

Kovacs, Agnes; Moran, Patrick

Working Paper

Temptation and commitment: Understanding the demand for illiquidity

IFS Working Papers, No. W19/18

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Kovacs, Agnes; Moran, Patrick (2019) : Temptation and commitment: Understanding the demand for illiquidity, IFS Working Papers, No. W19/18, Institute for Fiscal Studies (IFS), London,
<https://doi.org/10.1920/wp.ifs.2019.1819>

This Version is available at:

<https://hdl.handle.net/10419/223262>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Temptation and Commitment: Understanding the Demand for Illiquidity

IFS Working Paper W19/18

Agnes Kovacs
Patrick Moran

Temptation and Commitment: Understanding the Demand for Illiquidity^{*}

Agnes Kovacs[†]

Patrick Moran[‡]

June 30, 2019

Abstract

The vast majority of household wealth in the U.S. is held in illiquid assets, primarily housing, making households vulnerable to unexpected income shocks. To rationalize this preference for illiquidity, we build a life-cycle model where households are tempted to consume their liquid wealth but can use illiquid housing as a savings commitment device. The importance of temptation and commitment is identified using data on consumption, liquid assets, and housing wealth over the life-cycle. Our model matches observed portfolio choices and gives rise to a high demand for illiquid housing partially driven by the need for commitment. Preference for illiquidity has important implications for the consumption response to unexpected income shocks. Our model is able to replicate the recent empirical evidence that MPCs remain high in response to large shocks, a finding that cannot be explained by current heterogeneous agent models, but that has great significance for fiscal stimulus targeting.

Keywords: commitment; housing; liquidity; consumption

JEL classification: D11; D14; D91; E21; R21

^{*}We are grateful to Orazio Attanasio, Taha Choukhmane, Andreas Fagereng, Andrea Ferrero, Mike Keane, Hamish Low, Costas Meghir, Peter Neary, Cormac O'Dea, Carlo Pizinelli, Akos Valentinyi, Gianluca Violante, and seminar participants at Oxford, Yale, 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 University of Manchester and Institute for Fiscal Studies (agnes.kovacs@manchester.ac.uk)

[‡]University of Oxford and Institute for Fiscal Studies (patrick.donnellymoran@economics.ox.ac.uk)

1 Introduction

A large fraction of households hold almost zero liquid assets, despite owning substantial wealth in illiquid housing. In the United States, roughly 70 percent of households own homes, while 30 percent of homeowners hold essentially zero liquid assets, according to the Panel Study of Income Dynamics (PSID). Households’ overwhelming preference for illiquidity is puzzling for at least two reasons. First, by concentrating their wealth in housing, households limit their ability to respond to adverse income shocks. Second, the desire for illiquidity arises even though housing has a lower risk-adjusted return than liquid assets such as equities.

In this paper, we study households’ observed preference for illiquidity, evaluating competing theories as to why households keep the vast majority of their wealth in housing. The housing literature traditionally focuses on three distinct roles that drive demand for housing: the investment role, the collateral role, and the utility role.¹ In contrast, an alternative strand of literature suggests a fourth option: the role of housing as a savings commitment device. According to this theory, households may have difficulty saving in liquid assets due to issues of self-control, but can alleviate this problem by “locking away” their wealth in illiquid housing. The main contribution of this paper is to evaluate the importance of this theory by estimating a structural life-cycle model of consumption, housing, and wealth accumulation that includes these four potential drivers of housing demand. We find that it is not possible to rationalize the large share of homeowners holding almost zero liquid assets using only the three traditional roles of housing. Instead, households have a demand for illiquidity that stems from their desire for commitment. We find that the alternative view of illiquidity that we highlight is consistent with recent empirical evidence on the consumption response to large income shocks, evidence that is difficult to rationalize with a traditional model, but that has important implications for the optimal design of targeted fiscal stimulus.

Understanding the preference for illiquidity is interesting not only from a theoretical point of view, but also for answering longstanding macroeconomic questions on fiscal stimulus, stabilization, and redistribution. There exists a growing literature in macroeconomics that argues that “asset rich but cash poor” households are important in explaining aggregate consumption behavior in heterogeneous agent models.² In these models, large consumption fluctuations are driven by households whose wealth is locked in illiquid assets that are difficult to access for consumption smoothing purposes. As we present a new model to explain why households keep the vast majority of their wealth in housing, we evaluate whether the alternative view of illiquidity that we highlight generates different

¹See for example Cocco (2005), Fernandez-Villaverde and Krueger (2011), Kaplan and Violante (2014), or recently Gorea and Midrigan (2017) and Sommer and Sullivan (2018).

²See for instance Kaplan and Violante (2014), Carroll et al. (2017), Kaplan, Moll, and Violante (2018), Auclert, Rognlie, and Straub (2018), Luetticke (2018), and Bayer et al. (2019).

implications for consumption behavior relative to a traditional model. We find that our model is able to match the recent empirical evidence that the marginal propensity to consume (MPC) remains high even for households that receive large transitory income shocks, a finding that cannot be explained by a traditional heterogeneous agent model. Moreover, our model matches the widely recognized empirical evidence that MPCs are highest for households with low liquid assets. Taken together, these two findings suggest that large and targeted fiscal stimulus payments may be more effective in boosting aggregate consumption than previously believed, a finding that we confirm by simulating a series of budget-equivalent fiscal stimulus policies using our estimated model.

We begin by proposing a structural model of consumption, housing, and wealth accumulation. We allow households to invest in two instruments: liquid assets and illiquid housing, where housing provides direct utility, serves as collateral for mortgages, and provides tax advantages. In addition, households are able to borrow using fixed-rate, fully-amortizing mortgages. The most novel feature of our model is that we allow for the possibility that households may be tempted to consume their liquid assets, following Gul and Pesendorfer (2001), but that they can mitigate this problem by purchasing illiquid housing as a savings commitment device. Temptation represents the idea that households face self-control problems that make it difficult to achieve their optimal savings plan, due to the possibility of instantaneous gratification that is hard to resist. Using a simplified version of our model, we demonstrate that this implies a desire for illiquidity, as households can reduce temptation by locking away their wealth in housing. As a result, housing acts as a commitment device that helps households accumulate wealth.

The key challenge of our analysis is to identify the importance of temptation and commitment relative to the other potential factors that could drive households to keep the vast majority of their wealth in illiquid form. While it is easy to quantify the financial returns to housing using historical data, it is much more difficult to measure the utility and commitment benefits separately. For this reason, we structurally estimate key parameters of our model related to housing utility and temptation costs, as well as time-preferences. We do so by matching the mean life-cycle profiles of consumption, liquid assets, net housing wealth, and share of homeowners with zero liquid assets.³ We find that temptation and commitment are important in explaining the observed wealth accumulation patterns of U.S. households.⁴ In our model, identification comes from the ability of the temptation model to match these four life-cycle profiles in a way that cannot

³More specifically, estimation is performed using the Method of Simulated Moments, with data from the Panel Study of Income Dynamics (PSID) between 1999 and 2015.

⁴This finding confirms the importance of a growing literature that studies the interaction between public policy and welfare when households suffer from problems of self-control. Some nice examples include Krusell, Kuruscu, and Smith (2010) who study optimal savings subsidies, Nakajima (2012) who looks at the welfare effects of credit cards, and Schlafmann (2016) who studies the consequences of mortgage market regulation. Due to the difficulty of measuring self-control problems, all of these papers calibrate the importance of temptation.

be obtained by the standard model (where we turn off temptation) using any combination of the other parameters. Especially important for the identification of temptation is the share of homeowners with little or no liquid assets over the life-cycle. We find that a standard model with high housing utility is unable to match the large share of these households, as it creates a stronger precautionary savings motive for homeowners, who become more averse to losing their home and therefore store large liquid buffers.

This paper contributes to a growing literature that tries to understand households' observed preference for illiquidity. We build upon the idea that self-control problems may make it difficult to accumulate liquid wealth, therefore households keep the vast majority of their wealth in illiquid form (Laibson (1997), Harris and Laibson (2001), and Angeletos et al. (2001)). While there exist many papers that estimate the relevance of self-control problems in a controlled laboratory environment,⁵ there are very few that attempt to estimate this using observed life-cycle patterns of portfolio choice. Most notably, Laibson et al. (2017) estimates the importance of self-control using a structural life-cycle model with hyperbolic discounting. Our analysis differs from theirs in that we identify the importance of self-control problems and commitment relative to the other potential factors that could drive households to keep the vast majority of their wealth in illiquid assets. Most importantly, we jointly estimate the strength of self-control and housing utility, whereas Laibson et al. (2017) choose to calibrate housing utility externally. It is important to estimate housing utility and self-control jointly because either parameter could explain households' decision to keep the vast majority of their wealth in housing.

