

NBER WORKING PAPER SERIES

TEMPTATION AND INCENTIVES TO WEALTH ACCUMULATION

Orazio Attanasio
Agnes Kovacs
Patrick Moran

Working Paper 28938
<http://www.nber.org/papers/w28938>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2021

We are grateful to Taha Choukmane, 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 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.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Orazio Attanasio, Agnes Kovacs, and Patrick Moran. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Temptation and Incentives to Wealth Accumulation
Orazio Attanasio, Agnes Kovacs, and Patrick Moran
NBER Working Paper No. 28938
June 2021
JEL No. D11,D14,D91,E21,R21

ABSTRACT

We propose a rich model of household behavior to study the effect of two important policies: mortgage interest tax deduction and mandatory mortgage amortization. These policies have attracted some controversy, first because they are conceived to increase overall saving, an objective that the literature does not agree they can achieve, and second because they incentivize illiquid savings and may thus increase the share of ‘wealthy hand-to-mouth’ households. We build a life-cycle model where housing may act as a commitment device to counteract present biases arising from temptation. We show that the model matches several empirical facts, including the large share of wealthy hand-to-mouth households. We evaluate the effect of the two policies and find that they increase wealth accumulation by 7 and 10% respectively. Our results demonstrate that these policies not only induce portfolio re-balancing, as emphasized by the previous literature, but also increase savings by making commitment more accessible.

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

Patrick Moran
Department of Economics
University of Copenhagen
Oster Farimagsgade 5
Building 26
Copenhagen, Denmark 1353
patrick.moran@econ.ku.dk

Agnes Kovacs
The University of Manchester
Department of Economics
Oxford Road
Manchester M139PL
United Kingdom
agnes.kovacs@manchester.ac.uk

1 Introduction

Many governments, policy makers, and even academics have worried about the level of savings of households in the U.S. and other countries. Numerous commentators have judged the level of savings to be ‘too low’. Such low levels of wealth accumulation could jeopardize the sustainability of living standards after retirement, as well as the ability of households to smooth over temporary fluctuations to income. Moreover, even when households accumulate wealth, they often concentrate their wealth in illiquid form. Indeed, a large share of households have almost zero liquid assets, despite holding a sizable amount of illiquid wealth, primarily in housing but also in retirement accounts. The fact that wealthy consumers and in particular homeowners do not have liquid savings implies that they have a limited ability to absorb adverse income shocks. These consumers, who have been labeled as ‘wealthy hand-to-mouth’ (WH2M) households, have featured prominently in a large number of recent macroeconomic models and policy discussions.

Numerous countries have implemented policies intended to stimulate household savings, in part as answer to the worries mentioned above. However, most of these policies are designed to incentivize saving in *illiquid* form. For instance, the U.S. and other countries have introduced large tax subsidies to housing and retirement saving. Further, many countries have implemented policies that compel homeowners to accumulate wealth in housing via mandatory mortgage payments. These policies appear puzzling from the perspective of traditional public finance and macroeconomic models. In these models, such policies may simply induce a reshuffling from liquid to illiquid assets, without increasing aggregate savings. Further, such policies may generate more wealthy hand-to-mouth households, thus increasing consumption sensitivity.

The main goal of this paper is to evaluate policies designed to incentivize illiquid wealth accumulation through the lens of a model of individual preferences that generates a demand for commitment. We make three sets of contributions. First, we develop a two-asset model of household behavior that generates present biases, demand for commitment, and in turn wealthy hand-to-mouth behavior. Second, we take this model to the data and find that it is consistent with several important empirical facts, including recent evidence on consumption and saving behavior. Third, we use this model to evaluate the impact of very visible and controversial policies that incentivize illiquid wealth accumulation, such as tax subsidies to housing and mandatory mortgage amortization.

Modeling demand for commitment. We develop a two-asset life-cycle model with self-control problems that make it difficult for households to save. In particular, we introduce a temptation element into preferences, using the approach proposed by Gul and Pesendorfer (2001, 2004). Temptation represents the idea that households may find it difficult to save in liquid assets, due to the possibility of instantaneous gratification

that is hard to resist. Temptation generates a demand for saving commitment devices that may allow households to mitigate self-control problems.

In our model, households are allowed to save in two instruments: liquid assets and illiquid housing. We focus on housing because it constitutes the vast majority of illiquid wealth for most households.¹ We assume that housing provides direct utility, gives realistic tax advantages, and serves as collateral for fixed-rate fully-amortizing mortgages. As a result of temptation, housing may act as a savings commitment device for two different reasons. First, housing is illiquid and allows households to lock away their wealth. Second, housing is purchased with amortizing mortgages that force households to make regular mortgage payments and gradually accumulate wealth in the form of home equity.

Temptation preferences are not the only ones that give rise to self control problems, present biases, and possibly a demand for commitment. Another strand of the literature that focuses on self-control problems makes use of 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, are dynamically consistent, which allows for a simple recursive formulation and makes this approach very convenient compared to hyperbolic discounting. As a result of dynamic consistency, temptation preferences even allow for straightforward welfare analysis, something that has been problematic under hyperbolic discounting. Finally, temptation preferences are derived from fundamental axioms that generalize standard preferences, making an analysis of policies and policy recommendations based on such preferences well grounded in theory.

The behaviour implied by temptation is consistent with both experimental evidence and observed consumption behavior. Toussaert (2018) finds evidence of temptation in a lab experiment where individuals voluntarily choose to restrict their choice set to avoid costly temptation. Further, Kovacs, Low, and Moran (2020) exploit consumption data to empirically test for the presence of a temptation motive in a life-cycle model. Using data from the Consumer Expenditure Survey, they reject the hypothesis that the parameter that captures temptation is zero. This result justifies our use of temptation preferences.

Taking the model to the data. We calibrate the model to fit U.S. data on consumption and wealth accumulation. We first set the strength of temptation based on the results from Kovacs, Low, and Moran (2020), who use two different estimation strategies to recover the degree of temptation and find broadly similar results. We then calibrate the remaining preference parameters by targeting the life-cycle profiles of consumption, wealth accumulation, and the share of wealthy hand-to-mouth households.

We find that the model obtains a good fit of the data, including the large share of

¹In the Survey of Consumer Finances, the median (mean) U.S. household has \$31,000 (\$72,500) in net housing wealth, but less than \$1,000 (\$34,500) in retirement accounts (Kaplan and Violante, 2014).

wealthy hand-to-mouth households. This is true regardless of whether we set the liquid asset returns to be high or low, reflecting either equities or bonds. This is an important benefit of modeling a demand for illiquidity based on commitment: the presence of WH2M consumers does not depend on large excess returns to housing, but rather on the presence of sizable transaction costs, which generates commitment. In contrast, traditional macroeconomic models of WH2M behavior require the strong assumption that all liquid assets, including equities, deliver lower risk-adjusted returns than housing.²

Further, the view of illiquidity that we highlight is consistent with a growing experimental literature which finds that (i) households often save in low-return illiquid accounts despite the availability of liquid accounts with higher or equal returns and (ii) the ability to save in illiquid accounts helps households achieve a higher saving rate (Thaler and Benartzi, 2004; Ashraf, Karlan, and Yin, 2006; Beshears et al., 2020). Again, this type of behavior is difficult to rationalize within traditional models, but has important implications for our understanding of WH2M behavior and saving subsidies.

Finally, we find that our model is consistent with empirical evidence on the Marginal Propensity to Consume (MPC), including recent findings that have been tough to explain. For instance, recent evidence shows that the MPC declines only slowly with shock size (Kueng, 2018; Fuster, Kaplan, and Zafar, 2018; Fagereng, Holm, and Natvik, 2021). This evidence is difficult to rationalize in traditional two-asset models, where the consumption response to large income shocks is negligible, but is well explained by our framework. The slow decline of MPCs by shock size may affect the optimal design of targeted fiscal stimulus policies, which we discuss briefly.

Implications for policy. Having established the consistency of the model with a variety of facts, as well as its ability to explain the puzzling evidence mentioned above, we use it to analyse the impact of two important and widely debated policies: the tax deductability of mortgage interest payments, which effectively constitutes a sizeable subsidy to home ownership, and mandatory mortgage amortization, which forces homeowners to gradually pay down their mortgage and accumulate wealth in the form of home equity.

Temptation preferences and the demand for commitment (and illiquidity) have important implications for saving behavior and wealth accumulation and therefore provide a new lens to assess the above policies. The opportunity to purchase a savings commitment device (in our case, housing) enables households to accumulate more wealth over the life-cycle. As a result, a subsidy to the commitment device may increase overall wealth

²This is a somewhat controversial assumption. For instance, Robert Shiller remarks, “It would perhaps be smarter, if wealth accumulation is your goal, to rent and put money in the stock market, which has historically shown much higher returns than the housing market.” Similarly, a large number of studies document that housing delivers a lower risk-adjusted return than equities, even when accounting for imputed rent and other benefits to homeownership (Flavin and Yamashita, 2002; Goetzmann and Spiegel, 2002; Piazzesi, Schneider, and Tuzel, 2007).

accumulation by enabling more households to access commitment. Further, mandatory amortization policies may boost household savings by helping homeowners commit to save in the future via regular mortgage payments.

There is a large debate about the effect of the mortgage interest tax deduction on household saving behavior.³ The arguments in this debate share many similarities with the debate surrounding the provision of tax incentives to save in illiquid retirement accounts. There are two main opposing views: some studies claim that subsidies to retirement accounts increase the average return on household savings and thus increase overall wealth accumulation (e.g. Venti and Wise, 1986, 1987), while others argue that such incentives simply result in portfolio re-allocation from liquid to illiquid assets and do not substantially increase overall savings (e.g. Gale and Scholz, 1994; Attanasio and DeLeire, 2002). Yet none of these papers consider the effect of making commitment more affordable, which could stimulate savings in the spirit of Laibson (1997). Using our calibrated model, we assess the relative importance of these different effects. We find evidence of some substitution from liquid to illiquid assets as a result of housing subsidies: every additional dollar of housing wealth reduces liquid assets by roughly 33 cents. However, the increase in housing wealth is not perfectly offset by a reduction in liquid assets, thus overall wealth accumulation increases by around 7% by the time of retirement.

There is a similar debate about the effect of mandatory amortization policies. While mandatory amortization requirements have recently been introduced by Sweden and the Netherlands, there is concern that these policies may simply result in a re-balancing of household portfolios away from liquid assets and into illiquid housing. However, if mortgage payments serve as a commitment device, mandatory amortization may result not only in simple portfolio re-balancing, but also additional wealth accumulation. Based on our counterfactual simulations, we find some portfolio re-balancing from liquid to illiquid assets, but it is far from one-to-one. Overall, mandatory amortization policies increase wealth accumulation by around 10% by the time of retirement. Every additional dollar of housing wealth only reduces liquid assets by roughly 60 cents. This finding is consistent with recent empirical evidence showing that mandatory amortization results in increased wealth accumulation, as households accumulate more housing wealth, but do not reduce their non-housing assets (Bernstein and Koudijs, 2020).

Finally, there is concern that the above policies may increase households' responsiveness to income shocks, which would have detrimental effects for macroeconomic stability (Svensson, 2019, 2020). This may be the case if such policies encourage illiquid wealth accumulation at the expense of liquid wealth accumulation. We use our calibrated model to evaluate this concern. Overall, our model predicts that housing subsidies and manda-

³See for instance Rosen (1985), Poterba (1992), Gervais (2002), Chambers, Garriga, and Schlagenhaut (2009), Floetotto, Kirker, and Stroebel (2016), Sommer and Sullivan (2018), and Nakajima (2020).

tory amortization have very little effect (less than 0.02 percentage point) on the average MPC. The economic intuition for this result is that these policies have little effect on the share of WH2M households. This is because households' desire for illiquidity does not come simply from a desire to exploit higher returns, but also from a desire to obtain commitment. This model prediction comes in sharp contrast with traditional macroeconomic models, such as Kaplan and Violante (2014), where policies that incentivize illiquid wealth accumulation would have a large effect on the share of WH2M households.

