NBER WORKING PAPER SERIES

TEMPTATION AND COMMITMENT:
A MODEL OF HAND-TO-MOUTH BEHAVIOR

Orazio Attanasio
Agnes Kovacs
Patrick Moran

Working Paper 27944
http://www.nber.org/papers/w27944

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2020, Revised October 2023

An earlier version of this paper was circulated as Kovacs and Moran (2018) “Temptation, Commitment, and Hand-to-Mouth Consumers”. We are grateful to Taha Choukhmane, Andreas Fagereng, Andrea Ferrero, Greg Kaplan, Mike Keane, Hamish Low, Costas Meghir, Ben Moll, Peter Neary, Cormac O’Dea, Carlo Pizinelli, Akos Valentinyi, Gianluca Violante, and seminar participants at Oxford, Yale, Manchester, Reserve Bank of New Zealand, Statistics Norway, Bank of England, and Institute for Fiscal Studies for helpful comments. We would also like to thank two referees and the editor for useful comments. We would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work (http://dx.doi.org/10.5281/zenodo.22558). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Temptation and Commitment: A Model of Hand-to-Mouth Behavior
Orazio Attanasio, Agnes Kovacs, and Patrick Moran
NBER Working Paper No. 27944
October 2020, Revised October 2023
JEL No. D11,D14,D91,E21,R21

ABSTRACT

This paper presents a model of consumption behavior that explains the presence of ‘wealthy hand-to-mouth’ consumers using a mechanism that differs from those analyzed previously. We show that a two-asset model with temptation preferences generates a demand for commitment and thus illiquidity, leading to hand-to-mouth behavior even when liquid assets deliver higher returns than illiquid assets. This preference for illiquidity has important implications for consumption behaviour and for fiscal stimulus policies. Our model matches the recent empirical evidence that MPCs remain high even for large income shocks, suggesting a larger response to targeted fiscal stimulus than previously believed.

Orazio Attanasio
Department of Economics
Yale University
87 Trumbull Street
New Haven, CT 06511
and Institute for Fiscal Studies,
FAIR, BREAD and CEPR
and also NBER
orazio.attanasio@yale.edu

Patrick Moran
Federal Reserve Board
Washington, DC 20551
patrick.e.donnellymoran@frb.gov

Agnes Kovacs
30 Aldwych
King's Business School
Bush House
London WC2B 4BG
United Kingdom
agnes.kovacs@kcl.ac.uk

1 Introduction

A large fraction of households in the United States and other developed countries have almost zero liquid assets, despite holding substantial wealth in illiquid assets, primarily housing but also retirement accounts. Data from the Survey of Consumer Finances indicate that roughly 20% of US households own a house, but have liquid assets worth less than two weeks of income. The fact that wealthy consumers and in particular homeowners do not have liquid savings implies that they have limited ability to absorb adverse income shocks and therefore, as documented in many studies, have a relatively large Marginal Propensity to Consume (MPC) out of temporary shocks. These consumers have been labeled as *Wealthy Hand-to-Mouth* (*WH2M*) households.

As stressed in Kaplan and Violante (2022), explaining the behaviour of these households is an important challenge, as their behavior determines a large fraction of aggregate consumption. Modeling *Hand-to-Mouth* behaviour and understanding its origin is important both to identify the main drivers of individual behaviour and to predict the aggregate reaction to a number of policies, including the so-called *stimulus packages* that have been implemented in response to several recent recessions. In this paper, we present a model of individual preferences and consumption behaviour that can account for the presence of wealthy hand-to-mouth consumers and a relatively high MPC for certain sectors of households.

We make two contributions to the existing literature. First, to explain the presence of wealthy hand-to-mouth consumers, we propose a model of preferences that generates present biases because of the presence of self-control problems, which, in turn, generate a demand for commitment and, as a consequence, an incentive to hold savings in illiquid form. In particular, the model of preferences we use and embed in a rich life cycle framework, is that proposed by Gul and Pesendorfer (2001, 2004), who introduce an element of *temptation* in individual preferences. Because of temptation, individuals may want to avoid saving in liquid assets because of the possibility of instantaneous gratification that might be hard to resist. Temptation, therefore, creates a demand for commitment devices that may allow households to mitigate self-control problems. Housing may act as such a savings commitment device, not only because it is illiquid and allows households to lock away their wealth, but also because amortizing mortgages force households to make regular mortgage payments and accumulate wealth in the form of home equity.

Second, we show that in addition to the ability to explain the large share of *WH2M* households with a high MPC, our model generates patterns of consumption behaviour in response to changes in income that are in line with the available evidence. For instance, our model is consistent with recent empirical evidence, which we review below, showing that the consumption response to income shocks declines only slowly with shock size.

This evidence obviously has important implications for the optimal design of targeted fiscal stimulus policies.

Temptation preferences are attractive from a theoretical point of view, as choices emerging from this axiomatization are time consistent and allow for internally-consistent welfare analysis. Moreover, the behaviour implied by these preferences has been found to be consistent with both experimental evidence and observed consumption decisions, as discussed in Kovacs, Low, and Moran (2021). For our exercise, we calibrate the parameters of the model using data on consumption, assets, and housing from the Panel Study of Income Dynamics and Survey of Consumer Finances. We find that our model fits the data well, matching the large share of wealthy hand-to-mouth households.

The view of illiquidity that we highlight helps improve our understanding of observed consumption behavior. For instance, our model can explain recent empirical evidence showing that the MPC declines relatively slowly with shock size. While historically it has been difficult to study how the MPC varies by shock size, since most stimulus payments are small, there is a growing empirical literature that studies this question using new sources of variation and better quality data. The early empirical literature suggested that MPCs were small for large shocks (Hsieh, 2003). However, new empirical evidence has overturned this finding (Kueng, 2018), and shown that MPCs remain sizeable even in response to large income shocks (Kueng, 2018; Aladangady et al., 2018; Fuster, Kaplan, and Zafar, 2020).[1] This new empirical evidence is difficult to rationalize in traditional two-asset models, where the consumption response to large income shocks is negligible. We show that our model can accommodate it.

In the last part of the paper, we also briefly summarize the implications of temptation and commitment for fiscal stimulus targeting. We evaluate the effect of fiscal stimulus payments directed towards different segments of the income distribution. We find that the largest change in aggregate spending is achieved when larger stimulus payments are targeted towards the bottom 20% of the income distribution, rather than smaller payments to the entire population. This result is driven by the fact that these households have a higher than average MPC, which declines only slowly with the size of the stimulus payment. In contrast, stimulus payments have historically been rather small and targeted to a large proportion of the population. In the US, for instance, stimulus checks were given to 80-85% of households in both 2008 and 2020.

Our contribution is not the only one that provides an explanation of the presence of *WH2M* consumers. Indeed, Kaplan and Violante (2022) review different alternative explanations. One of the first hypotheses discussed in the literature was in an influential

[1]Further note that Bunn et al. (2018) and Fagereng, Holm, and Natvik (2021) also find slow-declining MPCs by shock size. These studies, however, can not distinguish between durable and nondurable consumption, hence are less relevant for our case.

paper by Kaplan and Violante (2014) (KV from now on), who considered a model with two assets: one which is liquid and delivers a meager return (such as bank deposits) and one which is illiquid and yields an attractive return. KV show that such a model is able to deliver a substantial fraction of wealthy hand-to-mouth consumers that are willing to withstand some fluctuations in consumption to exploit the returns on the illiquid asset.

We view our approach as complementary to KV; the mechanism that generates *WH2M* is similar. In both settings, many wealthy individuals choose to keep a large share of their wealth in illiquid form. The main difference between the two approaches is that our model generates *WH2M* with an assumption about preferences, while KV make one about the structure of returns. In our model, temptation increases the demand for illiquid assets but not for liquid ones, without requiring an assumption about the relative returns of these assets. In contrast, the key assumption put forward by KV that generates a demand for illiquidity, as noted by the authors, is a positive gap in returns between illiquid and liquid assets. Clearly, the two drivers for the presence of *WH2M* can co-exist as they may both provide incentives to keep a large fraction of individual wealth in illiquid rather than liquid assets.

In reality, there are many types of liquid assets, some of which give high returns (such as shares) and some of which give low returns (such as cash). In calibrating their model, KV bundle all liquid assets, including cash and shares directly held by households. They then document that the average return on the liquid bundle (of which shares constitute a small fraction) is low relative to that on the illiquid bundle (housing and pension funds). In other words, KV's assumption about asset returns reflects this classification but is driven by specific portfolio choices. In our context, these choices on the allocation of savings among different assets are driven by preferences: even if there exists a high-return liquid asset, households in our model, as in reality, would still choose to allocate the majority of their wealth towards illiquid savings because of the commitment motive. Obviously, there could be other reasons not to hold high-return liquid assets, such as aggregate shocks that might be difficult to smooth with assets like shares, which are not considered explicitly in this context. Nevertheless, our model provides an attractive story to explain observed asset holdings.

In our baseline specification, we explain hand-to-mouth behaviour given the availability of a wide variety of assets, including high-return liquid ones, which we calibrate using the risk-adjusted return to publicly-traded equities.[2] We explicitly model housing, given that housing constitutes the vast majority of illiquid wealth for the average household. However, our results do not require any assumption on the structure of returns. Our

[2]Indeed, many studies show that publicly-traded equities deliver higher risk-adjusted returns than housing, even when accounting for imputed rent and other benefits to homeownership. See for instance Flavin and Yamashita (2002), Goetzmann and Spiegel (2002), and Piazzesi, Schneider, and Tuzel (2007).

model can generate the high concentration of household wealth in illiquid form that we observe in reality, regardless of whether the return on other liquid assets is higher or lower than that on illiquid assets. The demand for illiquidity in our model is generated to a large extent by a demand for commitment, in addition to financial returns and other factors, such as the utility benefit to housing, which we model directly.

One could consider commitment as a form of return on illiquid assets, above and beyond the financial return, so that one could re-frame the KV model in terms of preferences.[3] In our model, 'excess returns' on illiquid assets are built in terms of utils from the value of commitment. More generally, *WH2M* behavior arises from a sufficiently large excess return on illiquidity. Such excess return can be of a financial nature (as in the basic KV model), or be connected to utility flows from illiquid assets (housing but also cars or art), or to a commitment value in preferences.

Temptation preferences are not the only ones that give rise to self control problems, present biases, and, possibly, a demand for commitment devices. Another strand of the literature that focuses on self-control problems, which may make it difficult for households to save in liquid assets, uses a model with hyperbolic discounting, as developed by Strotz (1956), Phelps and Pollak (1968), Laibson (1997), Harris and Laibson (2001), and Angeletos et al. (2001). Temptation preferences, however, unlike hyperbolic discounting, do not generate time inconsistent choices, making equilibria easier to define and allowing for straightforward welfare analysis.

Other potential explanations for the presence of *WH2M* include heterogeneity in preferences (Aguiar, Bils, and Boar, 2020) and heterogeneity in financial sophistication (Bhutta, Blair, and Dettling, 2021). We view these approaches as complementary to those based on present-biased behavior. One important benefit of temptation and commitment, however, is that it is also consistent with recent empirical evidence showing that households desire illiquidity even when it does not deliver excess financial returns or other utility flows (Ashraf, Karlan, and Yin, 2006; Cho and Rust, 2017; Beshears et al., 2020; Vihriälä, 2021). We demonstrate that our model is able to generate a similar demand for illiquidity, in line with this empirical evidence, which plays an important role in generating *WH2M* behavior.[4]

The rest of the paper proceeds as follows. First we develop a life-cycle model of hand-

[3]KV considered versions of their model where: (a) transaction costs are expressed in terms of utility, and (b) the excess return on illiquid assets is a preference flow (e.g., utility flow from owner-occupied housing).

[4]Whether the demand for commitment devices is large or small is an empirical question. While large scale commitments (such as housing and retirement accounts) are widespread, it is not obvious that their prevalence is generated primarily by a need for commitment. Furthermore, they may crowd out demand for private small scale commitments, as discussed in Laibson (2015). Furthermore, recent evidence for the existence of demand for commitment, which we discuss in Appendix A.1, has become available. Carrera et al. (2021), for example, provide a nice summary of studies finding demand for commitment.

to-mouth behaviour driven by temptation preferences (Section 2). We demonstrate that temptation generates a desire for illiquidity and a demand for commitment (Section 3). We calibrate the model using US data (Section 4). Finally, we evaluate the implications for consumption behavior and the optimal design of fiscal policy (Section 5).

2 A Model with Temptation Preferences

We develop a model of household behavior with temptation preferences, building upon a traditional life-cycle model of consumption and saving decisions with uninsurable income risk. In this model, households save for two reasons: to maintain consumption following adverse income shocks (the precautionary motive) and following retirement (the life-cycle motive). We extend this model by allowing households to save in either liquid assets or illiquid housing, where housing gives flow utility, serves as collateral, and provides realistic tax advantages. This generates a third incentive for households to accumulate wealth (the housing motive). Housing transactions incur significant costs, thus making housing illiquid.

Households live for T years and work during the initial W years. Households maximize their present discounted lifetime utility, which depends on both nondurable consumption and housing services. Households have access to two investment assets: liquid assets and illiquid housing. All households are born as renters, but they have the possibility to purchase housing once they have accumulated a downpayment.

Temptation Preferences. Households with standard preferences have no demand for commitment, as more choice and more flexibility are always weakly beneficial. To allow for the possibility that households desire commitment, we incorporate temptation preferences by Gul and Pesendorfer (2001, 2004). According to these preferences, it may be difficult to save due to the temptation to spend for short term gratification. This generates a desire for commitment, as households would like to reduce temptation by locking away their wealth in illiquid form. More specifically, households want to maximize the sum of their expected discounted lifetime utility:

$$\max \mathbb{E}_t \sum_{t=0}^{T} \beta^t U(c_t, h_t, \tilde{c}_t, \tilde{h}_t) \tag{1}$$

subject to a budget constraint we define later. Here, c_t and h_t are the chosen levels of nondurable consumption and housing, while $\tilde{c}_t$ and $\tilde{h}_t$ are the most tempting consumption and housing alternatives each period. The key feature of temptation preferences is that utility, $U()$, depends not only on actual consumption and housing decisions, but also on the most tempting consumption and housing alternatives available in the choice set each period. We define the utility function as:

$$U(c_t, h_t, \tilde{c}_t, \tilde{h}_t) = u(c_t, h_t) - \lambda\Big[u(\tilde{c}_t, \tilde{h}_t) - u(c_t, h_t)\Big] \tag{2}$$

where the felicity function, $u()$, is concave and increasing in both c_t and h_t. We define the most tempting feasible alternative as that which maximizes current period felicity:

$$\big[\tilde{c}_t, \tilde{h}_t\big] = \arg\max_{c_t, h_t} u(c_t, h_t), \tag{3}$$

subject to the current period budget constraint, which will be defined later. The term in square brackets in equation (2) represents the utility cost of temptation. The cost of temptation is proportional to the felicity you could enjoy if you only cared about the present $u(\tilde{c}_t, \tilde{h}_t)$, minus the felicity you actually enjoy given your current consumption and housing decisions $u(c_t, h_t)$. The parameter λ captures the degree of temptation. This model nests standard preferences, as temptation disappears when $\lambda = 0$.

Temptation preferences have been applied in a number of macroeconomic models (Amador, Werning, and Angeletos, 2006; Krusell, Kuruscu, and Smith, 2010; Nakajima, 2012; Schlafmann, 2020). That said, temptation is not the only way to generate present-biased households and create demand for commitment devices. The main alternative is to use dynamically inconsistent preferences in the form of hyperbolic discounting (Strotz, 1956; Laibson, 1997; Harris and Laibson, 2001). There are two main advantages of temptation over hyperbolic-discounting. First, temptation preferences are dynamically consistent, which makes welfare analysis straightforward.[5] Second, the importance of temptation preferences can be measured directly from observed consumption decisions using a linearized Euler equation, as shown by Kovacs, Low, and Moran (2021).

Assets. Households who wish to save can invest in two types of assets: a fully liquid financial asset, a_t, or less-liquid housing, h_t. The financial asset, a_t, yields a certain return r in each period. We abstract from the idea of return risk in our model, therefore we calibrate our model using risk-adjusted returns.

Households can adjust housing in any period. We assume that housing exists on a discrete grid with k different sizes: $h^k \in \{h^1, h^2, ..., h^k\}$. The price of each house, $p_t(h^k)$, depends on its size and is determined relative to the index price, $\bar{p}_t$:

$$p_t(h^k) = g(h^k)\bar{p}_t$$

where $0 < g(h^k) \leq 1$, $g'(h^k) > 0$ and $g''(h^k) < 0$. House prices grow at a constant rate,

[5] Fang and Silverman (2009) discuss the difficulty of welfare analysis under hyperbolic discounting. For an example of welfare analysis with temptation, see Kovacs and Moran (2021), who evaluate the welfare effects of financial liberalization when there is a trade-off between flexibility and commitment.