In contrast, the macroeconomic literature usually explains households' preference for illiquidity by relying upon the assumption that illiquid assets deliver excess returns relative to all available liquid assets. Most prominently, Kaplan and Violante (2014) assume a 3.7% gap in returns between illiquid and liquid assets. In contrast, there exists a wide body of literature showing that the most prevalent form of illiquid assets, housing, delivers worse risk adjusted returns than liquid assets such as equities, even when accounting for imputed rents and other benefits to homeownership (Flavin and Yamashita (2002), Goetzmann and Spiegel (2002), and Piazzesi, Schneider, and Tuzel (2007)). We therefore set ourselves the challenge of explaining why so many homeowners hold zero liquid assets, despite the fact that housing has lower risk-adjusted returns than equities.⁶ More specifically, we calculate the risk-adjusted returns to housing and stocks during the past 65 years and use these returns when calibrating our model. We find that our model with temptation and commitment can match the targeted life-cycle moments while relaxing

⁵See Toussaert (2018) or Houser et al. (2018) for example.

⁶Put differently, we are investigating whether housing utility has the same effects as financial returns, as we allow for the possibility that housing delivers substantial utility benefits. We find that high housing utility is not able to generate a large fraction of homeowners with no liquid wealth, as it makes households more averse to losing their home, therefore they keep a buffer stock of liquid assets. In other words, there exists a "housing smoothing" motive in addition to the standard "consumption smoothing" motive.

the assumption in Kaplan and Violante (2014) that illiquid assets deliver higher returns than all available liquid assets.

After establishing that temptation is crucial to successfully match the observed data from the PSID, we study the implications for consumption behavior. More specifically, we look at the marginal propensity to consume in response to exogenous and transitory income shocks. We compare our model results to the empirical MPC literature, not only to validate the model’s predictive power, but also to assess whether the alternative view of illiquidity that we highlight is important for understanding consumption dynamics. We find that our model obtains a good match of (1) the average annual MPC, (2) the slow decline of MPCs by wealth, and (3) the slow decline of MPCs by shock size. While there exists a wide variety of traditional heterogeneous agent models that are able to match finding (1), Kaplan and Violante (2014) was the first to match findings (1) and (2). In turn, our framework is the first to match findings (1), (2), and (3). We find that this has important implications for the optimal design of fiscal stimulus targeting, as large stimulus payments still induce a large consumption response.

Our model is consistent with a recent empirical literature that finds that MPCs decline slowly with shock size. More specifically, our model predicts that households that receive a shock of \$1,000 have an annual MPC of 0.51, whereas households that receive a shock of \$5,000 still have an MPC of 0.44. This is driven not only by the mechanical effect of temptation, but also by sizable housing adjustment costs. While historically it has been difficult to study how MPCs vary by shock size (as most stimulus payments are small), there exists a growing empirical literature that studies this question using new sources of variation and better quality data. The early empirical literature suggested that MPCs were small for large shocks (see for instance Hsieh (2003)), however, new empirical evidence shows that MPCs remain large even in response to sizable income shocks (Fagereng, Holm, and Natvik (2019) and Kueng (2018)). For instance, Kueng (2018) finds a quarterly MPC of just under 0.3 in response to an average payment of \$4,600.⁷ In contrast, this new empirical evidence is difficult to rationalize in traditional heterogeneous agent models. While Kaplan and Violante (2014) give a good explanation for the average MPC in response to small income shocks, they predate the recent empirical findings and predict that the consumption response to large income shocks is almost zero.

Finally, given this evidence on the consumption response to income shocks, we evaluate the implications for the optimal design of fiscal stimulus. More specifically, targeted fiscal stimulus may be more effective than previously believed, due to the sizeable consumption response to large income shocks. We find that targeting households in the bottom 20% of the income distribution results in the largest aggregate consumption response, based on

⁷In addition, he replicates the results from Hsieh (2003) and shows that the small and insignificant finding was a result of the nonclassical measurement error in the income data, which attenuated the estimates. Moreover, the estimates in Hsieh (2003) are of consumption elasticities rather than MPCs.

a series of simulations where we implement budget-equivalent fiscal stimulus policies.⁸ In contrast, most governments have historically relied upon small fiscal stimulus payments to a large proportion of the population, which our model shows to be suboptimal. During the Great Recession in the U.S., for instance, the federal government gave stimulus payments to approximately 80-85% of households, with an average payment of \$600-\$1,200.

The rest of the paper proceeds as follows. First we develop a life cycle model of consumption, saving, and housing, where we add the novel feature that households suffer from temptation, but can mitigate this problem using housing (Section 2). To develop intuition, we use a simplified version of our model to demonstrate that this implies a desire for illiquidity (Section 3). We then return to the full model and discuss calibration and estimation (Section 4). The results from estimation show that temptation and commitment are important in fitting the wealth accumulation behavior of U.S. households (Section 5). Moreover, the model obtains a good fit of the empirical MPC literature and suggests a more important role for targeted fiscal stimulus (Section 6). We conclude by discussing next steps for further research (Section 7).

2 Model

In this section, we develop a life-cycle model with temptation preferences following Gul and Pesendorfer (2001). We build upon an otherwise standard model of consumption, housing, and wealth accumulation where households face idiosyncratic income risk during their working life and therefore accumulate wealth both to smooth consumption over income shocks and to smooth consumption over retirement.

We build upon the standard life-cycle model of consumption and saving, where there exists income uncertainty, liquidity constraints, and retirement. These assumptions give an incentive for households to save for two reasons: to finance consumption after retirement (the life-cycle motive) and to support consumption when negative income shocks hit (the precautionary motive). Next, we allow households to invest in two instruments: liquid assets or housing, where housing provides direct utility, serves as collateral for mortgages, and provides tax advantages, thus generating a third incentive for households to accumulate wealth (the housing motive). Housing transactions incur significant costs, thus making housing illiquid.

Finally, our model contains large and realistic tax benefits to homeownership. Specifically, our model matches the U.S. tax code by (1) allowing households to deduct all mortgage interest payments from their taxable income and (2) levying no capital gains tax on housing. It is important to include these benefits to housing as they may be important in explaining households wealth accumulation behavior.

⁸We study the response to income targeting as most governments have comprehensive information on their citizens' income, but not their liquid assets.

2.1 Model Structure

In this model, households live for T periods as adults, of which W periods are spent as workers. 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 devices because they are ex-post fully committed to their ex-ante choices. In order to generate demand for commitment, households have to exhibit some sort of present-biased behavior. In this section, we incorporate the temptation preferences of Gul and Pesendorfer (2001) that represent preferences for immediate gratification.

More formally, households with temptation preferences, similarly to those with standard preferences, want to maximize the sum of their expected, discounted lifetime utility, which can be written as:

$$\max \mathbb{E}_t \sum_{t=0}^T \beta^t U_t. \quad (1)$$

In contrast to standard preferences, the per period utility function under temptation preferences depends not only on the chosen consumption bundle, but also on the most tempting consumption bundle in the feasible choice set:

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

where the felicity function u is a concave function that is increasing both in c_t and h_t and is specified later. c_t and h_t are the chosen level of nondurable consumption and housing status, while \tilde{c}_t and \tilde{h}_t are the most desirable nondurable consumption and housing status. More specifically, households may be tempted to maximize their current period utility instead of maximizing their discounted lifetime utility. In particular, they may wish to spend all of their available liquid resources on nondurable consumption and housing, since that is the most tempting alternative of all. Therefore the most tempting alternative, $(\tilde{c}_t, \tilde{h}_t)$ maximizes their immediate utility, rather than lifetime utility:

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

where \mathcal{A}_t represents the budget set of the households, to be defined later. The term in square brackets in equation (2) represents the temptation motive of the households. It is the utility cost of not choosing the most tempting consumption alternative: the difference

between the temptation values of the most tempting and of the chosen consumption bundles. When exposed to temptation, households can decide to exercise self-control or succumb to temptation. If they exercise self-control they have to pay the utility cost of temptation resistance. On the other hand, if households succumb to temptation then the cost of self-control becomes zero and the utility function simplifies to its standard form.

Assets.

