ \$1000	2.5	10.0	11.7	7.1	8.3	12.0
$w \leq$ \$5000	8.0	20.8	21.5	19.3	19.7	21.8
$w \leq$ \$10000	13.4	23.5	23.7	23.3	23.2	24.1
$w \leq$ \$50000	34.9	45.9	46.1	45.2	45.3	47.5
$w \leq$ \$100000	48.0	57.7	57.9	57.6	57.5	59.7
Wealth, Top 10% share	40.3	62.0	62.4	60.9	60.9	64.0
Wealth, Top 1% share	5.6	19.2	19.6	18.0	17.7	21.0
Gini coefficient, total wealth	0.6	0.8	0.8	0.8	0.8	0.8

Table 5: Baseline two-asset model and sensitivity with respect to the liquid and illiquid rates of return. b denotes liquid wealth and $w = a + b$ total wealth (or net worth). Column (1) baseline one-asset model. Column (2): baseline two-asset model. Columns (3)-(4): two-asset model solved with a different rate of return on the liquid asset. Columns (3)-(4): two-asset model solved with a different rate of return on the illiquid asset.

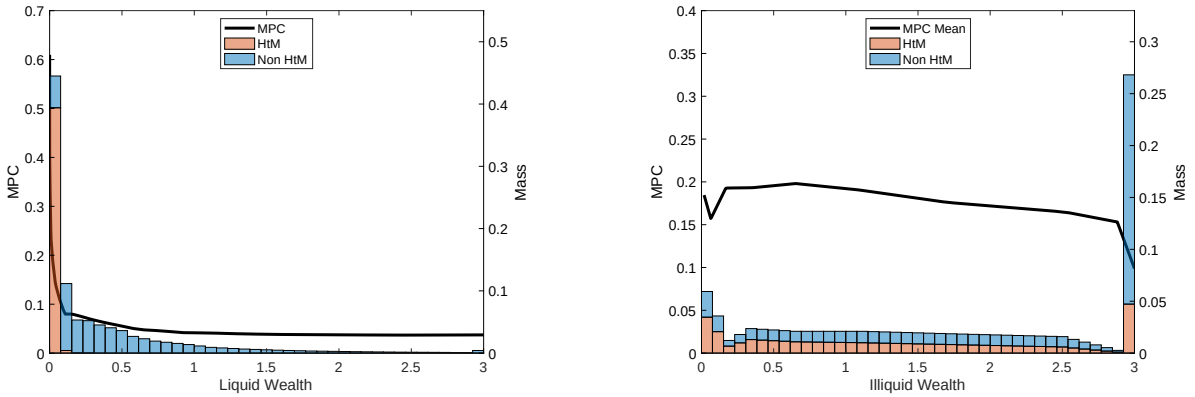


Figure 4: The MPC (left scale) and the household wealth distribution (right scale) for the baseline two asset model as a function of liquid (left panel) and illiquid (right panel) wealth. Annual labor income normalized to 1.

pared to the one-asset model, we report the results of the decomposition in Equation (4) in Panel C. We implement the decomposition by computing the average MPC in the two asset model as a function of net worth w, i.e. $\mathfrak{m}_0(w,y) = \int_{a+b=w} \mathfrak{m}_0(a,b,y)d\mu(a,b,y)$. The shapes of the consumption function and the wealth distribution play similar roles. Because of the presence of a high-return asset, the intertemporal saving motive is strong and the model asks for a lower discount factor to match total net worth, a force that increases the MPC at all wealth levels. In addition, the larger share of hand-to-mouth consumers with low liquid wealth shifts the distribution toward the steep region of the MPC function.

Figure 4 confirms that what accounts for the bulk of cross-sectional variation in MPCs is liquid wealth (left panel) in line with much of the recent empirical evidence discussed in the Introduction. The MPC function is instead quite flat as a function of illiquid wealth (right panel) and net worth, because there is a sizable share of wealthy HtM types even among households with high levels of illiquid wealth (up to around three times annual income).

The statistics on the wealth distribution in panel D reveal that the two-asset model does not suffer from the missing middle problem. Median net worth is 1.1, which is much closer to the empirical target of 1.56. The reason is that in the two-asset models households in the middle of the distribution hold the bulk of their wealth in illiquid assets, which have only a small effect on the MPC.

4.2 Alternative Calibrations of the Two Asset Model

We now examine the robustness of our findings for the two-asset model with respect to some of the key parameters of the model. In these experiments, whenever we change parameters, we recalibrate only the discount rate ρ to match the same target for total wealth, unless otherwise noted.

Rates of Return. Changing the rates of return on liquid and illiquid wealth has a significant impact on the MPCs, as illustrated in Columns (3)-(6) of Table 5. Raising the liquid rate or reducing the illiquid rate narrows the return gap between the two assets, which leads to fewer HtM households, lowering the average MPC.[14]

Rebalancing Frequency. Table E.7 in the Appendix illustrates that decreasing the frequency of rebalancing opportunities from an average of once per month to an average of once per quarter (Column 2) or once per year (Column 3) has only a small effect on the average MPC. However, these calibrations with less frequent rebalancing opportunities is more successful in matching the mean level of liquid wealth in the data (0.56). Households hold a larger buffer of liquid wealth both because opportunities to rebalance arrive less frequently and because households that are accumulating wealth must hold their savings in the liquid account for a longer duration on average before transferring them to their illiquid account. The model with rebalancing opportunities on average once per year also features a slightly larger average MPC, despite having fewer hand-to-mouth households, because the less frequent adjustment opportunities means that the wealthy hand-to-mouth households who are waiting to withdraw resources have a higher MPC than in the baseline calibration (38% vs 30%).

Rebalancing Cost. Varying the transaction cost from \$0 to \$3,000, while keeping the arrival rate of rebalancing opportunities at its baseline value of 3, has only very small effects on the average MPC because there are offsetting effects on the distribution and the shape of the consumption function. See Columns (4)-(7) in Table E.7 in the Appendix.

4.3 Tackling the Large Return Differential

The baseline two-asset model can simultaneously match the evidence on MPCs and wealth targets only with a rate of return differential of around 8% per year. A reasonable question is whether a return differential of this magnitude is plausible, and if not, whether there are extensions to the model that could yield the same MPC and wealth statistics as with a smaller return difference.

For the United States since 1950, Jordà, Knoll, Kuvshinov, Schularick, and Taylor (2019) estimate a gap of 5-6% between the real financial return on stocks and housing (which are the biggest components of illiquid wealth) versus the return on T-Bills. However, these calculations do not account for the potentially large tax advantages from investing in housing due to deductibility of

[14]The MPC decompositions in this section are performed relative to the baseline two-asset model in column (2).

	(1) Baseline 2-asset	(2) φ=0.01	(3) φ=0.05	(4) φ=0.01	(5) φ=0.015
Rebalance arrival rate	3.00	3.00	3.00	3.00	3.00
Quarterly MPC (%)	16.1	23.8	46.7	18.1	21.1
Annual MPC (%)	41.2	49.7	63.9	42.5	44.3
Quarterly PHtM MPC(%)	24.3	24.8	51.9	23.9	23.6
Quarterly WHtM MPC (%)	29.8	47.8	75.9	40.3	50.7
Mean MPC at Mean Wealth (%)	7.0	10.4	37.0	6.5	13.3
Prob. HtM status at year t and year t+1	0.9	0.9	0.9	0.9	0.9
Panel A: Calibrated Variables					
Annualized discount factor	0.937	0.937	0.937	0.964	0.973
Annualized liquid return	-2.0	-2.0	-2.0	-2.0	-2.0
Annualized illiquid return	6.2	6.2	6.2	3.2	2.2
Rebalance cost ($)	1456.8	1456.8	1456.8	2400.0	2700.0
Panel B: Targeted Statistics					
Mean total wealth	4.1	4.1	4.1	4.1	4.1
Share hand-to-mouth	39.5	48.6	62.9	39.0	40.4
Share poor hand-to-mouth	13.2	18.0	24.0	13.0	13.3
Panel C: Decomposition					
Gap with Baseline MPC		7.7	30.5	2.0	5.0
Effect of MPC function		0.9	5.0	-1.1	-1.5
Distributional Effect		2.0	15.1	0.1	0.2
Interaction		4.7	10.4	3.0	6.3
Panel D: Wealth Statistics					
Mean liquid wealth	0.2	0.2	0.1	0.2	0.2
Median total wealth	1.1	1.0	0.7	1.0	1.0
Median liquid wealth	0.1	0.0	0.0	0.1	0.0
$b \leq$ \$1000	33.0	44.3	60.2	36.0	38.2
$b \leq$ \$5000	54.7	59.7	70.2	55.3	56.4
$b \leq$ \$10000	66.1	69.8	78.8	66.2	67.1
$w \leq$ \$1000	10.0	14.2	23.4	10.3	10.5
$w \leq$ \$5000	20.8	22.4	26.0	24.0	24.5
$w \leq$ \$10000	23.5	24.6	27.5	29.1	30.0
$w \leq$ \$50000	45.9	47.2	51.0	46.4	47.2
$w \leq$ \$100000	57.7	58.8	61.3	58.0	58.4
Wealth, Top 10% share	62.0	63.5	65.8	60.9	60.6
Wealth, Top 1% share	19.2	20.2	21.6	16.2	15.2
Gini coefficient, total wealth	0.8	0.8	0.8	0.8	0.8

