UNIVERSITÀ DEGLI STUDI DI TORINO
ALMA UNIVERSITAS TAURINENSIS
Founded in 1404

DEPARTMENT OF
ECONOMICS AND STATISTICS
WORKING PAPER SERIES

Quaderni del Dipartimento di Scienze
Economico-Sociali e Matematico-Statistiche

ISSN 2279-7114

CONSUMPTION INSURANCE, EARNINGS RISK AND ILLIQUID HOUSING WEALTH

Claudio Campanale

Luis Franjo

Working paper No. 98 - March 2025

Consumption Insurance, Earnings Risk and Illiquid Housing Wealth*

Claudio Campanale[a] and Luis Franjo[b]

[a]University of Turin and CERP

[b]Fundamentos del Análisis Económico (FAE), Universidad de Alicante

March 11, 2025

Abstract

Households appear to smooth consumption in the face of income shocks much more than implied by the standard incomplete market model with one fully liquid asset and permanent plus temporary income shocks. However, it is well known that for most households, illiquid housing represents the most important form of wealth holdings. Moreover, the last decade has witnessed the development and estimation of richer models of earnings dynamics. In this paper we extend the basic SIM model to include a second illiquid asset and a more complex earnings process based on recent empirical estimates. We show that under the assumed earnings process with lower persistence that increases with age, the insurance coefficient against persistent shocks increases by about 20 percentage points compared to the baseline permanent shocks model, overshooting its empirical counterpart in Blundell, Pistaferri and Preston (2008). The presence of illiquid housing reduces it by about 4 percentage points aligning the model more closely to the data. We conclude that both housing and a richer specification of income risk are important for understanding insurance against shocks, with the latter playing a quantitatively more important role.

Keywords: Consumption insurance coefficients, housing, earnings dynamics.

JEL Classification: D15, E21

*We would like to thank Orazio Attanasio, Joao Cocco, Antonia Díaz, Paolo Sodini and Francesco Turino for their helpful comments and suggestions. Luis Franjo acknowledges financial support from Consellería de Innovación, Universidades, Ciencia y Sociedad Digital de la Generalitat Valenciana (CIPROM/2021/060).

1 Intorduction

Microeconomic evidence on individual consumption growth shows a large degree of idiosyncratic volatility, the observable sign of imperfect risk sharing. Macroeconomic models with heterogeneous agents inherently feature imperfect risk sharing at least qualitatively, but if they are to be credible tools for analyzing economic phenomena and assess economic policies they need to be able to match quantitatively the extent to which consumption is insured against income shocks that is observed in the data. In an influential paper Blundell, Pistaferri and Preston (2008), henceforth BPP, provided an empirical measure of the household consumption response to earnings shocks of different persistence. They did that by estimating, so called insurance coefficients, defined as the fraction of the variance of the shock that is not passed to consumption growth. In practice if consumption is perfectly insured it does not co-vary with the shock and the insurance coefficient is one. At the opposite end of the spectrum, if there is no insurance at all, shocks are translated one-to-one into consumption growth and the insurance coefficient is zero. Blundell et al. (2008) estimated these coefficients and found that about 65 percent of permanent shocks is passed to consumption growth while only about 5 percent of temporary shocks translates into consumption movements. After these estimates became available Kaplan and Violante (2010) constructed a basic life cycle self insurance model with one asset and permanent and transitory shocks and concluded that the model does quite well at matching insurance against temporary shocks but falls short with respect to permanent shocks, generating insurance of only 7 to 22 percent of the shock, depending on the tightness of the borrowing constraint.

The current paper builds upon the framework in Kaplan and Violante (2010) but extends it in two ways. First, it recognizes the fact that a large part of households' wealth is tied in the form of housing. Housing on the one hand is illiquid, hence housing wealth can only be used to insurance non durable consumption at a cost. On the other hand, it provides collateral against which to borrow if the need arises. Second, it considers a richer process for earnings dynamics, based on the work of Karahan and Ozkan (2013), that allows for shocks that are persistent but not fully permanent, with persistence that declines with age. To be more precise about the model, we consider life-cycle agents that go through the two stages of working life and retirement. During working life they receive earnings from working in the market, during retirement they receive a pension benefit that is tied to average working life earnings following a progressive formula. At each point in time households make decisions about whether to rent or own their home, about the size of the house and about consumption of non durable goods and saving in a liquid asset. We introduce three frictions in the model. First a transaction cost on house sales. Second a borrowing constraint: households can only borrow against their house and only up to a certain fraction of the house value. Finally we assume that a household can tap home equity only by paying a refinancing cost. We calibrate the model so that it matches the main features of home ownership and wealth, both liquid and illiquid, over working life. We then simulate it and study its implications for the insurance coefficients.