$1 + r^H$, over time, representing a fixed gross return on the housing asset. Therefore the initial index price determines all other house prices for each time period:

$$\bar{p}_t = (1 + r^H)\bar{p}_{t-1}. \qquad \forall t \quad \text{given} \quad \bar{p}_1 \tag{4}$$

Buying or selling a house incurs a fixed cost, which is a fraction F of the house price:

$$F p_t(h^k)$$

We assume that housing markets are segmented, therefore renters are only permitted to live in the smallest available house (h^1).[6] We assume that the cost of renting is proportional to the price of this unit:

$$rent_t = \eta p_t(h^1)$$

where η represents the rental scale. Note that we abstract from the option to save in retirement accounts, as we want to focus our attention on the choice between liquid assets and housing. This seems like a reasonable simplification, as most households hold very little in retirement assets. For instance, Kaplan and Violante (2014) note that the median household keeps only $950 in their retirement account, according to the Survey of Consumer Finances. In Section A.9.2 we consider sensitivity to this assumption.

Mortgages. The most widely used mortgage contract in the US is the fully-amortizing fixed-rate mortgage.[7] Therefore, we assume that mortgages are of this kind with regular required mortgage payments that force households to gradually build wealth in the form of home equity. As a result, housing may act as a commitment device not only because of its illiquidity, but also because of the regular mortgage payments, mp_t, in every period. The mortgage balance for households who buy a house at time t is

$$m_{t+1} = (1 - \psi)p_t(h_t)(1 + r^M) \tag{5}$$

where ψ is the down-payment households choose to pay, which is required to be at least 10% of the home's actual value, so $\psi^{\min} = 0.1$. On the other hand, the law-of-motion for existing mortgages is

$$m_{t+1} = (m_t - mp_t)(1 + r^M) \tag{6}$$

[6]For evidence of segmentation in housing markets, see Greenwald and Guren (2019).

[7]These mortgages accounted for approximately two-thirds of mortgage origination in the US during the 2000's (Amromin et al., 2018). The prevalence of these mortgage features began with the passage of the National Housing Act in 1934, which created the Federal Housing Administration (FHA). By offering to insure mortgages, the FHA was able to insist on fixed-rate mortgages with constant-level fully-amortizing payment plans (Wiedemer and Baker, 2012). The Dodd-Frank Financial Reform Bill of 2010 reaffirmed these standards by introducing the concept of a "qualified" mortgage that requires fixed rate mortgages to have regular amortizing payments.

where mp_t represents the required mortgage payment at time t. We assume that mortgages are fully-amortizing with constant-level payment plans, as is the case for the vast majority of mortgages in the US. Therefore households must make equal mortgage payments, mp_t, every year that they own the house until they pay off the mortgage. We assume that all mortgage debt must be paid off by age W when households retire. Thus, households make fixed repayments each year based on the following formula:

$$mp_t = \frac{(1+r^M)^s}{\sum_{j=1}^{s}(1+r^M)^j} m_t \tag{7}$$

where the required payment depends on $s = W - t + 1$, which is the number of periods until retirement. If there exists a positive mortgage balance, $m_t > 0$, at the time a house is sold, the value of the house is used to repay the mortgage and the remaining home equity goes to the household. In order to reduce the state space, we assume that all households are required to pay off their mortgage by the time of retirement, thus we impose the condition that $m_{W+1} = 0$.

If households receive large negative income shocks such that they cannot make their mortgage payment, they are forced to default on their mortgage. In this situation, households must sell their home and repay the remaining mortgage debt.[8]

Income. Each household i receives idiosyncratic labor income, $y^l_{i,t}$, in every period before retirement, $t \leq W$, which is assumed to evolve according to the following:

$$\ln y^l_{i,t} = g_t + z_{i,t} \tag{8}$$

where g_t is a deterministic age profile approximated by a third-order age-polynomial, and $z_{i,t}$ is an idiosyncratic shock to log income described by an AR(1) Markov process:

$$\begin{aligned} z_{i,t} &= \rho z_{i,t-1} + \varepsilon_{i,t} \\ \varepsilon_{i,t} &\sim N(0, \sigma^2_\varepsilon) \\ \varepsilon_{i,0} &\sim N(0, \sigma^2_0). \end{aligned} \tag{9}$$

Note that we let the initial variance of the income innovations, $\varepsilon_{i,0}$, to be different from the subsequent periods' in order to account for initial heterogeneity in income at age 22 in the data.

Transfers. We assume that the government supports low income households by providing cash transfers to those below a certain income level. The government transfer guarantees that no household receives less than this income floor, y_{min}. As a result, household income

[8] If mortgage debt is larger than the house value plus transaction costs, then the remaining debt is written off and the government provides a minimum consumption floor. This modeling choice ensures that households never experience infinite negative utility.

can be expressed as:

$$y_{i,t} = \max\left\{y_{min}, y^l_{i,t}\right\} \tag{10}$$

Taxes and Pensions. We incorporate a number of realistic features into our model, which are important if the model is going to have a chance to fit observed life-cycle profiles of consumption and wealth accumulation. More specifically, we include progressive income taxation, large and realistic tax benefits to homeownership, and social-security based retirement. We build progressive income taxation into the model following Keane and Wasi (2016), who assume a nonlinear tax function:

$$\tau(y_{i,t}, a_{i,t}) = e^{\tau_1 + \tau_2 \log(y_{i,t} + ra_{i,t} - \tau_d)} \tag{11}$$

where the parameters τ_1 and τ_2 determine the progressivity of the aggregate tax schedule. These parameters are estimated based on income and tax data from the Current Population Survey, therefore $\tau(y_{i,t}, a_{i,t})$ represents the sum of federal, state, and municipal taxes, plus mandatory social security contributions. Taxes are levied on both labor income, $y_{i,t}$, and capital gains, $ra_{i,t}$, although it is important to note that capital gains to owner-occupied housing are not taxed in our model, thus providing a tax benefit to homeownership.

In addition, τ_d represents the deduction which is subtracted from income before the tax is applied. We define τ_d to be the greater of either the standard deduction, τ_d^{standard}, or mortgage interest payments. This allows our tax schedule to incorporate the mortgage interest tax deduction, a second large subsidy to homeownership in the US. This results in an after-tax income for households given by the following equation:

$$\widetilde{y}_{i,t} = y_{i,t} - \tau(y_{i,t}, a_{i,t})$$

Following retirement at age W, households get a progressive social security-style pension determined by the following rule:

$$\widetilde{y}_{i,t} = \max\left\{\text{SS Income Floor, Annual PIA}(y_W)\right\} \tag{12}$$

where Annual PIA(y_W) is the annual social security benefit (the primary insurance amount) received upon retirement, based on average indexed monthly earnings (AIME), which we approximate based on the last working period income, y_W.[9] We calibrate the social security income floor and primary insurance amount based on US legislation from

[9] In reality, to calculate AIME, the worker's wage during the years of employment is first expressed in today's dollars, and then the wages of the highest 35 years are summed up. This sum is then divided by 420, (12*35), in order to get the real average monthly earnings.

2015.[10]

Functional Form. Turning to the choice of functional form for the felicity function, u, we follow Attanasio et al. (2012) and let homeownership affect the felicity function flexibly. This is important as we do not have a strong prior on whether housing utility is additive or multiplicative, and therefore we want a very flexible functional form that includes both options:

$$u(c_t, h_t) = n_t \left(\frac{\left(\frac{c_t}{n_t}\right)^{1-\gamma}}{1-\gamma} \exp\Big[\theta\phi(h_t, n_t)\Big] + \mu\phi(h_t, n_t) - \chi I_{h_t \neq h_{t-1}} \right) \tag{13}$$

where n_t is the exogenously given equivalence scale capturing the evolution of household composition over the life-cycle, γ is the risk aversion parameter, θ and μ are housing preference parameters, and $\phi(h_t, n_t)$ represents the benefit of owning house h_t with family size n_t. Housing affects immediate utility both directly and via the marginal utility of consumption. The direct effect represented by $\mu\phi(h_t, n_t)$ makes the utility function non-homothetic in consumption and housing. We will later calibrate the importance of μ and θ in explaining observed demand for housing.

The utility benefit of housing depends on the size of the house, h, which exists on a discrete grid with k values: $h^k \in \{h^1, h^2, ..., h^k\}$. We assume a segmented housing market by only allowing the smallest house, h^1, to be rented, which also provides lower utility than owning the same unit. In addition, the utility benefit of housing, $\phi(h_t, k_t)$, increases with the size of the house, h_t, and decreases with the size of the family.

$$\phi(h_t, k_t) = \ln\left(\frac{h_t}{n_t}\right) \tag{14}$$

Whenever a household adjusts housing (i.e. when $I_{h_t \neq h_{t-1}}$ equals one in equation (13)), it suffers a utility cost, χ.[11] The utility cost, besides the financial transaction cost, plays an important role in our model, as it increases the illiquidity of housing, thus making housing more useful as a commitment device.

Budget Constraint. Households start each period t with liquid assets a_t, after-tax income $\widetilde{y}_t$, housing h_{t-1}, and mortgage balance m_t. Households then decide on consumption c_t and housing h_t. If they buy a home, they also decide on the level of the mortgage

[10] The PIA is a piecewise linear function with two break points. Currently, the PIA is computed as 90% of AIME up to breakpoint 1, 32% of AIME up to breakpoint 2, and 15% of AIME up to the social security wage base.

[11] Here we think of the non-monetary cost of changing homes, like finding new schools, setting up new utility providers, facing stress etc.

m_{t+1}. The intertemporal budget constraint therefore can be written as follows:

$$a_{t+1} = (1+r)\begin{cases} a_t + \widetilde{y}_t - c_t - \mathbb{I}_t^{own} mp_t - (1-\mathbb{I}_t^{own}) rent_t \\ \qquad \text{if no housing adjustment} \\ \\ a_t + \widetilde{y}_t - c_t - \left[(1+F)p_t(h_t) - \dfrac{m_{t+1}}{(1+r^M)}\right] \\ \qquad + \left[(1-F)p_t(h_{t-1}) - m_t\right] \\ \qquad \text{if housing adjustment} \end{cases} \tag{15}$$

The intertemporal budget constraint implicitly defines the most tempting feasible alternative, $\left[\tilde{c}_t, \tilde{h}_t\right]$, where all the resources available to a consumer at time t are exhausted. Households are always tempted to spend all of their available resources on either consumption or housing, maximising utlity *in that period using all available resources.* However, to be used in period t, these resources need to be in liquid form. As a result, the most tempting alternative crucially depends on how much of the available housing equity can be liquidated *at time t*. As apparent from (15), our assumption is that households can liquidate their housing equity in full. Households might be tempted to sell their home and use that income to buy a bigger house or to consume more. More realistic schemes, where it is only possible to extract a portion of house equity could be considered but are not the focus of this paper.

Recursive Formulation. We define the recursive formulation as follows:

$$V_t(\Omega_t) = \max\left\{V_t^0(\Omega_t), V_t^1(\Omega_t)\right\} \tag{16}$$

where $V_t^0(\Omega_t)$ and $V_t^1(\Omega_t)$ are the value functions conditional on not adjusting and adjusting housing. We define the vector of state variables $\Omega_t = \{a_t, z_t, m_t, h_{t-1}\}$. Those who choose not to adjust in period t solve the following dynamic problem:

$$V_t^0(\Omega_t) = \max_{\{c_t, a_{t+1}\}} U(c_t, h_t, \tilde{c}_t, \tilde{h}_t) + \beta \mathbb{E}_t V_{t+1}(\Omega_{t+1}), \tag{17}$$

subject to:

$$\begin{aligned} a_{t+1} &= (1+r)\left[a_t + \widetilde{y}_t - c_t - \mathbb{I}_t^{own} mp_t - (1-\mathbb{I}_t^{own}) rent_t\right] \\ \widetilde{y}_t &= \begin{cases} \max\left\{y_{min}, exp(g_t + z_t)\right\} - \tau, & \text{if } t \leq W \\ \text{SS Benefit}(y_W), & \text{if } t > W \end{cases} \\ z_t &= \rho z_{t-1} + \varepsilon_t \quad \text{and} \quad c_t > 0 \end{aligned} \tag{18}$$

Those who choose to adjust housing in period t solve the following dynamic problem:

$$V_t^1(\Omega_t) = \max_{\{c_t, h_t, m_{t+1}, a_{t+1}\}} U(c_t, h_t, \tilde{c}_t, \tilde{h}_t) + \beta \mathbb{E}_t V_{t+1}(\Omega_{t+1}), \tag{19}$$

subject to:

$$\begin{aligned}
a_{t+1} &= (1+r)\Big[a_t + \widetilde{y}_t - c_t - (1+F)p_t(h_t) + \frac{m_{t+1}}{(1+r^M)} + (1-F)p_t(h_{t-1}) - m_t\Big] \\
m_{t+1} &\leq (1-\psi^{\min})p_t(h_t)(1+r^M) \\
\tilde{y}_t &= \begin{cases} \max\Big\{y_{min}, exp(g_t + z_t)\Big\} - \tau, & \text{if } t \leq W \\ \text{SS Benefit}(y_W), & \text{if } t > W \end{cases} \\
z_t &= \rho z_{t-1} + \varepsilon_t \quad \text{and} \quad c_t > 0
\end{aligned} \tag{20}$$

3 Temptation Generates Demand for Illiquidity

In this section, we demonstrate one of the key implications of our model of temptation and commitment. Namely, our model generates a demand for illiquidity that is absent in a traditional two-asset model, but which is consistent with a growing experimental and quasi-experimental literature. The demand for illiquidity plays an important role in helping our model generate a large share of wealthy hand-to-mouth households.

3.1 Insights from a Simplified Model

To highlight the most important implications of temptation and commitment, we focus on a simplified version of our model, where housing is strictly dominated in a traditional model. We simplify our model by assuming that (i) housing and liquid assets give identical returns, (ii) housing does not enter the utility function, (iii) labor income is deterministic, and (iv) there is only one house size. Table 1 presents the parameter restrictions in the simplified model. We later relax these assumptions in Section 4.

In this simplified model, households with standard preferences ($\lambda = 0$) have no demand for housing. Home-ownership comes with sizeable transaction costs, yet delivers no benefits in terms of either utility or returns. This is demonstrated in Figure 1, which presents the life-cycle profiles of assets and consumption in the model without temptation. The left panel presents asset accumulation, which reaches a peak at age 65 when the household retires. The household saves only in liquid assets (hence net wealth and liquid wealth coincide) and never purchases a home. The right panel presents income and consumption over the life-cycle. We see that income rises in a hump shape, before dropping at the time of retirement. Despite this hump-shaped income process, the household is able to perfectly smooth consumption between the early 30s and the end of life.

Table 1: Parameters in the Simplified Model

Parameter		Value
k	Housing options	1
θ	Housing preference (MU of consumption)	0
μ	Housing preference (non-homotheticity)	0
z	Idiosyncratic shock to log income	0
r	Net return on liquid asset	0.021
r^H	Net return on housing	0.021

Note: This table presents the parameter assumptions that we use to simplify our model in Section 3, relative to the full model that we calibrate in Section 4.

In contrast, households with temptation preferences demand housing, despite the fact that housing delivers no direct financial or utility benefits. In the left panel of Figure 2, we see that households with temptation preferences begin to accumulate liquid assets relatively late in life. This is because liquid wealth accumulation is difficult in the presence of temptation, owing to the disutility of deviating from your most tempting consumption alternative each period. Further, households purchase housing despite the presence of sizeable transaction costs. This is driven by the fact that housing provides commitment, which helps households accumulate wealth for two reasons. First, households are able to reduce temptation by locking their wealth in illiquid housing; second, households are able to bind themselves to save in the future by committing to regular mortgage payments that build up wealth in the form of home equity.[12]

Households are able to decrease the utility cost of temptation by locking away their wealth in housing. In other words, temptation and commitment generate a preference for illiquidity. As a result, households spend a significant fraction of their lives holding no liquid wealth despite owning housing.

The right panel of Figure 2 shows the implications of household portfolio decisions for consumption. Since tempted households do not accumulate liquid wealth early in life, consumption closely tracks income until shortly before home purchase.[13] After buying a

[12]This has been highlighted by Shiller (2014), who says: "One nice thing about investing in a house is that you're committed to a mortgage payment. So if you don't take out a home equity line of credit or do something like that, you will accumulate wealth." Recent empirical evidence on the effect of mortgage payments on wealth accumulation is consistent with our model predictions (Bernstein and Koudijs, 2020).