Table 6: Two-asset model with temptation and self-control. b denotes liquid wealth and $w = a + b$ total wealth (or net worth). Column (1): baseline two-asset model. Columns (2)-(5) two-asset model with Gul-Pesendorfer preferences solved with different temptation parameter.

mortgage interest and favorable capital gains treatment, nor the tax advantages that arise from investing in equities inside qualified retirement accounts. Moreover, in the United States many 401(k) plans include a generous employer matching contribution (often around 50%) which further raises the effective rate of return. Finally, owner-occupied housing yields a direct utility flow, which in Kaplan and Violante (2014) we estimated to be equivalent to a financial return of 3-4% per year. Including any of these features in the model can lower the nominal return gap, without changing the implied consequences for savings and spending patterns.

An alternative modification to reduce the rate of return differential is to add a commitment value of holding illiquid assets. This approach is similar in spirit to including a direct utility flow because it amounts to modifying preferences in such a way that households desire to hold more illiquid assets as fraction of total wealth, for a given financial return gap. To assess the quantitative reach of this mechanism, in Table 6 we introduce temptation and self-control, like in Section 3.2.1, into the two-asset model, along the lines proposed by Attanasio, Kovacs, and Moran (2020). Details of the model are in Appendix C.2.

When we recalibrate the discount rate to keep mean total wealth constant at 4.1, but hold the return differential and adjustment cost fixed at their baseline values, adding even a small amount of temptation raises the aggregate MPC substantially. For example, with a temptation parameter of $\varphi = 0.01$, the average quarterly MPC rises from 16% to 24% (Column 2), and with a temptation parameter of $\varphi = 0.05$, the MPC rises to 47% (Column 3). However, as in the one-asset versions of the model, these larger MPCs arise from having a much larger fraction of hand-to-mouth households than in the data. The decomposition relative to the baseline two-asset model confirms the majority of the higher MPC is due to the wealth distribution rather than the consumption function.

We can bring the two-asset model with temptation in line with the evidence on hand-to-mouth households by also recalibrating the return gap and adjustment cost to target the fraction of wealthy HtM (41%) and poor HtM (14%) households. Columns 3 and 4 report the results of these calibrations with temptation parameters of $\varphi = 0.01$ and $\varphi = 0.015$. Both versions are successful in matching these targets, and both yield even higher average MPCS than in the baseline model, with much lower return differences. For example, with $\varphi = 0.015$ (Column 5), the return gap is 4%, compared with 8% in the baseline, and the average MPC is 21%. In order to match the hand-to-mouth targets with a lower return difference, the models with temptation require larger adjustment costs than in the baseline model.

	(1) 1A Baseline	(2) E[a] = 0.56	(3) β Het	(4) IES Het	(5) Spender-Saver	(6) 2A Baseline
MPC (%), out of -$5000	8.3	22.5	28.3	28.6	20.2	44.9
MPC (%), out of -$500	5.0	14.9	19.5	22.7	17.2	26.6
MPC (%), out of $500	4.6	14.0	18.6	20.8	16.8	16.1
MPC (%), out of $5000	3.6	11.3	15.4	14.5	13.2	12.3

Table 7: Quarterly MPC for positive and negative windfalls of different sizes across various models. Column (1): baseline one-asset model. Column (2): one-asset model calibrated to median net worth. Columns (3) and (4): one-asset model with preference heterogeneity. Column (5): spender-saver model. Column (6): baseline two-asset model.

5 Other Concepts of MPCs

So far, we have focused on the quarterly MPC out of an unexpected $500 windfall. In this section, we extend the analysis to (i) MPCs out of unexpected income shocks of different sizes, both positive and negative, (ii) MPCs at different horizons, including MPCs out of the news of future windfalls, and (ii) MPCs out of unexpected illiquid injections (or illiquid capital gains) in the two asset models.

5.1 Sign and Size Asymmetries

To explore sign and size asymmetries, we compute MPCs out of - $500 and+$5,000.[15] Table 7 summarizes these results. Because of the concavity in the consumption function, the MPC out of small windfalls is bigger than the MPC out of larger windfalls. Concavity also implies that MPCs out of income losses are larger than for income gains of the same absolute magnitude.

In the baseline one-asset model these effects are not very strong because the asymmetries are significant only in regions where the consumption function is very concave, where there is only a small share of households in the stationary distribution (Column 1). But in models with a larger share of households in regions of the state space where the consumption function is steep, such as the one-asset model calibrated to mean liquid wealth calibration, one asset models with ex-ante heterogeneity and two-asset models, there is stronger size and sign asymmetry (Columns 2-7).

So far we have focused on MPCs out of a fixed dollar amount as opposed to a fixed fraction of each household's own earnings. We adopted this measure as a baseline because there

[15]Simulating MPCs out of large negative shocks have the problem that they lead to an empty budget set for many poor households, so one has to make ad-hoc assumptions about what happens in that case.

exists reliable quasi-experimental evidence which can be used to assess different models and because of the size-asymmetries documented above. However, the average MPC out of an unexpected common percentage change in income, is also interesting because it corresponds to the partial-equilibrium impact of transitory aggregate income shock with even incidence across all households. We computed the aggregate MPC out of 1% increase in earnings for each households (mean earnings is \$67,000 so this corresponds to a mean windfall of \$670, which is close to \$500). For the one-asset model, the aggregate MPC out of this shock (computed as the total change in consumption divided by the total change in income) is 2.7% , which is close to the income-weighted MPC out of a \$500 gain, 2.7%. For the two-asset model, the aggregate MPC out of the 1% earnings shock is 13.1%. In both models the aggregate MPC out of the 1% earnings shock is lower than the MPC out of a \$500 gain because higher income households have lower MPCs.

5.2 Intertemporal MPCs

In general equilibrium versions of heterogeneous agent models, the entire time path of MPCs matters for the extent of shock propagation (see Auclert, Rognlie, and Straub (2018)). We compute the time profile of MPCs from horizon $t = -4$ (one year before the shock) to $t = +4$ (one year after the shock). The quarterly MPC out of a \$500 gain at horizons $t = -\tau$ should be interpreted as the fraction of the windfall consumed by households upon receiving the news that a windfall will be received in τ quarters' time. The MPC at horizon $t = +\tau$ is the fraction consumed τ quarters after receipt of an unexpected windfall. Appendix D describes how to compute MPCs at different horizons.

Figure 5 reports the time profile of MPCs across several models. In the baseline one-asset model the profile is very flat because of the low share of hand-to-mouth households. In one-asset models with discount factor heterogeneity, the profile is more tent-shaped: MPCs out of news are much smaller than actual MPCs, as are lagged MPCs. This is also true in versions of the model with IES heterogeneity (not displayed). These dynamics are driven by impatient or high IES households. As pointed out by Auclert, Rognlie, and Straub (2018), the tent-shape is particularly extreme in the spender-saver model. In this model, the time profile is an average of the MPCs for the savers, which closely mimics the shape of the profile in the baseline one-asset model, and the MPCs for the spenders, who have a strong spending response on impact, but not at other leads or lags.

The time path for the two-asset model shows a high response at impact because of the large share of hand-to-mouth households, as well as a sizable reaction to news about future windfalls.