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

In addition, households can put their wealth into housing at any point in time. 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 :

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

where $0 < g(h^k) \leq 1$, $g'(h^k) > 0$ and $g''(h^k) < 0$. House prices grow at a constant rate, R^H , over time, representing a fixed gross return on the housing asset, therefore the initial index price determines all other house prices for each time period:

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

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

$$Fp_t(h^k).$$

When households do not own a home, they are renters of the smallest house, h^1 . We assume that the cost of renting is proportional to the price of this unit:

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

where η represents rental scale.

Mortgages.

The most widely used mortgage contract in the U.S. is the fully-amortizing fixed-rate mortgage.⁹ Therefore we assume that mortgages are of this kind with regular required

⁹These 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

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 illiquidity, but also because of the regular mortgage payments mp_t every period.

The mortgage balance for households who buy a house at time t is

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

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

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

where mp_t represents the required mortgage payment at time t . We assume that mortgages are fully-amortizing with 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.

In our baseline model, we assume that households are not allowed to increase the size of their mortgage using cash-out refinancing or home equity loans. While this makes housing more illiquid than in reality, it is consistent with evidence that homeowners are often not allowed to extract home equity when they need it most (i.e. when their income falls). DeFusco and Mondragon (2018) highlight that employment documentation requirements prevent many homeowners from refinancing when they lose their job. Moreover, Gorea and Midrigan (2017) find that there exists substantial frictions that prevent homeowners from extracting home equity. They highlight that payment-to-income ratios often bind on households that have experienced negative income shocks. For these reasons, we believe it is worthwhile to assume that when household income falls, households will have to rely on

fully-amortizing payment plans (Wiedemer and Baker, 2012). The Dodd-Frank Financial Reform Bill of 2010 reaffirmed these standards by introducing the concept of a “qualified” mortgage that requires fixed rate mortgages to have fully amortizing payments.

their liquid asset buffer, rather than extract home equity.¹⁰ Similarly, we also assume that households are not allowed to adjust their mortgage balance through early prepayment. In reality, most households are allowed to adjust their mortgage balance, though lenders often charge large prepayment penalties in order to discourage this behavior (Wiedemer and Baker, 2012).

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

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

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:

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

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 munic-

¹⁰In addition, we have experimented with allowing home equity withdrawal in our model and find that the standard model still has difficulty matching the large share of homeowners holding zero liquid assets.

ipal 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. In our case, we allow τ_d to be the sum of the standard deduction τ_d^{standard} and mortgage interest payments made in that period. This allows our tax schedule to incorporate the mortgage interest tax deduction, a second large subsidy to homeownership in the U.S. 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:

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

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

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.

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 estimate 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 plays an important role in our model, as it increases the illiquidity of housing, thus making housing more useful as a commitment device.

Budget Set.

In order to close the model we need to define the budget set, \mathcal{A}_t , which is the constraint households face when they only optimize for the current period, period t (i.e. they spend all their available resources by setting $a_{t+1} = 0$). Tempted households take into account their budget set whenever they evaluate their most tempting alternatives.

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

Recursive Formulation.

In order to solve the problem we define the following recursive formulation:

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

where $V_t^0(\Omega_t)$, and $V_t^1(\Omega_t)$ are the value functions conditional on not adjusting and adjusting housing. Those who choose not to adjust in period t solve the following dynamic

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

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:

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

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:

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

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	Return on liquid asset	2.10
r^H	Return on housing	2.10

Note: This table presents the assumptions that we impose to simplify our model in Section 3, relative to the full model that we estimate in Section 4. A full list of parameters is given in Appendix A.1.


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 either utility or returns. This is demonstrated in Figure 1, which presents the life-cycle profiles for a representative household. 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.

Our simulation results for tempted households are shown in Figure 2. Households with temptation preferences purchase homes despite having to pay the sizeable transaction costs. This is a rational choice of households with temptation preferences, since they not only buy housing for its future return but also for its commitment value due to its illiquidity. Keeping their savings in the illiquid housing asset decreases their cost of temptation and at the same time allows them to accumulate greater wealth for retirement. As a result, the negative effect of housing transactions cost on the demand for housing is offset by the positive effect of the commitment benefit of housing.

Panel (a) of Figure 2 shows that tempted households begin to accumulate liquid assets quite late in their life, at around age 55. The reason is that accumulating wealth in liquid form is costly in the presence of temptation: households have to exercise self-control since otherwise, they optimize for the current period only and spend their liquid assets immediately. By contrast, tempted households buy homes relatively early in their life,


Figure 1: Lifecycle Profiles for Standard Households



Note: $R = R^H = 1.021$, $p_1^{\max} = 5$

at around age 35. As a result of the temptation and commitment motives in our model, households spend a significant part of their lives as wealthy hand-to-mouth: they hold no liquid wealth while owning a sizeable illiquid, housing asset.

Figure 2: Lifecycle Profiles for Tempted Households



Note: $R = R^H = 1.021$, $p_1^{\max} = 5$

Panel (b) of Figure 2 shows the implications of households' asset portfolio decisions for their consumption, relative to their labor income. Since tempted households do not accumulate liquid wealth at the beginning of their lives, their consumption coincides with their labor income up to the point when they invest in housing. This implies that the downpayment requirement for mortgages has an immediate effect on their consumption when they buy their homes. This is why consumption drops significantly for one period during the early 30s. After buying the home, consumption follows labor income closely: the difference between the two is the per period mortgage repayment. 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 when their labor income is high. As a result, facing decreasing labor income after age 55, households' consumption declines.

3.2 Availability of Housing

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 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 therefore try to 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.

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.

In this model, the type of asset choice does not impact the amount of asset accumulation.

4 Data and Calibration

In this section, we start with some descriptive statistics about household portfolios in the Panel Study of Income Dynamics (PSID) and document life-cycle patterns of consumption and wealth accumulation. We then calibrate our model in two stages. In the first stage, we calibrate standard parameters based either on the existing literature or on our direct estimates from the data. In the second stage, we use the Method of Simulated Moments to estimate the rest of the model parameters by minimizing the distance between key moments in the data and the model.

4.1 The PSID Data

We estimate our model using the 1999-2015 Panel Study of Income Dynamics (PSID). While the PSID has been collecting information on income and demographics ever since the study began in 1968, the survey received a large overhaul in 1999 when they added detailed questions on asset holdings and consumption expenditure. We, therefore, focus our analysis on 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.

For our baseline specification, we focus on households with a head between 22 and 65

years old, with non-missing information on key demographics such as age, education, and state. We do not select our sample based on the working status of the household head or spouse. Whenever there is a change in family composition we drop that observation and treat the household as a new unit starting with the subsequent observation. We focus on an oversample of low-income families in the PSID, as we believe that these households will be the households that are most likely to hold very low liquid assets and therefore are the most interesting to study.¹⁴ We therefore include households from both the core sample of the PSID (which is representative of the U.S. population), as well as households from the Survey of Economic Opportunity (which purposefully oversamples the poor). Furthermore, we drop the top 5% of households based on income in order to focus our analysis on middle and low-income households. Finally, 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). As a result, the share of homeowners with essentially zero liquid assets is slightly higher in our sample than in the Survey of Consumer Finances, as we document below.¹⁵

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.


We compute liquid assets as the sum of all household bank account deposits (checking and savings accounts) and directly held public 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.

The patterns of life-cycle nondurable consumption and wealth accumulation in the

¹⁴Furthermore, an additional benefit of oversampling the poor is that our results for wealth accumulation will not be driven by the rich, who may behave differently due to access to better investment opportunities.

¹⁵For a comprehensive survey of the share of homeowners with zero liquid assets, see Kaplan, Violante, and Weidner (2014).

Figure 4: Mean Life-Cycle Profiles, PSID




Note: All the variables are measured at 2015 prices.

U.S. is shown in Figure 4, between ages 22 and 65. In the first part of the Figure, we see the hump-shaped consumption profile with a peak at around age 50.¹⁶ While in the second part of the Figure, we plot the evolution of wealth accumulation both in the form of liquid assets and illiquid housing. The purple line shows net liquid wealth, the thick blue line shows net housing wealth. Apart from very young ages, households' wealth is concentrated in illiquid housing: by the age of 40 housing wealth account for about 70% of the average U.S. households' wealth. This overwhelming dominance of housing wealth only starts declining later in life, close to retirement, reaching roughly 60% by age 65.

In Figure 5 we plot the share of households who are homeowners but hold little or no liquid wealth. Following Kaplan, Violante, and Weidner (2014), we identify these households in the PSID as those who own homes, but whose liquid wealth is less than two weeks of income. Kaplan, Violante, and Weidner (2014) call these households 'wealthy hand-to-mouth', which is an appealing description: these households are wealthy as they own a house, yet they are living hand-to-mouth with essentially zero liquid wealth that can be used for consumption smoothing.