[13]Temptation causes households to accumulate their downpayment quickly prior to house purchase. Consistent with this finding, Haurin, Hendershott, and Wachter (1997) find that households save the majority of their downpayment in the year prior to homeownership. Similarly, Charles and Hurst (2002) find that the vast majority of mortgage applicants have liquid wealth less than 10% of their predicted house value at the time of mortgage application.

Figure 1: Life-Cycle Asset and Consumption Profiles without Temptation

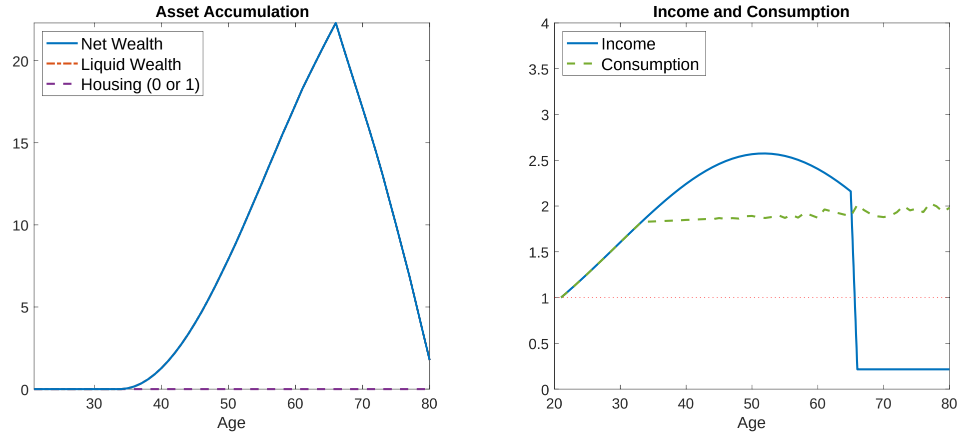

Note: This figure shows the life-cycle profiles of assets, income, and consumption in a model without temptation, where housing delivers equal return to liquid assets and provides no direct utility benefit.

Figure 2: Life-Cycle Asset and Consumption Profiles with Temptation

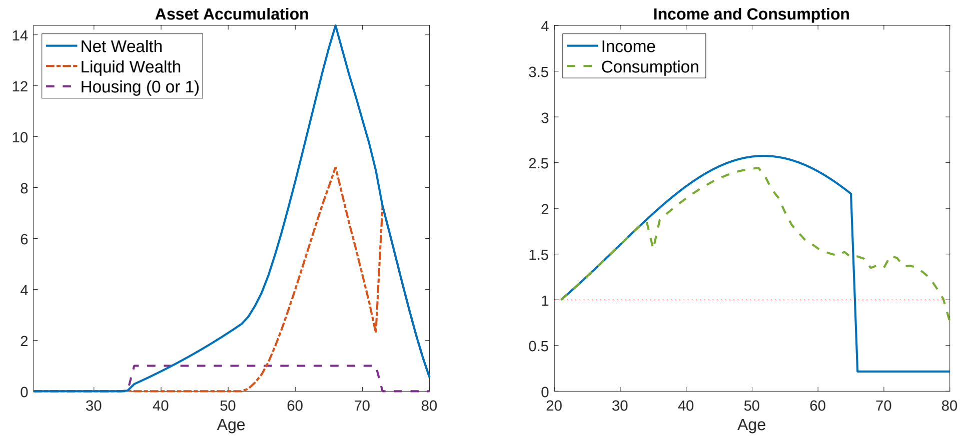

Note: This figure shows the life-cycle profiles of assets, income, and consumption in a model with temptation, where housing delivers equal return to liquid assets and provides no direct utility benefit.

house, consumption continues to follow income closely: the difference between the two is equal to the mandatory mortgage repayment each period. After age 55, when households start accumulating liquid wealth, consumption drops steadily. This is the consequence of temptation: households do not accumulate much wealth for retirement early in life, and therefore consumption declines when nearing retirement.

4 Calibration and Model Fit

We divide our model parameters into two sets. The first set ($\mathbb{P}$) contains parameters that we assign based on external evidence. The second set (Γ) contains parameters that we internally calibrate using the Method of Simulated Moments (MSM), which minimizes the distance between key moments in the model and the data.

4.1 External Parameter Values

In this section we discuss the parameters that are assigned based on external evidence. The complete list of parameter values can be found in Table A.1.

Asset Returns. As mentioned in the introduction, our model does not require specific assumptions about the rates of return on different types of assets in order to generate a desire for illiquidity. In our baseline specification, we allow households to invest directly in a high return liquid asset. This differs from previous models of *WH2M* behavior, which require a positive gap in returns between illiquid and liquid assets, and therefore assume that households cannot invest directly in the high-return liquid asset, but instead can only access it as part of a larger bundle of liquid assets, most of which give low returns.[14] One way to interpret differences in assumptions is that in previous models, the demand for illiquid assets is generated directly by a difference in returns, while in our case it is generated by a desire fore commitment.

We calibrate the model using average real risk-adjusted returns in the US between 1950 and 2016. We use data from the S&P 500 index and the Case-Shiller house price index, where the latter is augmented with housing service flows, maintenance costs, and home insurance. Housing service flows are measured using imputed rents, following Piazzesi, Schneider, and Tuzel (2007). Given that equities and housing have different levels of risk, we compute risk-adjusted returns using the same method as Kaplan and Violante (2014). The full details can be found in Appendix A.3.

Table 2: Real Asset Returns

	Mean	St.Dev.	Risk-adj. Mean
Equities	8.24	16.82	5.40
Housing	2.34	5.06	2.10

Note: Stock: S&P 500. Housing: Case-Shiller index augmented by average housing service flow, maintenance cost, and home insurance.

Table 2 shows the average real returns of equities and housing in the US. While equities deliver higher mean returns than housing, they also exhibit higher variance. Nonetheless, when looking at risk-adjusted returns, we see that equities deliver a higher risk-adjusted return than housing. The average real risk-adjusted return for equities is 5.40% and for

[14] That said, like other contributions in the literature, our model does not try to explain all aspects of portfolio choice and ignores many features of asset returns. For instance, there could be other reasons not to hold high-return liquid assets, such as the correlation between asset returns and (aggregate) earnings, which might make liquid wealth an inadequate insurance mechanism. Nevertheless, our model provides an attractive story to explain households' desire for illiquidity, even when high return liquid assets exist.

housing it is 2.10%. These results are consistent with a wide body of literature including Flavin and Yamashita (2002), Goetzmann and Spiegel (2002), and Piazzesi, Schneider, and Tuzel (2007) who show that equities deliver higher risk-adjusted returns than housing, even when accounting for imputed rents and other benefits to ownership.

Based on the above findings, we set ourselves the challenge of rationalizing the large share of wealthy hand-to-mouth households, despite the presence of a high-return liquid asset. We set a real risk-adjusted return of $r^H = 2.1\%$ for housing and $r = 5.4\%$ for liquid assets in our baseline model. It is important to note, however, that we do not require this assumption. We assess sensitivity to an alternative calibration ($r^H = r = 0.021$) and find that it has little effect on our main results in Appendix A.9.[15]

House Sizes and House Prices. Determining the set of available house sizes in the data is difficult because we only observe house prices for those that have chosen to purchase a house. Nonetheless, we calculate the distribution of house prices for households between age 20 and 30. We use this distribution of house prices to define the different house sizes, and these sizes are then kept constant over time. However, the price of each house will evolve over time following equation (4). In our model, we set the maximum house price (size) at 8 times average income at age 22, corresponding to the 90th percentile of observed house prices for the age group 20-30 in the data, and we set the minimum price to be equal to average income. We allocate the remaining points on the house size grid to a logarithmic scale, following Nakajima and Telyukova (2020). We assume that there are $k = 7$ different house sizes available. We impose the same house size structure on the models with and without temptation. While it is difficult to obtain a realistic estimate of adjustment costs, following Attanasio et al. (2012), we impose a 5% fixed cost of moving, F, representing the cost of real estate agents, lawyers, surveyors, and moving companies. This is consistent with empirical evidence showing that transaction costs for housing are usually at least 5% of the asset value (OECD, 2011). Finally, we assume that capital gains to owner-occupied housing are not taxable.

Mortgages. We calculate the average mortgage rate over the period between 1950 and 2016 based on the 30-year fixed rate mortgage. The average mortgage rate is 4.1%, and therefore we calibrate the net mortgage rate, $r^M = 0.041$, which is two percentage points higher than the risk-adjusted return on housing. We assume that each household can borrow up to 90% of the value of its home, hence the minimum down-payment requirement, ψ, is set to be 10%. Following the US tax code, we assume that mortgage interest payments are fully deductible from taxable income.

[15]There is some controversy about how to best compute imputed rents. In our baseline, we adopt the balance sheet approach following Piazzesi, Schneider, and Tuzel (2007). Jordà et al. (2019) develop an alternative method to compute imputed rents. Based on their imputation, housing and equities in the US have delivered roughly equivalent returns since the 1950s. For further discussion see Appendix A.3.1. We take a flexible stance on imputed rents by including housing in the utility function.

Income and Taxes. We calibrate income over the life-cycle in two steps. First, we use the estimated parameters for the stochastic component of income from Choukhmane (2019). We then estimate a third-order age polynomial on income in order to approximate the deterministic part of labor income. We approximate the annual income floor, y_{min}, by using the federal minimum wage in 2015.[16] For the parameters of the non-linear tax function we use the estimation results by Keane and Wasi (2016) and convert them to 2015 units. We parametrize the progressive social security-style pension based on US legislation from 2015. All income parameters are listed in Table A.1 in the Appendix.

Demographics. In the model, we account for changes in household composition over the life-cycle by assuming an exogenous and deterministic life-cycle profile for household composition. This is performed using the equivalence scale n_t which enters into the utility function. To calculate the equivalence scale n_t, we follow the OECD methodology using PSID data. This methodology assigns weight 1 to the first adult in the household, weight 0.7 to the second adult and each subsequent person aged 14 and over, and weight 0.5 to each child aged under 14. We then average by age in order to construct n_t.

Initial wealth. We assume zero initial housing wealth. We set the initial liquid wealth distribution to match the distribution for 22-25-year-old households in the SCF.

Prices. All variables in the model are expressed in 2015 prices. Where necessary, exogenous parameters from the existing literature are adjusted to represent 2015 prices.

4.2 Calibrated Parameters

The remaining model parameters are internally calibrated using the Method of Simulated Moments to match aggregate statistics of consumption and wealth accumulation. The remaining model parameters are $\Gamma = \{\lambda, \beta, \gamma, \mu, \theta\, \chi\}$, which represent temptation, time preferences, risk aversion, the housing preference (taste) parameters respectively, and the utility cost of changing home. This second stage takes the first stage calibrated parameters fixed, $\hat{\mathbb{P}}$, while choosing Γ to minimize some measure of the distance, f, between the empirical moments, m^e, and the simulated moments, $m^s(\hat{\mathbb{P}}, \Gamma)$:

$$f(\hat{\mathbb{P}}, \Gamma) \equiv [m^s(\hat{\mathbb{P}}, \Gamma) - m^e] \tag{21}$$

We choose to target the mean life-cycle profiles of four variables: consumption, liquid assets, net housing wealth, the share of homeowners with zero liquid asset and the share of renters with zero liquid assets. To focus on the working life, we target the mean of each variable between ages 25 and 65, giving 205 moment conditions. We also target the average homeownership rate and the average share of homeowners who move in a given

[16] Assuming 50 weeks of work at an hourly minimum wage of \$7.25 in 2015 results in an annual labor income of \$14,500, which we consider as the income floor.

year, giving two additional moments. Altogether we target $N_m = 207$ moment conditions to calibrate the five parameters in Γ.

In order to capture the fact that these targeted moments vary substantially in both scale and volatility, we use a weighting matrix, W, to create our scalar-valued final distance function, f^{W}, equal to the weighted sum of squared deviations of simulated moments from their corresponding empirical counterparts:

$$f^{\mathrm{W}}(\hat{\mathbb{P}}, \Gamma) \equiv f(\hat{\mathbb{P}}, \Gamma) \cdot W^{-1} \cdot f(\hat{\mathbb{P}}, \Gamma)' \tag{22}$$

where W is a diagonal N_m x N_m matrix that includes the variance of the targeted moments along the main diagonal. In effect, this means that our MSM approach places more weight on moments that are more precisely estimated in the data.[17]

Pinning down the structural parameters requires that each structural parameter in Γ has an independent effect on at least one targeted moment in $m^s(\hat{\mathbb{P}}, \Gamma)$. More formally, our model is identified if the mapping from structural parameters Γ to targeted moments $m^s(\hat{\mathbb{P}}, \Gamma)$ is full rank near the true Γ. In Section 5 we discuss the way in which structural parameters impact targeted moments.

Our MSM results are based on simulations for 1,000 households for two scenarios each. In the first scenario, we allow λ to be greater than or equal to zero, and we call this the model with temptation based on the finding that λ is significantly different than zero. In the second scenario, we set parameter λ to be zero and we call this the model without temptation. As a result, we can compare the ability of these two models to match the empirical patterns of household consumption and portfolio allocation together with their calibrated parameters.

We target the mean life-cycle profiles of consumption, liquid assets, net housing wealth, the share of homeowners with no liquid wealth, and the share of renters with no liquid wealth between ages 25 and 65. We also target the average homeownership rate and the average share of homeowners who move in a given year. Data comes from the Survey of Consumer Finances (SCF) and the Panel Study of Income Dynamics (PSID) waves covering the period between 1999 and 2015. We use both data sets given that we want detailed information on both asset holdings (in the SCF) and nondurable consumption (in the PSID). Detailed information on the sample is contained in Appendix A.4. In this section, we describe each targeted moment in turn.

Consumption. Using data from the PSID, we compute real nondurable consumption following the classification in Blundell, Pistaferri, and Saporta-Eksten (2016). With this

[17]We choose to use the diagonal weighting matrix rather than the full variance-covariance matrix as many authors have found that the full variance-covariance matrix leads to biased estimates in small samples. See Altonji and Segal (1996) for example.

classification, our consumption measure covers approximately 70 percent of consumption expenditure on nondurable goods and services.

Liquid assets. Using data from the SCF, we measure liquid assets as the sum of cash, bank account deposits, directly held mutual funds, stocks, bonds and T-bills. Note that our model does not account for the existence of credit cards, hence we do not consider negative liquid asset holdings. For this reason, we also exclude credit cards from our definition of liquid assets in the data.

Net housing wealth. Using data from the SCF, we measure net housing wealth as the total value of housing assets, residential and nonresidential real estate net of mortgages and home equity loans.

Wealthy Hand-to-Mouth. We define the wealthy hand-to-mouth as those homeowners who have liquid assets less than two weeks of income. This is an appealing description: these households are wealthy as they own a house, yet live hand-to-mouth with essentially zero liquid wealth that can be used for consumption smoothing (Kaplan, Violante, and Weidner, 2014). Note that this definition is different from the one given by KV as we only consider housing as illiquid asset (while KV includes retirement accounts, life insurance policies, CDs, and saving bonds as well).

Poor Hand-to-Mouth. We define the poor hand-to-mouth as renters who have liquid assets less than two weeks of income. These households are poor as they do not own any illiquid housing, and they also live hand-to-mouth with essentially zero liquid wealth that can be used for consumption smoothing.

Homeownership. Using data from the SCF, we measure homeownership as the mean homeownership rate between ages 25 and 65.

Movers. Using data from the PSID, we define movers as the average share of homeowners that move for non-work reasons each year. We exclude moves for work since that is outside the scope of our model.

4.2.1 Calibrated Parameter Values

Table 3 presents the results for the calibrated preference parameters: temptation (λ), time preference (β), risk aversion (γ), taste for housing (μ, θ), and the disutility of moving (χ). The first column shows the calibrated parameters for the temptation model. The second column shows the calibrated parameters for the model without temptation.