The rest of the paper proceeds as follows. First we develop a two-asset model of hand-to-mouth behaviour driven by temptation preferences (Section 2). We demonstrate that temptation generates a desire for illiquidity and a demand for commitment devices (Section 3). We calibrate the model using U.S. data (Section 4) and establish its consistency with both targeted (Section 5) and untargeted moments (Section 6). Finally, we use our calibrated model to study the effect of policies designed to incentivize illiquid wealth accumulation (Section 7).

2 A Model with Temptation Preferences

We develop a life-cycle model of consumption and housing with temptation preferences. In this model, households face uninsurable income risk and save for two reasons: to maintain consumption following adverse income shocks (the precautionary motive) and following retirement (the life-cycle motive). Households can save in either liquid assets or illiquid housing. Housing gives flow utility, serves as collateral, and may provide tax advantages, therefore generating a third incentive 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. While all households are born as renters, they have the possibility to purchase housing which comes in discrete sizes, offers a utility benefit, and serves as collateral for mortgages.

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, subject to a budget constraint we define later:

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

where c_t and h_t are the chosen level 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 the most tempting consumption and housing alternatives available in the choice set each period. We define the utility function as:

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

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:

$$[\tilde{c}_t, \tilde{h}_t] = \arg \max_{c_t, h_t \in \mathcal{A}_t} u(c_t, h_t), \quad (3)$$

where \mathcal{A}_t represents the current period choice set, 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 is not the only way to generate present-bias behavior and possibly generate demand for commitment. The main alternative is to use dynamically inconsistent preferences as formally introduced by Phelps and Pollak (1968) based on ideas proposed in (Strotz, 1956) and later studied by (Laibson, 1997; Harris and Laibson, 2001). These preferences relax the standard assumption on time discounting and introduce “hyperbolic discounting.” Under hyperbolic discounting, the discount rate from today’s perspective between today and tomorrow is larger than the discount rate between two consecutive dates in the future. In contrast, preference proposed by Gul and Pesendorfer (2001) generate present-bias behavior by relaxing the standard assumption on instantaneous utility – more specifically, by allowing utility to depend not only on actual choices but also the most tempting alternative available in the choice set.

There are three main advantages of temptation preferences over hyperbolic-discounting. First, modeling hyperbolic discounting poses a number of conceptual problems, as one needs to specify how individuals perceive their future selves. In the literature, different selves are modeled as playing games the outcomes of which inform dynamic life-cycle

choices. In contrast, temptation allows for a simple recursive formulation, as seen in the Bellman equations at the end of the model description.

Second, as temptation preferences are dynamically consistent, they make welfare analysis straightforward.⁴ In contrast, it is impossible to perform internally consistent welfare analysis under hyperbolic discounting, as the present self and the future self disagree with each other (see for instance Fang and Silverman, 2009, who describe the difficulty of welfare analysis under hyperbolic discounting).

Third, temptation preferences are consistent with experimental evidence showing a demand for commitment even when individuals expect to be able to resist temptation (Toussaert, 2018). In contrast, individuals with hyperbolic discounting only desire commitment if they expect to succumb to the tempting alternative.

Assets. Households who wish to save can invest in two types of assets: liquid financial assets, a_t , or partially illiquid housing, h_t . The financial asset, a_t , yields a certain return r in each period. We abstract away 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: $h^k \in \{h^1, h^2, \dots, h^k\}$. The price of each house $p_t(h^k)$ depends on its size and is determined relative to the index price \bar{p}_t :

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

where for simplicity we assume a linear mapping between house sizes and prices.⁵ House prices grow at a constant rate, $1 + r^H$, over time, representing a fixed gross return on housing, therefore the initial index price determines all other house prices for each period:

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

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

$$F p_t(h^k).$$

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

⁴Examples of welfare analysis under temptation include Krusell, Kuruscu, and Smith (2010), Nakajima (2012), Schlafmann (2020), and Kovacs and Moran (2020).

⁵We previously experimented with different mappings from house sizes to prices, but found that it had little effect on our results, therefore we chose a linear mapping for simplicity.

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

proportional to the price of this unit:

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

where η represents rental scale.

For simplicity, we assume that house prices are determined exogenously and do not adjust in response to policy. While this assumption seems reasonable for regions of the U.S. where land is abundant and housing supply is highly elastic, such as Texas, we should note that this assumption is unlikely to hold in places where land is more scarce, such as New York City. That said, it would be straightforward to extend our model to include equilibrium house prices and partially inelastic housing supply.⁷

In addition, 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. As Kaplan and Violante (2014) note, the median household keeps only \$950 in their retirement account, according to the Survey of Consumer Finances. In Section A.6.2, we consider sensitivity to this assumption.

Mortgages. The most widely used mortgage contract in the U.S. is the fully-amortizing fixed-rate mortgage.⁸ We assume that mortgages are of this kind, with regular required mortgage payments that force households to gradually build up wealth in the form of home equity. As a result, housing may act as a commitment device not only because of its relative illiquidity, but also because of mandatory mortgage payments mp_t every period. The mortgage balance for households who buy a house at time t is

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

where ψ is the maximum loan-to-value (LTV) ratio. Conditional on not moving homes,

⁷If we were to incorporate equilibrium house prices in our model, we would expect that this would alter the quantitative but not qualitative results in Section 7. Of course, the quantitative magnitude of this effect would vary depending on the local housing supply elasticity. Interestingly, the direction of this effect may differ for housing subsidies and mandatory amortization: the former would increase house prices while the latter may slightly reduce house prices.

⁸These mortgages accounted for approximately two-thirds of mortgage origination in the U.S. 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 Act of 2010 reaffirmed these standards with the "qualified mortgage" regulation, which introduced strong incentives for banks to offer mortgages that are fully-amortizing. For more discussion of the Dodd-Frank Act, see DeFusco, Johnson, and Mondragon (2020).

mortgage balances evolve according to the following law of motion:

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

where mp_t represents the required mortgage payment at time t . We assume that mortgages require constant-level payment plans, as is the case for the vast majority of mortgages in the United States, 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:

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

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. As households are required to pay off their mortgages by the time of retirement, the terminal condition, $m_{W+1} = 0$, is satisfied.

While the model features a realistic amortization schedule, we assume for simplicity that home equity withdrawal is not permitted. As a result of this assumption, households can only liquidate their housing wealth by moving homes, conditional on paying the housing transaction cost.⁹ The fact that housing is partially illiquid means that it does not provide perfect commitment, although it still serves as a savings commitment device.

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 are forced to sell their home and repay their remaining mortgage debt.¹⁰

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

$$\ln y_{i,t} = g_t + z_{i,t} \quad (8)$$

where g_t is a deterministic age profile approximated by a third-order age-polynomial,

⁹This simplifying assumption allows us to keep track of mortgage balances using a discrete state variable which represents the time since home purchase. That said, it may be interesting to relax this assumption and explicitly allow for home equity withdrawal. We expect that doing so would alter the quantitative but not qualitative results, as housing would still be partially illiquid and would require regular amortization payments, thus providing partial commitment.

¹⁰In the case where 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.

while $z_{i,t}$ is an idiosyncratic shock to log income that is described by an AR(1) Markov process:

$$\begin{aligned} z_{i,t} &= \rho z_{i,t-1} + \varepsilon_{i,t} \\ \varepsilon_{i,t} &\sim N(0, \sigma_\varepsilon^2) \\ \varepsilon_{i,0} &\sim N(0, \sigma_0^2). \end{aligned} \tag{9}$$

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.

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:

$$\tau(y_{i,t}, a_{i,t}) = e^{\tau_1 + \tau_2 \log(y_{i,t} + ra_{i,t} - \tau_d)} \tag{10}$$

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 United States. This results in an after-tax income for households given by the following equation:

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

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

$$\tilde{y}_{i,t} = \max \left\{ \text{SS Income Floor, Annual PIA}(y_W) \right\} \tag{11}$$

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 .¹¹ We calibrate the social security income floor and primary insurance amount based on U.S. legislation from 2015.¹²

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

$$u(c_t, h_t) = n_t \left(\frac{\left(\frac{c_t}{n_t}\right)^{1-\gamma}}{1-\gamma} \exp \left[\theta \phi(h_t, n_t) \right] + \mu \phi(h_t, n_t) - \chi I_{h_t \neq h_{t-1}} \right) \quad (12)$$

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 $\phi(h_t, k_t)$ 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: $h^k \in \{h^1, h^2, \dots, h^k\}$. 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.

$$\phi(h_t, k_t) = \ln \left(\frac{h_t}{n_t} \right) \quad (13)$$

Whenever a household adjusts housing (i.e. when $I_{h_t \neq h_{t-1}}$ equals one in equation (12)), it suffers a utility cost, χ .¹³ The utility cost, besides the financial transaction cost, plays an important role in our model, as it increases the illiquidity of housing, thus making

¹¹In reality, to calculate AIME, the worker's wage during the years of employment is first expressed in today's dollars, 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.

¹²The 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.

¹³Here we think of the non-monetary cost of changing homes, like finding new schools, setting up new utility providers, facing stress etc.

housing more useful as a commitment device.

Budget Set. We define the budget set, \mathcal{A}_t , as follows:

$$\mathcal{A}_t = \begin{cases} x_t \in R^+ : x_t \leq a_t + \tilde{y}_t - \mathbb{I}_t^{own} mp_t - (1 - \mathbb{I}_t^{own}) rent_t \\ \text{if no housing adjustment} \\ x_t \in R^+ : x_t \leq a_t + \tilde{y}_t - \left[(1 + F)p_t(h_t) - \frac{m_{t+1}}{(1 + r^M)} \right] \\ \quad + \left[(1 - F)p_t(h_{t-1}) - m_t \right] \\ \text{if housing adjustment} \end{cases} \quad (14)$$

This determines the choice set for households each period, which in turn pins down the most tempting alternative each period. Households will always be tempted to spend all of their available resources on either consumption or housing.

Recursive Formulation. We define the recursive formulation as follows:

$$V_t(\Omega_t) = \max \left\{ V_t^0(\Omega_t), V_t^1(\Omega_t) \right\} \quad (15)$$

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:

$$V_t^0(\Omega_t) = \max_{\{c_t, a_{t+1}\}} U(c_t, h_t, \tilde{c}_t, \tilde{h}_t) + \beta \mathbb{E}_t V_{t+1}(\Omega_{t+1}), \quad (16)$$

subject to:

$$\begin{aligned} a_{t+1} &= (1 + r) \left[a_t + \tilde{y}_t - c_t - \mathbb{I}_t^{own} mp_t - (1 - \mathbb{I}_t^{own}) rent_t \right] \\ \tilde{y}_t &= \begin{cases} \exp(g_t + z_t), & \text{if } t \leq W \\ \text{SS Benefit}(y_W), & \text{if } t > W \end{cases} \\ z_t &= \rho z_{t-1} + \varepsilon_t \quad \text{and} \quad c_t > 0 \end{aligned} \quad (17)$$

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

$$V_t^1(\Omega_t) = \max_{\{c_t, h_t, m_{t+1}, a_{t+1}\}} U(c_t, h_t, \tilde{c}_t, \tilde{h}_t) + \beta \mathbb{E}_t V_{t+1}(\Omega_{t+1}), \quad (18)$$

subject to:

$$\begin{aligned}
a_{t+1} &= (1+r) \left[a_t + \tilde{y}_t - c_t - (1+F)p_t(h_t) + \frac{m_{t+1}}{(1+r^M)} + (1-F)p_t(h_{t-1}) - m_t \right] \\
m_{t+1} &\leq \psi p_t(h_t)(1+r^M) \\
y_t &= \begin{cases} \exp(g_t + z_t), & \text{if } t \leq W \\ \text{SS Benefit}(y_W), & \text{if } t > W \end{cases} \\
z_t &= \rho z_{t-1} + \varepsilon_t \quad \text{and} \quad c_t > 0
\end{aligned} \tag{19}$$

3 Key Model Insights

In this section, we demonstrate two implications of our model that differ from those of the standard model. First, our model generates a demand for illiquidity that is absent in a standard model. Second, the availability of housing helps households save for retirement. To better highlight the implications of temptation and commitment, we focus on a simplified version of our model in this section.