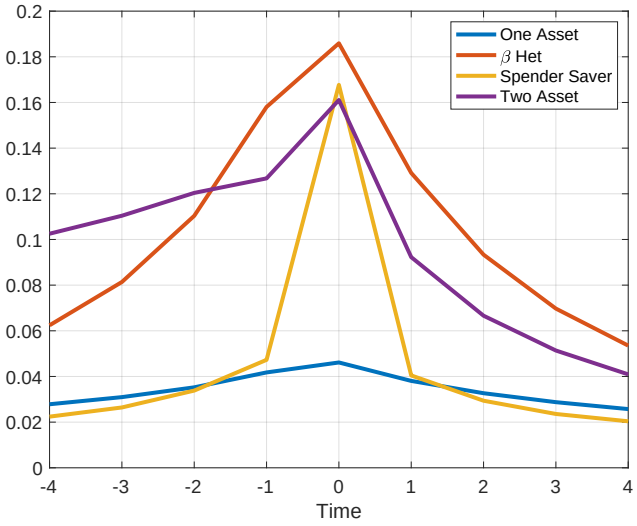


Figure 5: Intertemporal MPCs for four models. At $t = 0$ we observe the impact MPC; at $t > 0$, the lagged MPC out of a windfall received at $t = 0$; at $t < 0$, the MPC out of the news of a windfall that will be received at $t = 0$.

The reason for this latter effect is the low calibrated discount factor, which raises the MPCs of unconstrained households, who are the ones that respond to future windfalls. Comparing these anticipation effects in the model with corresponding empirical evidence requires caution because in the model future windfalls are received with certainty, whereas in reality there is always some degree of uncertainty around whether the windfalls will ultimately be received.

In Section 2.2, we determined that in the one-asset model the annual MPC is much smaller than four times the quarterly MPC. Figure 5 clearly illustrates that this conclusion holds across models, because the time path of the intertemporal MPCs is –sharply, in some economies– declining over time after $t = 0$.

5.3 MPC out of Illiquid Wealth

In the two-asset model, we can also compute the MPC out of a small injection of illiquid wealth. Conceptually, the closest empirical counterpart of this MPC would be an estimated MPC out of changes in housing wealth or out of unrealized capital gains/losses on equity in retirement accounts or private businesses. To put our analysis in the context of recent empirical contributions, we note that Carroll, Otsuka, and Slacalek (2011) estimate an average quarterly MPC out of housing wealth around 2% percent. Mian, Rao, and Sufi (2013) estimate it at 1.5% percent, and uncover a great deal of heterogeneity with respect to income and leverage. Di Maggio, Kermani, and Majlesi (2020) analyze unrealized capital gains in stock market wealth and conclude that the literature, including their own research, finds values between 0 and 2.5% for the average quar-

terly MPC. In our model, the average quarterly MPC out of an illiquid windfall of $500 is 1.4% and the MPC out of an illiquid windfall of $5,000 is 2.4%. As in the data, the average illiquid MPC is roughly an order of magnitude smaller than the liquid MPC.

6 Conclusions

Marginal propensities to consume (MPCs) are the most important feature of household spending behavior for macroeconomics. They play an essential role in determining the effects of changes in aggregate demand, the strength of the fiscal multiplier, the transmission mechanism of monetary policy and the portfolio share of wealth held in risky assets. Empirical evidence suggests that in response to an unexpected windfall, the MPC is large on average and displays a wide degree of heterogeneity across households. But canonical representative agent models fail in accounting for this evidence.

One-asset heterogeneous agent models with idiosyncratic income risk and incomplete markets generate an average MPC that is an order of magnitude bigger than in representative agent models, but still significantly smaller than in the data. Allowing for ex-ante heterogeneity across households or certain behavioral biases in preferences can raise the average MPC to be in line with the data. However, this comes at the cost of misrepresenting key features of the wealth distribution. Most importantly, these models vastly overstate the fraction of households with low wealth.

Two-asset heterogeneous agent models with a liquid and illiquid asset can resolve this tension when they feature a large enough return gap between the two assets. Theory and data suggest a number of plausible justifications for such return differentials.

References

ACHDOU, Y., J. HAN, J.-M. LASRY, P.-L. LIONS, AND B. MOLL (2017): "Income and wealth distribution in macroeconomics: A continuous-time approach," Discussion paper, National Bureau of Economic Research.

AGUIAR, M. A., M. BILS, AND C. BOAR (2020): "Who are the Hand-to-Mouth?," Discussion paper, National Bureau of Economic Research.

ANDERSEN, A. L., A. S. JENSEN, N. JOHANNESEN, C. T. KREINER, S. LETH-PETERSEN, AND A. SHERIDAN (2021): "How do households respond to job loss? Lessons from multiple high-frequency data sets," .

ATTANASIO, O., A. KOVACS, AND P. MORAN (2020): "Temptation and commitment: Understanding hand-to-mouth behavior," Discussion paper, National Bureau of Economic Research.

AUCLERT, A. (2019): "Monetary policy and the redistribution channel," *American Economic Review*, 109(6), 2333–67.

AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2018): "The intertemporal keynesian cross," Discussion paper, National Bureau of Economic Research.

AYDIN, D. (2022): "Consumption response to credit expansions: Evidence from experimental assignment of 45,307 credit lines," *American Economic Review*, 112(1), 1–40.

BENHABIB, J., AND A. BISIN (2018): "Skewed wealth distributions: Theory and empirics," *Journal of Economic Literature*, 56(4), 1261–91.

BILBIIE, F. O. (2020): "The new Keynesian cross," *Journal of Monetary Economics*, 114, 90–108.

BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): "Consumption inequality and partial insurance," *American Economic Review*, 98(5), 1887–1921.

BUNN, P., J. LE ROUX, K. REINOLD, AND P. SURICO (2018): "The consumption response to positive and negative income shocks," *Journal of Monetary Economics*, 96, 1–15.

CAMPBELL, J. Y., AND N. G. MANKIW (1989): "Consumption, income, and interest rates: Reinterpreting the time series evidence," *NBER Macroeconomics Annual*, 4, 185–216.

CARROLL, C. D. (2001): "A theory of the consumption function, with and without liquidity constraints," *Journal of Economic Perspectives*, 15(3), 23–45.

CARROLL, C. D., M. B. HOLM, AND M. S. KIMBALL (2021): "Liquidity constraints and precautionary saving," *Journal of Economic Theory*, 195, 105276.

CARROLL, C. D., AND M. S. KIMBALL (1996): "On the concavity of the consumption function," *Econometrica: Journal of the Econometric Society*, pp. 981–992.

——— (2001): "Liquidity constraints and precautionary saving," .

CARROLL, C. D., M. OTSUKA, AND J. SLACALEK (2011): "How large are housing and financial wealth effects? A new approach," *Journal of Money, Credit and Banking*, 43(1), 55–79.

CARROLL, C. D., J. SLACALEK, K. TOKUOKA, AND M. N. WHITE (2017): "The distribution of wealth and the marginal propensity to consume," *Quantitative Economics*, 8(3), 977–1020.

COMMAULT, J. (2017): "How does consumption respond to a transitory income shock? reconciling natural experiments and structural estimations," Discussion paper, Working Paper.

DE NARDI, M., AND G. FELLA (2017): "Saving and wealth inequality," *Review of Economic Dynamics*, 26, 280–300.

DEATON, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59(5), 1221–48.

DI MAGGIO, M., A. KERMANI, AND K. MAJLESI (2020): "Stock market returns and consumption," *The Journal of Finance*, 75(6), 3175–3219.

EPSTEIN, L. G., AND S. E. ZIN (1991): "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis," *Journal of Political Economy*, 99(2), 263–286.

FAGERENG, A., M. B. HOLM, AND G. J. J. NATVIK (2019): "MPC Heterogeneity and Household Balance Sheets," *Journal of Monetary Economics*.

FUSTER, A., G. KAPLAN, AND B. ZAFAR (2021): "What would you do with $500? Spending responses to gains, losses, news, and loans," *The Review of Economic Studies*, 88(4), 1760–1795.

GANONG, P., D. JONES, P. NOEL, F. GREIG, D. FARRELL, AND C. WHEAT (2020): "Wealth, race, and consumption smoothing of typical income shocks," *NBER Working Paper*, (w27552).