Our calibration strategy results in a value of λ of around 0.15. This estimate is statistically significantly different from zero, which confirms the importance of temptation. We interpret $\frac{\lambda}{1+\lambda}$ as the degree of relative temptation that measures the importance of

Table 3: Internally calibrated parameters

PARAMETER		Temptation Model	No Temptation Model
Temptation	λ	0.149 (0.007)	-
Time preference	β	0.967 (0.004)	0.938 (0.009)
Risk aversion	γ	2.143 (0.049)	2.379 (0.218)
Housing utility (separable)	μ	0.249 (0.010)	0.533 (0.048)
Housing utility (non-separable)	θ	0.002 (0.001)	0.174 (0.025)
Utility cost of housing adjustment	χ	0.899 (0.091)	0.339 (0.089)
Goodness of Fit	$f^{\mathrm{W}}(\hat{\mathbb{P}}, \hat{\Gamma})$	16.803	39.310

Note: In the temptation model we estimate parameter λ in the range of $[0, 1]$. In the model without temptation we set $\lambda = 0$. Goodness of fit is defined by equation (22).

temptation relative to consumption, in consumption utility terms.[18] Hence our results indicate that the utility cost of temptation is roughly one-sixth of the utility benefit of consumption. The value of λ that we obtain from our calibration strategy is between the estimates of Bucciol (2012) who finds a value of 0.05 and Kovacs, Low, and Moran (2021) and Kovacs and Moran (2021) who find λ to be closer to 0.3.[19] That said, in a previous version of this paper, we set λ based on the results from Kovacs, Low, and Moran (2021) and found that it only resulted in higher MPCs and had little impact on the other results.

The model with temptation yields an annual discount factor, β, of 0.96 and a risk aversion parameter, γ, of 2.14, consistent with most of the macroeconomic literature. The additive utility benefit of housing, μ, is roughly 0.25, while the non-separable utility benefit of housing, θ, is 0.002. It is worth noting that the calibrated positive value of θ

[18]Note that equation (2) can be easily rewritten as $U(\cdot) = (1+\lambda)u(c_t, h_t) - \lambda u(\tilde{c}_t, \tilde{h}_t)$, hence $\frac{\lambda}{1+\lambda}$ can be viewed as the relative importance of temptation in the utility function.

[19]The latter two papers exploit additional data to pin down temptation via its implications for the consumption Euler equation, something which is outside the scope of the present paper.

implies that consumption and housing are complements (however this complementarity is weak). This is in line with the results from Attanasio et al. (2012) who calibrate a similar utility function for housing, albeit with two types of homes, and find that consumption and housing are complements.[20] We find that the utility cost of moving, χ, is roughly 0.89 in the temptation model. For an easier interpretation, we calculate the consumption equivalence of parameter χ by expressing it as the amount of additional consumption that households require to be indifferent between moving and not moving homes. An average homeowner aged 25, for example, has to face a utility cost that is equivalent to an additional $\$11,700$ of consumption if it decides to buy a bigger house.[21]

The model without temptation yields a lower impatience parameter, β, of 0.94 and a slightly higher risk aversion parameter, γ, of 2.37. Moreover, without temptation the model features significantly higher housing taste parameters, μ and θ, of 0.53 and 0.17. This is because the model without temptation requires strong discounting of future utility (and large utility benefits to housing) to try to explain the large share of the hand-to-mouth, while matching the substantial housing wealth accumulation over the life-cycle. Compared to the model with temptation, the model without temptation delivers a lower utility cost of moving, χ, of 0.33. The difference between χ in these two models reflects the fact that illiquidity is partially desirable in the temptation model, as housing wealth that is more illiquid results in less temptation, whereas illiquidity is strictly undesirable in the standard model.

4.3 Model Fit

Baseline Model with Temptation. We first analyze the performance of the calibrated model with temptation. Figure 3 shows the simulated life-cycle moments from the temptation model (the solid blue line) and the targeted moments from the SCF (the black dots). In general, the temptation model obtains a good fit of the life-cycle profiles of liquid and illiquid asset accumulation and the hand-to-mouth status. Notice that, apart from very young ages, household wealth is concentrated in illiquid housing: by age 40, housing wealth accounts for about 70% of the average US household's wealth in both the model and the data.

We find that the model with temptation obtains a good fit of the share of wealthy and poor hand-to-mouth households over the life-cycle. In the temptation model, WH2M behaviour is driven not only by the tax and utility benefits of housing, but also the

[20]Since μ is positive, an increase in housing has a direct positive effect on utility. In addition, $\theta > 0$ implies Edgeworth complementarity of consumption and homeownership, as the cross derivative of utility with respect to consumption and housing is positive: $\frac{\partial^2 u(c_t,h_t)/\partial c_t}{\partial h_t} = \theta\phi'(h_t,n_t)c_t^{-\gamma}\exp(\theta\phi(h_t,n_t)) > 0$. Thus in the model with temptation, housing and consumption are complements, whereas in the model without temptation, housing and consumption are (weak) complements.

[21]Naturally, the consumption equivalence is different depending on the age of households, their current housing status, and their next period housing status.

Figure 3: Fit of the Temptation Model

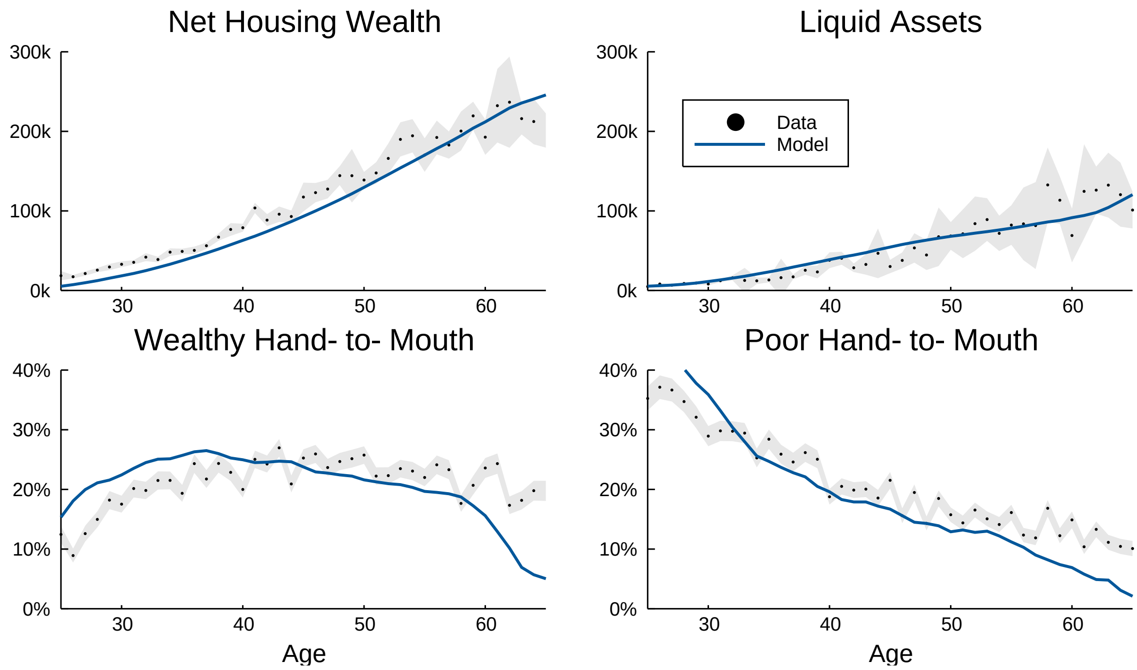

Note: This figure shows the life-cycle moments from the temptation model (solid blue line) and the SCF (black dotted line). The moments from the SCF are shown with bands of 1.96 standard deviations around the mean. Figure A.2a in the Appendix shows the life-cycle paths of nondurable consumption in the simulated model and in the PSID. Table A.5 in the Appendix shows the fit of the aggregate moments that we target: the average homeownership rate and the average share of homeowners moving each period.

fact that households are willing to withstand fluctuations in consumption to obtain the commitment benefit of illiquidity. The age profile for the wealthy hand-to-mouth is hump-shaped with a peak around age 40, when roughly 25% of households own a home but hold essentially zero liquid wealth. The age profile for the poor hand-to-mouth is declining over the life-cycle. It is mainly driven by two factors: increasing homeownership and increasing liquid wealth accumulation over the life-cycle. The one place where the model fails to fit the data is during the years immediately prior to retirement. This is driven by a simplifying assumption in the modeling of retirement benefits.[22]

As noted by Parker (2017), hand-to-mouth status may be either situational (i.e. a result of poor income shocks or temporary illiquidity, as in Kaplan and Violante (2014)) or reflective of persistent household traits (i.e. preferences or behavioral characteristics). Our model of temptation and commitment implies a relatively high persistence of hand-to-mouth status. In our calibrated model, there is roughly a 60% probability that a household who is wealthy hand-to-mouth in one year remains wealthy hand-to-mouth two years later. This is consistent with empirical evidence from the PSID showing that hand-to-mouth status is highly persistent in the data (Aguiar, Bils, and Boar, 2020). In their review, Kaplan and Violante (2022) report that this level of persistence (or even

[22]We assume that social security benefits depend only on income at the end of the work life. As a result, precautionary motives just prior to retirement are stronger in the model than in the data.

higher) is also observed in the two asset models with illiquid high-returns assets.[23]

Model without Temptation. We next analyze the performance of the model where we turn off temptation ($\lambda = 0$) and re-calibrate the remaining preference parameters. We document that our two-asset model without temptation cannot fit the large share of wealthy and poor hand-to-mouth households when housing does not deliver higher risk-adjusted returns than liquid assets.

Figure 4: Fit of the Model without Temptation

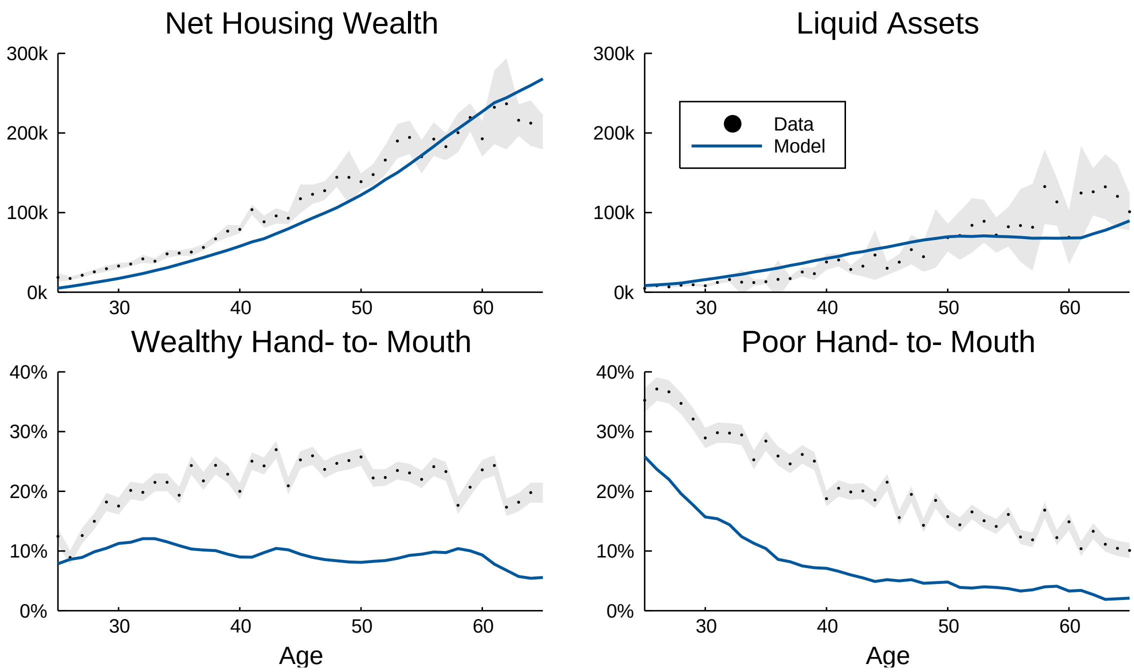

Note: This figure shows the life-cycle moments from our no temptation model (solid blue line) and the SCF (black dotted line), where we impose $\lambda = 0$. The moments from the SCF are shown with bands of 1.96 standard deviations around the mean. Figure A.2b in the Appendix shows the life-cycle paths of nondurable consumption in the simulated model and in the PSID. Table A.5 in the Appendix shows the fit of the aggregate moments that we target: the average homeownership rate and the average share of homeowners moving each period.

Figure 4 shows the simulated life-cycle moments from the model without temptation (the solid blue line) and the targeted moments from the SCF (the black dotted line). We observe that the model without temptation obtains a relatively good match of the life-cycle profile of housing and liquid wealth accumulation. However, in contrast to the temptation model, in the model without temptation there is no combination of the preference parameters that can match the large share of wealthy and poor hand-to-mouth households without failing to match other targeted moments. This is despite the fact that the model is relatively flexible and allows for varying levels of impatience, risk aversion, and housing taste.

High impatience is unable to explain the large share of hand-to-mouth households without compromising the models fit elsewhere. For instance, if β were lowered below

[23] Kaplan and Violante (2022) report yearly persistence rates between 0.7 and 0.9 depending on the parametrization of the model.

our calibrated value in the model without temptation ($\hat{\beta} = 0.93$), we would obtain a better fit of the wealthy and poor hand-to-mouth, but a worse fit of the life-cycle profile of wealth accumulation.

Higher housing taste (i.e. μ or θ) is unable to explain the large share of the wealthy hand-to-mouth. This is because higher housing taste makes homeowners more averse to losing their homes, thus generating a stronger precautionary savings motive for homeowners. This "housing smoothing" motive results in a smaller share of homeowners holding zero liquid assets.

These considerations indicate that the empirical moments we target are able to identify the parameters of interest and, in particular, the strength of the temptation motive in our model and to distinguish it empirically from other parameters that affect the demand for total or illiquid wealth, such as the discount factor and taste for housing.

To provide further evidence on the identifiability of the temptation parameter, in Appendix A.6, we consider the life-cycle profile for the following variables: net housing wealth, liquid wealth, the proportion of wealth hand-to-mouth, and the proportion of poor hand-to-mouth consumers. We plot these profiles as observed in the data and in two different versions of the model with different values for the discount factor and the temptation parameter. We find that changes in these parameters move these moments substantially and in ways that are not correlated, indicating that they are able to identify them separately.

5 Implications for Consumption Behavior

Understanding households' preference for illiquidity is interesting not only from a theoretical perspective, but also for understanding the consumption response to income shocks and fiscal stimulus payments. A growing macro literature argues that households with low liquid assets are important in explaining aggregate consumption behavior.[24] In this section, we study the implications of the alternative view of illiquidity that we have developed.

First, we evaluate the ability of our model to match empirical evidence on heterogeneity in the marginal propensity to consume (MPC). Our most important finding is that the average MPC in our model declines relatively slowly with shock size. In other words, large income shocks still induce a sizeable increase in consumption. This result is consistent with a growing empirical literature, but inconsistent with previous models of hand-to-mouth behavior. To the best of our knowledge, our model is the first that is

[24] See for instance Kaplan and Violante (2014), Carroll et al. (2017), Kaplan, Moll, and Violante (2018), Auclert, Rognlie, and Straub (2018), Luetticke (2018), and Bayer et al. (2019).

able to replicate this empirical finding. In addition, our model fits well-known evidence that MPCs decline slowly with wealth, but quickly with liquid assets.

Finally, we assess the implications for the design of targeted fiscal stimulus payments, which depends crucially on MPC heterogeneity. We use our calibrated model to evaluate the trade-off associated with targeted fiscal stimulus payments. We find that fiscal stimulus is most effective when targeted towards households in the bottom 20% of the income distribution.

5.1 MPC Heterogeneity: Model meets Data

We use our calibrated model to study the consumption response to unanticipated and transitory income shocks. This is performed by simulating N households with randomly drawn initial conditions and income shocks. We then simulate consumption behavior under two alternative scenarios. In the the first scenario, households are given a one-time unanticipated shock of size x, while in the second scenario, household income is unchanged. We then compute the marginal propensity to consume, averaging across all households. The full details of this procedure can be found in Appendix A.7.

We begin by looking at the average MPC in response to an income shock of $1,000. We find that our model generates an average annual MPC of 0.26. This lies well within the standard range of annual MPCs estimated by the empirical literature. Carroll, Slacalek, Tokuoka, and White (2017) give a comprehensive summary of the empirical literature and report that average annual MPC estimates for the US range between 0.2 and 0.6.

5.1.1 Heterogeneity by Shock Size

To study the effect of shock size on consumption behavior, we experiment with changing the magnitude of the income shock. The results are contained in Table 4. We observe that the average MPC declines only slowly with the size of the income shock. For instance, changing the shock size from $1,000 to $10,000 only causes the average MPC to decline from 0.26 to 0.18.

Table 4: MPC Heterogeneity by Shock Size and H2M status

	SHOCK SIZE		
	$\$1,000$	$\$5,000$	$\$10,000$
Average MPC	0.26	0.20	0.18
Average MPC for PH2M	0.68	0.52	0.41
Average MPC for WH2M	0.55	0.41	0.35

Note: Each coefficient represents the average annual MPC in our calibrated model. The procedure used to compute MPCs is described in Appendix A.7.