More specifically, we simplify our model by assuming that there is only one size of housing to rent and buy; that housing does not enter the utility function; that labor income is deterministic; and that the returns on liquid assets and housing are the same. Table 1 presents the parameter restrictions imposed in the simplified model.

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.

3.1 Demand for Illiquidity

In this simplified model, households with standard preferences ($\lambda = 0$) have no demand for housing. Homeownership 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 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 drastically 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.

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.


In contrast, households with temptation 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 cost of deviating from your most tempting consumption alternative each period. As a result, households purchase housing despite the presence of sizeable transaction costs. This is driven by the fact that housing provides commitment, which helps households accumulating 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.¹⁴

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

¹⁴This 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.”

no liquid wealth despite owning housing.

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.

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.¹⁵ After buying a 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, therefore consumption declines when nearing retirement.

The demand for illiquidity that we highlight is consistent with experimental evidence showing that individuals often choose to save in illiquid accounts even when there is a liquid asset that gives equivalent returns. For instance, Thaler and Benartzi (2004) find that many individuals are willing to commit in advance to allocate future income increases to an illiquid retirement account. Ashraf, Karlan, and Yin (2006) conduct a field experiment in the Philippines and find that many individuals choose to save in an illiquid savings account that restricts access to their funds but provides no further benefits. Further, Beshears et al. (2020) conduct a lab experiment in the US and find that (i) many individuals choose to allocate their funds to an illiquid saving account that offers equivalent returns to a liquid savings account and (ii) demand for the illiquid saving account actually increases as the degree of illiquidity is increased.


¹⁵Temptation causes households to accumulate their downpayment quickly prior to home 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.

3.2 Access to Commitment Increases Wealth Accumulation

In this section, we study how the availability of housing impacts the wealth accumulation of households. We demonstrate that in a standard model, the presence of housing has no impact on wealth accumulation, as housing delivers identical returns to liquid assets in our simplified model. In contrast, in a temptation model, the presence of housing enables households to accumulate greater wealth for retirement, as illiquid housing enables households to “lock away” their wealth and therefore mitigate the effects of temptation.

To see this, we look at patterns of wealth accumulation, imposing different assumptions on the availability of housing and household preferences. We therefore simulate households under four scenarios: with and without housing and with and without temptation. This allows us to observe the difference in wealth accumulation when housing is not available (i.e. when a savings commitment device is not available).

Figure 3: The Change in Wealth when Housing is Available



Note: This figure shows the difference in net wealth when housing is available versus when housing is unavailable. The red line shows the difference for a tempted household, while the blue line with round markers shows the difference for a standard household. The duration of homeownership is measured for a household that owns housing and all other simulations are compared based on age.

Figure 3 presents our results from these simulations using our simplified model. The solid red line shows the difference in net wealth for a tempted household that has access to housing, relative to an identical household that does not have access to housing. The line is increasing over the duration of home ownership, indicating that the presence of housing changes the savings behavior of the tempted households. After buying a home, households are required to repay the mortgage in each period, which acts as a self-imposed commitment device that forces them to accumulate home equity. When housing is not available, households have no way to mitigate their temptation, therefore they accumulate

less wealth by the time of retirement.

It is interesting to note that the red line begins to decline towards the end of the 30 year window that we consider, as households near retirement. This is because tempted households without access to housing decide to accumulate liquid assets rapidly immediately prior to retirement. When the tempted household does not have access to a commitment device, they suffer a temptation cost for every period that they hold liquid assets, therefore they try to accumulate wealth for retirement as late and as quickly as possible. They catch up with the tempted household with housing, although this effect is only partial.

In contrast, the blue line with round markers in Figure 3 shows the difference in wealth accumulation for a standard household that has access to housing, relative to an identical household that does not have access to housing. The line is horizontal, indicating that the presence of housing does not change the savings behavior of the standard household. In this model, the type of asset choice does not impact the amount of asset accumulation.

Again, experimental evidence shows that access to commitment allows individuals to achieve higher wealth accumulation. Thaler and Benartzi (2004) find that access to a commitment device results in a large increase in saving rates, with the average saving rate going from 3.5 percent to 13.6 percent over 40 months among compliers. Ashraf, Karlan, and Yin (2006) find that the availability of an illiquid saving account results in an 81% increase in average bank account balances one year later for individuals who are offered the illiquid account compared to those who are not offered it.

In addition, these results extend to the context of mortgage amortization payments, which in our framework also serves as a commitment device. Bernstein and Koudijs (2020) study the introduction of a mandatory amortization policy in the Netherlands and find that amortization payments result in a near one-for-one increase in net wealth accumulation, as households increase their housing wealth but do not reduce their liquid assets. In addition, Bäckman and Khorunzhina (2020) study the elimination of mandatory amortization in Denmark (the opposite of the policy in the Netherlands) and find that the elimination of mandatory amortization results in a large increase in consumption, which implies lower saving rates.

4 Calibration

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 Externally Set Parameters

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

Temptation Parameter. We set $\lambda = 0.28$ following the parameter estimate obtained by Kovacs, Low, and Moran (2020). These authors develop two complementary strategies to estimate the strength of temptation: first, a semi-structural Euler equation approach using observed consumption data, and second, a simulation based approach targeting the life-cycle profiles of consumption and wealth accumulation. The authors apply these two estimation strategies to the U.S. using data from the Consumer Expenditure Survey (CEX) and find similar estimates of λ in both cases. We opt to use the point estimate from the Euler equation method, as this semi-structural approach requires fewer assumptions about the economic environment faced by households.¹⁶

By assuming that households suffer from temptation, our model implies a desire for illiquidity which is consistent with a growing experimental literature, as discussed in the previous section.

Asset Returns. Traditional macroeconomic models generate wealthy hand-to-mouth households by assuming that illiquid assets deliver large excess returns relative to all available liquid assets. However, this is a somewhat controversial assumption, especially when looking at specific assets such as housing and equities. There is a wide body of empirical research showing that housing delivers lower risk-adjusted returns than equities, even when accounting for imputed rents and other benefits of homeownership. This poses a challenge for traditional models of hand-to-mouth behavior. If a high-return liquid asset such as equities were added to the model of Kaplan and Violante (2014), the wealthy hand-to-mouth would disappear.

We set ourselves the challenge of rationalizing the large share of WH2M households despite the presence of a high-return liquid asset (i.e. equities). We thus assume that r is given by the real risk-adjusted return on equities. This assumption is not crucial to our main results, but it does highlight an important benefit of modeling a demand for illiquidity based on commitment: the presence of WH2M does not require $r > r^H$. However, our main results do not depend on this assumption, as we demonstrate in Section A.6.1 where we ignore the presence of high return liquid assets.

We compute real risk-adjusted returns to housing and equities in Appendix A.2. We explicitly account for imputed rents in housing using the balance sheet approach, following Piazzesi, Schneider, and Tuzel (2007) and Kaplan and Violante (2014). We find a real risk-adjusted return of 2.1% for housing and 5.4% for equities, therefore we set the returns

¹⁶For instance, the Euler equation approach is invariant to alternative assumptions about the asset structure of the economy, as explained by Kovacs, Low, and Moran (2020).

in our model to be $r^H = 0.021$ and $r = 0.054$ respectively. This calibration is consistent with a wide body of research including Flavin and Yamashita (2002), Goetzmann and Spiegel (2002), and Piazzesi, Schneider, and Tuzel (2007) who show that housing delivers lower risk-adjusted returns than equities, even when accounting for imputed rents and other benefits to homeownership.

Housing Transaction Costs. Following Attanasio et al. (2012), we assume that moving homes requires households to pay a transaction cost F equal to 5% of the value of the house. F represents the cost of real estate agents, lawyers, surveyors, and moving companies, and is consistent with empirical evidence showing that transaction costs for housing are usually at least 5% of the asset value (OECD, 2011). The presence of large transaction costs is in contrast to Kaplan and Violante (2014), who assume a much smaller transaction cost of \$1,000. The low transaction cost in Kaplan and Violante (2014) helps their model generate a large share of WH2M households. However, this low transaction cost would be unreasonable for a model that explicitly focuses on housing.¹⁷

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.

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%, 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 we set the maximum loan-to-value (LTV) ratio $\psi = 0.9$.¹⁸

Income and Taxes. We calibrate income over the life-cycle in two steps. First, we

¹⁷As we later demonstrate, the two-asset model without temptation has difficulty generating a large share of WH2M households when there is a large transaction cost, even when allowing for potentially very large utility benefits to housing. This holds in both the baseline calibration where $r > r^H$ and an alternative calibration where $r = r^H$ (see Appendix A.6.1).

¹⁸We do not study optimal LTV requirements, however, a nice paper by Schlafmann (2020) evaluates the welfare effect of alternative LTV requirements when households value commitment.

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. 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. Following the U.S. tax code in 2015, we assume that mortgage interest payments are fully deductible from taxable income and that capital gains to owner-occupied housing are not taxed. We parametrize the progressive social security-style pension based on U.S. 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 PSID.

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 Internally Calibrated Parameters

The remaining parameters in the model, $\Gamma = \{\beta, \gamma, \mu, \theta, \chi\}$, the the impatience parameter, the risk aversion parameter, the housing preference (taste) parameters respectively, and the utility cost of changing home, are internally calibrated by the Method of Simulated Moments such that the model matches aggregate statistics of consumption and wealth accumulation. 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, $m^s(\hat{\mathbb{P}}, \Gamma)$.

$$f(\hat{\mathbb{P}}, \Gamma) \equiv [m^s(\hat{\mathbb{P}}, \Gamma) - m^e] \quad (20)$$

We choose to target the mean life-cycle profiles of four variables: consumption, liquid assets, net housing wealth, and the share of homeowners with zero liquid assets. To focus on the working life, we target the mean of each variable between ages 25 and 65, giving 164 moment conditions. We also target the average homeownership rate and the average share of homeowners who move in a given year, giving two additional moments. Altogether we target $N_m = 166$ 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:

$$f^W(\hat{\mathbb{P}}, \Gamma) \equiv f(\hat{\mathbb{P}}, \Gamma) \cdot W^{-1} \cdot f(\hat{\mathbb{P}}, \Gamma)' \quad (21)$$

where W is a diagonal $N_m \times N_m$ 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.¹⁹

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 set λ to be 0.28 following Kovacs, Low, and Moran (2020) and we call this the model with temptation. 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.

Data comes from the Panel Study of Income Dynamics (PSID) waves 1999 to 2015. Detailed information on the sample is contained in Appendix A.3. In this section, we describe each targeted moment in turn.

Consumption. We compute real nondurable consumption following the classification in Blundell, Pistaferri, and Saporta-Eksten (2016). With this classification our consumption measure covers approximately 70 percent of consumption expenditure on nondurable goods and services.