GELMAN, M. (2021): "What drives heterogeneity in the marginal propensity to consume? Temporary shocks vs persistent characteristics," *Journal of Monetary Economics*, 117, 521–542.

GELMAN, M., S. KARIV, M. D. SHAPIRO, D. SILVERMAN, AND S. TADELIS (2020): "How individuals respond to a liquidity shock: Evidence from the 2013 government shutdown," *Journal of Public Economics*, 189, 103917.

GOLOSOV, M., M. GRABER, M. MOGSTAD, AND D. NOVGORODSKY (2021): "How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income," Discussion paper, National Bureau of Economic Research.

GUL, F., AND W. PESENDORFER (2001): "Temptation and Self-control," *Econometrica*, 69(6), 1403–1435.

HAVRANEK, T., AND A. SOKOLOVA (2020): "Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say "probably not"," *Review of Economic Dynamics*, 35, 97–122.

HELPMAN, E. (1981): "Optimal spending and money holdings in the presence of liquidity constraints," *Econometrica: Journal of the Econometric Society*, pp. 1559–1570.

JAPPELLI, T., AND L. PISTAFERRI (2010): "The consumption response to income changes," *Annual Review of Economics*, 2(1), 479–506.

——— (2014): "Fiscal policy and MPC heterogeneity," *American Economic Journal: Macroeconomics*, 6(4), 107–36.

JOHNSON, D. S., J. A. PARKER, AND N. S. SOULELES (2006): "Household expenditure and the income tax rebates of 2001," *American Economic Review*, 96(5), 1589–1610.

JORDÀ, Ò., K. KNOLL, D. KUVSHINOV, M. SCHULARICK, AND A. M. TAYLOR (2019): "The rate of return on everything, 1870–2015," *The Quarterly Journal of Economics*, 134(3), 1225–1298.

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary policy according to HANK," *American Economic Review*, 108(3), 697–743.

KAPLAN, G., AND G. L. VIOLANTE (2014): "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 82(4), 1199–1239.

KAPLAN, G., G. L. VIOLANTE, AND J. WEIDNER (2014): "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 2014(1), 77–138.

KEKRE, R., AND M. LENEL (2021): "Monetary policy, redistribution, and risk premia," Discussion paper, National Bureau of Economic Research.

KIMBALL, M. S. (1990): "Precautionary saving and the marginal propensity to consume," Discussion paper, National Bureau of Economic Research.

KRUEGER, D., K. MITMAN, AND F. PERRI (2016): "Macroeconomics and household heterogeneity," in *Handbook of Macroeconomics*, vol. 2, pp. 843–921. Elsevier.

KRUSELL, P., B. KURUŞÇU, AND A. A. SMITH JR (2002): "Time Orientation and Asset Prices," *Journal of Monetary Economics*, 49(1), 107–135.

——— (2010): "Temptation and Taxation," *Econometrica*, 78(6), 2063–2084.

KUENG, L. (2018): "Excess sensitivity of high-income consumers," *The Quarterly Journal of Economics*, 133(4), 1693–1751.

LAIBSON, D. (1997): "Golden eggs and hyperbolic discounting," *The Quarterly Journal of Economics*, 112(2), 443–478.

LAIBSON, D., P. MAXTED, AND B. MOLL (2021): "Present bias amplifies the household balance-sheet channels of macroeconomic policy," Discussion paper, National Bureau of Economic Research.

LEWIS, D. J., D. MELCANGI, AND L. PILOSSOPH (2019): "Latent heterogeneity in the marginal propensity to consume," *FRB of New York Staff Report*, (902).

MIAN, A., K. RAO, AND A. SUFI (2013): "Household balance sheets, consumption, and the economic slump," *The Quarterly Journal of Economics*, 128(4), 1687–1726.

MISRA, K., AND P. SURICO (2014): "Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs," *American Economic Journal: Macroeconomics*, 6(4), 84–106.

NAKAJIMA, M. (2017): "Assessing bankruptcy reform in a model with temptation and equilibrium default," *Journal of Public Economics*, 145, 42–64.

PARKER, J. A., AND N. S. SOULELES (2019): "Reported effects versus revealed-preference estimates: Evidence from the propensity to spend tax rebates," *American Economic Review: Insights*, 1(3), 273–90.

PARKER, J. A., N. S. SOULELES, D. S. JOHNSON, AND R. MCCLELLAND (2013): "Consumer spending and the economic stimulus payments of 2008," *American Economic Review*, 103(6), 2530–53.

PATTERSON, C. (2019): "The matching multiplier and the amplification of recessions," *Unpublished Manuscript, Northwestern University*.

PAVONI, N., AND H. YAZICI (2017): "Optimal life-cycle capital taxation under self-control problems," *The Economic Journal*, 127(602), 1188–1216.

SHAPIRO, M. D., AND J. SLEMROD (2003): "Consumer response to tax rebates," *American Economic Review*, 93(1), 381–396.

——— (2009): "Did the 2008 tax rebates stimulate spending?," *American Economic Review*, 99(2), 374–79.

WERNING, I. (2015): "Incomplete markets and aggregate demand," Discussion paper, National Bureau of Economic Research.

WOLF, C. (2020): "The missing intercept: A demand equivalence approach," Discussion paper, MIT.

ZELDES, S. P. (1989): "Consumption and liquidity constraints: an empirical investigation," *Journal of political economy*, 97(2), 305–346.

Online Appendix

The Marginal Propensities to Consume in Heterogeneous Agent Models

Greg Kaplan and Giovanni L. Violante

This Appendix is organized as follows. Section A describes the estimation of the income process and how we computed some key statistics of the wealth distribution. Section B describes the annual calibration of the baseline model, lays out its continuous time version and derives the MPC under certainty and no borrowing constraints. Section C lays out the household problem under temptation-self control and under present bias. Section E contains additional Tables.

A Data

A.1 Panel Study of Income Dynamics

We use data from the Panel Study of Income Dynamics (PSID) on total annual household labor income for households with heads aged 25 to 65 from 1968 to 2008. We drop households with annual labor income less than \$7,250 in 2016 dollars, which correspond to 1,000 hours per year at \$7.25 per hour (or part-time employment at the ongoing minimum wage). We remove age and year effects in a first stage by regressing household labor income on a full set of year and age dummies and we construct the empirical counterparts to $m_{2,d}$ using the residuals from this regression. The resulting moments are shown in the first column of Table A.1.

Let $\log y_t^{ann}$ be annual labor income in year t, and let annual income growth at lag d be

$$\Delta_d \log y_t^{ann} = \begin{cases} \log y_t^{ann} & \text{if } d = 0 \\ \log y_{t+d}^{ann} - \log y_t^{ann} & \text{if } d > 0 \end{cases}$$

Define cross-sectional moments of annual income growth of order j at lag d as

$$m_{j,d} = \mathbb{E}\left[(\Delta_d \log y_t^{ann})^j\right]$$

Lag (d)	$m_{2,d}$	$m_{4,d}$	κ_d
0	0.504	0.930	3.65
1	0.142	0.220	10.90
2	0.207	0.369	8.57
3	0.235	0.410	7.42
4	0.280	0.544	6.96
5	0.295	0.557	6.39
6	0.335	0.694	6.19
7	0.352	0.729	5.87
8	0.383	0.838	5.72
9	0.398	0.885	5.59
10	0.422	0.967	5.43

Table A.1: Empirical moments of annual income growth at different lags. Source: PSID 1968-2008

and the kurtosis of income growth at different lags as

$$\kappa_d = \frac{m_{4,d}}{(m_{2,d})^2}$$

Table A.1 reports the empirical estimates for the cross-sectional moments that we use in estimation.

A.2 Estimated Income Processes

A.2.1 Discrete Time

We model the discrete-time quarterly income process y_t as follows:

$$\log y_s = \begin{cases} z_t + \varepsilon_t & \text{with probability } \lambda_\varepsilon, \quad \varepsilon_t \sim \mathcal{N}\left(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2\right) \\ z_t & \text{with probability } 1-\lambda_\varepsilon. \end{cases}$$

$$z_t = \begin{cases} \phi_z z_{t-1} + \eta_t & \text{with probability } \lambda_\eta, \quad \eta_t \sim \mathcal{N}\left(-\frac{\sigma_\eta^2}{2}, \sigma_\eta^2\right) \\ \phi_z z_{t-1} & \text{with probability } 1-\lambda_\eta \end{cases} \tag{A.1}$$

We define annual income y^{anm} as the sum of the four quarterly income within the year and, based on this definition, we are able to construct the model counterparts of all the empirical moments in Table A.1.