This result is consistent with a growing empirical literature that finds that the average MPC remains large in response to large income shocks. Kueng (2018) studies the consumption response to large and anticipated payments of the Alaska Permanent Fund and also finds a large MPC. In response to an average payment of \$4,600, Kueng (2018) finds a quarterly MPC of just under 0.3.[25] In addition, Aladangady et al. (2018) study the consumption response to the Earned Income Tax Credit and also find a large MPC in response to this large and anticipated income shock. This quasi-experimental evidence has been corroborated by survey evidence that find similar results. For instance, Fuster, Kaplan, and Zafar (2020) find a large MPC in response to large income shocks.

Large MPCs in response to large income shocks cannot be explained by traditional two-asset models. For instance, while the model of Kaplan and Violante (2014) obtains a good fit of the average MPC out of small income shocks, it predates these recent empirical findings and cannot rationalize sizeable MPCs in response to large income shocks.[26] Section 5.2 explores the reasons for these differences.

5.1.2 Heterogeneity by Asset Holdings

There is a large literature showing that liquidity is an important determinant of MPC heterogeneity, most importantly Kaplan and Violante (2014). We assess the predictive power of our model relative to existing literature. First, we simply look at MPCs by hand-to-mouth status. Table 4 reports average the MPC for wealthy and poor hand-to-mouth households. Compared to the whole population, hand-to-mouth households exhibit a higher average MPC for any shock size considered in the analysis. The average MPC is highest for the poor hand-to-mouth, but only slightly lower for the wealthy hand-to-mouth. This results confirm the importance of liquidity as a main determinant of MPC heterogeneity.

We further investigate the relationship between wealth and MPCs in Appendix A.8. We find that the average MPC in our model declines slowly with net wealth and quickly with cash-on-hand. This is consistent with empirical evidence by Jappelli and Pistaferri (2014) and Fagereng, Holm, and Natvik (2021), as well as the model of hand-to-mouth households by Kaplan and Violante (2014). In addition, we find that the average MPC

[25]While the early empirical literature suggested that MPCs were small for large shocks (Hsieh, 2003), new empirical evidence has overturned this finding. More specifically, Kueng (2018) replicates the analysis in Hsieh (2003) and shows that (i) the small and insignificant consumption response was a result of a nonclassical measurement error in income data attenuating the estimates, and (ii) the estimates in Hsieh (2003) are of consumption elasticities rather than MPCs.

[26]Note that historically it has been difficult to measure the consumption response for large shocks, as most governmental stimulus payments are small. As a result, most of the empirical work on MPCs focuses on the consumption response to small income shocks. See for instance Shapiro and Slemrod (2003), Johnson, Parker, and Souleles (2006), Shapiro and Slemrod (2009), Parker et al. (2013) or Misra and Surico (2014) who all analyze the consumption response of households to the 2001 and the 2008 fiscal stimulus, with average payments of \$600 and \$1,200.

in our model remains elevated even for households with a high degree of liquidity, though of course less so than for households with less cash-on-hand. The results are contained in Table A.7. This implication of our model is consistent with recent empirical evidence which shows that the average MPC stays relatively high even for highly liquid agents (Lewis, Melcangi, and Pilossoph, 2021, Kueng, 2018, Olafsson and Pagel, 2018, Fagereng, Holm, and Natvik, 2021 and McDowall, 2020).

5.2 MPC Decomposition

In our model, households exhibit high MPCs for two reasons. First, there is the effect of illiquidity: households keep the vast majority of their wealth in illiquid form, thus restricting their ability to smooth consumption over transitory income shocks. Second, there is the mechanical effect of temptation: households experience increased temptation ($\tilde{c}$) in response to a positive income shock, therefore they consume more today. In this section, we decompose the relative importance of these two mechanisms.

We decompose the mechanical effect of temptation using the following procedure. For each household $i \in [1, 2, ..., N]$ and each period of working life $t \in [1, 2, ...W]$, we simulate two versions of household i starting at time t, where one receives an unanticipated and transitory income shock of x dollars, while the other does not. This is performed using our baseline parameter estimates for the temptation model with $\lambda = 0$, thus turning off the mechanical effect of temptation. We assume that the state variables at the time of the shock, $\Omega_{i,t}$, are identical for the two versions of the same household and are based on our baseline parameter estimates for the temptation model:[27]

$$\widetilde{\text{MPC}}(x) = \frac{1}{N}\frac{1}{W}\sum_{i=1}^{N}\sum_{t=1}^{W}\frac{c_{i,t}(\text{shock} = x, \lambda = 0, \Omega_{i,t}) - c_{i,t}(\text{shock} = 0, \lambda = 0, \Omega_{i,t})}{x} \tag{23}$$

Figure 5 presents the results from our model decomposition. The striped black bar shows the effect of illiquidity on consumption behavior when we turn off temptation, while the white bar shows the mechanical effect of temptation. We find that both mechanisms are important in generating large MPCs. The effect of illiquidity explains roughly 55-60% of the overall consumption response to an unanticipated and transitory income shock, while the mechanical effect of temptation explains the remainder. Even when households receive an income shock of $10,000, the effect of illiquidity remains substantial.

Why does illiquidity have a large effect even when households receive an income shock

[27] In other words, we assume that households suffer from temptation prior to the realization of the unanticipated and transitory income shock. This ensures that turning off temptation does not alter the distribution of liquid assets, housing, and hand-to-mouth households prior to the shock.

Figure 5: MPC Decomposition

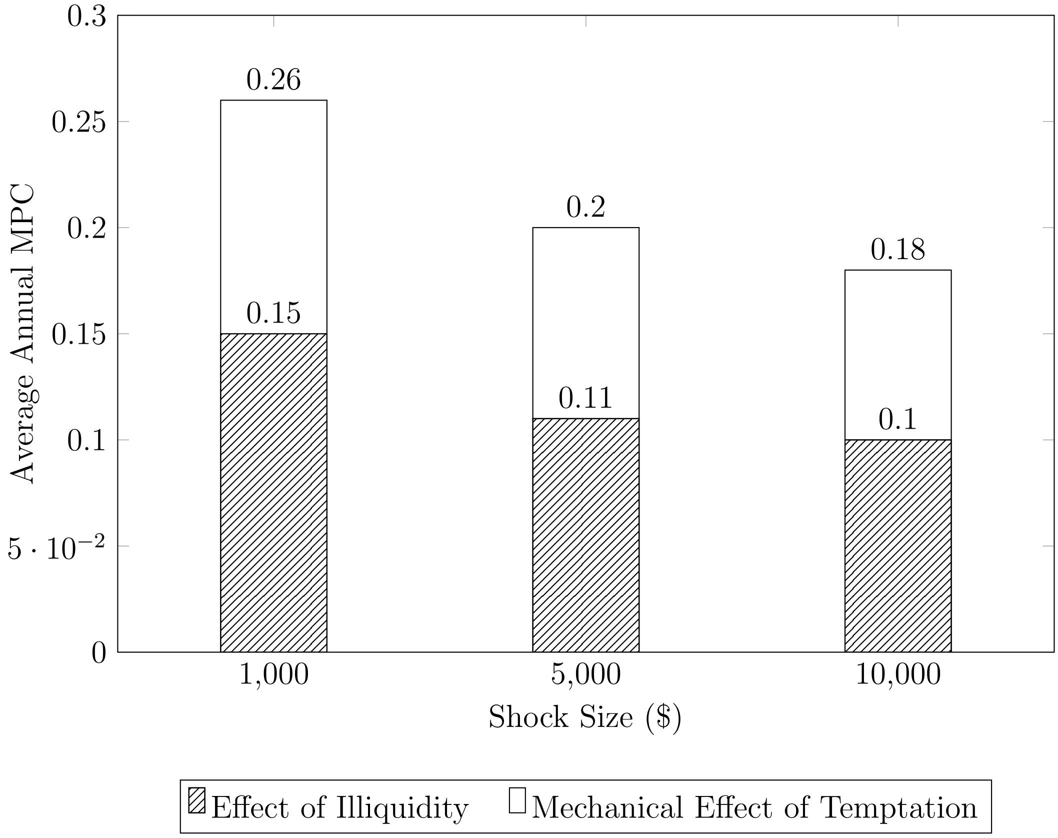

Note: This figure shows the mechanisms driving consumption behavior in response to an unanticipated and transitory income shock. The striped black bar shows the effect of illiquidity, e.g. the MPC absent temptation ($\widetilde{\text{MPC}}(x)$). The white bar shows the mechanical effect of temptation ($\text{MPC}(x) - \widetilde{\text{MPC}}(x)$).

of $10,000? This is because a shock of $10,000 is insufficient to induce many households to pay the housing transaction cost and adjust their housing stock. In our model, households face both financial and non-financial adjustment costs to housing. As a result of these adjustment costs, many households are liquidity constrained in the short-term, even if they have significant housing wealth, therefore they respond strongly to unexpected and transitory income shocks, even when they are relatively large. As a result, the MPC declines only slowly with the size of the shocks.

In contrast, Kaplan and Violante (2014)'s model implies a very small MPC (almost zero) in response to large income shocks. This is because their model has a relatively small, additive adjustment cost of $1,000. While the simple structure of adjustment costs is useful to convey the main ideas behind the model, it introduces an important tension that prevents it from fitting some aspects of the data.[28] Given a certain level of adjustment costs, consumers will have a high MPC when facing small shocks and a low MPC when facing large shocks, as they would be willing to bear the adjustment cost.

[28]Kaplan and Violante (2014) perform extensive sensitive analysis with different adjustment costs, while Kaplan, Moll, and Violante (2018) use different specifications.

Increasing the adjustment costs, on the other hand, might prevent the use of illiquid asset all together. In a model with temptation, larger adjustment costs may actually increase demand for housing, whereas in a traditional model, larger adjustment costs always reduce demand for housing. An important benefit of temptation, therefore, is that it allows the model to accommodate a realistic housing adjustment cost, yet still explain the presence of wealthy hand-to-mouth households.

5.3 Targeted Fiscal Stimulus

As noted in the previous section, hand-to-mouth households have the largest consumption response to transitory income shocks. In addition, the average consumption response declines relatively slowly with shock size. These two findings may have important implications for the design of fiscal stimulus policies, as they suggest that large and targeted fiscal stimulus payments could be very effective in boosting aggregate consumption. In contrast, most governments have historically relied upon small fiscal stimulus payments given to a large proportion of the population. In this section, we use our estimated model to study the efficiency of targeted stimulus.

We study the consumption response to alternative stimulus targeting policies by varying the fraction of households that receive a one time unanticipated stimulus payment from the government, where the government uses an income based targeting approach.[29] We focus on budget equivalent policies, for instance, giving \$500 to all households or \$1,000 to the bottom half of the income distribution. We simulate N households using our model and then compare their baseline consumption to a counterfactual simulation where the same households (with the same income shocks) are given a one time unanticipated stimulus payment at age t. We assume that all households between the ages of 22 and 65 are eligible for stimulus payments, therefore we repeat this exercise for all t within this age range and then aggregate our results. We then report the fraction of aggregate stimulus payments that are consumed within one year after disbursal.

The solid line in Figure 6 shows the aggregate one year marginal propensity to consume (MPC) to budget equivalent fiscal stimulus policies that target different fractions of the income distribution. At one extreme, all households are given a stimulus payment of \$500, while at the other extreme, the bottom 2% of households in the income distribution are given a stimulus payment of \$25,000. At either extreme, under 30% of stimulus payments are consumed within the year of disbursal. We observe that the consumption response gradually rises as the government moves from a policy that distributes stimulus to all households to a policy that targets the bottom 20% of the income distribution. At the

[29] We study the response to income targeting as most governments have comprehensive information on residents' income, but not liquid assets. Of course, Fagereng, Holm, and Natvik (2021) might be able to help Norway perform targeted stimulus based on liquid assets, which may be even more effective.

optimum, when \$2,500 is given to each household in the bottom 20% of the distribution, we observe that 50% of aggregate stimulus is consumed within one year.

Figure 6: Income Targeted Stimulus Payments

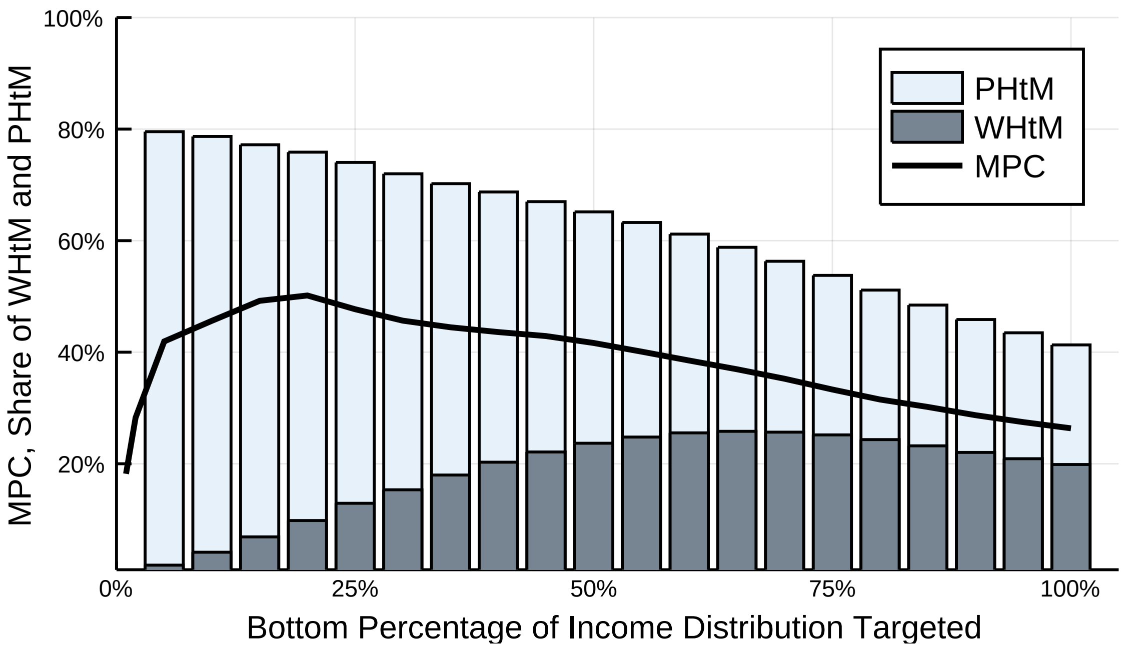

Note: This figure shows results from our baseline temptation model.The solid line represents the aggregate one year consumption response (MPC) to budget equivalent fiscal stimulus policies that target different fractions of the income distribution. The light (dark) bars show the fraction of poor hand-to-mouth (wealthy hand-to-mouth) households over the income distribution.

These results imply that fiscal stimulus can produce a much larger consumption response when it is heavily targeted towards households in the lowest quintile of the income distribution. In contrast, during the Great Recession, most governments that provided stimulus payments decided to give payments to a large fraction of the population, with very little targeting. For instance, under the Economic Stimulus Act of 2008, the US government gave tax rebates to approximately 80-85% of households, with an average stimulus payment of \$600 to \$1,200.

To better understand why it is optimal to target households at the bottom of the income distribution, Figure 6 also shows the fraction of poor hand-to-mouth (light blue bars) and wealthy hand-to-mouth (dark blue bars) households over the income distribution. The share of poor hand-to-mouth households is highest at the bottom of the income distribution and declines slowly with income. The share of wealthy hand-to-mouth households on the other hand is relatively low at the bottom of the income distribution and increases slowly with income. Targeted fiscal stimulus allows the government to reach the households with lower income and liquidity and thus higher MPCs. However, as seen in Table 4, there is a trade-off between providing larger stimulus payments to a subset of households and providing smaller stimulus payments to all households. Overall, we find the optimal policy to be targeting households in the bottom quintile of the income distribution, as this is the point where the benefits of further income-targeting are outweighed by the costs of giving larger stimulus payments. That said, we note that income-based targeting is an imperfect mechanism, and an even larger aggregate effect could be obtained by targeting households based on both income and assets, as this would allow governments to directly target H2M households.