Liquid assets. We compute liquid assets as the sum of bank account deposits and publicly traded stock. We believe that publicly traded stock are essentially liquid because there are very low transaction costs on these assets, thus it would be easy for households to sell their stock position for consumption smoothing purposes.

Net housing wealth. We measure net housing wealth as the reported value of the household's main residence minus all mortgage debt on this home. We exclude net wealth from other real estate.

¹⁹We 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.

Wealthy Hand-to-Mouth. We target the share of homeowners with liquid assets less than two weeks of income, which we refer to as the wealthy hand-to-mouth. 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).

Homeownership. We target the average homeownership rate between ages 25 - 65.

Movers. We target 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 2 presents the results for the calibrated preference parameters: 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, where we turn off temptation.

The model with temptation yields an annual discount factor (β) of 0.96 and a risk aversion parameter (γ) of 2.41, consistent with most of the macroeconomic literature. The additive utility benefit of housing (μ) is roughly 0.3, while the non-separable utility benefit of housing (θ) is 0.16. It is worth noting that the calibrated positive value of θ implies that consumption and housing are complements. 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 0.43 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 \$6,000 of consumption if it decides to buy a bigger house.²¹

The model without temptation yields both a lower impatience parameter of ($\beta = 0.91$) and a lower risk aversion parameter ($\gamma = 1.88$). Moreover, without temptation the model

²⁰Since μ is positive, an increase in housing has a direct positive effect on utility. In addition, it is worth noting that a positive θ implies Edgeworth complementarity of consumption and home ownership, as the cross derivative of utility with respect to consumption and housing is positive

$$\frac{\partial u(c_t, h_t)/\partial c_t}{\partial h_t} = \theta \phi'(h_t, n_t) c_t^{-\gamma} \exp(\theta \phi(h_t, n_t)) > 0$$

Thus in the model with temptation, housing and consumption are complements, whereas in the model without temptation housing and consumption are (weak) complements.

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

Table 2: Internally calibrated parameters

PARAMETER		Temptation Model	No Temptation Model
Time preference	β	0.959 (0.000)	0.914 (0.000)
Risk aversion	γ	2.415 (0.052)	1.884 (0.019)
Housing utility (separable)	μ	0.291 (0.000)	0.591 (0.004)
Housing utility (non-separable)	θ	0.160 (0.007)	0.282 (0.003)
Utility cost of housing adjustment	χ	0.437 (0.037)	0.887 (0.039)
Goodness of Fit	$f^W(\hat{\mathbb{P}}, \hat{\Gamma})$	5.667	10.059

Note: In the temptation model we set $\lambda = 0.28$. In the model without temptation we set $\lambda = 0$. Goodness of fit is defined by equation (21).

features significantly higher housing taste parameters ($\mu = 0.59$ and $\theta = 0.28$). This is because the model without temptation requires strong discounting of future utility and greater tolerance to risk to try to explain the large share of the wealthy 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 higher value of utility cost of moving ($\chi = 0.88$).

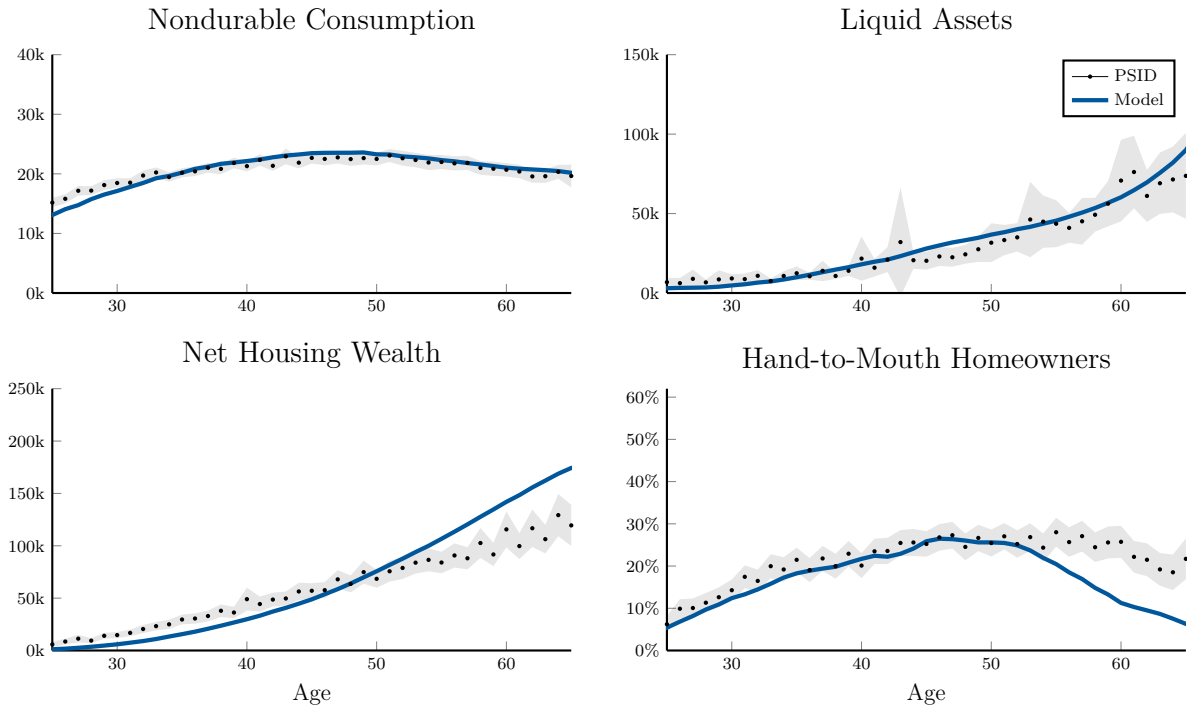
5 Model Fit

We find that the baseline model with temptation generates a good match of the life-cycle profiles that we target, including the striking illiquidity of household portfolios and the large share of wealthy hand-to-mouth households. In contrast, we show that a restricted version of the model without temptation is unable to match key features of the data.

Baseline Model with Temptation. We first analyze the performance of their calibrated model with temptation. Figure 4 shows the simulated life-cycle moments from the temptation model (the solid blue line) and the targeted moments from the PSID (the black dotted line). In general, the temptation model obtains a good fit of the life-cycle

profile of consumption and liquid asset accumulation. The model obtains a good fit of net housing wealth up until age 50, although after this point, net housing wealth starts diverging a bit as households approach retirement. The divergence of net housing wealth arises from the simplifying assumption in our model that all mortgages must be paid off by age 65, when households are forced to retire. In reality, some households continue to work and/or maintain positive mortgage balances after age 65, which helps explain this difference. Notice that, apart from very young ages, household’s wealth is concentrated in illiquid housing: by the age of 40, housing wealth accounts for about 70% of the average U.S. household’s wealth. This overwhelming dominance of housing wealth only starts declining later in life, close to retirement, reaching roughly 60% by age 65.

Figure 4: Fit of Baseline Model with Temptation



Note: This figure shows the life-cycle moments from the temptation model (solid blue line) and the PSID (black dotted line). The moments from the PSID are shown with bands of 1.96 standard deviations around the mean. Table A.4 in the Appendix shows the fit of the other moments that we target: the average homeownership rate and the average share of homeowners moving each period.

Most importantly, the temptation model obtains a good fit of the wealthy hand-to-mouth. In the temptation model, wealthy hand-to-mouth behaviour is mainly driven by the fact that households are willing to withstand fluctuations in consumption in order to obtain the commitment benefit of illiquid assets. The age profile for the wealthy hand-to-mouth is hump-shaped with a peak around age 45, when roughly 25% of households own a home but hold essentially zero liquid wealth. These findings are very similar to what Kaplan, Violante, and Weidner (2014) document using the U.S. Survey of Consumer Finances: they find a hump-shaped life-cycle profile of the wealthy hand-to-mouth with a peak of around 22% at age 40.

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 a roughly 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, Bills, and Boar, 2020).

Restricted 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 the wealthy hand-to-mouth when housing does not deliver higher risk-adjusted returns than liquid assets.

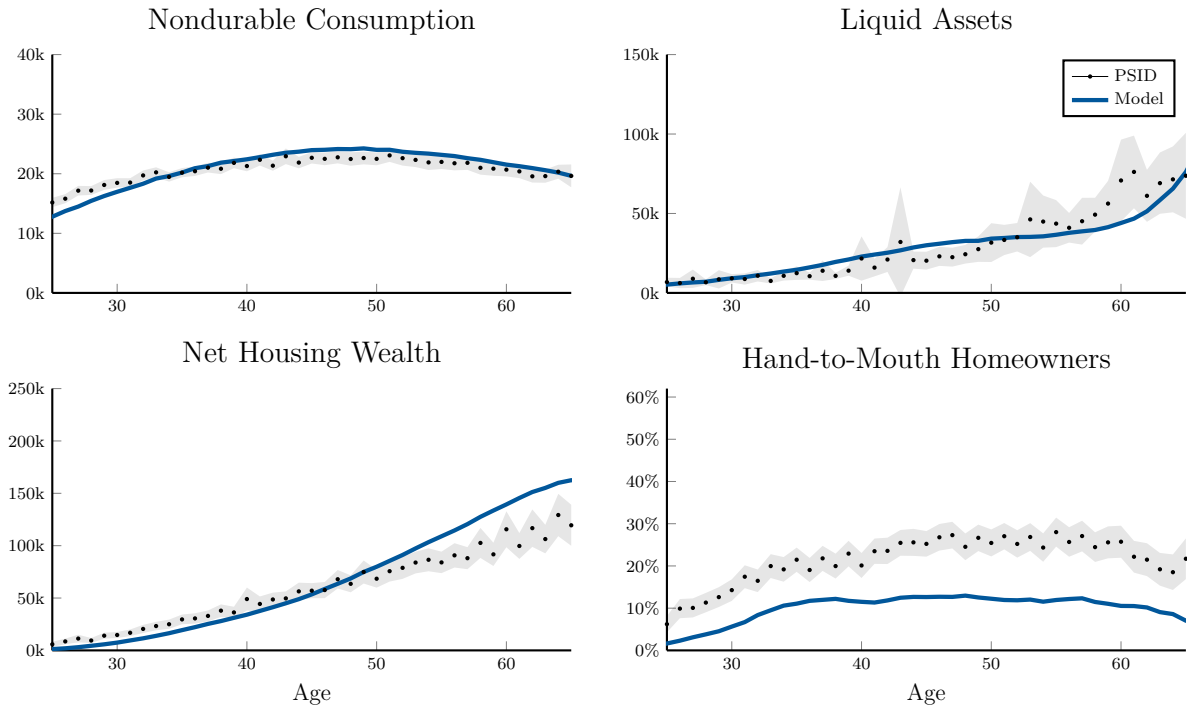
Figure 5 shows the simulated life-cycle moments from the model without temptation (the solid blue line) and the targeted moments from the PSID (the black dotted line). We observe that the restricted model without temptation obtains a relatively good match of the life-cycle profile of consumption, liquid assets, and net housing wealth. However, the restricted model fails to match the large share of hand-to-mouth homeowners, obtaining roughly half of the hand-to-mouth homeowners observed in the data. This is despite the fact that the model is relatively flexible and allows for impatience, different types of housing taste, and varying degrees of risk aversion. This tells us that there is no combination of impatience, housing taste, and risk aversion that is able to match the large share of wealthy hand-to-mouth households without failing to match some other targeted moments.