With $\lambda_\epsilon, \lambda_\eta$ fixed exogenously, we require three moments to estimate the three parameters $\left(\phi_z, \sigma_\eta^2, \sigma_\epsilon^2\right)$. Note that the set of moments $\{m_{2,d}\}$ for $d = 1 \ldots D$ contains the identical infor-

Process	ϕ_z	σ^2_η	λ_η	ϕ_u	σ^2_ε	λ_ε	σ^2_{FE}
Quarterly							
Baseline	0.988	0.0439	0.250		0.6376	0.250	
Shocks Arrive Quarterly	0.988	0.0108	1		0.2087	1	
Estimated Shock Arrival Rates	0.987	0.0516	0.237		1.6243	0.073	
Krueger, Mitman and Perri (2017 formula)	0.988	0.0108	1		0.0494	1	
Annual							
Baseline	0.953	0.0422			0.0494		
No transitory shocks	0.953	0.0422			0		
High persistence ($\phi_z^{ann} = 0.995$)	0.995	0.0043			0.0688		
With fixed effects	0.916	0.0445			0.0479		0.180
Continuous							
Baseline	0.009	0.134	0.250	0.347	0.652	0.250	
Estimated Shock Arrival Rates	0.012	0.239	0.060	0.347	1.28	0.063	

Table A.2: Parameter estimates of various statistical models for income dynamics

mation to the auto-covariance function out to D lags. We express the data in this way since it is more convenient for extending the estimation strategy to the case where λ_ϵ and λ_η are also estimated.

In our baseline specification, we choose $(m_{2,0}, m_{2,1}, m_{2,5})$ as our moments to match. The first row of Table A.2 shows our baseline estimates in which we assume that income shocks arrive on average once per year, $\lambda_\epsilon = \lambda_\eta = 0.25$. In the second row, we shows corresponding estimates when the income shocks arrive every quarter, $\lambda_\epsilon = \lambda_\eta = 1$. In the third row of Table A.2 we estimate the shock arrival rates $\lambda_\epsilon, \lambda_\eta$ alongside the other parameters of the income process. This requires two additional moments. To find moments that identify these parameters, we note that the main effect of lowering the arrival rates below 1 is that it induces excess kurtosis into the distribution of annual income growth, more a given variance of income growth.

The second and third columns of Table A.1 report $m_{4,d}$ and κ_d out to ten lags. Note that log income itself does not display much excess kurtosis, but annual income growth is very leptokurtic, with the degree of leptokurtosis declining as the lag length increases. We add κ_1 and κ_5 as the additional moments to identify $\lambda_\epsilon, \lambda_\eta$. The estimates are reported in the third row of Table A.2. They suggest that persistent shocks arrive on average close to once per year, but that transitory shocks are much less frequent and much larger on average than implied by the more restrictive model.

The fourth row of Table A.2, labeled "Krueger, Mitman and Perri (2017) formula" constructs

the quarterly estimates by applying the following formulas to annual estimates :

$$\begin{aligned}
\phi_z &= (\phi_z^{ann})^{0.25} \\
\sigma_\varepsilon^2 &= (\sigma_\varepsilon^{ann})^2 \\
\sigma_\eta^2 &= \left(\frac{1-\phi_z^2}{1-(\phi_z^{ann})^2}\right)\left(\sigma_\eta^{ann}\right)^2
\end{aligned}$$

The annual estimates upon which these are based are shown in the bottom panel of Table A.2. These are constructed by reinterpreting y_t as annual income and estimating the parameters by matching the moments $(m_{2,0}, m_{2,1}, m_{2,5})$.

In the next section of the table we report estimates for income processes estimated at annual frequency. We first exclude transitory shocks, interpreting them as measurement error. Next, we estimate a version where we restrict the AR(1) component to have high persistence $\phi_z^{ann} = 0.995$, and a version where we include an individual-specific fixed effect. For this latter model we add the moment to $m_{2,10}$ to identify the additional parameter (the variance of the fixed effect σ_{FE}^2).

A.2.2 Continuous Time

We model the continuous time income process y_t as follows:

$$\begin{aligned}
\log y_t &= z_t + u_t \\
dz_t &= -\phi_z z_t + \eta_{it} dJ_{\eta,t} \\
du_t &= -\phi_u u_t + \varepsilon_t dJ_{\varepsilon,t}
\end{aligned}$$

where dJ_η is a Poisson process with arrival rate λ_η and dJ_ε is a Poisson process with arrival rate λ_ε. The innovations are given by

$$\begin{aligned}
\eta_{it} &\sim N\left(0, \sigma_\eta^2\right) \\
\varepsilon_{it} &\sim N\left(0, \sigma_\epsilon^2\right)
\end{aligned}$$

Note that since income is a flow, there is no natural concept of purely transitory shock in continuous time. For consistency with the discrete time formulation, so that the two versions match the same data moments, in our baseline model we restrict $\phi_u = \frac{1}{2}\log 2$ which implies a half-life of two quarters. This is broadly consistent with a discrete time annual formulation in which a transitory shocks lasts for one year. In our baseline model we restrict $\lambda_\eta = \lambda_\varepsilon = 0.25$ as in the discrete time model.

The parameter estimates for the continuous time income process are reported in the bottom

panel of Table A.2. We also report estimates for the version where we estimate the shock arrival rates. As in the quarterly discrete time model, when the shock arrival rates are estimated we find them to be larger and less frequent than when restricted to arrive on average once per quarter.

A.3 Survey of Consumer Finances and Wealth Statistics

To compute moments of the wealth distribution, we first select all households in the 2019 *Survey of Consumer Finances*, without any age restriction. Then as explained we drop the top 5% of the wealth distribution.

Our definition of household labor income includes wage and salary income plus social security income. It excludes other business income, other government transfers, as well as interests, dividends and capital gains. Mean household labor income is $\$67,132$ and median income is $\$54,266$.

Our definition of net worth is the baseline definition of the SCF for total net worth (variable NETWORTH). See the document Networth Flow chart.pdf in https://www.federalreserve.gov/econres/scfindex.htm. It includes all financial assets (bank accounts, CDs, mutual funds, retirement accounts, and directly held stocks and bonds), vehicles, housing wealth and private business equity net of all types of unsecured and secured debt. Mean wealth is $\$275,665$ and median wealth is $\$103,380$.

Our definition of net liquid and illiquid wealth follows Kaplan and Violante (2014) and Kaplan, Violante, and Weidner (2014). Net liquid wealth includes bank accounts and directly held mutual funds, stocks and bond net of credit card debt. In terms of the SCF variables: FIN - CDS - SACVBND - CASHLI - OTHMA - RETLIQ - (OTHLOC + CCBAL + ODEBT).

Our definition of net illiquid wealth is residual, i.e. net worth minus net liquid wealth. The biggest items among financial assets are retirement accounts, among non-financial assets are housing and business equity. The biggest components on the liability side are mortgages. In terms of the SCF variables net illiquid wealth is: (CDS + SACVBND + CASHLI + OTHMA + RETLIQ) + NFIN - MRTHEL - RESDBT - INSTALL - ODEBT.

B One-Asset Models

B.1 Annual calibration

As in the baseline quarterly calibration, we set $\gamma = 1$, the credit limit to zero, $\delta = 1/50$ so that the expected adult life span is 50 years, and the real interest rate $r = 0.01$. Table A.2 reports the annual value for variances and correlation coefficient estimated to match the same annual

covariances restrictions as for the baseline calibration. The discount factor is calibrated internally to match a ratio of mean net worth to mean annual household labor income ratio of 4.1. We obtain an annualized value of 0.980 for the effective discount factor β, i.e. virtually the same value as in the quarterly calibration. This is reassuring, since annual and quarterly calibrations should replicate exactly the same set of moments.