We find a more important role for stimulus targeting than the existing theoretical literature. For instance, while the model of Kaplan and Violante (2014) implies a similar aggregate response when stimulus payments are given to the entire population, they find that the optimal policy is to give payments to the bottom half of the income distribution. This difference is driven by the above trade-off between targeting households with high MPCs and giving larger payments. Their model requires relatively small transaction costs ($1,000 in their preferred calibration) to explain the presence of wealthy hand-to-mouth households, therefore there is a rapid decline in the MPC based on size of stimulus payment, as larger stimulus payments induce more households to pay this cost and put their wealth in the illiquid asset. In contrast, in our model we have a housing transaction cost of 5% of the value of the home, as well as a utility cost that we estimate. Therefore fewer households are willing to adjust housing due to a stimulus payment, unless that payment is very large. As a result, our model is consistent with the recent empirical evidence showing a gradual decline in MPCs based on shock size, thus suggesting a more important role for targeted fiscal stimulus.

6 Conclusion

In this paper, we integrate the idea of temptation preferences, proposed by Gul and Pesendorfer (2001), into a life-cycle model with incomplete markets. We show that this model is able to explain the existence of the wealthy hand-to-mouth by emphasizing the role of illiquid housing as a savings commitment device. We document the model's ability to generate consumption behaviour that is in line with observed choices. Specifically, our model is able to match the large share of the wealthy hand-to-mouth despite the fact that housing delivers higher returns than equities, which is hard to reconcile with traditional life-cycle models. Moreover, the model rationalizes households' overwhelming preference for illiquidity.

Using the Method of Simulated Moments, we internally calibrate the preference parameters of our model with and without temptation. Parameters are pinned down by the life-cycle patterns of households consumption and portfolio compositions. Crucially, we target the observed large fraction of wealthy hand-to-mouth households. The model without temptation is not able to explain the existence of the wealthy hand-to-mouth, even though it has great flexibility that allows it to have different types of housing taste (separable and nonseparable), utility cost of housing adjustment, risk attitude, and impatience. In contrast, the model with temptation generates the observed life-cycle profiles, including the share of households who own housing but no liquid wealth.

Households' illiquidity has important consequences for aggregate consumption behavior, which in turn crucially affects the design of fiscal stimulus policies. We find that targeted fiscal stimulus is more powerful than previously believed: targeting households at the bottom of the income distribution results in the largest aggregate consumption response. This result is driven by the finding that the average MPC declines slowly with shock size, a result that is consistent with a number of recent empirical studies (Kueng, 2018; Aladangady et al., 2018; Fuster, Kaplan, and Zafar, 2020), but which has been difficult to rationalize using traditional models of wealthy hand-to-mouth behavior.

A Appendix

A.1 Empirical Evidence shows Demand for Commitment

There is a growing empirical literature that documents a demand for commitment in both experimental and quasi-experimental settings. This evidence is broadly consistent with the model in this paper, but not traditional explanations of wealthy hand-to-mouth behavior. To give an example, Ashraf, Karlan, and Yin (2006) conduct a field experiment where households are offered both a liquid saving account and a partially illiquid saving account that does not permit withdrawals before a certain date or before a certain savings goal is achieved. The authors find that the partially illiquid account is taken up by a large share of households and leads to increased savings, despite the fact that this account gives no additional returns and thus would be strictly dominated in a traditional model. More recently, Beshears et al. (2020) conduct a lab experiment and study how individuals divide money between liquid and illiquid accounts. The authors ask individuals to allocate money between various accounts, which feature equal rates of return but different degrees of illiquidity. The authors find that illiquid accounts with higher early-withdrawal penalties attract more deposits. This preference for illiquidity would be difficult to explain in traditional models of WH2M behavior.

Both of the above experimental studies demonstrate demand for *pure* commitment, i.e., contracts that tie an individual's hands without providing any other financial or utility benefits. While demand for pure commitment has been documented in numerous lab and field experiments, Laibson (2015) notes that there are surprisingly few examples of pure commitment arising in real-world markets. Of course, housing is not a pure commitment, since there are many other incentives to homeownership above and beyond the potential role of commitment. Various theories have been put forward to explain this phenomenon (see e.g., Laibson (2018) and Ericson and Laibson (2019)). One explanation offered by these authors, which we find particularly compelling, is that many types of commitment may already be embedded in large-scale social institutions, such as housing markets and retirement systems, and that the commitment features in these institutions may crowd out demand for most small-scale commitment technologies offered by private companies. That said, this question remains an active line of inquiry, and not one which we attempt to settle.[30]

Recent evidence has shown various examples of individuals choosing to restrict their choice set in real-world markets, even when there are no monetary or other utility benefits

[30]Another reason why demand for commitment in the lab may be higher than in real markets is risk. Pure commitment contracts reduce the ability to self-insure, thus demand for pure commitments may be low even if the desire for commitment is high. Our model captures this effect.

to doing so. For example, Cho and Rust (2017) study credit card borrowing decisions in Korea and find that many individuals choose to voluntarily reduce their credit card borrowing limits as a form of commitment. Further, Vihriälä (2021) studies a natural experiment in Finland where households are given the option to freely suspend mortgage payments for up to a year. While standard theory suggests that all households should opt for greater flexibility, he finds that a large share of households choose to forgo this option, including households for whom the option is especially salient, as well as households who are already at their credit constraint. Both of these studies are consistent with the present model where households may choose to restrict their choice set to reduce temptation.

Finally, one related question is whether households are willing to pay for commitment. While the early literature struggled to find such evidence, more recent studies have documented a substantial willingness to pay for commitment. For instance, Casaburi and Macchiavello (2019) show in a field experiment that Kenyan farmers are willing to incur sizable costs to receive infrequent payments as a commitment device. Further, Schilbach (2019) documents in a field experiment in India that cycle-rickshaw drivers are willing to forgo large monetary payments in order to set incentives for themselves to remain sober. That said, one benefit of housing compared to *pure* commitments is that housing carries additional benefits above and beyond illiquidity, and the existence of these tied inducements may also help explain its important role in household portfolios.

A.2 Model Parameters

Table A.1: External Parameter Values

Parameter	Definition	Value
	Timing	
T	number of years as adult	59
W	number of years as worker	44
	Asset Returns, Prices	
r	stock return	0.054
r^H	housing return	0.021
r^M	mortgage interest rate	0.040
η	rental scale	0.03
F	fixed cost of moving	0.05
ψ	down-payment requirement	0.10
$p_1^{\max}$	initial house price	$250,000
	Income Process	
ρ	stochastic process: income persistence	0.90
σ_ε^2	stochastic process: std. dev. income shock	0.05
σ_0^2	stochastic process: std. dev. initial income	0.184
g_0	deterministic process: constant	6.391
g_1	deterministic process: *age*	0.256
g_2	deterministic process: age^2	-0.045
g_3	deterministic process: age^3	0.002
y_{min}	income floor	$14,500
	Taxes and Social Security	
τ_1	income tax function: constant	-4.034
τ_2	income tax function: progressivity	1.226
τ_d	income tax function: deduction	$6,116
	social security: income floor	$10,998[1]
	social security: PIA bend points	[$816, $4,917]
	social security: wage base	$118,500

[1] Supplemental Security Income is $8,796 for individuals and $13,200 for couples. From the 2015 Bureau of Labor Statistics Report we know that about half of the population is married (50.2%) and the other half is single, therefore average households get $10,998 as SS income.

A.3 Asset Returns

In this section, we calculate the real risk-adjusted returns of housing and publicly traded equitites. We start with the consumption-based pricing equation, which expresses asset returns in terms of prices and dividends:

$$r_{t+1} = \frac{p_{t+1} + d_{t+1} - p_t}{p_t} \tag{A.1}$$

where r_{t+1} is the net return on the asset between periods t and $t+1$, p_t is the price of the asset in period t, while d_{t+1} is the dividend in period $t+1$. We use this pricing formula to calculate the return on housing. Households who invest in housing in period t enjoy housing service flows between periods t and $t+1$, but also pay the costs related to homeownership over the same period. More explicitly, we can write the return on housing similarly to equation (A.1) as

$$r^{\mathrm{h}}_{t+1} = \frac{p_{t+1} + s_{t+1} - c^{\mathrm{m}}_{t+1} - c^{\mathrm{i}}_{t+1} - p_t}{p_t} \tag{A.2}$$

where p_t is the price of the house in period t, while s_{t+1} and c_{t+1} are the housing service flow and the costs that arise between periods t and $t+1$. Maintenance cost is denoted by c^{m}, and the cost of home insurance by c^{i}. Note that we implicitly assume that depreciation is roughly equal to the maintenance cost.

In what follows we measure aggregate house prices by the Case-Shiller house price index,[31] while we use data from the Bureau of Economic Analysis (BEA) to calculate the average housing service flow. We follow the approach of Kaplan and Violante (2014) to calibrate the size of different ownership-related costs. Housing service flow and related costs are all proportional to the value of the house. Given that these costs are relatively constant over time in terms of the value of the house, we will in the rest of the paper use constant fractions of changing house value in order to calculate these variables. Under these conditions equation (A.2) can be rewritten as

$$r^{\mathrm{h}}_{t+1} = \frac{p^{\mathrm{h}}_{t+1} + (s - c^{\mathrm{m}} - c^{\mathrm{i}} - 1)p^{\mathrm{h}}_t}{p^{\mathrm{h}}_t} \tag{A.3}$$

where s, c^m and c^i are the housing service flows and different costs relative to the value of the house.

We compute imputed rents using the balance sheet approach, following Piazzesi, Schneider, and Tuzel (2007) and Kaplan and Violante (2014), among others. We use housing gross value added at current dollars from the BEA to approximate the housing service flow and use residential fixed assets at current dollars to approximate the housing stock.[32] The average of gross housing value added over residential fixed assets between 1950 and 2016 is around 8%.

Following Kaplan and Violante (2014), we set the maintenance cost at 1% and the insurance cost at 0.35% of the value of housing. In Figure A.1 we plot the calculated

[31] The Case-Shiller house price index is available at `http://www.econ.yale.edu/~shiller/data.htm`.

[32] Gross value added can be found in Table 7.4.5, "Housing Sector Output, Gross Value Added and Net Value Added" in National Income and Product Accounts (NIPA) of the BEA. Residential fixed assets can be found in Table 1.1, "Current-Cost Net Stock of Fixed Assets and Consumer Durable Goods" of the Fixed Asset Tables of the BEA.

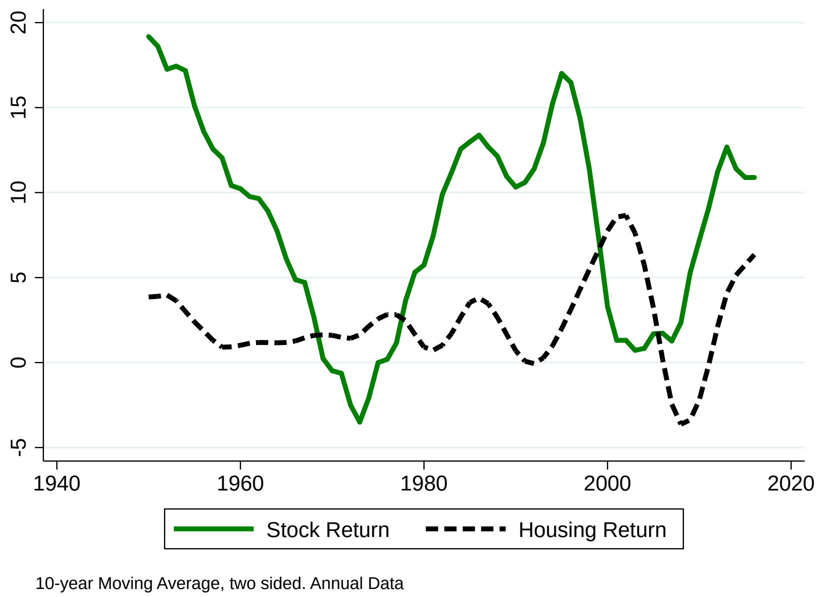

Figure A.1: Real Returns

real return on housing together with the returns on the S&P 500 between 1950 and 2016. The most important thing to notice is that stock returns are in general much higher than the return on housing. There was only a short period of time in the '70s and a couple of years in the early '20s when stocks underperformed housing.

A part of these return differences can obviously be interpreted as reflecting differences in the riskiness of these assets. To allow for this, we calculate the risk-adjusted returns. Following Kaplan and Violante (2014) in order to calculate the risk-adjusted returns on the three assets, we subtract the variance of the return from the expected return of the asset:

$$r^i_{adj} = E(r^i) - var(r^i) \tag{A.4}$$

where superscript i refers to the type of the asset, i.e. 3 Months T-Bill[33], S&P 500, and housing. Since we are using the variance as a measure of riskiness, we cannot generate a similar graph of risk-adjusted returns as in Figure A.1. Instead, we have the average, risk-adjusted real returns over the period between 1950 and 2016, which is 0.69% for the T-bill, 5.40% for the stocks, and 2.10% for the housing asset as seen in Table A.2.

Table A.2 presents our results. We find that stock deliver substantially higher returns than housing, even when accounting for imputed rents and differential volatility in returns. In our baseline results, we find that stock have delivered a risk-adjusted real return of 5.40% per year, while housing has delivered a risk-adjusted real return of 2.10%. It is important to note that most of the real return in housing comes from imputed rents. This is consistent with Shiller (2015), who notes that there has been essentially no real growth

[33]The 3 Month T-Bill times series is downloaded from the database of the Federal Reserve Bank of St. Louis (Fred).

Table A.2: Real Asset Returns

	Mean	St.Dev.	Risk-adj. Mean
Treasury Bill	0.74	2.12	0.69
Stock	8.24	16.82	5.40
Housing (Capital Gains)	0.70	5.06	0.01
Housing (Capital Gains + Imputed Rents)	2.34	5.06	2.10

Note: This table shows our baseline return calculations. Stock returns include both capital gains and dividend income from the S&P. Housing capital gains come from the Case-Shiller index. Imputed rents include the imputed rental income net of maintenance costs and home insurance, as described in Equation A.3.

in house prices during the 20th century. That said, even when accounting for imputed rents, we find that housing has delivered substantially lower returns than equities.

For robustness, we also compute the Sharpe ratios for stocks and housing. The Sharpe ratio measures the expected value of the excess return of the asset relative to the standard deviation of the excess return. The higher the Sharpe ratio, the more attractive the asset, as it delivers better excess returns relative to its riskiness. We find a Sharpe ratio of 0.45 for stock and 0.30 for housing, including both capital gains and imputed rents. This leads further credibility to the result that housing yields a lower risk-adjusted return than stocks.

A.3.1 Sensitivity to Imputed Rents

We compute imputed rents using the above described balance sheet approach, following Piazzesi, Schneider, and Tuzel (2007) among others. Our results are consistent with a wide body of empirical literature including Flavin and Yamashita (2002), Goetzmann and Spiegel (2002), and Piazzesi, Schneider, and Tuzel (2007) who show that housing delivers worse risk-adjusted returns than stock, even when accounting for imputed rents and other benefits to homeownership. That said, Jordà et al. (2019) develop an alternative method to compute imputed rents. In this subsection, we explore sensitivity to alternative imputed rent calculations.

Using the data provided by Jordà et al. (2019), we calculated average risk-adjusted returns to housing and equities in the US.[34] For comparability with our previous return calculations, we focus on the sample period between 1950 and 2016.

Table A.3 presents our results using data from Jordà et al. (2019). We find that stock

[34] We thank the authors for providing their data at http://www.macrohistory.net/data/.

deliver risk-adjusted returns that are very similar to housing. The risk-adjusted return to stock is 5.98% whereas the risk-adjusted return to housing (including both capital gains and imputed rents) is 5.80%. This result is driven by the measure of imputed rents constructed by Jordà et al. (2019). When we look at just the capital gains to housing, we find a risk-adjusted real return of 0.65%, only slightly larger than our results when using the Case-Shiller index.

Table A.3: Real Asset Returns – Sensitivity to Alternative Data

	Mean	St.Dev.	Risk-adj. Mean
Stock	9.03	17.47	5.98
Housing (Capital Gains)	0.78	3.52	0.65
Housing (Capital Gains + Imputed Rents)	5.93	3.60	5.80

Note: This table shows alternative return calculations for the US between 1950 and 2016. Data on stock, housing, and imputed rents come from Jordà et al. (2019).

Based on our reading of the literature, where the vast majority of studies find that stock deliver higher risk-adjusted returns than housing, we decide to adopt that result in our baseline calibration. That said, even if stock and housing delivered roughly similar risk-adjusted returns, this would still pose a challenge for the traditional returns-based explanation of wealthy hand-to-mouth behavior, as put forward by Kaplan and Violante (2014). If liquid and illiquid assets delivered equivalent risk-adjusted returns in their model, then the wealthy hand-to-mouth would disappear. For this reason, in Appendix A.9.1 we evaluate the sensitivity of our model calibration to the assumption that $r = r^H$ and find that it has little effect on our core results.