¹⁶Note that the PSID doesn't cover total consumption expenditures. Later, in the matching exercise, we account for this fact by scaling down the consumption measure implied by our model.

Figure 5: Share with Housing but no Liquid Wealth



Note: This variable is measured at 2015 prices.

As seen in this figure, the age profile for the ‘wealthy hand-to-mouth’ is humped-shaped with a peak around age 45, when roughly 25% of our sample owns a home but holds 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 also find a hump-shaped life-cycle profile with a peak of around 22% at age 40. Note that the average homeownership rate in our sample (up until retirement) is 53%,¹⁷ which implies that the fraction of households with no liquid wealth conditional on being a homeowner is even higher, in our sample 30%. Figure 5 highlights the importance of housing in wealth accumulation, as well as the high proportion of households, who are wealthy but keep their wealth only in illiquid form (the wealthy hand-to-mouth).

These four life-cycle profiles are key moments we want our model to match.

4.2 Model Calibration

Next, we show the calibration of the standard, exogenous parameters, \mathbb{P} , based on the existing literature or on direct estimation from the data, we then describe the Method of Simulated Moments (MSM) estimation of the remaining parameters, Γ . The complete list of parameter values can be found in Table A.1 in Appendix A.1.

4.2.1 First-Stage Parameters

Asset Returns.

We calculate the average, risk-adjusted real returns over the period between 1950 and 2016 using data from the Federal Reserve Bank of St. Louis. Details of the computation

¹⁷While a homeownership rate of 53% in our sample is slightly lower than what is observed in the SCF, this makes sense for two reasons. First, our sample only includes working age households, while homeownership is traditionally even higher among retired households. And second, our sample includes an oversample of low income households who are less likely to be homeowners.

can be found in Appendix A.2, while the results are reported below, in Table 2. The average risk-adjusted real return is 0.69% for T-bills, 5.40% for stocks, and 2.10% for housing.

Table 2: Real Asset Returns

	Mean	St.Dev.	Risk-adj. Mean
T-Bill	0.74	2.12	0.69
Stock	8.24	16.82	5.40
Housing	2.34	5.06	2.10

Note: T-Bill: 3-month treasury bill rate, Stock: S&P 500, Housing: Case-Shiller index augmented by average housing service flow, maintenance cost and home insurance.

Most importantly, our calculated returns show that there exists a liquid asset (stock) that delivers substantially higher risk-adjusted returns than housing. This is consistent with a wide body of literature including Flavin and Yamashita (2002), Goetzmann and Spiegel (2002), and Piazzesi, Schneider, and Tuzel (2007) who show that housing delivers worse risk adjusted returns than stock, even when accounting for imputed rents and other benefits to homeownership. In our model we calibrate a risk-adjusted return of 5.4% for liquid assets ($R = 1.054$) and 2.1% for housing ($R^H = 1.021$). This parametrization relaxes the key assumption of macroeconomic models (for instance, Kaplan and Violante (2014)) that assume a high excess return on illiquid assets in order to generate high demand for illiquidity. This makes it more difficult for our model to match the large share of homeowners holding essentially zero liquid assets, thus requiring the model to match these moments using either housing utility or commitment benefit.

Housing.

We allow seven different sizes of housing ($k = 7$), with different prices, as described above. We set the price of the largest available home initially to \$250,000, which, thanks to the positive return on housing (R^H) continuously increases over the life-cycle.¹⁸ Following Attanasio et al. (2012), we impose a 5 % fixed cost of moving, F , representing the cost of real estate agents, lawyers, surveyors, and moving companies. This is consistent with empirical evidence showing that transaction costs for housing are usually at least 5% of the asset value (OECD, 2011).

¹⁸We tried to estimate p_1^{max} , which always converged roughly to this value. Eventually, we decided to calibrate this parameter, given that we noticed negligible demand for the biggest home in the model.

Mortgage rate.

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 gross mortgage rate, $R^M = 1.041$, which is two percentage points higher than the risk-adjusted return on housing. We assume that each household can borrow up to 90% of the value of its home, hence the minimum downpayment requirement, ψ is set to be 10%.

Utility function.

The curvature parameter, γ , is calibrated to match findings in Blundell, Browning, and Meghir (1994) and in Attanasio and Weber (1995). It corresponds to an inverse elasticity of intertemporal substitution of 2.0. The remaining parameters in the utility function are calibrated in Section 4.2.2 by using the Method of Simulated Moments. We find that estimating the remaining utility parameters within the model is of crucial importance when we evaluate competing theories as to why households keep the vast majority of their wealth in housing.

Initial wealth and Income.

We assume zero initial housing endowments but calibrate the initial liquid wealth distribution in order to match the distribution for 22-25-year-old households in the PSID. We calibrate income over the life-cycle in two steps. First, we use the estimated parameters for the stochastic component of income from Choukhmane (2019). We then estimate a third-order age polynomial on income in order to approximate the deterministic part of labor income. 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. Income after retirement is not subject to any risk and is a function of households' last working period income. All these parameters are listed below, in Table 3, while in Appendix A.4 we plot the before-tax income profiles from the model and from the data.

Demographics.

In the model, we account for changes in households' composition over the life-cycle by the equivalence scale measure. Calculating these, we adapt OECD weighting, which assigns weight 1 to the first adult in the household, weight 0.7 to the second and each subsequent person aged 14 and over and weight 0.5 to each child aged under 14.

Table 3: Income Process Parameters

Stochastic Income Component			Deterministic Income Component			
ρ	σ_ε^2	σ_0^2	constant	age	age ²	age ³
0.9	0.05	0.184	6.391	0.256	-0.045	0.002

Income Tax Function			Social Security		
τ_1	τ_2	τ_d	SS inc. floor	PIA bend points	SS wage base
-4.034	1.226	\$6,116	\$10,998 ¹	[\$816, \$4,917]	\$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.

Prices.

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

4.2.2 Second-Stage Parameters

The remaining parameters in the model, $\Gamma = \{\lambda, \beta, \chi, \mu, \theta\}$, the degree of temptation, the impatience parameter, the utility cost of changing home, and the housing preference (taste) parameters respectively, are estimated 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 chooses Γ to minimizes 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 four life-cycle profiles shown in Figure 4 and Figure 5. More specifically, we target the mean life cycle paths of four variables between ages 22-65: consumption, liquid assets, net housing wealth, and the share of households who own homes but hold no liquid assets (176 moment conditions). We also target the average homeownership rate in the our sample (1 moment condition). We choose to match this additional moment on homeownership rate based on the example set by Attanasio et al. (2012). Altogether we target $N_m = 177$ moment conditions to estimate 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 estimation routine places more weight on moments that are more precisely estimated in the data.¹⁹

Identification of 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 estimation routine is based on simulations for 1,000 households for two scenarios each. In the first scenario, we let λ be different from zero and we call this the unrestricted or temptation model. In the second scenario, we restrict parameter λ to be zero and we call this the restricted or standard model. 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 estimated parameters.

5 Results and Model Fit

In this section, we report the parameter estimates that we obtain when we estimate the model using the Method of Simulated Moments. We then evaluate how well the two models match targeted moments from the PSID. While the temptation model obtains a good fit of the targeted life-cycle moments, the standard model (with $\lambda = 0$) fails to match the share of homeowners with zero liquid assets. The intuitive explanation for this result is that while high impatience (low β) can generate a large share of homeowners with zero liquid assets, this results in asset holdings that are implausibly low relative to the data. Meanwhile high taste for housing (μ and θ) is ineffective in increasing the share of homeowners with zero liquid assets, as this makes households more averse to losing their home, thus encouraging households to keep a buffer stock of liquid assets. In other words, there exists a “housing smoothing” motive in addition to the standard “consumption smoothing” motive.

Parameter Estimates

Table 4 presents the MSM estimates for the five parameters of interest: temptation (λ), time preference (β), taste for housing (μ, θ), and the disutility of moving (χ). The first

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.

column shows the parameter estimates from the temptation model, where λ is allowed to take on any non-negative value during the estimation routine. The second column shows the parameter estimates from the standard model, where we impose the restriction that $\lambda = 0$, thus turning off temptation.

It is interesting to note four differences between the unrestricted and restricted parameter estimates. First, we estimate that $\lambda = 0.16$ in the unrestricted model, meaning that households suffer from temptation in this model. The λ that we obtain is between the estimates of Bucciol (2012) and Kovacs and Low (2019) who find values of 0.05 and 0.28 respectively.²⁰ Our result is also consistent with Toussaert (2018) who finds evidence that individuals suffer from temptation using a clever lab experiment.