High impatience is unable to explain the large share of wealthy hand-to-mouth households without compromising model fit elsewhere. For instance, if β were lowered below our calibrated value in the model without temptation ($\hat{\beta} = 0.91$), we would obtain a better fit of the wealthy hand-to-mouth, but a worse fit of the life-cycle profile of wealth accumulation.

Greater tolerance towards risk (i.e. lower γ) reduces the degree to which households are worried about consumption fluctuations. However, it also reduces the degree to which households accumulate wealth for a down-payment in order to purchase housing, thus reducing the homeownership rate. As a result, a lower value of γ slightly decreases the share of the wealthy hand-to-mouth, while simultaneously increasing the share of poor hand-to-mouth households.

Higher housing taste (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

Figure 5: Fit of the Restricted Model without Temptation



Note: This figure shows the life-cycle moments from our no temptation model (solid blue line) and the PSID (black dotted line), where we impose $\lambda = 0$. The moments from the PSID are shown with bands of 1.96 standard deviations around the mean.

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.

6 Consistency with Empirical Evidence on Consumption

We next evaluate our model’s ability to match untargeted empirical evidence on the Marginal Propensity to Consume (MPC). We believe that this is important given that the view of illiquidity that we highlight differs from the returns-based explanation offered by the traditional macroeconomic literature.

We find that our model is consistent with existing empirical evidence on the average MPC and heterogeneity in the MPC based on liquid asset holdings. This result is similar to Kaplan and Violante (2014) and therefore we do not dwell on it.

More importantly, however, our model also generates realistic MPC heterogeneity based on shock size. In particular, we find that the average MPC predicted by our model declines relatively slowly with shock size. Large income shocks still induce a sizeable increase in consumption. This finding is consistent with a growing empirical literature, which we review below, but inconsistent with traditional models of wealthy hand-to-mouth behavior such as Kaplan and Violante (2014).

6.1 MPC Heterogeneity

We use our baseline model to study the consumption response to unanticipated and transitory income shocks. The details of this procedure can be found in Appendix A.5.1. We find that the model generates an average annual MPC of 0.43 in response to a \$1,000 income shock. This lies well within the standard range of annual MPCs estimated by both the empirical and theoretical literature. For instance, Carroll, Slacalek, Tokuoka, and White (2017) give a comprehensive summary of the empirical literature and report that average annual MPC estimates for the U.S. range between 0.2 and 0.6. Similarly, Kaplan and Violante (2014) report an average annual MPC of 0.36 in response to a \$1,000 income shock, only slightly lower than our baseline result.²²

In addition, when evaluating heterogeneity across households, we find that the average MPC in our model declines slowly with total wealth, but quickly with liquid wealth. This is consistent with a large body of empirical research, such as 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). The results are contained in Appendix A.5.2.

6.1.1 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 3. 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 0.43 to 0.29.

Table 3: MPC Heterogeneity by Shock Size

	SHOCK SIZE		
	\$1,000	\$5,000	\$10,000
MPC	0.43 (0.002)	0.37 (0.002)	0.29 (0.001)

Note: Each coefficient represents the average annual MPC in our calibrated model. The procedure used to compute MPCs is described in Appendix A.5.1.

This result is consistent with a growing empirical literature that finds that the average MPC remains large in response to large income shocks. Most similar to our analysis, Fagereng, Holm, and Natvik (2021) study the consumption response to large and unanticipated lottery winnings using administrative tax data from Norway. They find that the average MPC only gradually declines as the size of the shock increases: households

²²The authors report a quarterly MPC of 0.16 in their model, which leads to an annual MPC of 0.36.

that receive large payments still consume most of their payment shortly after receipt.

In addition, 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 0.3.²³ Finally, 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 finds similar results. For instance, Bunn et al. (2018) and Fuster, Kaplan, and Zafar (2018) both 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 large MPCs in response to large income shocks.²⁴ Section A.5.3 explores the reasons for these differences.

6.2 Interpretation and Implications

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 Appendix A.5.3, we decompose the relative importance of these two mechanisms. 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.

Our findings on MPC heterogeneity have important implications for the design of targeted fiscal stimulus policies. As noted previously, households with low cash-on-hand have the largest MPC in response to transitory income shocks. Further, the average MPC declines relatively slowly with shock size. Taken together, these two facts suggest that large and targeted fiscal stimulus payments could be very effective in boosting aggregate

²³While 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 nonclassical measurement error in income data attenuating the estimates, and (ii) the estimates in Hsieh (2003) are of consumption elasticities rather than MPCs.

²⁴Note that historically it has been difficult to measure 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) all analyze the consumption response of households to the 2001 and the 2008 fiscal stimulus, with average payments of \$600 and \$1,200.

consumption. We study the design of targeted fiscal stimulus policies in Appendix A.5.4 and find that fiscal stimulus can produce a much larger consumption response when it is heavily targeted towards households in the bottom of the income distribution. In contrast, most governments have historically relied upon small fiscal stimulus payments given to a large proportion of the population.

7 How Does Policy Affect Wealth Accumulation?

The view of illiquidity that we highlight has important implications for saving behavior and wealth accumulation. As we demonstrated in Section 3, the ability to purchase a savings commitment device (in our case, housing) allows households to accumulate more wealth over the life-cycle. This occurs for two reasons: first, households are able to reduce temptation by locking their wealth in illiquid form; second, households are able to bind their future self to save by committing to regular mortgage payments that build up wealth in the form of home equity.

In this section, we study two important policies that alter households' incentives to accumulate wealth in housing, and which thus may affect their ability to obtain commitment. First, we evaluate the effect of housing subsidies, such as the mortgage interest tax deduction and capital gains exemption. Second, we study the impact of mandatory mortgage amortization policies, which force households to gradually pay off their mortgage and accumulate wealth in the form of home equity. In both cases, there is substantial controversy surrounding (a) whether these policies may even increase household saving and (b) whether these policies may generate more wealthy hand-to-mouth households.²⁵

7.1 Housing Subsidies

There is a long-standing debate about the effect of tax benefits that incentivize wealth accumulation in illiquid assets such as housing or individual retirement accounts. In the United States and many other countries, a large number of tax advantages have been given to housing and retirement accounts. Some studies have argued that such tax incentives may result in an increase in overall savings (Venti and Wise, 1986, 1987), while others have argued that such incentives may only result in portfolio re-allocation from liquid to illiquid assets (Gale and Scholz, 1994; Attanasio and DeLeire, 2002). In addition, a growing literature argues that tax benefits to housing may distort household portfolio allocations, resulting in increased accumulation of housing wealth at the expense of reduced accumulation of non-housing assets.²⁶ Yet none of these papers consider the

²⁵For this reason, we focus our current analysis on the positive effects of these policies. That said, it would be easy to extend our analysis to the normative effects of these policies, given that temptation preferences allow for straightforward welfare analysis due to time consistency.

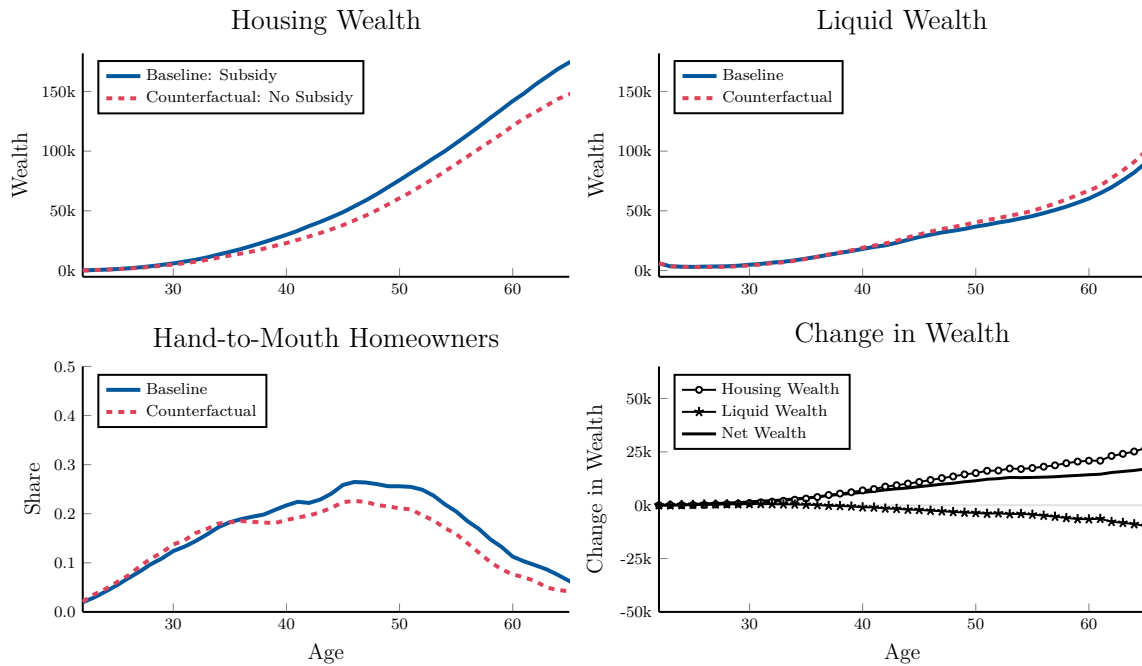
²⁶See for instance Rosen (1985), Poterba (1992), Gervais (2002), Chambers, Garriga, and Schlagenhaut (2009), Floetotto, Kirker, and Stroebel (2016), and Nakajima (2020).

effect of making commitment more affordable, which could be an important channel through which governments may stimulate household saving.

Using our calibrated model, we evaluate the effect of housing subsidies on both portfolio allocation and overall wealth accumulation. We focus on two large subsidies to homeownership that are relevant in the U.S. context: the mortgage interest tax deduction and the tax exemption of capital gains on owner-occupied housing. We compare our baseline model, where both tax benefits exist, to a counterfactual model, where both tax benefits are eliminated.

Figure 6 shows the life-cycle profiles of wealth accumulation in the baseline model with housing subsidies (the solid blue line) relative to the counterfactual model without housing subsidies (the dashed pink line). We find evidence of some portfolio re-balancing from liquid to illiquid assets, however, this is not a one-to-one effect. As a result, overall wealth accumulation is increased due to the presence of housing subsidies.

Figure 6: Effect of Housing Subsidies



The first two panels of Figure 6 show the life-cycle profiles of net housing wealth and liquid wealth in the models with and without housing subsidies. We find that housing subsidies increase housing wealth accumulation at the cost of reduced liquid wealth accumulation. Every additional dollar of housing wealth reduces liquid assets by roughly 33 cents by the time of retirement.

The third panel of Figure 6 shows the share of hand-to-mouth homeowners in the two different models. We see that the share of hand-to-mouth homeowners is reduced only slightly by eliminating the tax benefits to housing. This is an important finding

which differs substantially from a traditional model. In a model where wealthy hand-to-mouth households are generated by the assumption that illiquid assets deliver excess returns relative liquid assets, such as Kaplan and Violante (2014), we would expect that a reduction in tax benefits to illiquid assets would have a large effect on the share of hand-to-mouth households, since it reduces the gap in returns between illiquid and liquid assets. In contrast, in our model, hand-to-mouth homeowners exist not due to excess returns on housing, but rather due to a desire to lock away wealth in illiquid form.