A.4 Data: Sample and Definitions

We calibrate our model using two different data sets, which we describe in turn. First, we get consumption data fron the Panel Study of Income Dynamics (PSID) waves 1999 to 2015. While the PSID has collected information on income and demographics since 1968, the survey received a large overhaul in 1999 with the addition of detailed questions on household expenditure. We therefore use the modern PSID, which contain information on consumption. We then construct wealth data from the Survey of Consumer Finances (SCF) waves 1998 to 2016. We follow the same sample selection routine as Kaplan and Violante (2014), and rely upon the code which they have very kindly made available.[35]

[35] We thank the authors for providing their data at `https://gregkaplan.me/s/otherreplication_whtm_brookings.zip`.

Sample Selection. We focus on households with a head between 25 and 65 years old with non-missing information on age, education, and state. We do not select our sample based on the working status of the household head or spouse. We include households from both the core sample of the PSID as well as households from the Survey of Economic Opportunity. To reduce the influence of measurement error, we drop observations with extremely high assets, for instance, observations with a total net worth higher than $20 million, following the criteria of Blundell, Pistaferri, and Saporta-Eksten (2016). In addition, both in the PSID and in the SCF, we drop the top 5% of households by income due but not the top 5% of the distribution.

Consumption. Using the PSID, we compute real nondurable consumption following the classification in Blundell, Pistaferri, and Saporta-Eksten (2016). Prior to 1999, the PSID collected data on very few components of consumption, namely food, rent, and child care. The coverage was greatly increased starting in 1999, to include many other components of nondurable consumption and services including transportation, utilities, gasoline, car maintenance, health expenditures, education, and childcare. In total, this allows the PSID to cover approximately 70% of consumption expenditure on nondurable goods and services. While additional categories such as clothing and entertainment were added to the survey in 2005, we exclude these categories to keep the consumption series consistent over time.

Assets and Income. Using the SCF, we compute different asset types (liquid vs housing), hand-to-mouth status and income measures following the definitions and stata code provided by KV. We make two modifications to their asset definitions in order to construct measures in the data that are more comparable to those in our simulated model. First, we exclude credit card balances from our measure of liquid assets as our model does not account for the existence of credit cards. Second, we exclude retirement accounts, life insurance policies, CDs, and saving bonds from our measure of illiquid assets as our model does not account for the existence of such illiquid assets. In fact, the only illiquid asset that we consider in our model is housing. Other than these changes, we define income, poor and wealthy hand-to-mouth households exactly the same way as KV.

A.5 Model Fit

Table A.5 presents the model fit of the average homeownership rate and share of homeowners moving each year. Both the model with temptation and the model without temptation obtain a good fit of both targeted moments.

Table A.4: Internally calibrated parameters

PARAMETER		Temptation Model	Temptation Model Fix λ	No Temptation Model
Temptation	λ	0.149 (0.007)	0.280	-
Time preference	β	0.967 (0.004)	0.975 (0.003)	0.938 (0.009)
Risk aversion	γ	2.143 (0.049)	2.634 (0.061)	2.379 (0.218)
Housing utility (separable)	μ	0.249 (0.010)	0.435 (0.016)	0.533 (0.048)
Housing utility (non-separable)	θ	0.002 (0.001)	0.149 (0.013)	0.174 (0.025)
Utility cost of housing adjustment	χ	0.899 (0.091)	0.880 (0.112)	0.339 (0.089)
Goodness of Fit	$f^{W}(\hat{\mathbb{P}}, \hat{\Gamma})$	16.803	19.680	39.310

Note: In the temptation model we estimate parameter λ in the range of $[0, 1]$. In the model without temptation we set $\lambda = 0$. Goodness of fit is defined by equation (22).

Figure A.2: Nondurable Consumption

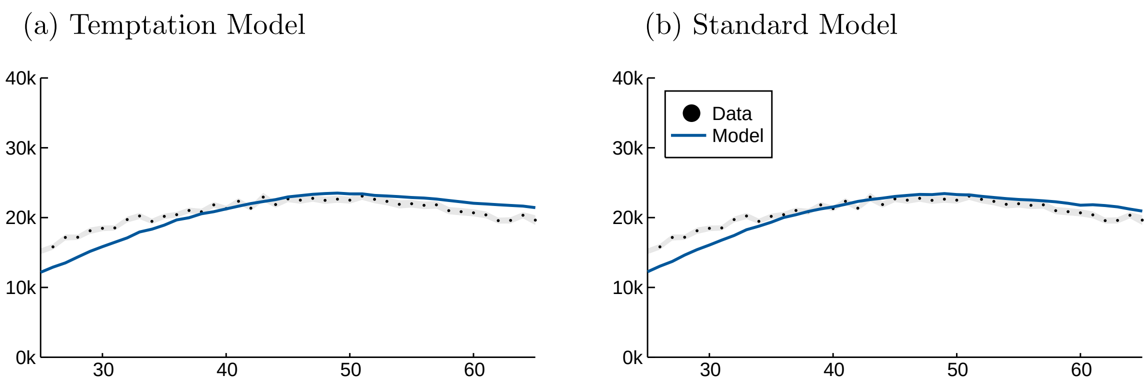

Note: This figure shows the life-cycle consumption profiles from our two different models (solid blue line) and the PSID (black dotted line). In subfigure (a) we allow λ to vary, while in subfigure (b) we impose $\lambda = 0$. The moments from the PSID are shown with bands of 1.96 standard deviations around the mean.

Table A.5: Additional Targeted Moments

	Data	Temptation Model	No Temptation Model
Homeownership rate	0.65	0.69	0.76
Share of movers	0.068	0.041	0.074

Note: This table shows fit of the average homeownership rate and share of movers each period.

A.6 Sensitivity of Targeted Moments to Parameter Changes

In this section, we check how certain moments used in the calibration vary with different values of the temptation parameter λ in order to provide some evidence on the identification of the temptation parameter. To provide a reference, we also perform the same exercise for the discount factor β, which might affect similar variables. If changes in these two parameters induced correlated changes on a set of targeted moments this might indicate that the two parameters could be difficult to disentangle using this set of moments.

Figure A.3: Sensitivity of the Life-Cycle Moments to a Decrease in β

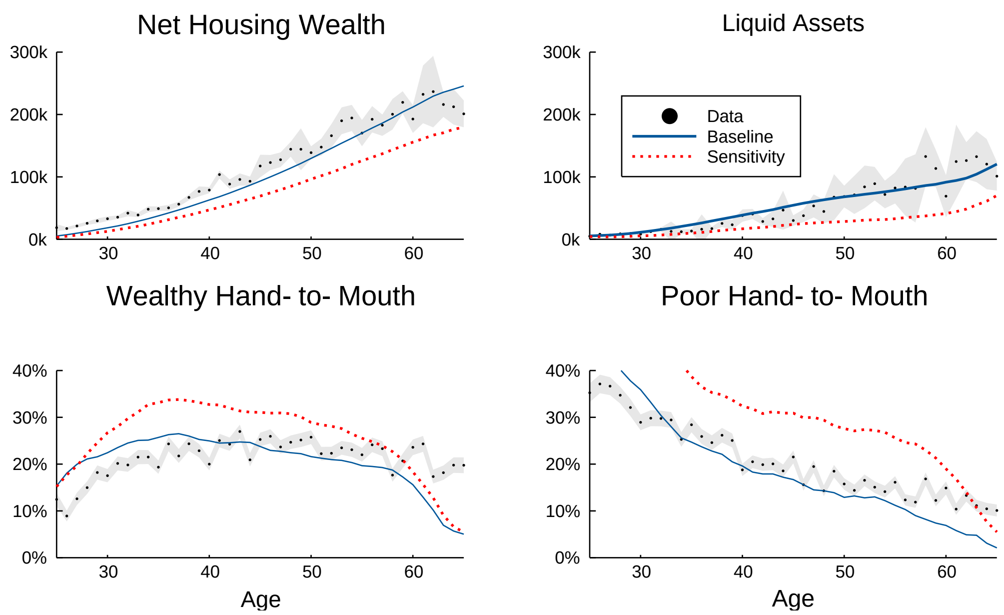

Figure A.3 plots the life cycle profile for net housing wealth, liquid asset, and the proportion of wealthy and poor *H2M* consumers in the data, in the baseline calibrated model and in a model with a lower level of β. In Figure A.4, we perform the same exercise, but lowering λ. In Table A.6, we report how homeownership rates and the share of movers change with changes in λ and β.

Figure A.4: Sensitivity of the Life-Cycle Moments to a Decrease in λ

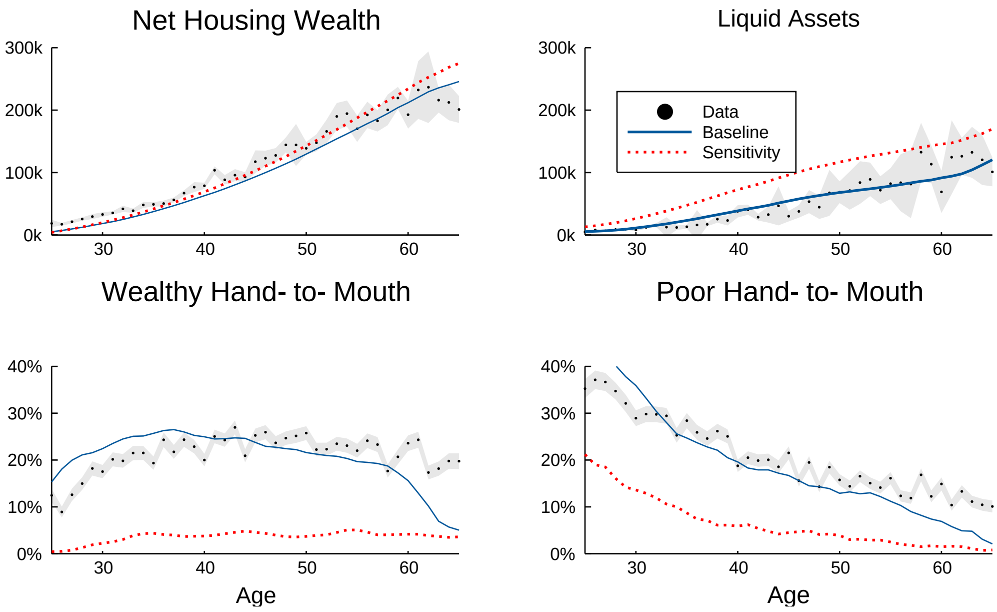

Table A.6: Sensitivity of the Aggregate Moments to a Decrease in β and λ

	Data	Baseline Model	Sensitivity to β	Sensitivity to λ
Homeownership rate	0.65	0.69	0.55	0.72
Share of movers	0.068	0.041	0.047	0.040

We note that changes in the temptation parameter have large impacts on liquid asset holding relative to the impact on housing wealth, while impatience impacts liquid and housing wealth holdings similarly; homeownership rate is very sensitive to impatience but not much to temptation. Finally, the proportion of wealthy *H2M* change dramatically with changes in λ but is not affected much by changes in β. To summarize, these moments change substantially when changing these two parameters and in ways that are not perfectly correlated, indicating that they are able to pin down and identify separately these parameters.

A.7 MPC Caclulations

To calculate the consumption response to an unanticipated and transitory income shock of size x, we first generate N households, each with a different series of randomly drawn income shocks and initial heterogeneity. We simulate the behavior of each household, producing the state variable $\Omega_{i,t}$ and consumption behavior $c_{i,t}(\text{shock} = 0, \Omega_{i,t})$ for each household and each time period. Next, we simulate counterfactual consumption behavior in response to an unanticipated and transitory income shock at time t, conditional on state $\Omega_{i,t}$, producing $c_{i,t}(\text{shock} = x, \Omega_{i,t})$ for all households and all time periods. We set the size of the transitory income shock x to be \$1,000 in our baseline simulation, though we also show results for shock of \$5,000 and \$10,000, respectively. The annual marginal propensity to consume is computed as follows:

$$\mathrm{MPC}_{i,t}(x) = \frac{c_{i,t}(\text{shock} = x, \Omega_{i,t}) - c_{i,t}(\text{shock} = 0, \Omega_{i,t})}{x} \tag{A.5}$$

$$\mathrm{MPC}(x) = \frac{1}{N}\frac{1}{W}\sum_{i=1}^{N}\sum_{t=1}^{W}\mathrm{MPC}_{i,t}(x) \tag{A.6}$$

A.8 MPC Heterogeneity: The Importance of Liquidity

To evaluate the relationship between wealth and consumption behavior, we group simulated households into quartiles based on net wealth and cash-on-hand. We then estimate the average MPC in each of these categories. This is performed in a regression framework, allowing us to control for age, similar to Jappelli and Pistaferri (2014):

$$\mathrm{MPC}_{i,t} = \beta_0 + \sum_{j=2}^{4}\gamma_j \mathrm{CashQ}^j_{i,t} + \sum_{j=2}^{4}\delta_j \mathrm{WealthQ}^j_{i,t} + \sum_{j=23}^{65}\psi_j \mathrm{Age}^j_{i,t} + \epsilon_{i,t} \tag{A.7}$$

If low wealth is important in generating large MPCs, then we would expect to see a rapid decline of MPCs as we move from the lowest to the highest wealth quartiles. In contrast, if low cash-on-hand is more important in generating large MPCs, we would expect to see a rapid decline of MPCs as we move from the lowest to the highest quartile of cash-on-hand.

The results are presented in Table A.7. We find that the average MPC declines very quickly with cash-on-hand. Households in the lowest quartile of cash-on-hand have the highest average MPC. Households in the second quartile have an average MPC that is 0.14 lower. This is even more pronounced in the third and fourth quartiles. In short, low liquidity is an important determinant of high MPCs.

Once we have controlled for cash-on-hand, net wealth has less effect on the average MPC. The average MPC is 0.1 smaller when moving from the lowest to second lowest net wealth quartile. Moreover, we see that even the richest households still have a large average MPC: households in the top wealth quartile have an MPC that is just as large as households in the bottom quartile, once we control for cash-on-hand.

It is also important to notice that even highly liquid households in our model (who has high levels of liquid wealth) exhibit high average MPCs. Households in the top cash-on-hand quartile have an average MPC of around 0.18. It is consistent with the recent

Table A.7: MPC Heterogeneity by Household Type

		MPC Coefficient	MPC Standard Error
CASH-ON-HAND			
Quartile			
	2^{nd}	-0.136^{***}	(0.005)
	3^{rd}	-0.359^{***}	(0.006)
	4^{th}	-0.491^{***}	(0.007)
NET WEALTH			
Quartile			
	2^{nd}	-0.114^{***}	(0.005)
	3^{rd}	-0.071^{***}	(0.006)
	4^{th}	-0.095^{***}	(0.008)
Constant		0.765^{***}	(0.010)

Note: MPCs are based on a $1,000 transitory income shock. We control for age using Equation A.7. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

empirical evidence that points to puzzlingly high MPCs even for highly liquid agents (eg. Lewis, Melcangi, and Pilossoph (2021), Kueng (2018), Olafsson and Pagel (2018), Fagereng, Holm, and Natvik (2021) and McDowall (2020)).

These results are consistent with a wide body of empirical evidence that finds that the average MPC declines only slowly, if at all, with net wealth, while it declines quickly with cash-on-hand. Jappelli and Pistaferri (2014) show that households in the top quintile of cash-on-hand have an avereage MPC that is 0.44 lower than that of households in the bottom quintile, affirming the importance of cash-on-hand for MPCs.[36] Similarly, Fagereng, Holm, and Natvik (2021) find that net wealth is unimportant in explaining MPC heterogeneity, once controlling for liquid wealth.[37]

The reason behind these empirical observations is that households might have substantial wealth, but if it is kept in illiquid form, it cannot be used easily for consumption-smoothing purposes. As a result, wealth is a less important determinant of MPCs than cash-on-hand. This reaffirms the importance of modeling household illiquidity (using a

[36] Jappelli and Pistaferri (2014) use Italian survey data to study the consumption response to unexpected transitory income shocks. They exploit the survey question from the 2010 Italian Survey of Household Income and Wealth, which asks households how much of an unexpected transitory income change they would spend.

[37] Fagereng, Holm, and Natvik (2021) study the consumption response to winning the lottery in Norway. This study is unique in the quality of its data: the authors use administrative tax data from Norway, which contains rich information on household income and asset holdings.

two asset model) in order to study consumption behavior in response to transitory income shocks. In this regard, our model delivers similar results to Kaplan and Violante (2014), who also find that liquid wealth is more important than total wealth in explaining MPC heterogeneity. In contrast, these empirical results are almost impossible to justify using a traditional heterogeneous agent model with only one asset. For instance, Jappelli and Pistaferri (2014) study whether an Ayiagari model with heterogeneous households and a standard calibration is able to replicate the slow decline of MPCs by wealth. They find that this requires implausibly impatient households: β has to be 0.6 or lower.