Second, the unrestricted model yields an annual discount factor of $\beta = 0.95$, a value that is consistent with the literature and the most common calibration of β in macroeconomic models. In contrast, the restricted model yields a very low impatience parameter of $\beta = 0.91$. This is because the standard model requires strong discounting of future utility in order to explain the large share of homeowners with zero liquid assets.

Table 4: Estimated parameters from MSM

PARAMETER		Temptation Model $\lambda \geq 0$	Standard Model $\lambda = 0$
Temptation	λ	0.16	-
Impatience	β	0.95	0.91
Utility cost of housing adjustment	χ	0.57	0.07
Housing utility (non-homothetic)	μ	0.44	0.17
Housing utility (MU of consumption)	θ	0.24	0.04

Third, we find that the utility cost of moving χ is 0.57 in the temptation model. For an easier interpretation, we calculate the consumption equivalence of parameter χ , by expressing it as additional consumption that households require to be indifferent between moving and not moving homes. An average renter, for example, has to face a utility cost that is equivalent to an additional \$42,000 of consumption if it decides to buy

²⁰ Of these two papers, only Kovacs and Low (2019) estimate the importance of temptation in the presence of housing. This is performed using an estimated consumption Euler equation. We not only use a different estimation strategy, but also go one step further by studying households' observed preference for illiquidity and the implications for consumption behavior.

a house.²¹ In contrast, the restricted model delivers a close-to-zero estimate of χ at 0.07. The low value of χ makes housing more liquid, thus allowing households to carry smaller buffer stocks of liquid assets for consumption smoothing purposes. The difference between χ in these two models reflects the fact that illiquidity is partially desirable in the temptation model, as housing wealth that is more illiquid results in less temptation, whereas illiquidity is strictly undesirable in the standard model.

Fourth, we estimate higher taste for housing (μ, θ) in the temptation model than in the standard model. This may be driven in part by the fact that housing adjustment costs are so low in the standard model. In the standard model, where it is easy to move homes, it is sufficient to have a low $\mu = 0.17$ and $\theta = 0.04$ in order to match the homeownership rate observed in the data.

Finally, it is worth noting that we estimate consumption and housing to be complements in both the temptation model and the standard model, although the strength of this complementarity is much stronger in the temptation model, where $\theta = 0.24$ relative to $\theta = 0.04$. This is in line with the results from Attanasio et al. (2012) who estimate a similar utility function for housing, albeit with two types of homes, and find that consumption and housing are complements.²²

Model Fit

Figure 6 shows the simulated life-cycle moments from the unrestricted temptation model (the solid line) and the empirical moments from the PSID (the dashed line). We observe that the model obtains a very good fit of the life-cycle profiles of consumption and liquid asset holdings. Most importantly, the model obtains a good fit of the share of households who simultaneously own homes while holding essentially zero liquid assets. The model is able to achieve this result through a combination of temptation and taste for housing. In addition, the model obtains a good fit of both consumption and liquid wealth. While the fit of net housing wealth is very good prior to age 50, it diverges slightly as households approach retirement. In the PSID, net housing wealth is \$130k at age 65, whereas in our model net housing wealth is \$170k. This difference arises from the assumption in our model that all mortgages must be paid off by age 65, when households are forced to

²¹Naturally, the consumption equivalence is different depending on the age of households, their current housing status, and their next period housing status. In Figure A.4 in the Appendix, we plot the consumption equivalence for parameter χ over the whole life-cycle for renters.

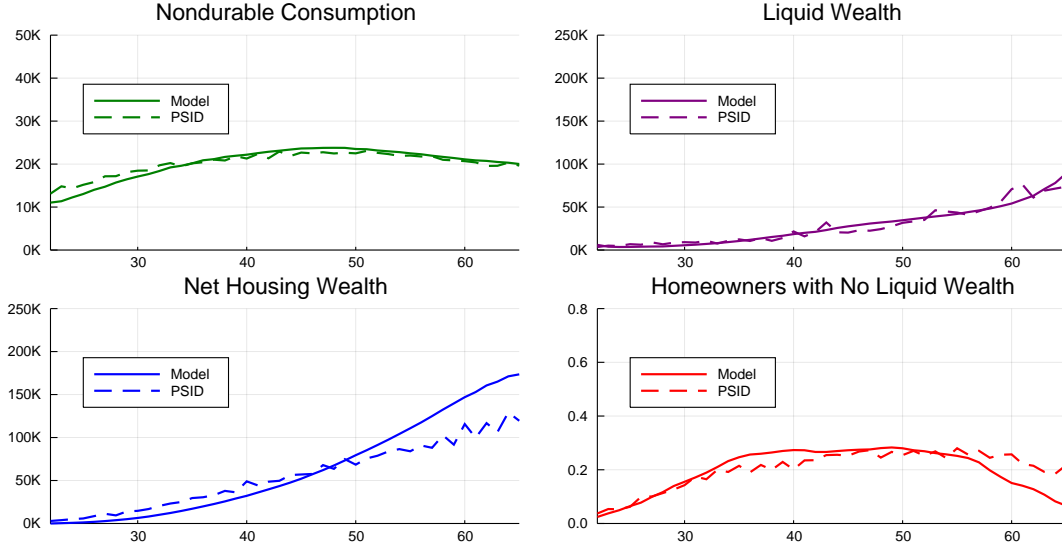
²²Since μ is positive, an increase in housing has a direct positive effect on utility. In addition, it is worth noting that a positive θ implies Edgeworth complementarity of consumption and home ownership, as the cross derivative of utility with respect to consumption and housing is positive

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

Thus in the unrestricted model, housing and consumption are complements, whereas in the restricted model housing and consumption are (weak) compliments.

retire. In reality, some households continue to work and/or maintain positive mortgage balances after age 65, which helps explain this difference.

Figure 6: Fit of the Temptation Model




Note: This figure shows the targeted life-cycle moments from our model (the solid line) and the PSID (the dashed line). Given our interest in households with zero liquid assets, we use the PSID which provides an oversample of low income households in the U.S. The average homeownership rate is 53% in the PSID and 48% in the model.

In order to gauge the importance of temptation and housing taste in matching the life-cycle profiles that we target, we now analyze the performance of the standard model, where we turn off temptation by imposing the restriction that $\lambda = 0$ and then estimate the remaining parameters. Figure 7 shows the results for the standard model. In general, the model obtains a good match of the life-cycle profile of consumption, as well as a relatively good match of liquid assets and net housing wealth. After age 50, the model predicts slightly lower liquid assets than is observed in the data. In addition, prior to age 50, the model slightly under predicts housing wealth accumulation. Most importantly, when we look at the share of households with housing but no liquid assets, we see that the standard model obtains a very poor fit of the data. There is no combination of model parameters that work to match the consumption and portfolio allocations simultaneously, even though the model is quite flexible and allows for different types of housing taste, utility cost of housing adjustment and impatience.

Identification

Identification of the model parameters requires that each parameter has an independent effect on at least one targeted moment. We find that the share of households that own housing but hold almost zero liquid assets is especially important to identify the

Figure 7: Fit of the Standard Model



Note: This figure shows the targeted life-cycle moments from our model (the solid line) and the PSID (the dashed line). Given our interest in households with zero liquid assets, we use the PSID which provides an oversample of low income households in the U.S. The average homeownership rate is 53% in the PSID and 48% in the model.

importance of temptation (λ). This is because no combination of the other parameters is able to match this moment without failing to match other targeted moments.

We first consider how the standard model could obtain a better fit of the share of homeowners with zero liquid assets. There exist four options: it could decrease the time preference parameter (β), increase taste for housing (μ or θ), or decrease the utility cost of housing adjustment (χ). While some of these options would give a better fit of the wealthy hand-to-mouth, they would all involve sacrifice along one of the other dimensions.

To understand why, we need an intuitive understanding of the role of each of these parameters. Time preference (β) impacts the level of wealth accumulation, taste for housing (μ and θ) impacts the share of wealth held in housing, and the housing transaction cost (χ) impacts the degree to which housing is illiquid. Of these four parameters, we find that β and χ have the most impact on the share of homeowners with zero liquid assets.²³ In contrast, μ and θ are much less effective in generating homeowners with zero liquid assets. This is because high taste for housing makes homeowners more averse to losing their home, thus generating a strong precautionary savings motive for homeowners. This “housing smoothing” motive results in a smaller share of homeowners holding zero liquid assets.