Finally, the fourth panel of Figure 6 shows the difference in wealth accumulation in the model with housing subsidies relative to the model without housing subsidies. We find that overall wealth accumulation is increased by roughly 7% by the time of retirement when there are tax-benefits to housing. This demonstrates that although tax-benefits to housing do result in distortion, as households accumulate more housing and less liquid assets, these subsidies also increase overall wealth accumulation by making the commitment device more affordable.²⁷

7.2 Mandatory Amortization

There is a growing debate about the effect of mandatory amortization policies on household spending and wealth accumulation. While mandatory amortization requirements have recently been introduced by Sweden and the Netherlands, there exists concern that these policies may simply result in a re-balancing of household portfolios away from liquid assets and into illiquid housing, potentially resulting in reduced resilience to income shocks (Svensson, 2019, 2020).²⁸ But if mortgage payments serve as a commitment device, then amortization might not result in simple portfolio re-balancing, but may also help households accumulate wealth.

We use our calibrated model to study the effect of mandatory amortization policies on household wealth accumulation and to quantify the role of commitment. To do so, we consider two different policy regimes. In the baseline, we assume that all mortgages are fully-amortizing, thus mortgagors are required to gradually pay back their mortgage using a periodic debt repayment plan that ends at the time of retirement. In the counterfactual, we assume that all mortgages are interest-only (IO), with a balloon payment covering the mortgage principal at the time of retirement. The latter option implies that homeowners

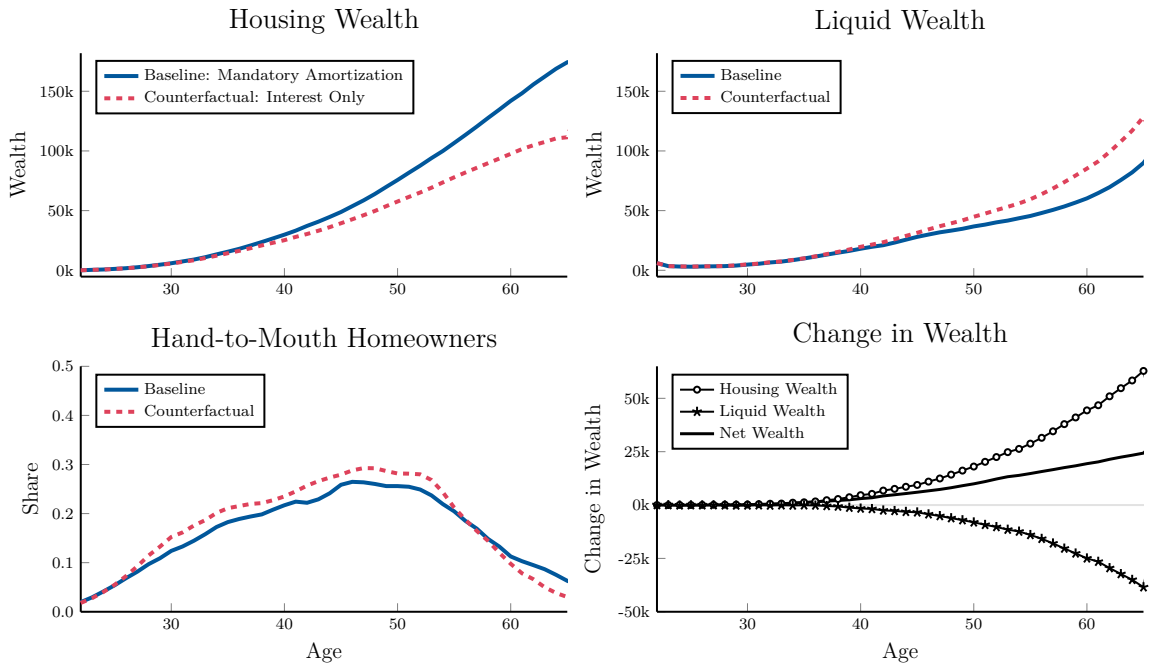
²⁷In future work, it may be interesting to evaluate other ways that governments can make commitment devices more accessible to households. For instance, first-time homebuyer tax credits targeted towards low-income households may be an attractive way to help households onto the housing ladder (giving the benefit of commitment) without the regressive tax implications of the mortgage interest tax deduction. It would be straightforward to extend our model to study such a policy.

²⁸While some countries have explicit mandatory amortization policies, others have introduced similar but less explicit policies. For instance, the United States has created strong incentives against interest-only mortgages via the “qualified mortgage” provision in the Dodd-Frank Act. As a result, the vast majority of mortgages in the U.S. require regular amortization payments.

will only be required to make small interest payments on their mortgage during the course of their working-life.²⁹ This may potentially increase wealth accumulation, as it allows households to invest in the high-return liquid asset using the money that they otherwise would have been required to accumulate in the form of home equity. However, interest-only mortgages may instead decrease wealth accumulation if amortization payments serve as a savings commitment device.

Figure 7 shows the life-cycle profiles of wealth accumulation in the baseline model with mandatory amortization (the solid blue line) relative to the counterfactual model with interest-only mortgages (the dashed pink line). Overall, we find evidence of portfolio re-balancing from liquid to illiquid assets, but this is not a one-to-one effect, therefore overall wealth accumulation is increased.

Figure 7: Effect of Mandatory Amortization



The first two panels of Figure 7 show the life-cycle profiles of net housing wealth and liquid wealth in the two different models. Mandatory amortization increases housing wealth accumulation at the cost of slightly reduced liquid wealth accumulation. Overall, we find that every additional dollar of housing wealth reduces liquid assets by roughly 60 cents by the time of retirement. The third panel of Figure 7 shows the share of the wealthy hand-to-mouth in the two different models. We see that the share of the wealthy hand-to-mouth is very similar across the two different models, and in fact is very slightly higher in the model with interest-only mortgages. This suggests that mandatory amortization policies are not an important determinant of hand-to-mouth behavior. Finally, the

²⁹More specifically, the mortgage payment given by equation (7) is modified to be $mp_t = m_t r^M$ during the working life and $mp_t = m_t$ at the time of retirement.

fourth panel shows the difference in wealth accumulation in the model with mandatory amortization relative to the model with interest-only mortgages. We find that mandatory amortization increases net wealth by around 10% by the time of retirement, relative to the model with interest-only mortgages.

The reason that wealth accumulation is increased is because mortgage amortization serves as a form of commitment. By purchasing a home using an amortizing mortgage, households not only lock-away their current wealth, but they also commit to accumulate wealth in the future in the form of home equity. This finding is consistent with recent empirical evidence from the Netherlands, showing that mandatory amortization results in increased wealth accumulation, as households accumulate more home equity due to this policy, but do not reduce their non-housing assets (Bernstein and Koudijs, 2020).

More specifically, Bernstein and Koudijs (2020) study the effect of mortgage amortization on wealth accumulation by exploiting a macroprudential policy change in the Netherlands that restricted the availability of interest-only mortgages for first-time home buyers starting in January 2013. These authors find that mandatory amortization results in a near 1-for-1 rise in net worth, with no observable change in non-housing savings, at least in the four years for which data is available. Our results are consistent with their findings. Moreover, our calibrated model allows us to evaluate the potential long-term effects of such a policy on household wealth accumulation. We find that the effect on wealth accumulation builds over time, with the largest increase in wealth occurring at age 65, when households in our model finish paying off their mortgage.³⁰

7.3 Housing Subsidies, Mandatory Amortization, and MPCs

Finally, we analyze the effect of these two policies on households' responsiveness to income shocks. This is an important consideration, as there is concern that these policies may reduce households' resilience to income shocks by encouraging substitution from liquid to illiquid forms of wealth accumulation. For instance, the previously mentioned work by Svensson (2019, 2020) argues that this concern may substantially undermine the macroprudential motives for which mandatory amortization policies have been implemented.

To evaluate the effect of housing subsidies and mandatory amortization on household sensitivity to income shocks, we compute the average MPC in response to a transitory income shock under three different policy regimes. Table 4 presents our results. The first row presents the average annual MPC in the baseline model. The second row presents the average MPC in the counterfactual model without housing subsidies. We find an average

³⁰In this paper, we do not explore the welfare implications of mandatory amortization. That said, the presence of temptation generates an interesting trade-off between flexibility and commitment, which is absent in a traditional model. Kovacs and Moran (2020) study the implications of this trade-off as it relates to financial liberalization which has given households greater access to home equity withdrawal.

MPC of 0.42 in response to a \$1,000 income shock, only 1 percentage point lower than in the baseline model with large and realistic housing subsidies as in the United States. The difference is even smaller in response to larger income shocks. These results are in line with the previous finding that housing subsidies are not an important determinant of wealthy hand-to-mouth behavior and therefore do not have an economically meaningful effect on the average MPC.

Table 4: Average MPC under Different Policy Regimes

	SHOCK SIZE		
	\$1,000	\$5,000	\$10,000
MPC in Baseline Model	0.43	0.37	0.29
MPC in Model without Housing Subsidies	0.42	0.37	0.29
MPC in Model without Mandatory Amortization	0.41	0.36	0.27

Note: Each coefficient represents the average annual MPC in our calibrated model. The procedure used to compute MPCs is described in Appendix A.5.1.

The final row of Table 4 presents the average MPC in the counterfactual model without mandatory amortization. We find that the average MPC is slightly lower without mandatory amortization, however, the difference is small and therefore unlikely to be economically meaningful. For instance, in response to a \$1,000 income shock, the MPC is 0.43 in the baseline model versus 0.41 in the counterfactual model without mandatory amortization.

Why do housing subsidies and mandatory amortization have such a small effect on the average MPC? The economic intuition for this result is that households' desire for illiquidity does not come simply from a desire to exploit better interest rates on different types of assets, but rather from a fundamental difficulty to accumulate wealth in liquid form, which thus generates a desire for illiquidity. For policy makers looking to reduce households' responsiveness to income shocks, the answer is not the elimination of housing subsidies and mandatory amortization. Instead, policy makers should focus on policies that provide temporary liquidity relief to households who have experienced income loss, such as the mortgage forbearance policy studied by Boar, Gorea, and Midrigan (2020).

8 Conclusion

A growing body of experimental evidence shows that households suffer from present bias that makes it difficult to save. At the same time, policy makers around the world have implemented a large number of incentives intended to encourage wealth accumulation in illiquid assets. And yet, the decision to incentivize *illiquid* rather than *liquid* wealth accumulation has remained controversial, due to concerns that these policies may sim-

ply encourage portfolio reallocation, without actually generating an increase in overall saving. Further, some economists have suggested that these policies may be actively detrimental if they increase the share of wealthy hand-to-mouth households. This raises the question: can the government actually help households save by incentivizing *illiquid* wealth accumulation? And do such policies harm macroeconomic stability by increasing the sensitivity of consumption to income shocks?

In this paper, we answer these questions using a rich model of individual preferences with liquid and illiquid assets. Our contribution is related to two important studies: Krusell, Kuruscu, and Smith (2010) show that saving subsidies may be beneficial due to temptation in a one-asset framework, while Kaplan and Violante (2014) show the importance of illiquid assets in explaining household consumption and saving behavior. We bring together these two different strands of literature by incorporating temptation into a two-asset model of wealthy hand-to-mouth behavior, which we then bring to the data. This allows us to carefully evaluate the effect of housing subsidies and mandatory amortization on household portfolio allocation and wealth accumulation.

There are a number of potential extensions to our analysis that may be interesting and that would be straightforward to implement within our framework. For instance, we could evaluate the effect of downpayment assistance policies, such as the \$15,000 first-time homebuyer tax credit recently proposed by President Biden. We expect that downpayment assistance may be more effective than the mortgage interest tax deduction, as it is targeted directly towards the extensive margin of adjustment and may thus help more households obtain access to commitment. In addition, while we have focused on the positive effects of policies intended to encourage illiquid wealth accumulation (given the substantial controversy in the literature about the effect of these policies on household portfolio decisions), it would be straightforward to use our framework to assess the normative effects of these policies. As we find that these policies increase savings without significantly distorting households' ability to smooth consumption, we expect that the welfare effect of these policies would be positive and substantial.