B.2 Continuous-time formulation

The continuous-time version of the household problem (1) is:

$$\begin{aligned} & \max_{\{c_t\}} \mathbb{E}_0 \int_0^\infty e^{-(\delta+\tilde{\rho})t} u\left(c_t\right) dt \\ & s.t. \\ \dot{b}_t &= \exp\left(y_t\right) + r b_t - c_t \\ b_t &\geq -\underline{b} \\ y_t &\sim F\left(y_t, y_{t-1}\right) \end{aligned} \tag{B.2}$$

In this formulation, $\delta > 0$ is the instantaneous death rate, $\tilde{\rho} > 0$ the discount rate, $\rho = \tilde{\rho} + \delta$, and $\dot{b}_t$ represents savings. The corresponding HJB equation is:

$$\begin{aligned} \rho v\left(b, y\right) = \max_c \quad & u\left(c\right) + v_b\left(b, y\right) \dot{b} + \mathcal{A}(y) v\left(b, y\right) \\ & \text{subject to} \\ & \dot{b} = r b + y - c \\ & b \geq 0 \end{aligned}$$

where $\mathcal{A}$ is the infinitesimal generator of the income process. The continuous time equivalent of the income process in (A.1) is:

$$\begin{aligned} y_t &= z_t + \varepsilon_t dJ_{\varepsilon t}, \\ dz_t &= -\left(1-\phi\right) z_{t-1} + \eta_t dJ_{\eta t}, \text{ with } \eta_t \sim \mathcal{N}\left(0, \sigma_\eta\right) \\ \varepsilon_t &\sim \mathcal{N}\left(0, \sigma_\varepsilon\right) \end{aligned} \tag{B.3}$$

where $J_{\varepsilon t}$ and $dJ_{\eta t}$ are jump processes with arrival rate λ_ε and λ_η respectively. To estimate the parameters of the income process, we time-aggregate in order to match the same set of annual moments described above. Table A.2 reports the point estimates of the parameters of the income process, expressed quarterly for ease of comparison with the discrete time counterpart.

MPC in continuous time To define and compute the MPC in the continuous time version of the model we follow Achdou, Han, Lasry, Lions, and Moll (2017). In continuous time, the MPC is defined over an interval τ as:

$$\mathfrak{m}_\tau(b,y) = \frac{\partial C_\tau(b,y)}{\partial b} \simeq \frac{C_\tau(b+x,y) - C_\tau(b,y)}{x}, \tag{B.4}$$

where

$$C_\tau(b,y) = \mathbb{E}_0\left[\int_0^\tau c(b_t,y_t)dt | b_0 = b, y_0 = y\right].$$

The conditional expectation $C_\tau(b,y)$ can be conveniently computed using the Feynman-Kac formula. This formula establishes a link between conditional expectations of stochastic processes and solutions to partial differential equations. Applying the formula, we have $C_\tau(b,\log y) = K(b,\log y,0)$, where $K(b,y,t)$ satisfies the partial differential equation on $[0,\tau]$

$$c(b,y) + K_b(b,y,t)\dot{b}(b,y) + K_y(b,y,t)\left[-(1-\phi)z\right] + \mathcal{A}(y)K(b,y,t) \tag{B.5}$$

with terminal condition $\Gamma(b,y,\tau) = 0$, where $\mathcal{A}$ is the infinitesimal generator of the income process.

B.3 MPC Under Certainty and No Borrowing Constraints

The budget constraint of the household problem (1) is:

$$c_t = Rb_t + y_t - b_{t+1}$$

Iterating forward, we obtain:

$$c_0 + \frac{1}{R}c_1 + \frac{1}{R^2}c_2 + ... = Rb_0 + \sum_{t=0}^{\infty}\left(\frac{1}{R}\right)^t y_t.$$

Using the household Euler equation between t and $t+1$

$$c_{t+1} = (\beta R)^{\frac{1}{\gamma}} c_t$$

to substitute c_t at every t on the left hand side as a function of c_0, we arrive at:

$$c_0 + \frac{1}{R}c_0(\beta R)^{\frac{1}{\gamma}} + \frac{1}{R^2}c_0\left[(\beta R)^{\frac{1}{\gamma}}\right]^2 + ... = Rb_0 + \sum_{t=0}^{\infty}\left(\frac{1}{R}\right)^t y_t$$

and collecting terms on the left hand side:

$$c_0 \left[\frac{1}{1 - R^{-1} (\beta R)^{\frac{1}{\gamma}}} \right] = Rb_0 + \sum_{t=0}^{\infty} \left(\frac{1}{R} \right)^t y_t$$

which proves that $\mathfrak{m}_0^* = 1 - R^{-1} (\beta R)^{\frac{1}{\gamma}}$.

C Models with Behavioral Biases in Preferences

C.1 Temptation and Self-Control: Discrete Time One-Asset Model

In this Appendix we describe the model of temptation and self-control that we solve in Section 3.2.1. We assume that in each period the agent is tempted to consume its entire wealth and the temptation utility function is the same as for actual consumption. The household problem can then be written in recursive form as:

$$\begin{aligned} v(b,y) &= \max_{b' \geq 0} \left\{ u(c) + \beta \mathbb{E} v\left(b', y'\right) \right\} + \varphi \left[u(c) - u(Rb + y) \right] \\ &\quad \text{subject to} \\ c + b' &= Rb + y, \quad b' \geq 0 \end{aligned}$$

The parameter $\varphi \geq 0$ measures the strength of the temptation. When $\varphi = 0$, the model collapses to the model without temptation. When φ is very large, the agents gives in to temptation and consumes all its cash in hand every period. In this case the MPC out of additional income is one.

The first-order condition for this problem is

$$u_c(c) = \beta R \mathbb{E} \left[\left(1 - \frac{\varphi}{1+\varphi} \frac{u_c(Rb' + y')}{u_c(c')} \right) u_c(c') \right]. \tag{C.6}$$

This first-order condition can be interpreted as a modified Euler equation, with an endogenous discount factor. For example, with log preferences $u(c) = \log(c)$ the endogenous discount factor becomes

$$\beta \left(1 - \frac{\varphi}{1+\varphi} \left(\frac{c'}{Rb' + y'} \right) \right)$$

which makes it clear that, for a given value of ϕ, households who consume a higher fraction of their wealth act as if they are more impatient. These are typically poorer households, and so with this preference formulation, the effective discount factor tends to be lower for household with lower wealth. In the limit, as households become hand-to-mouth, $c' = Rb' + y'$ and their discount factor becomes $\beta \frac{1}{1+\varphi} < \beta$.

C.2 Temptation and Self-Control: Continuous Time Two-Asset Model

In this Appendix we describe the model of temptation and self-control that we solve in Section 4.3. In continuous time, the recursive formulation of the two-asset model with temptation and self-control can be written as follows.

$$\begin{aligned}\rho v\left(a,b\right)=\max_{c}\left(1+\varphi\right)u\left(c\right)+\left[\partial_{b}v\left(a,b\right)+\varphi\partial_{b}\hat{v}\left(a,b\right)\right]\left(r^{b}b+y-c\right)\\+\left[\partial_{a}v\left(a,b\right)+\varphi\partial_{a}\hat{v}\left(a,b\right)\right]\left(r^{a}a\right)+\mathcal{A}\left[v+\varphi\hat{v}\right]\left(a,b\right)\\+\lambda\left[v^{*}\left(a,b\right)-v\left(a,b\right)\right]+\lambda\varphi\left[\hat{v}^{*}\left(a,b\right)-\hat{v}\left(a,b\right)\right]\\-\varphi\delta\hat{v}\left(a,b\right)\end{aligned}$$

where $\mathcal{A}$ is the infinitesimal generator of the income process and

$$v^{*}\left(a,b\right)=\max\left\{v\left(a,b\right),\max_{a'+b'\leq a+b-\kappa}v\left(a',b'\right)\right\}$$

The function $\hat{v}(a,b)$ is the temptation value function, which solves

$$\begin{aligned}\hat{\rho}\hat{v}\left(a,b\right)=\max_{c}u\left(c\right)+\partial_{b}\hat{v}\left(a,b\right)\left(r^{b}b+y-c\right)\\+\partial_{a}\hat{v}\left(a,b\right)r^{a}a+\mathcal{A}\hat{v}\left(a,b\right)\\+\lambda\left[\hat{v}^{*}\left(a,b\right)-\hat{v}\left(a,b\right)\right]\end{aligned}$$

where

$$\hat{v}^{*}\left(a,b\right)=\max\left\{\hat{v}\left(a,b\right),\max_{a'+b'\leq a+b-\kappa}\hat{v}\left(a',b'\right)\right\}$$

The key assumption is that $\hat{\rho}>>\rho$, so that the household is tempted to act according to a preference specification that discounts the future at a much higher rate. In our simulations we set the quarterly value for $\hat{\rho}=90\%$. The first-order condition for consumption satisfies

$$\left(1+\varphi\right)u'\left(c\right)=\partial_{b}v\left(a,b\right)+\varphi\partial_{b}\hat{v}\left(a,b\right).$$

When $\rho=\hat{\rho}$ or $\varphi=0$, the model collapses to the standard model without temptation.