In the standard model, we find that the best fit comes from a low β and low χ ,

²³More specifically, a low β implies that households do not accumulate much wealth, therefore increasing the share of wealthy hand-to-mouth households. In addition, a low χ implies that housing is relatively liquid and that households can easily downgrade homes for consumption smoothing purposes. This reduces the incentive for households to keep a liquid asset buffer.

although it is not worthwhile to lower these parameters any further, as then the standard model would fail along other dimensions. For instance, if β were to be lowered below our estimate of 0.91, we would get a better fit of the wealthy hand to mouth, but a worse fit of the liquid asset profile. Similarly if χ were to be lowered further, we would get more wealthy hand to mouth households, but implausibly low liquid asset holdings. In contrast, μ and θ behave very differently. If μ or θ were higher, we would change the share of wealth held in housing versus liquid assets, but we would have very little impact on the share of homeowners with zero liquid assets.

In contrast, the temptation model is able to obtain a good fit of the data for three reasons. First, the temptation parameter λ makes it difficult to hold liquid assets, but does not make it difficult to hold housing, provided that housing is sufficiently illiquid. This allows the model to obtain a better fit of homeowners holding zero liquid assets than a model with high housing taste alone.²⁴ Second, tempted households still care about saving for retirement, but they seek to accumulate their retirement savings relatively quickly. This is because households suffer temptation costs from holding liquid assets, therefore if a household accumulates liquid assets quickly during the final years before retirement, it suffers less temptation costs than if it gradually accumulates liquid assets over many years. Third, temptation impacts the relationship between housing adjustment costs (f and χ) and demand for housing. In the next section, we will discuss how there exists a desire for illiquidity that allows large adjustment costs to fit the data.

Housing Adjustment Costs


In our model, households face both a financial cost and a utility cost to adjust housing. While we calibrate the financial cost as 5% of the value of the house, we estimate the utility cost along with the other parameters. We find that in a standard model, a higher χ always makes housing less desirable, whereas in the temptation model, a positive value of χ can actually increase the demand for housing, as greater illiquidity makes housing more effective as a savings commitment device. In other words, there exists χ has both a cost effect and commitment effect that work in opposite directions.

Figure 8 shows net housing wealth as a function of χ , the utility cost of moving. The red dashed line shows the effect of χ on net housing wealth in the standard model, while the solid blue line shows that in the temptation model.

In the standard model, there exists a negative relationship between χ and housing wealth. In contrast, in the temptation model, this relationship is humped shaped, with an increase in χ initially causing an increase in net housing wealth. This is driven by the fact that the cost and commitment effects of χ , defined above, work in opposite directions. For

²⁴While both temptation and housing taste impact the share of wealth held in housing, high housing taste also makes homeowners more averse to losing their home if their income falls, resulting in more homeowners holding a liquid asset buffer.

Figure 8: The Importance of Housing Adjustment Cost χ



Note: This figure shows the level of simulated net housing wealth at age 40 for different levels of parameter χ both under standard (red dashed line) and temptation (blue solid) models.

$\chi \leq 0.4$, an increase in χ motivates wealth accumulation, as the commitment effect helps households accumulate wealth in housing. Beyond that threshold, a further increase in χ has a negative effect on wealth accumulation, indicating that the cost effect dominates. Figure 8 shows that χ can go as high as 1 and yet still generate the same level of housing wealth as there would be without any utility cost. Consequently, any χ between 0 and 1 (*ceteris paribus*) implies an increased demand for illiquidity, as within this range the benefit from commitment dominates the utility cost of moving. Our estimated value of $\chi = 0.57$ lies within this range.

Discussion

The results in this section confirm that high housing taste and high impatience are not sufficient to explain the large fraction of homeowners who choose to hold essentially zero liquid assets. We therefore find that temptation and commitment are important to obtain a good fit of the targeted moments.

It is worth noting that we obtain this result even though our model contains large and realistic tax benefits to homeownership. Specifically, our model matches the U.S. tax code by (1) allowing households to deduct all mortgage interest payments from their taxable income and (2) levying no capital gains tax on housing. While these tax benefits certainly increase demand for housing, they do not explain the decision for homeowners to hold zero liquid assets. We therefore find that the traditional motives for housing wealth accumulation are insufficient to match households' observed portfolio choices.

Households' portfolio decisions have strong implications for their consumption behavior. In the next section, therefore we analyze the out of sample fit of our model by looking at households' marginal propensities to consume (MPC) out of unexpected, transitory

income shocks.

6 Consumption Behavior

In this section, we study the consumption response to transitory and unanticipated income shocks. We compare our model results to the empirical MPC literature, not only to validate the model’s predictive power, but also to assess whether the alternative view of illiquidity that we highlight is important for understanding consumption dynamics. We find that our model obtains a good match of (1) the average annual MPC, (2) the slow decline of MPCs by wealth, and (3) the slow decline of MPCs by shock size.

MPC heterogeneity has important implications for the optimal design of fiscal stimulus. With a better understanding of MPC heterogeneity, policy makers can use this information to give stimulus payments to households that are going to consume them more quickly, thus boosting aggregate demand during a recession. In the second half of this section, we study targeted fiscal stimulus and evaluate the tradeoffs to stimulus targeting. We find that fiscal stimulus is the most effective when the government targets households in the bottom 20% of the income distribution.

6.1 MPC Heterogeneity

To evaluate the performance of our model relative to the empirical evidence, we study the consumption response to a transitory and unexpected windfall income shock in our model. Details of this procedure can be found in Appendix A.3.

Average MPC

We find that our model generates an average annual MPC of 0.51 in response to a \$1,000 windfall income shock. This lies slightly on the upper side of the empirical literature, yet still well within the standard range of estimated MPCs. For instance, Carroll et al. (2017) give a comprehensive summary of the existing empirical literature and reports that average annual MPC estimates for the U.S. range between 0.2 and 0.6. In addition, Jappelli and Pistaferri (2014) and Fagereng, Holm, and Natvik (2019) also report average MPCs that are similar in magnitude. In our model, households exhibit high MPCs for two reasons. First, in order to avoid the cost of temptation, households keep the vast majority of their wealth in illiquid form, thus restricting their ability to consumption smooth over transitory income shocks. Second, there exists a mechanical effect of temptation: when households receive a positive income shock, they face increased temptation, therefore they consume more today. In the next two subsections, we will explore variation in MPCs along various different dimensions.

Slow declines of MPC with net wealth

We find that the average MPC in our model declines relatively slowly with net wealth, while it declines quickly with cash-on-hand, a finding that is consistent with a wide body of empirical evidence. This reaffirms the importance of using a two-asset model (with both liquid and illiquid assets) to study consumption behavior.

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

We find that cash-on-hand is the most important determinant of MPC heterogeneity in our model. This can be seen in Table 5, which presents the results from our regression. We see that household in the lowest quartile of both cash-on-hand and net wealth have an average MPC of 0.77. Moreover, MPCs decline very quickly for households with greater cash-on-hand holdings. For instance, households in the second quartile of cash-on-hand have an average MPC that is 0.23 lower than households in the bottom quartile. Moreover, households in the top quartile have an MPC that is 0.66 lower than households in the bottom quartile. We therefore see almost no response in consumption for households in the top quartile of cash-on-hand.

In contrast, once we have controlled for cash-on-hand, net wealth is almost entirely unimportant in explaining MPC heterogeneity in our model. Table 5 shows that households in the lowest wealth quartile have an MPC of 0.77, whereas households in the second wealth quartile have an MPC of 0.65. Moreover, we see that even the richest households still have a large MPC: households in the top wealth quartile have an MPC that is just as large as households in the bottom quartile.

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. For instance, Fagereng, Holm, and Natvik (2019) find that net wealth is unimportant in explaining MPC heterogeneity, once controlling for liquid wealth.²⁵ Similarly, 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

²⁵Fagereng, Holm, and Natvik (2019) 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.

Table 5: MPC Heterogeneity by Household Type

		MPC	
		Coefficient	Standard Error
CASH-ON-HAND			
Quartile			
	<i>2nd</i>	-0.229***	(0.006)
	<i>3rd</i>	-0.431***	(0.006)
	<i>4th</i>	-0.668***	(0.007)
NET WEALTH			
Quartile			
	<i>2nd</i>	-0.124***	(0.005)
	<i>3rd</i>	-0.018***	(0.006)
	<i>4th</i>	0.014*	(0.008)
Constant		0.776***	(0.012)

Note: MPCs are based on a \$ 1,000 transitory income shock, where we use a regression framework to control for age. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

quintile, affirming the importance of cash-on-hand for MPCs.²⁶

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.