A Appendix

A.1 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
Utility Parameters		
λ	temptation parameter	0.28
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	loan-to-value constraint	0.90
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: <i>age</i>	0.256
g_2	deterministic process: <i>age</i> ²	-0.045
g_3	deterministic process: <i>age</i> ³	0.002
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.2 Asset Returns

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

$$r_{t+1} = \frac{p_{t+1} + d_{t+1} - p_t}{p_t} \quad (\text{A.1})$$

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 home ownership over the same period. More explicitly, we can write the return on housing similarly to equation (A.1) as

$$r_{t+1}^h = \frac{p_{t+1} + s_{t+1} - c_{t+1}^m - c_{t+1}^i - p_t}{p_t} \quad (\text{A.2})$$

with 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, in the rest of the paper we use constant fractions of changing house value in order to calculate these variables. Under these conditions equation (A.2) can be rewritten as

$$r_{t+1}^h = \frac{p_{t+1}^h + (s - c^m - c^i - 1)p_t^h}{p_t^h} \quad (\text{A.3})$$

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

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

³²Gross 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.

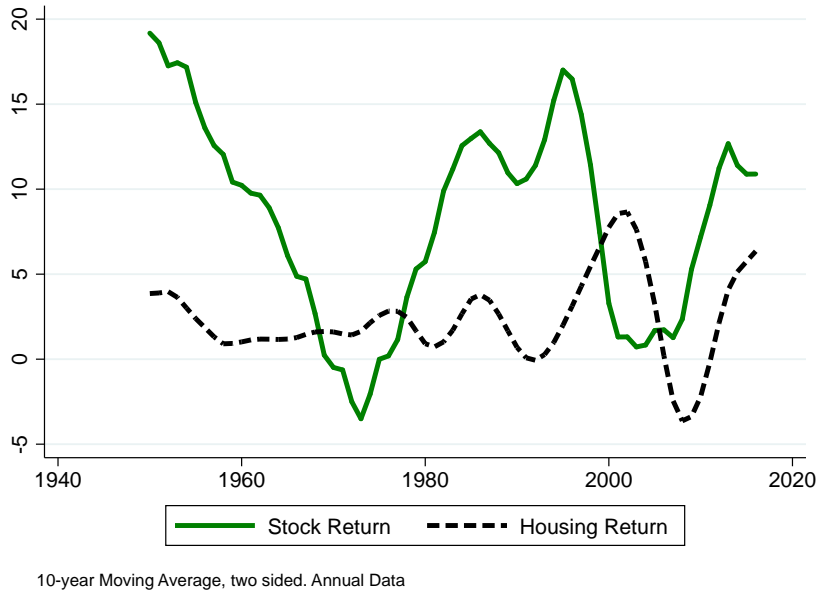


Figure A.1: Real Returns

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 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 seventies and a couple of years in the early twenties 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.

$$r_{adj}^i = E(r^i) - var(r^i) \quad (\text{A.4})$$

where superscript i refers to the type of the asset, i.e. 3 Months T-Bill³³, S&P500 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, while 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 re-

³³The 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.

turns. 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 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.

A.2.1 Sensitivity to Imputed Rents

We compute imputed rents using the balance sheet approach described above, 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 United States.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 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) 5.80%. This result is driven by the measure of imputed rents

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

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 United States 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.6.1 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.3 Data: Sample and Definitions

We calibrate our model using the Panel Study of Income Dynamics 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 to the best of our knowledge is the only large scale U.S. panel to contain information on income, consumption, and wealth accumulation.

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, we drop the top 5% of households by income due to evidence that the PSID performs well at measuring the wealth holdings of the bottom 95%, but not the top 5% of the distribution (Bosworth, Anders et al., 2008; Pfeffer et al., 2016).³⁵

Consumption. 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 percent 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. We compute liquid assets as the sum of bank account deposits and directly held stock. We believe that publicly traded stock are essentially liquid because there are very low transaction costs on these assets, thus it would be easy for households to sell their stock position for consumption smoothing purposes. We measure net housing wealth as the reported value of the household’s main residency minus all mortgage debt on this home. We exclude net wealth from other real estate, net business wealth, and IRA/annuity wealth, as we want to focus our analysis on owner-occupied housing. We also exclude credit cards debt from our analysis, as the PSID did not collect information on this variable for the full period of our sample.

Movers. We target the average share of homeowners that move for non-work reasons each year. This is computed using detailed information on the timing and reasons for moving homes, collected by the PSID. We include all moves due to consumptive reasons (expansion of housing, contraction of housing, better neighborhood, better school, etc.) and ambiguous reasons (desire to save money, neighbors moving away, unknown). We exclude moves for work reasons (to take another job, to get nearer to work) since these are outside the scope of our model.

A.4 Model Fit

Table A.4 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.

³⁵As a result, the share of the wealthy hand-to-mouth in our sample is slightly higher than in the Survey of Consumer Finances. In our sample, the share of such households reach a peak of 25% at age 45, whereas in the Survey of Consumer Finances this number is 22% at age 45.

Table A.4: Additional Targeted Moments

	PSID	Temptation Model	No Temptation Model
Homeownership rate	0.57 (0.006)	0.50	0.53
Share of movers	0.068 (0.004)	0.070	0.056

Note: This table shows fit of the average homeownership rate and share of movers each period. The average homeownership rate in our sample is 57%, slightly lower than the homeownership rate observed in the SCF. This is driven by two factors. First, our sample includes only working age households, who have less homeownership than retirees. Second, the PSID oversamples low income households.

A.5 Consistency with Empirical Evidence on Consumption

A.5.1 Calculation of the Marginal Propensity to Consume

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:

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

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

A.5.2 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):

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

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.5. 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 0.18 lower. This is even more pronounced in the third and fourth quartiles. In short, low liquidity is an important determinant of high MPCs.

Table A.5: MPC Heterogeneity by Household Type

		MPC	
		Coefficient	Standard Error
CASH-ON-HAND			
Quartile			
	2 nd	−0.181***	(0.006)
	3 rd	−0.368***	(0.006)
	4 th	−0.595***	(0.007)
NET WEALTH			
Quartile			
	2 nd	−0.062***	(0.005)
	3 rd	−0.093	(0.007)
	4 th	−0.055***	(0.008)
Constant		0.821***	(0.012)

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

In contrast, once we have controlled for cash-on-hand, net wealth has very little effect on the average MPC. The average MPC is only 0.06 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.

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 average 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 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 Aiyagari 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.5.3 MPC Decomposition

As mentioned in Section 6.2, large MPCs may be driven by either illiquidity or the mechanical effect of temptation. In this subsection, we decompose the relative importance of these two different channels. To do so, we use the following procedure. For each household $i \in [1, 2, \dots, N]$ and each period of working life $t \in [1, 2, \dots, W]$, 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.³⁸ This allows us to isolate

³⁶Jappelli 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.

³⁷Fagereng, 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.


³⁸In 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.

the mechanical effect of temptation.

$$\widetilde{\text{MPC}}(x) = \frac{1}{N} \frac{1}{W} \sum_{i=1}^N \sum_{t=1}^W \frac{c_{i,t}(\text{shock} = x, \lambda = 0, \Omega_{i,t}) - c_{i,t}(\text{shock} = 0, \lambda = 0, \Omega_{i,t})}{x} \quad (\text{A.8})$$

Figure A.2 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.

Figure A.2: 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 ($\text{MPC}(x) - \widetilde{\text{MPC}}(x)$).

Why does illiquidity have a large effect even when households receive an income shock 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, there are substantial adjustment costs to housing, both financial and nonfinancial. 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.

In contrast, Kaplan and Violante (2014) find a very small MPC (almost zero) in response to large income shocks. This is because their model has a small, additive adjustment cost of \$1,000. Moreover, if they were to include a larger adjustment cost, the wealthy hand-to-mouth would disappear. An important benefit of temptation is that it is able to accommodate a large and realistic adjustment cost,³⁹ yet still explain the presence of wealthy hand-to-mouth households. 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.

A.5.4 Targeted Fiscal Stimulus

We use our estimated model to study the design of targeted fiscal stimulus policies. In particular, we evaluate the consumption response to alternative stimulus policies that give one time unanticipated stimulus payments to varying fractions of the population. We focus on budget equivalent policies where the government uses income based targeting to stimulate aggregate consumption.⁴⁰ For instance, the government may give \$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.⁴¹ We then report the fraction of aggregate stimulus that is consumed within one year after disbursement.

Figure A.3 shows the aggregate one year consumption response 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, just under 40% of stimulus payments are consumed within the year of disbursement. 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 optimum, when \$2,500 is given to each household in the bottom 20% of the distribution, we observe that 68%

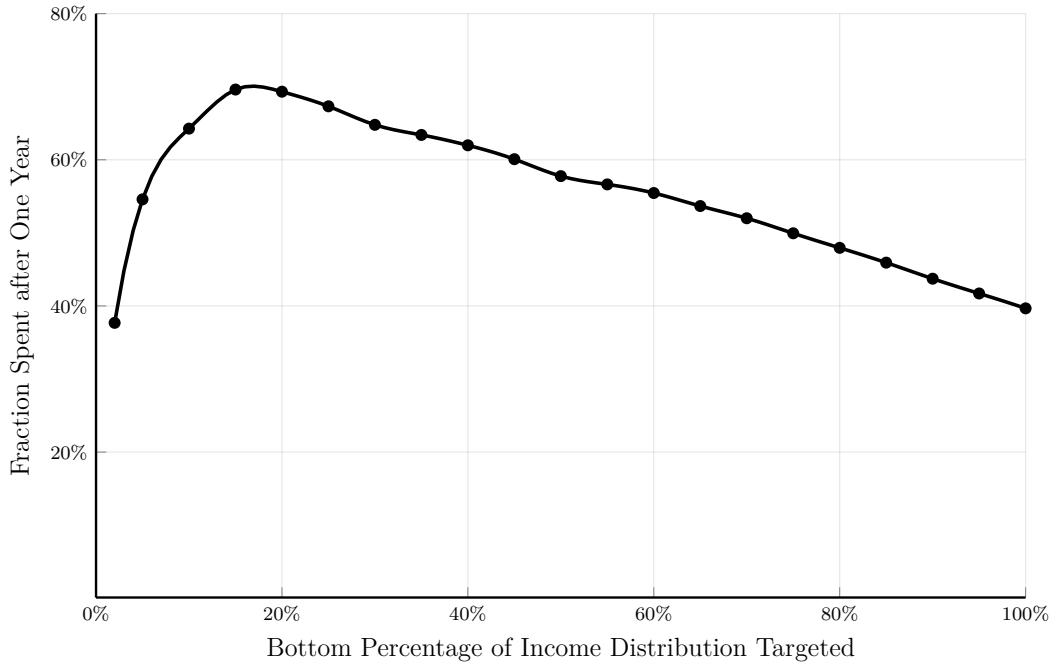
³⁹In our case, 5% of the value of the home (OECD, 2011), as well as the utility cost of moving.

⁴⁰We 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.

⁴¹We choose $N = 1,000$ for each time period, therefore there are 44,000 households in our simulation.

of aggregate stimulus is consumed within one year.

Figure A.3: Income Targeted Stimulus Payments



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 U.S. government gave tax rebates to approximately 80-85% of households, with an average stimulus payment of \$600-\$1,200.

Targeted fiscal stimulus allows the government to reach households with higher MPCs, but there exists a trade-off, as larger stimulus payments induce households to save a larger fraction of their income in either housing or liquid assets. As a result, the consumption response observed in Figure 4 declines when the government targets households in the very bottom of the income distribution. For instance, in response to a stimulus payment of \$25,000, approximately 29% is saved in housing wealth. Nevertheless, a very large stimulus payment is needed to convince households to increase their investment in housing, due to the presence of sizable housing transaction costs.