C.3 Present-Bias

In this Appendix we describe the model of naive present bias that we solve in Section 3.2.2. First consider the problem of a household that does not suffer from present bias:

$$\begin{aligned}\rho\tilde{v}\left(b,y\right) &= \max_{c} u(c)+\tilde{v}_{b}\left(b,y\right)\dot{b}+\mathcal{A}\tilde{v}\left(b,y\right)\\ &\quad \text{subject to}\\ \dot{b} &= rb+y-c, \quad b\geq 0\end{aligned}$$

where $\mathcal{A}$ is the infinitesimal generator of the income process. The first order condition to this optimization problem for $b>0$ is

$$u_{c}(c)=\tilde{v}_{b}\left(b,y\right)$$

The solution to this problem defines a consumption function given by

$$\tilde{c}\left(b,y\right)=\min\{u_{c}^{-1}\left[\tilde{v}_{b}\left(b,y\right)\right],y\}$$

A household with naive present bias has a continuation value given by

$$v\left(b,y\right)=\zeta\tilde{v}\left(b,y\right) \quad \text{for} \quad \zeta<1.$$

So for $b>0$, consumption solves the first order condition condition

$$u_{c}(c)=v_{b}\left(b,y\right)=\zeta\tilde{v}_{b}\left(b,y\right).$$

With CRRA utility, this gives the consumption function

$$c\left(b,y\right)=\min\{\zeta^{-\frac{1}{\gamma}}\tilde{c}\left(b,y\right),y\}.$$

D MPCs at Different Horizons

To compute the MPC at different horizons, we proceed as follows. Recall the definition of the impact MPC in equation (2) in the main text. Let, for example, $t=1$ be the horizon of interest. Then, the MPC at horizon 1 out of a windfall income x is:

$$\mathrm{m}_{1}\left(x;b,y\right)=\frac{\int_{Y}\left[c\left(b'\left(b+x,y\right),y'\right)-c\left(b'\left(b,y\right),y'\right)\right]dF\left(y',y\right)}{x}$$

Iterating this procedure forward, one obtains $\mathrm{m}_{t}\left(x;b,y\right)$, for all $t>0$. The cumulative MPC until horizon T is simply the sum of the MPCs at each horizon $t=0,1,...,T$. The average (or aggre-

gate) MPC at horizon t is obtained by integrating the function $\mathfrak{m}_t(x;b,y)$ under the stationary distribution, i.e.

$$\bar{\mathfrak{m}}_t(x) = \int_{B\times Y} \mathfrak{m}_t(x;b,y)\, d\mu(b,y). \tag{D.7}$$

Finally, we are also interested in the MPC out of the news that a windfall of size x will be received in the future. For example, the MPC at horizon -1, i.e. out of the announcement that x will be paid next period, is:

$$\mathfrak{m}_{-1}(x;b,y) = \frac{c(x;b,y) - c(b,y)}{x} \tag{D.8}$$

where $c(b,y)$ is the solution to the Bellman equation corresponding to the optimization problem (1):

$$v(b,y) = \max_c u(c) + \beta \mathbb{E}\left[v(b',y')|y\right]$$

and $c(x;b,y)$ is the solution to the following Bellman equation, modified to account for the fact that the household expects x next period:

$$v(b,y) = \max_c u(c) + \beta \mathbb{E}\left[v(b'+x,y')|y\right] \tag{D.9}$$

and subject to the same set of constraints as (1).

E Additional Tables

	(1) Quarterly (Discrete)	(2) Annual (Discrete)	(3) Quarterly (Continuous)
Quarterly MPC (%)	4.6		3.0
Annual MPC (%)	14.6	14.3	11.5
Quarterly MPC of the HtM (%)	28.7		33.7
Share of HtM (%)	2.5	8.5	2.0
Annualized discount factor	0.980	0.980	0.985
Panel A: Decomposition			
Gap with Baseline MPC			
Effect of MPC Function			
Effect of Distribution			
Interaction			
Panel B: Wealth Statistics			
Mean wealth	4.1	4.1	4.1
Median wealth	1.3	1.2	1.6
$a \leq$ \$1000	2.5	8.6	2.5
$a \leq$ \$5000	11.6	15.6	8.0
$a \leq$ \$10000	18.5	21.9	13.4
$a \leq$ \$50000	40.3	42.7	34.9
$a \leq$ \$100000	51.9	54.1	48.0
Wealth, top 10% share	46.6	53.0	40.3

Table E.1: Baseline one-asset model and calibrations for different model frequency (annual and continuous time). Model frequency indicates the frequency at which consumption and saving decisions are made. See Table A.2 for details on the income process at different frequencies.

	(1) Baseline	(2) No transitory shocks	(3) Arrival rate estimated	(4) KMP	(5) Arrival rate quarterly	(6) High persistence	(7) FE het
Quarterly MPC (%)	4.6	5.9	4.7	5.1	4.7		
Annual MPC (%)	14.6	17.1	15.2	15.6	14.5	8.8	12.2
Quarterly MPC of the HtM (%)	28.7	24.6	23.5	32.3	36.4		
Share of HtM (%)	2.5	18.2	2.3	5.3	2.8	5.3	7.5
Annualized discount factor	0.980	0.981	0.980	0.980	0.980	0.988	0.984
Panel A: Decomposition							
Gap with Baseline MPC		1.3	0.1	0.5	0.1	-5.4	-2.0
Effect of MPC Function		-2.3	-0.4	-0.8	0.3	-3.9	-0.2
Effect of Distribution		3.7	0.6	1.4	-0.3	-1.9	-1.8
Interaction		-0.1	-0.0	-0.1	0.0	0.4	-0.0
Panel B: Wealth Statistics							
Mean wealth	4.1	4.1	4.1	4.1	4.1	4.1	4.1
Median wealth	1.3	1.2	1.3	1.1	1.2	0.9	1.6
$a \leq$ \$1000	2.5	17.3	2.9	5.6	2.7	5.2	7.2
$a \leq$ \$5000	11.6	23.6	14.9	17.2	9.2	9.4	12.5
$a \leq$ \$10000	18.5	27.8	21.9	24.6	16.1	15.3	17.4
$a \leq$ \$50000	40.3	43.3	41.1	44.0	41.2	45.5	36.3
$a \leq$ \$100000	51.9	53.1	52.0	54.7	53.7	61.2	48.7
Wealth, top 10% share	46.6	47.4	47.2	53.5	52.6	63.7	48.1

Table E.2: Baseline one-asset model and sensitivity analysis with respect to the statistical process for income dynamics. The columns correspond, respectively, to income processes whose parameters are in lines (1), (6), (3), (4), (2), (7) and (8) of Table A.2.

	(1) Baseline	(2) With Bequests	(3) No Death	(4) Annuities
Quarterly MPC (%)	4.6	4.6	4.4	4.9
Annual MPC (%)	14.6	14.6	14.7	15.8
Quarterly MPC of the HtM (%)	28.7	28.7	30.3	29.1
Share of HtM (%)	2.5	2.5	1.4	2.5
Annualized discount factor	0.980	0.980	0.975	0.961
Panel A: Decomposition				
Gap with Baseline MPC		0.0	-0.2	0.3
Effect of MPC Function		0.0	0.4	0.4
Effect of Distribution		0.0	-0.6	-0.0
Interaction		0.0	-0.0	-0.0
Panel B: Wealth Statistics				
Mean wealth	4.1	4.1	4.1	4.1
Median wealth	1.3	1.3	1.6	1.3
$a \leq$ \$1000	2.5	2.5	1.6	2.5
$a \leq$ \$5000	11.6	11.6	9.4	11.4
$a \leq$ \$10000	18.5	18.5	15.7	18.2
$a \leq$ \$50000	40.3	40.3	36.7	40.1
$a \leq$ \$100000	51.9	51.9	48.3	51.7
Wealth, top 10% share	46.6	46.6	42.0	46.5

Table E.3: Baseline one-asset model and sensitivity analysis with respect to survival rates and to assumptions on how assets of the deceased are distributed among the living. In the baseline, everyone starts with zero wealth. In the model with bequest, assets of the deceased are distributed equally to the newborn. In the model with no death ($\delta = 0$) households have an infinite horizon. The last specification has perfect annuity markets.