²⁶ Jappelli and Pistaferri (2014) use Italian survey data to study the consumption response to unexpected transitory income shocks. They exploit the survey question from the 2010 Italian Survey of Household Income and Wealth, which asks households how much of an unexpected transitory income change they would spend.

Slow declines of MPC with income shock size

Finally, we find that the average MPC declines relatively slowly with shock size. In other words, large income shocks still induce a large increase in consumption. While this result is consistent with a growing empirical literature, it cannot be explained by traditional heterogeneous agent models. To the best of our knowledge, our model is the first that is able to replicate this empirical finding.

Table 6 reports the average MPC in response income shocks of different sizes. We observe that the average MPC declines only slowly with the size of the shock. For instance, changing the shock size from \$1,000 to \$10,000 only causes the MPC to decline from 0.51 to 0.35.

Table 6: MPC Heterogeneity by Shock Size

	SHOCK SIZE		
	\$1, 000	\$5, 000	\$10, 000
MPC	0.51 (0.01)	0.44 (0.00)	0.35 (0.00)
Observations	43,876	43,392	42,892

Note: Each coefficients represent the average annual MPC after controlling for age in a regression.

This implication of our model is consistent with a growing empirical literature that documents the consumption response to large income shocks and finds that the average MPC remains relatively high in response to large shocks. Most similar to our analysis, Fagereng, Holm, and Natvik (2019) study the consumption response to large and unanticipated lottery payments using administrative tax data from Norway. They find that the average MPC only gradually declines as the size of the shock increases, therefore households that receive large payments still consume most of their payment quickly 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. For instance, Kueng (2018) finds a quarterly MPC of just under 0.3 in response to an average payment of \$4,600.

Large MPCs out of large, transitory income shocks cannot be explained by other heterogeneous agent models in the current literature. For instance, while the model of Kaplan and Violante (2014) obtains a very good fit for MPCs out of small income shocks, it predates these empirical findings and cannot rationalize large MPCs out of large income shocks. The reason our model is able to predict a slow decline of MPCs by shock size is because we have large and reasonable adjustment costs for housing (5% of the value of the house (OECD, 2011), plus the disutility of moving), whereas Kaplan and Violante

(2014) have very small adjustment costs for illiquid assets (\$1,000). Kaplan and Violante (2014) demonstrate that this small adjustment cost is necessary in order for their model to fit the large share of households with zero liquid assets. In contrast, our model does not require small adjustment costs to fit the large share of these households. Moreover, it is necessary to have reasonable adjustment cost to match the empirical evidence on the slow decline of MPCs by shock size.

This empirical phenomena has important implication for the optimal design of fiscal stimulus, as it suggests that large and targeted fiscal stimulus payments could be very effective in boosting aggregate consumption. We investigate this issue of optimal fiscal policy design further in the next section.

6.2 Targeted Fiscal Stimulus

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


We study the consumption response to alternative stimulus targeting policies by varying the fraction of households that receive a one time unanticipated stimulus payment from the government, where the government uses an income based targeting approach.²⁷ We focus on budget equivalent policies, for instance giving \$500 to all households or \$1,000 to the bottom 50% of the income distribution. Specifically we simulate 1000 households using our model and then compare their baseline consumption to a counterfactual simulation where the same households (i.e. 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 fiscal stimulus, therefore we repeat this exercise for all t within this age range and then aggregate our results. We then report the fraction of aggregate stimulus that is consumed within one year after disbursal by the government.

Figure 9 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

²⁷We study the response to income targeting as most governments have comprehensive information on their citizens' income, but not their liquid assets. Of course, Fagereng, Holm, and Natvik (2019) might be able to help Norway perform targeted stimulus based on liquid assets, which may be even more effective.

\$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% of aggregate stimulus is consumed within one year.

Figure 9: 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 engaged in fiscal stimulus decided to give stimulus 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 lower liquid assets and higher MPCs, but it is important to note that 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 6 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.

Finally, we find a more important role for stimulus targeting compared to the existing theoretical literature. For instance, while the model of Kaplan and Violante (2014) implies a very 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.²⁸ This difference is driven by the above trade-off 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.

7 Conclusion

In this paper, we integrate the idea of temptation preferences, proposed by Gul and Pe-sendorfer (2001), into a life-cycle model with incomplete markets. We show, that this model is able to explain households' overwhelming preference for illiquidity by emphasizing the role of illiquid housing as a savings commitment device. We document the model's ability to match several important features of MPCs, which are hard to reconcile with traditional life-cycle models.

Using the Method of Simulated Moments, we estimate our structural life-cycle model once with temptation and once when we shut down temptation. Our identification is based on the life-cycle patterns of households consumption and portfolio compositions. Crucially, we target the observed large fraction of U.S. households who own housing wealth but no liquid wealth (called 'wealthy hand-to-mouth' by Kaplan and Violante (2014)). The model without temptation is not able to explain the existence of households who own housing wealth but essentially zero liquid wealth, even though it has great flexibility that allows it to have different types of housing taste (direct and indirect via consumption), utility cost of housing adjustment and impatience. In contrast, the model with temptation generates the observed life-cycle profiles, including the share of

²⁸In their model, stimulus payments given to the entire population would imply that roughly 40% is 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.

households who own housing wealth but no liquid wealth.

Households' portfolio allocations have important consequences for aggregate consumption behavior and the optimal design of fiscal stimulus payments. We find that targeted fiscal stimulus is more powerful than previously believed: targeting households at the bottom 20% of the income distribution results in the largest aggregate consumption response, with households consuming approximately 70% of fiscal stimulus payments during the year of receipt. This result is largely driven by the finding that the average MPC declines slowly with shock size. While this is consistent with a number of recent empirical studies (Fagereng, Holm, and Natvik (2019), Bunn et al. (2018), and Kueng (2018)), there exist very few empirical papers that study the consumption response to large income shocks and it would therefore be interesting to see additional empirical research in this direction.

The results of this paper could be fruitfully applied in a number of ways. There already exists a growing literature within macroeconomics that studies the interaction between public policy and welfare when households suffer from problems of self-control. Some nice examples include Krusell, Kuruscu, and Smith (2010) who study optimal savings subsidies, Nakajima (2012) who looks at the welfare effects of credit cards, and Schlafmann (2016) who studies the consequences of mortgage market regulation when households suffer from temptation. Our estimation results reaffirm the importance of these studies.

A Appendix

A.1 Model Parameters

Table A.1: Static Annual Parameters

FIRST STAGE		VALUE
Timing		
T	number of years as adult	59
W	number of years as worker	44
Utility Parameters		
γ	risk aversion	2.0
Asset Returns, Prices		
r	stock return	0.054
r^H	housing return	0.021
r^M	mortgage interest rate	0.041
η	rental scale	0.03
F	fixed cost of moving	0.05
ψ	down-payment requirement	0.10
p_1^{\max}	initial house price	\$250,000
Income Process		
ρ	income persistence	0.90
σ_ε^2	std. dev. income shock	0.050
σ_0^2	std. dev. initial income	0.184
τ_1	income tax function, constant	-4.034
τ_2	income tax function	1.226
SECOND STAGE		
Utility Parameters		
β	impatience	0.95
λ	degree of temptation	0.16
θ	housing utility (MU of consumption)	0.24
μ	housing utility (non-homothetic)	0.44
χ	utility cost of housing adjustment	0.57

A.2 Housing Return Calculation

A key feature of our calibration is that housing does not deliver excess returns relative to liquid assets. In this section, we calculate the return of stock and housing. 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, while we use data from the Bureau of Economic Analysis (BEA) in order 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 use the 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.²⁹ The average of gross housing value added over residential fixed assets between 1950 and 2016 is around 8%.

Following Kaplan and Violante (2014), we set the maintenance cost at 1% and the

²⁹Gross 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

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.

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

where superscript i refers to the type of the asset, i.e. 3 Months T-Bill³⁰, 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 below.

³⁰The 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	Sharpe Ratio
T-Bill	0.74	2.12	0.69	-
Stock (S&P)	8.24	16.82	5.40	0.45
Housing (Case-Shiller)	2.34	5.06	2.10	0.30

We also report the Sharpe ratios for stocks and housing. The Sharpe ratio measures the expected value of the excess of the asset return over the T-bill return per unit of the standard deviation of the excess return. Therefore, the higher the value of the Sharpe ratio for a given risky asset, the more attractive is the asset, the more of its riskiness is compensated by its excess return. The Sharpe ratios confirm that housing yields a lower risk-adjusted return than stocks.