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 consumption response when stimulus payments are given to the entire population, they find smaller gains to targeted stimulus payments, and their optimal policy is to target the bottom half of the income distribution.⁴² This difference is driven by the above trade-off

⁴²In their model, stimulus payments given to the entire population would imply that roughly 40% is

between targeting households with high MPCs and giving larger payments. Their model requires very small transaction costs (\$1,000 in their preferred calibration) in order 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 realistic housing transaction cost of 5% of the value of the home, as well as a utility cost χ 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.

A.6 Sensitivity

In this section we evaluate the sensitivity of our results presented in Sections 4 and 6. We consider two modifications to our baseline model. First, we change the return structure in the model by assuming that housing and equities provide equivalent risk-adjusted returns ($r = r^H$). 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.6.1 Asset Returns

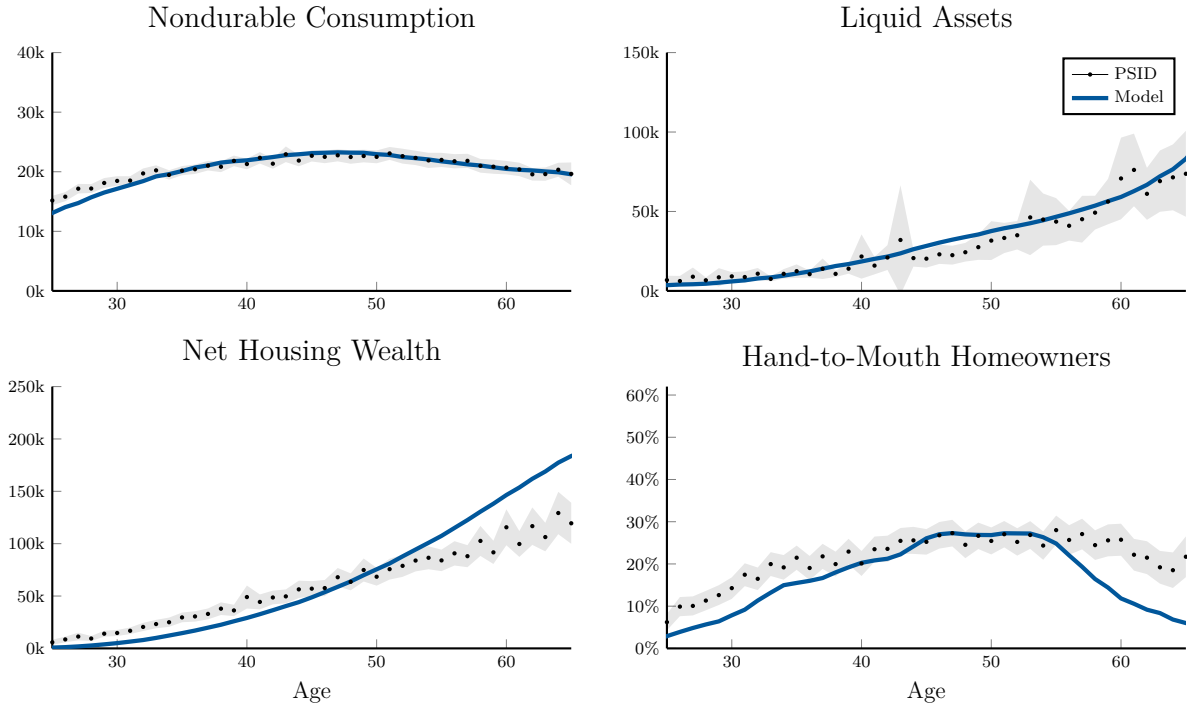
We consider a model where the two assets give equivalent risk-adjusted returns ($r = r^H$). It is important to note that this return structure would pose a challenge for the traditional returns-based explanation of wealthy hand-to-mouth behavior. If liquid and illiquid assets delivered equivalent returns 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. Figure A.4 shows the model fit when we re-calibrate the preference parameters of the model. We find that the model with equal returns obtains a good fit of the large share of wealthy hand-to-mouth households, as well as the other targeted moments, despite the fact that housing and liquid assets deliver equivalent returns.

The calibrated parameter results are reported in Table A.6. When the asset returns are equal (column 2), we find that that households are more patient and more risk averse

consumed within one year, while stimulus payments given to the bottom half would imply that 55% is consumed within one year. In contrast, we find that 70% is consumed within one year if stimulus payments are targeted towards the bottom 20% of the income distribution.

Figure A.4: Model Fit with Equal Returns



Note: This figure shows the life-cycle moments from our model with equal returns (solid blue line) and the PSID (black dotted line). The moments from the PSID are shown with bands of 1.96 standard deviations around the mean.

relative to the baseline calibration (column 1). More specifically, 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.7 shows the MPC by shock size. The average MPC in response to a \$1,000 shock is 0.40 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 0.27, only slightly lower than our baseline result of 0.29.

A.6.2 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

Table A.6: Model Parameters – Alternative Modeling Assumptions

PARAMETER	Baseline	Equal Returns ($r = r^H$)	Retirement Accounts
β	0.959	0.970	0.955
γ	2.415	2.795	2.596
μ	0.291	0.317	0.455
θ	0.160	0.267	0.295
χ	0.437	0.552	0.869

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.6.2. The temptation parameter is set at $\lambda = 0.28$ following Kovacs, Low, and Moran (2020).

Table A.7: MPC Heterogeneity – Alternative Modeling Assumptions

MPC	SHOCK SIZE		
	\$1,000	\$5,000	\$10,000
MPC in Baseline Model	0.43	0.37	0.29
MPC in Model with Equal Returns	0.40	0.33	0.27
MPC in Model with Retirement Accounts	0.43	0.38	0.28

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.6.2.

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:

$$y_{i,t}^{IRA} = \eta_1 * \text{IRA}(y_{i,W}) \quad (\text{A.9})$$

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:

$$\log[\text{IRA}(y_{i,W})] = \eta_2 * \log(y_{i,W}) \quad (\text{A.10})$$

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:

$$\eta_1 = \left[\sum_{t=W}^T \frac{s_t}{(1+r)^{t-W}} \right]^{-1} \quad (\text{A.11})$$

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.6 shows the parameter results. We find that the impatience parameter (β) and the risk aversion (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, κ).

The third row of Table A.7 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 0.28 in response to a \$10,000 shock, only slightly smaller than the MPC of 0.29 in our baseline model.

References

- Aguiar, Mark A., Mark Bilts, 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.
- 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.
- Attanasio, Orazio P. and Thomas DeLeire. 2002. "The Effect Of Individual Retirement Accounts On Household Consumption And National Saving." *Economic Journal* 112 (6):504–538.
- Bäckman, Claes and Natalia Khorunzhina. 2020. "Interest-only mortgages and consumption growth: Evidence from a mortgage market reform." *Available at SSRN 3533247* .
- 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.
- Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. "Consumption Inequality and Family Labor Supply." *American Economic Review* 106 (2):387–435.

- Boar, Corina, Denis Gorea, and Virgiliu Midrigan. 2020. “Liquidity constraints in the US housing market.”
- Bosworth, Barry, Sarah Anders et al. 2008. “Saving and Wealth Accumulation in the PSID, 1984-2005.” .
- 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.
- Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. 2017. “The Distribution of Wealth and the Marginal Propensity to Consume.” *Forthcoming, Quantitative Economics* .
- Chambers, Matthew, Carlos Garriga, and Don E Schlagenhauf. 2009. “Housing policy and the progressivity of income taxation.” *Journal of Monetary Economics* 56 (8):1116–1134.
- 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.
- Choukhmane, Taha. 2019. “Default Options and Retirement Saving Dynamics.” *mimeo* .
- DeFusco, Anthony A, Stephanie Johnson, and John Mondragon. 2020. “Regulating household leverage.” *The Review of Economic Studies* 87 (2):914–958.
- Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik. 2021. “MPC Heterogeneity and Household Balance Sheets.” *American Economic Journal: Macroeconomics* .
- 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.
- Floetotto, Max, Michael Kirker, and Johannes Stroebel. 2016. “Government intervention in the housing market: Who wins, who loses?” *Journal of Monetary Economics* 80:106–123.
- Fuster, Andreas, Greg Kaplan, and Basit Zafar. 2018. “What Would You Do with \$500? Spending Responses to Gains, Losses, News and Loans.” *NBER Working Paper* (24386).

- Gale, William G and John Karl Scholz. 1994. “IRAs and household saving.” *The American Economic Review* :1233–1260.
- Gervais, Martin. 2002. “Housing taxation and capital accumulation.” *Journal of Monetary Economics* 49 (7):1461–1489.
- 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 Pendorfer. 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 Fund.” *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 and Giovanni L. Violante. 2014. “A Model of the Consumption Response to Fiscal Stimulus Payments.” *Econometrica* 82 (4):1199–1239.
- 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. 2020. "Estimating Temptation and Commitment Over the Life-Cycle." *International Economic Review* forthcoming.
- Kovacs, Agnes and Patrick Moran. 2018. "Temptation, Commitment, and Hand-to-Mouth Consumers." *mimeo* .
- . 2020. "Breaking the Commitment Device: The Effect of Home Equity Withdrawal on Consumption, Saving, and Welfare." .
- 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.
- 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.
- . 2020. "Capital income taxation with housing." *Journal of Economic Dynamics and Control* 115:103883. URL <http://www.sciencedirect.com/science/article/pii/S016518892030052X>.
- 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* .
- 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 25MillionExperiment." *American Economic Journal: Macroeconomics* 9 (4) : 153 – –83.

- Pfeffer, Fabian, Robert F. Schoeni, Arthur Kinnickell, and Patricia Anderski. 2016. “Measuring Wealth and Wealth Inequality.” *Journal of Economic and Social Measurement* 41 (2):103–120.
- 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.
- Poterba, James M. 1992. “Taxation and housing: Old questions, new answers.” Tech. rep., National Bureau of Economic Research.
- Rosen, Harvey S. 1985. “Housing subsidies: Effects on housing decisions, efficiency, and equity.” In *Handbook of public economics*, vol. 1. Elsevier, 375–420.
- 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.
- Sommer, Kamila and Paul Sullivan. 2018. “Implications of US Tax Policy for House Prices, Rents, and Homeownership.” *American Economic Review* 108 (2):241–74. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20141751>.
- Strotz, Robert H. 1956. “Myopia and Inconsistency in Dynamic Utility Maximization.” *The Review of Economic Studies* 23:165–180.
- Svensson, Lars EO. 2019. “Amortization requirements, distortions, and household resilience: Problems of macroprudential policy II.” Tech. rep., Working Paper, Stockholm School of Economics.
- . 2020. “Macroprudential policy and household debt: What is wrong with Swedish macroprudential policy?” .
- Thaler, Richard H. and Shlomo Benartzi. 2004. “Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving.” *Journal of Political Economy* 102:S164–S187.

- Toussaert, Séverine. 2018. “Eliciting Temptation and Self-Control Through Menu Choices: A Lab Experiment.” *Econometrica* 86 (3):859–889.
- Venti, Steven F and David A Wise. 1986. “Tax-deferred accounts, constrained choice and estimation of individual saving.” *The Review of Economic Studies* 53 (4):579–601.
- . 1987. “IRAs and saving.” In *The effects of taxation on capital accumulation*. University of Chicago Press, 7–52.
- Wiedemer, J.P. and J.K. Baker. 2012. *Real Estate Finance, 9th ed.* Cengage Learning.