A.9 Sensitivity

In this section we evaluate the sensitivity of our results presented in Sections 4 and 5. We consider two modification to our baseline model. First, we change the return structure in the model by assuming that housing and equities provide equivalent risk-adjusted returns ($r = r^H$) following Jordà et al. (2019).[38] Second, we modify the asset structure of households by assuming that they receive annuitized disbursements from individual retirement accounts after retirement. Our core findings are robust to these modifications: both the calibrated parameters of the model and the MPC results are similar to our baseline results.

A.9.1 Asset Returns

First we consider a model where the two assets give equivalent risk-adjusted returns ($r = r^H$). It is important to note that this return structure would pose a challenge for the traditional returns-based explanation of wealthy hand-to-mouth behavior. If liquid and illiquid assets delivered equivalent returns in a model of the type proposed by Kaplan and Violante (2014), then the wealthy hand-to-mouth would disappear. In contrast, we find that the model with temptation and commitment doesn't require excess return on housing to explain the existence of wealthy hand-to-mouth households.

Table A.8 shows the model parameters and model fit when we re-calibrate the preference parameters under the assumption of equal returns. We find that the different return structure has little effect on our main results, whoever, the fit of the model (measured by $f^{\mathrm{W}}(\hat{\mathbb{P}}, \hat{\Gamma})$) is slightly worse than the fit of the baseline model. When the asset returns are equal (column 2), the temptation parameter hardly changes relative to the baseline calibration (column 1). Under this scenario, households are slightly more patient and more risk averse (β and γ are slightly higher than in our baseline calibration).

In addition, the alternative return structure has little impact on the main MPC results. Table A.9 shows the MPC by shock size. The average MPC in response to a \$1,000 shock

[38] There exists some disagreement in the literature about how to compute imputed returns to housing. For further discussion, see Appendix A.3.1.

Table A.8: Model Parameters – Alternative Modeling Assumptions

PARAMETER	Baseline	Equal Returns ($r = r^H$)	Retirement Accounts
λ	0.149	0.155	0.158
β	0.967	0.989	0.975
γ	2.143	2.501	2.094
μ	0.249	0.396	0.233
θ	0.002	0.202	0.003
χ	0.899	0.586	0.889
$f^{\mathrm{W}}(\hat{\mathbb{P}}, \hat{\Gamma})$	16.803	25.848	17.012

Note: This table shows our calibrated model parameters under alternative modeling assumptions. In the model with equal returns, we set $r = r^H = 0.021$. In the model with retirement accounts, we add retirement accounts as defined in Section A.9.2. Goodness of fit ($f^{\mathrm{W}}(\hat{\mathbb{P}}, \hat{\Gamma})$) is defined by equation (22).

is 0.22 in the model with equal returns (the second row), which is only slightly lower than in our baseline model (the first row). In addition, we find that the average MPC declines only gradually with respect to shock size. The average consumption response to a \$10,000 shock is 0.15, only slightly lower than our baseline result of 0.18.

Table A.9: MPC Heterogeneity – Alternative Modeling Assumptions

	SHOCK SIZE		
MPC	$\$1,000$	$\$5,000$	$\$10,000$
MPC in Baseline Model	0.26	0.20	0.18
MPC in Model with Equal Returns	0.22	0.17	0.15
MPC in Model with Retirement Accounts	0.28	0.21	0.18

Note: This table shows the average MPC in the temptation model. In the model with equal returns, we set $r = r^H = 0.021$. In the model with retirement accounts, we add retirement accounts as detailed in Section A.9.2.

A.9.2 Individual Retirement Accounts

In order to evaluate sensitivity to alternative modeling assumptions about the resources that are available to save for retirement, we extend our model by assuming that households have an individual retirement accounts (IRAs) from which they then receive annuitized disbursements after retirement. We make the simplifying assumption that households have no choice over the size of their retirement account or the timing at which they withdraw from their retirement account, i.e. retirement contributions are mandatory and must be converted to an annuity at the age of retirement (W). This assumption implies that households suffer zero temptation to consume their retirement account.

We require all households to purchase an annuity in the first year of retirement using the entirety of their retirement account. This ensures equal payments throughout the remainder of their life. During each year of retirement, households receive annuitized disbursements that depend on the replacement rate η_1 and the size of their IRA at retirement:

$$y_{i,t}^{IRA} = \eta_1 * \text{IRA}(y_{i,W}) \tag{A.8}$$

We assume that the size of the retirement account is a linear function of last working period income. This simplifying assumption allows us to include retirement accounts without the introduction of an additional state variable. The size of the retirement account is given by the following formula:

$$\log[\text{IRA}(y_{i,W})] = \eta_2 * \log(y_{i,W}) \tag{A.9}$$

The relationship between last period income and the size of the retirement account (η_2) is estimated using the PSID. Estimation is performed on a sample of households where the head is age 60 to 65 and currently employed. We find a value of $\hat{\eta}_2 = 0.99$.

Finally, the annuity is priced equal to its discounted value, giving the following replacement rate during retirement:

$$\eta_1 = \left[\sum_{t=W}^{T} \frac{s_t}{(1+r)^{t-W}} \right]^{-1} \tag{A.10}$$

We then re-calibrate the preference parameters of the model, under the assumption that households have individual retirement accounts as just described. The third column of Table A.8 shows the parameter results. We find that the temptation parameter, λ, the impatience parameter, β, and the risk aversion parameter, γ, hardly change compared to the baseline model. On the other hand, modifying the asset structure of households increases the housing preference parameters significantly (μ, θ).

The third row of Table A.9 shows the MPC by shock size for our alternative model with IRAs. We find that the average MPCs are almost identical in this alternative model, and there still exists a sizable consumption response to large income shocks. For instance, we find an MPC of 0.18 in response to a \$10,000 shock, identical to the MPC in our baseline model.

References

Aguiar, Mark A., Mark Bils, and Corina Boar. 2020. "Who Are the Hand-to-Mouth?" *NBER Working Paper* (26643).

Aladangady, Aditya, Shifrah Aron-Dine, David Cashin, Wendy Dunn, Laura Feiveson, Paul Lengermann, Katherine Richard, and Claudia Sahm. 2018. "High-frequency Spending Responses to the Earned Income Tax Credit." *FEDS Notes. Board of Governors of the Federal Reserve, Washington, DC* .

Altonji, Joseph G. and Lewis M. Segal. 1996. "Small-Sample Bias in GMM Estimation of Covariance Structures." *Journal of Business and Economic Statistics* 14 (3):353–366.

Amador, Manuel, Iván Werning, and George-Marios Angeletos. 2006. "Commitment vs. flexibility." *Econometrica* 74 (2):365–396.

Amromin, Gene, Jennifer Huang, Clemens Sialm, and Edward Zhong. 2018. "Complex mortgages." *Review of Finance* 22 (6):1975–2007.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. 2001. "The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives* 15 (3):47–68.

Ashraf, Nava, Dean Karlan, and Wesley Yin. 2006. "Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines." *The Quarterly Journal of Economics* 121 (2):635–672.

Attanasio, Orazio P., Renata Bottazzi, Hamish Low, Lars Neisham, and Matt Wakefield. 2012. "Modeling the Demand for Housing over the Lifecycle." *Review of Economic Dynamics* 15 (1):10–53.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2018. "The Intertemporal Keynesian Cross." *Mimeo* .

Bayer, Christian, Ralph Luetticke, Lien Pham-Dao, and Volker Tjaden. 2019. "Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk." *Econometrica* 87 (1):255–290. URL `https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA13601`.

Bernstein, Asaf and Peter Koudijs. 2020. "Mortgage Amortization and Wealth Accumulation." *mimeo* .

Beshears, John, James J Choi, Christopher Harris, David Laibson, Brigitte C Madrian, and Jung Sakong. 2020. "Which early withdrawal penalty attracts the most deposits to a commitment savings account?" *Journal of Public Economics* 183:104144.

Bhutta, Neil, Jacqueline Blair, and Lisa Dettling. 2021. "The Smart Money is in Cash? Financial Literacy and Liquid Savings Among U.S. Families." *Finance and Economics Discussion Series (FEDS)* .

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. "Consumption Inequality and Familiy Labor Supply." *American Economic Review* 106 (2):387–435.

Bucciol, Alessandro. 2012. "Measuring Self-Control Problems: A Structural Estimation." *Journal of the European Economic Association* 10(5):1084–1115.

Bunn, Philip, Jeanne Roux, Kate Reinold, and Paolo Surico. 2018. "The Consumption Response to Positive and Negative Income Shocks." *Journal of Monetary Economics* 96:1–15.

Carrera, Mariana, Heather Royer, Mark Stehr, Justin Sydnor, and Dmitry Taubinsky. 2021. "Who Chooses Commitment? Evidence and Welfare Implications." *The Review of Economic Studies* 89 (3):1205–1244. URL `https://doi.org/10.1093/restud/rdab056`.

Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N White. 2017. "The distribution of wealth and the marginal propensity to consume." *Quantitative Economics* 8 (3):977–1020.

Casaburi, Lorenzo and Rocco. Macchiavello. 2019. "Demand and Supply of Infrequent Payments as a Commitment Device: Evidence from Kenya." *American Economic Review* 109 (2):523–555.

Charles, Kerwin Kofi and Erik Hurst. 2002. "The Transition To Home Ownership And The Black-White Wealth Gap." *The Review of Economics and Statistics* 84 (2):281–297.

Cho, SungJin and John Rust. 2017. "Precommitments for financial self-control? Micro evidence from the 2003 Korean credit crisis." *Journal of Political Economy* 125 (5):1413–1464.

Choukhmane, Taha. 2019. "Default Options and Retirement Saving Dynamics." *mimeo* .

Ericson, Keith Marzilli and David Laibson. 2019. *"Intertemporal Choice"*, vol. 2. Elsevier.

Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik. 2021. "MPC Heterogeneity and Household Balance Sheets." *American Economic Journal: Macroeconomics* 13 (4):1–54. URL `https://www.aeaweb.org/articles?id=10.1257/mac.20190211`.

Fang, Hanming and Dan Silverman. 2009. "Time-inconsistency and welfare program participation: Evidence from the NLSY." *International Economic Review* 50 (4):1043–1077.

Flavin, Marjorie and Takashi Yamashita. 2002. "Owner-Occupied Housing and the Composition of the Household Portfolio." *American Economic Review* 92:345–362.

Fuster, Andreas, Greg Kaplan, and Basit Zafar. 2020. "What Would You Do with $500? Spending Responses to Gains, Losses, News, and Loans." *The Review of Economic Studies* 88 (4):1760–1795. URL `https://doi.org/10.1093/restud/rdaa076`.

Goetzmann, William N and Matthew Spiegel. 2002. "Policy implications of portfolio choice in underserved mortgage markets." *Low-income homeownership: Examining the unexamined goal* :257–274.

Greenwald, Daniel and Adam Guren. 2019. "Do credit conditions move house prices?" .

Gul, Faruk and Wolfgang Pesendorfer. 2001. "Temptation and Self-Control." *Econometrica* 69 (6):1403–1435.

———. 2004. "Self Control, Revealed Preference, and Consumption Choice." *Review of Economic Dynamics* 7:243–264.

Harris, Christopher and David Laibson. 2001. "Dynamic Choices of Hyperbolic Consumers." *Econometrica* 69:935–957.

Haurin, Donald R, Patric H Hendershott, and Susan M Wachter. 1997. "Borrowing Constraints and the Tenure Choice of Young Households." *Journal of Housing Research* 8 (2):137.

Hsieh, Chang-Tai. 2003. "Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fun." *American Economic Review* 93:397–405.

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

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

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor. 2019. "The Rate of Return on Everything, 1870–2015*." *The Quarterly Journal of Economics* 134 (3):1225–1298. URL `https://doi.org/10.1093/qje/qjz012`.

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

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

Kaplan, Greg and Giovanni L Violante. 2022. "The marginal propensity to consume in heterogeneous agent models." *Annual Review of Economics* 14:747–775.

Kaplan, Greg, Giovanni L. Violante, and Justion Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity* .

Keane, Michael P. and Nada Wasi. 2016. "Labor Supply: The Roles of Human Capital and the Extensive Margin." *The Economic Journal* 126 (592):578–617.

Kovacs, Agnes, Hamish Low, and Patrick Moran. 2021. "Estimating Temptation and Commitment Over the Life-Cycle." *International Economic Review* 62 (1):101–139.

Kovacs, Agnes and Patrick Moran. 2018. "Temptation, Commitment, and Hand-to-Mouth Consumers." *mimeo* .

———. 2021. "Breaking the Commitment Device: The Effect of Home Equity Withdrawal on Consumption, Saving, and Welfare." *C.E.P.R. Discussion Papers* (16634). URL `https://ideas.repec.org/p/cpr/ceprdp/16634.html`.

Krusell, Per, Burhanettin Kuruscu, and Anthony A. Smith. 2010. "Temptation and Taxation." *Econometrica* 78 (6):2063–2084.

Kueng, Lorenz. 2018. "Excess Sensitivity of High-Income Consumers." *The Quarterly Journal of Economics* 133 (4):1693–1751.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* :443–477.

———. 2015. ""Why don't present-biased agents make commitments?"." *American Economic Review Papers and Proceedings* 105 (5):267–272. URL `https://www.aeaweb.org/articles.php?doi=10.1257/aer.p20151084`.

———. 2018. "Private Paternalism, the Commitment Puzzle, and Model-Free Equilibrium." *AEA Papers and Proceedings* 108:1–21. URL `https://www.aeaweb.org/articles?id=10.1257/pandp.20181124`.

Lewis, Daniel J., Davide Melcangi, and Laura Pilossoph. 2021. "Latent Heterogeneity in the Marginal Propensity to Consume." *FRB of New York Staff Report* (902).

Luetticke, Ralph. 2018. "Transmission of Monetary Policy with Heterogeneity in Household Portfolios." Discussion Papers 1819, Centre for Macroeconomics (CFM). URL `https://ideas.repec.org/p/cfm/wpaper/1819.html`.

McDowall, Robert A. 2020. "Consumption Behavior Across the Distribution of Liquid Assets." *Working Paper* .

Misra, Kanishka and Paolo Surico. 2014. "Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs." *American Economic Journal: Macroeconomics* 6 (4):84–106.

Nakajima, Makoto. 2012. "Rising Indebtedness and Temptation: A Welfare Analysis." *Quantitative Economics* 3 (2):257–288.

Nakajima, Makoto and Irina A. Telyukova. 2020. "Home Equity in Retirement." *International Economic Review* 61 (2):573–616.

OECD. 2011. "Housing and the Economy: Policies for Renovation." *Economic Policy Reforms 2011: Going for Growth* .

Olafsson, Arna and Michaela Pagel. 2018. "The Liquid Hand-to-Mouth." *The Review of Financial Studies* 31 (11):4398–4446.

Parker, Johnathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland. 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103 (6).

Parker, Jonathan A. 2017. "Why Don't Households Smooth Consumption? Evidence from a $25 Million Experiment." *American Economic Journal: Macroeconomics* 9 (4):153–83.

Phelps, Edmund and Robert A. Pollak. 1968. "On Second-Best National Saving and Game-Equilibrium Growth." *The Review of Economic Studies* 35:185–199.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel. 2007. "Housing, Consumption and Asset Pricing." *Journal of Financial Economics* 83(3):531–569.

Schilbach, Frank. 2019. "Alcohol and Self-Control: A Field Experiment in India." *American Economic Review* 109 (4):1290–1322.

Schlafmann, Kathrin. 2020. "Housing, Mortgages, and Self-Control." *The Review of Financial Studies* URL `https://doi.org/10.1093/rfs/hhaa096`. Hhaa096.

Shapiro, Matthew D. and Joel 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: Papers and Proceedings* 99 (2):374–379.

Shiller, Robert J. 2014. *CNN Money* .

Shiller, Robert J. 2015. *Irrational exuberance: Revised and expanded third edition.* Princeton university press.

Strotz, Robert H. 1956. "Myopia and Inconsistency in Dynamic Utility Maximization." *The Review of Economic Studies* 23:165–180.

Vihriälä, Erkki. 2021. "Commitment in debt repayment: evidence from a natural experiment." *Available at SSRN 3792708* .

Wiedemer, J.P. and J.K. Baker. 2012. *Real Estate Finance, 9th ed.* Cengage Learning.