	(1) Baseline	(2) Het CRRA	(3) Het CRRA	(4) Het CRRA
Set of CRRA		$e^{-1},\ldots,e^{1}$	$e^{-2},\ldots,e^{2}$	$e^{-3},\ldots,e^{3}$
Quarterly MPC (%)	4.6	6.7	16.6	42.3
Annual MPC (%)	14.6	20.2	39.2	56.3
Quarterly MPC of the HtM (%)	28.7	30.6	46.0	88.6
Share of HtM (%)	2.5	4.0	19.6	39.8
Annualized discount factor	0.980	0.977	0.958	0.876
Panel A: Decomposition				
Gap with Baseline MPC		2.1	12.0	37.7
Effect of MPC Function		0.9	3.5	6.4
Effect of Distribution		1.0	5.6	10.5
Interaction		0.2	2.9	20.7
Panel B: Wealth Statistics				
Mean wealth	4.1	4.1	4.1	4.1
Median wealth	1.3	0.9	0.3	0.1
$a \leq$ \$1000	2.5	3.6	16.6	37.9
$a \leq$ \$5000	11.6	15.3	31.5	48.4
$a \leq$ \$10000	18.5	23.7	40.2	54.3
$a \leq$ \$50000	40.3	47.2	59.7	66.3
$a \leq$ \$100000	51.9	58.1	67.3	71.4
Wealth, top 10% share	46.6	55.7	66.6	73.1

Table E.4: One-asset model with heterogeneity in the curvature parameter γ of the CRRA utility function.

	(1) Baseline	(2) RA=1, IES=2	(3) RA=1, IES=0.25	(4) RA=8, IES=1	(5) RA=0.5, IES=1
Quarterly MPC (%)	4.6	5.3	5.7	3.4	4.6
Annual MPC (%)	14.6	16.4	17.6	11.9	14.6
Quarterly MPC of the HtM (%)	28.7	29.0	29.2	23.8	28.7
Share of HtM (%)	2.5	3.3	3.9	1.2	2.5
Annualized discount factor	0.980	0.986	0.954	0.951	0.980
Panel A: Decomposition					
Gap with Baseline MPC		0.7	1.1	-1.2	0.0
Effect of MPC Function		0.1	0.2	-0.0	0.0
Effect of Distribution		0.6	0.9	-1.6	0.0
Interaction		0.0	0.0	0.4	0.0
Panel B: Wealth Statistics					
Mean wealth	4.1	4.1	4.1	4.1	4.1
Median wealth	1.3	1.1	1.1	2.4	1.3
$a \leq$ \$1000	2.5	3.4	4.0	1.5	2.5
$a \leq$ \$5000	11.6	13.9	15.3	5.7	11.6
$a \leq$ \$10000	18.5	21.4	23.0	9.3	18.5
$a \leq$ \$50000	40.3	43.3	44.8	25.8	40.4
$a \leq$ \$100000	51.9	54.2	55.2	38.7	51.9
Wealth, top 10% share	46.6	48.6	49.5	35.3	46.6

Table E.5: One-asset model with Epstein-Zin preferences. RRA: coefficient of relative risk aversion (γ). IES: intertemporal elasticity of substitution $(1/\theta)$. See equation (4).

	(1) Baseline	(2) $\zeta = 0.9$	(3) $\zeta = 0.8$	(4) $\zeta = 0.7$
Quarterly MPC (%)	3.0	3.6	3.8	2.5
Annual MPC (%)	11.5	13.1	10.4	6.2
Quarterly MPC of the HtM (%)	33.7	21.7	14.4	8.1
Share of HtM (%)	2.0	6.6	17.1	19.6
Annualized discount factor	0.985	0.989	0.993	0.997
Panel A: Decomposition				
Gap with Baseline MPC		0.6	0.8	-0.5
Effect of MPC Function		-1.6	-7.6	-10.2
Effect of Distribution		1.4	2.0	0.9
Interaction		0.8	6.4	8.8
Panel B: Wealth Statistics				
Mean wealth	4.1	4.1	4.1	4.1
Median wealth	1.6	1.5	1.6	1.9
$a \leq$ \$1000	2.5	7.4	16.7	19.1
$a \leq$ \$5000	8.0	15.3	21.9	23.1
$a \leq$ \$10000	13.4	20.5	25.7	26.1
$a \leq$ \$50000	34.9	38.9	40.3	38.2
$a \leq$ \$100000	48.0	49.6	49.3	46.5
Wealth, top 10% share	40.3	39.4	37.7	34.8

Table E.6: One-asset model with preferences featuring present bias in consumption choices. The parameter $\zeta < 1$ measures the strength of the present bias (the baseline model without present bias features $\zeta = 1$).

	(1) Baseline 2-asset	(2) χ=1	(3) χ=0.25	(4) κ=0	(5) κ=3000
Rebalance arrival rate	3.00	1.00	0.25	3.00	3.00
Quarterly MPC (%)	16.1	14.9	16.3	14.4	15.3
Annual MPC (%)	41.2	37.7	42.3	32.7	40.4
Quarterly PHtM MPC(%)	24.3	24.1	23.9	31.0	24.5
Quarterly WHtM MPC (%)	29.8	25.9	37.7	18.7	28.6
Mean MPC at Mean Wealth (%)	7.0	8.0	9.2	8.8	6.7
Prob. HtM status at year t and year t+1	0.9	0.8	0.8	0.9	0.9
Panel A: Calibrated Variables					
Annualized discount factor	0.937	0.937	0.939	0.936	0.937
Annualized liquid return	-2.0	-2.0	-2.0	-2.0	-2.0
Annualized illiquid return	6.2	6.2	6.2	6.2	6.2
Rebalance cost ($)	1456.8	1456.8	1456.8	0.0	3000.0
Panel B: Targeted Statistics					
Mean total wealth	4.1	4.1	4.1	4.1	4.1
Share hand-to-mouth	39.5	35.2	30.7	50.3	36.8
Share poor hand-to-mouth	13.2	13.6	15.4	3.8	16.9
Panel C: Decomposition					
Gap with Baseline MPC		-1.2	0.2	-1.7	-0.8
Effect of MPC function		-0.0	3.0	-1.9	-0.4
Distributional Effect		-2.2	-3.5	2.6	-1.4
Interaction		1.1	0.7	-2.5	1.0
Panel D: Wealth Statistics					
Mean liquid wealth	0.2	0.3	0.4	0.1	0.3
Median total wealth	1.1	1.0	0.9	1.1	1.0
Median liquid wealth	0.1	0.1	0.1	0.0	0.1
$b \leq$ \$1000	33.0	27.1	25.1	42.6	31.0
$b \leq$ \$5000	54.7	51.1	47.7	68.3	52.3
$b \leq$ \$10000	66.1	63.0	59.1	79.7	63.1
$w \leq$ \$1000	10.0	10.4	11.7	2.7	12.8
$w \leq$ \$5000	20.8	21.5	24.2	15.0	26.6
$w \leq$ \$10000	23.5	24.7	27.8	16.2	31.0
$w \leq$ \$50000	45.9	46.7	47.8	47.4	46.8
$w \leq$ \$100000	57.7	58.7	59.4	57.8	58.9
Wealth, Top 10% share	62.0	62.2	62.9	62.0	62.9
Wealth, Top 1% share	19.2	18.7	19.5	19.5	19.4
Gini coefficient, total wealth	0.8	0.8	0.8	0.8	0.8

Table E.7: Two-asset baseline model and sensitivity with respect to the rebalancing frequency χ and the transaction cost κ.