A.3 MPC Calculations

In order to calculate the consumption response to a transitory and unexpected windfall income shock we generate N households, each with a different series of income shocks and initial heterogeneity. We simulate our model to obtain a baseline consumption path ($c_{i,t}^{\text{baseline}}$). Next, we simulate the same households (each with the same series of income shocks and initial heterogeneity) but now give these households an unexpected and transitory income shock at time j . This gives us an alternative consumption path ($c_{i,t}^{\text{alt}}$). We set the size of the transitory income shock to be \$1,000 at time j in our baseline simulation, while we also show results for shock sizes of \$5,000 and \$10,000. For household i , we then compute the MPC as the change in consumption at time j induced by the positive, transitory income shock, normalized by the size of the shock. Thus, the annual MPC is computed as

$$MPC_{i,j} = \frac{c_{i,j}^{\text{alt}} - c_{i,j}^{\text{baseline}}}{\text{ShockSize}}$$

for household i given a temporary shock to its income at time j . We repeat this procedure for all years over the life-cycle, $j = 1, 2, 3, \dots, T$.

A.4 Additional Figures and Tables

Figure A.2: Income Profiles




Figure A.3: Fit of the Standard Model with $\beta = 0.95$

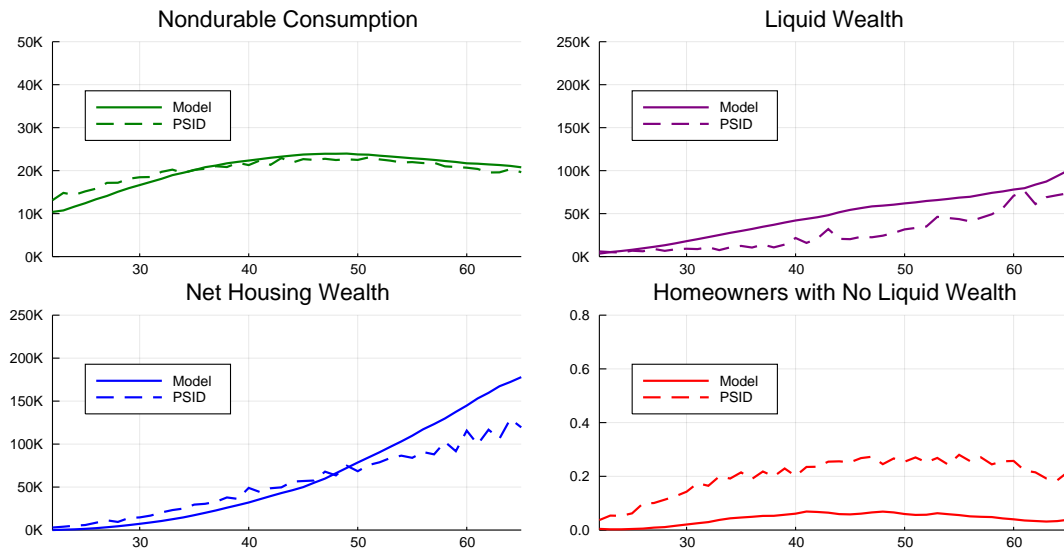


Figure A.4: Consumption Equivalence for χ if Renter




Table A.3: MPC Heterogeneity by Age Quartiles

	AGE QUARTILE			
	1 st	2 nd	3 rd	4 th
MPC	0.56 (0.00)	0.48 (0.00)	0.38 (0.00)	0.22 (0.00)
Observations	9,984	10,969	10,967	10,960

Note: Each coefficients represent a separate regression of MPCs within an age quartile.

References

- 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.
- 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 Guglielmo Weber. 1995. “Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey.” *Journal of Political Economy* 103(6):1121–1157.
- Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2018. “The Intertemporal Keynesian Cross.” *Mimeo* .
- Bayer, Christian, Ralph Luetticke, Lien Pham-Dao, and Volker Tjaden. 2019. “Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk.” *Econometrica* 87 (1):255–290. URL <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA13601>.
- Blundell, Richard, Martin Browning, and Costas Meghir. 1994. “Consumer Demand and the Life-Cycle Allocation of Household Expenditures.” *Review of Economic Studies* (61):57–80.
- Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. “Consumption Inequality and Family Labor Supply.” *American Economic Review* 106 (2):387–435.
- Bucciol, Alessandro. 2012. “Measuring Self-Control Problems: A Structural Estimation.” *Journal of the European Economic Association* 10(5):1084–1115.
- Bunn, Philip, Jeanne Roux, Kate Reinold, and Paolo Surico. 2018. “The Consumption Response to Positive and Negative Income Shocks.” *Journal of Monetary Economics* 96:1–15.
- Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. 2017. “The Distribution of Wealth and the Marginal Propensity to Consume.” *Forthcoming, Quantitative Economics* .

- Choukhmane, Taha. 2019. “Default Options and Retirement Saving Dynamics.” *mimeo* .
- Cocco, Joao F. 2005. “Portfolio Choice in the Presence of Housing.” *The Review of Financial Studies* 18 (2):535–567.
- DeFusco, Anthony and John Mondragon. 2018. “No job, no money, no refi: Frictions to refinancing in a recession.” *No Money, No Refi: Frictions to Refinancing in a Recession (August 2018)* .
- Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik. 2019. “MPC Heterogeneity and Household Balance Sheets.” *CESifo Working Paper* (7134).
- Fernandez-Villaverde, Jesus and Dirk Krueger. 2011. “Consumption and Saving over the Life-Cycle: How Important are Consumer Durables.” *Macroeconomic Dynamics* 15 (5):725–770.
- Flavin, Marjorie and Takashi Yamashita. 2002. “Owner-Occupied Housing and the Composition of the Household Portfolio.” *American Economic Review* 92:345–362.
- 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.
- Gorea, Denis and Virgiliu Midrigan. 2017. “Liquidity Constraints in the U.S. Housing Market.” *NBER Working Paper* (23345).
- Gul, Faruk and Wolfgang Pesendorfer. 2001. “Temptation and Self-Control.” *Econometrica* 69 (6):1403–1435.
- Harris, Christopher and David Laibson. 2001. “Dynamic Choices of Hyperbolic Consumers.” *Econometrica* 69:935–957.
- Houser, Daniel, Daniel Schunk, Joachim Winter, and Erte Xiao. 2018. “Temptation and Commitment in the Laboratory.” *Games and Economic Behavior* 107:329–344.
- 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.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. “Monetary Policy According to HANK.” *American Economic Review* 108 (3):697–743.

- Kaplan, Greg and Giovanni L. Violante. 2014. “A Model of the Consumption Response to Fiscal Stimulus Payments.” *Econometrica* 82 (4):1199–1239.
- Kaplan, Greg, 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 and Hamish Low. 2019. “Present Bias, Temptation and Commitment Over the Lif-Cycle: Estimating and Simulating Gul-Pesendorfer Preferences.” *mimeo* .
- 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.
- Laibson, David, Peter Maxted, Andrea Repetto, and Jeremy Tobacman. 2017. “Estimating Discount Functions with Consumption Choices over the Lifecycle.” .
- Luetticke, Ralph. 2018. “Transmission of Monetary Policy with Heterogeneity in Household Portfolios.” Discussion Papers 1819, Centre for Macroeconomics (CFM). URL <https://ideas.repec.org/p/cfm/wpaper/1819.html>.
- Nakajima, Makoto. 2012. “Rising Indebtedness and Temptation: A Welfare Analysis.” *Quantitative Economics* 3 (2):257–288.
- OECD. 2011. “Housing and the Economy: Policies for Renovation.” *Economic Policy Reforms 2011: Going for Growth* .
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel. 2007. “Housing, Consumption and Asset Pricing.” *Journal of Financial Economics* 83(3):531–569.
- Schlafmann, Kathrin. 2016. “Housing, Mortgages, and Self Control.” *mimeo* .
- Sommer, Kamila and Paul Sullivan. 2018. “Implications of US Tax Policy for House Prices, Rents, and Homeownership.” *American Economic Review* 108 (2):241–247.
- Toussaert, Séverine. 2018. “Eliciting Temptation and Self-Control Through Menu Choices: A Lab Experiment.” *Econometrica* 86 (3):859–889.
- Wiedemer, J.P. and J.K. Baker. 2012. *Real Estate Finance, 9th ed.* Cengage Learning.