There are four main findings that emerge from the quantitative analysis of the model. First we find that introducing housing in the model reduces the insurance coefficients by a sizable but not big amount: about 4 percentage points for the insurance coefficients against persistent or permanent shocks and 1 percentage point for the insurance coefficients against purely temporary shocks. Second, similar results are obtained when we consider variations of the benchmark model with housing where we change the extent of frictions by changing the size of the selling and refinancing costs. In this case we show that by moving from an extreme case with both costs reduced to zero to the opposite extreme where costs are set so high as to reduce refinancing and selling virtually to zero the variation in the insurance coefficients against the permanent shock is almost 4 percentage points, again a figure that is sizable but not big. Third, changing the labor earnings process from the original permanent plus temporary formulation to the one in Karahan and Ozkan (2013) has a much more dramatic effect, at least as far as insurance against persistent shocks is concerned. In this case in fact, both in the model with housing and in the model without housing the insurance coefficients are about 20 percentage points higher in the model with reduced persistence of shocks. Taken together the previous findings lead to the conclusion that in order to check whether the self-insurance model can match the data on consumption insurance the imperfect liquidity of part of households wealth does play a role, but this role is potentially dwarfed by the one played by the characteristics of the earnings process, hence a careful specification of the latter is more important. Finally, consistent with the existing quantitative literature we find that for the one asset model the insurance coefficients against the persistent shocks are below the empirical estimates in Blundell et al. (2008) when this shock is actually permanent and above them with the more recent and flexible earnings processes estimated in the literature. Adding housing to the model then leads those insurance coefficients to be more aligned to the estimates in the latter case, while moving them further away in the former. Hence, we can conclude that the preferred model is one in which both lower persistence, especially at young ages, and illiquid wealth are considered.

The present paper adds to the literature on the response of consumption to income shocks. Most closely related are those papers that look at insurance coefficients against shocks to non financial income averaging across the whole population. From an empirical perspective after the initial work of Blundell et al. (2008), Blundell, Pistaferri and Saporta-Eksten (2016) estimated consumption insurance coefficients with respect to wage shocks in two earner households. Following those contributions a number of quantitative papers have tested the ability of different versions of the self-insurance model to explain the data. First, Kaplan and Violante (2010) tested the most basic life-cycle SIM model with permanent plus temporary shocks and found that it substantially falls short of matching insurance coefficients against permanent shocks. Later several modifications to the base model have been proposed. Cerletti and Pijoan-Mas (2012) considered a model with durable and non durable goods and the impact on insurance coefficients for non durable goods when adjustments in the consumption bundle are possible. De Nardi, Fella and Paz-Pardo (2020) considered a version of the model with a more complex labor income dynamics featuring non constant persistence and non Gaussian shocks estimated non parametrically and found that in this case the model somewhat over-estimated insurance against persistent shocks. Campanale and Sartarelli (2024) pointed out that in order to test the insurance properties of the basic SIM model, we need to consider the whole profile of wealth accumulation over the working life rather than its average alone and found that when the former is correctly matched insurance coefficients are higher and not too far below those estimated in Blundell et al. (2008). Finally Wu and Krueger (2021) showed that a two earner version of the model where labor supply can be adjusted at both the intensive and extensive margin can match the corresponding insurance coefficients as estimated in Blundell et al. (2016) as well as the mechanism – saving or labor supply changes – through which consumption is smoothed. All of the above mentioned papers assume that savings occurs through a single fully liquid assets. The present paper then contributes to that literature by assessing the insurance properties of the SIM model when a large part of household wealth is instead tied in illiquid housing, hence it is usable for non durable consumption smoothing only by paying a substantial transaction cost. Moreover by considering both the traditional permanent plus temporary earnings shock process and more flexible formulations that allow for less persistence in shocks we can assess the relative importance of the liquidity of wealth and of the exact specification of the earnings process when testing the ability of the SIM model to explain consumption insurance in the data.

The current paper is also more loosely related to the literature that looks at the ability of quantitative consumption-saving models to match the marginal propensity to consume estimated in micro data. Among the literature in this area we can cite the works of Aguiar, Bils and Boar (2024), Kaplan and Violante (2014) and Carroll, Slacalek, Tokuoka and White (2023). Part of this literature has also included a non liquid asset in the menu of available savings instruments, however unlike the current paper it is based on the traditional permanent plus temporary formulation of earnings shocks. It is also different from the current work since it is generally focused on households close to the borrowing constraint and mainly to the response of consumption to temporary shock with a view towards analyzing macroeconomic stimulus policies. Finally the present paper is also related to recent work by Sodini, Van Nieuwerburgh, Vestman and von Lilienfeld-Toal (2023) that used quasi experimental evidence in the Swedish house market to identify in the data the role of housing collateral in smoothing income shocks.

The rest of the paper is organized in the following way. Section 2 presents a detailed description of the model. Section 3 presents the calibration and section 4 describes the results of the quantitative analysis of the model. Finally in section 5 some brief conclusions are outlined.

2 The Model

We consider a stationary partial equilibrium economy. Time is discrete and life is finite, lasting up to a maximum of T years, of which the first T_R are spent working and the remaining represent retirement. Given the structure of the economy, we will refer to time and age indifferently. The probability of surviving up to period t conditional on being alive at the end of period $t-1$ is denoted by p_t. In each period the agent derives utility from two goods, a non-durable consumption good c_t and housing services that we denote with s_t. Leisure is not valued, hence we suppose that the agent supplies her unitary endowment of time inelastically to the market. Given this description, the per period flow of utility is given by

$$u(c_t, s_t) = \frac{(c_t^{\sigma} s_t^{1-\sigma})^{1-\alpha} - 1}{1-\alpha} \tag{1}$$

where σ represents the weight of non-durable consumption in utility and α is the coefficient of relative risk aversion.[1]

Housing services can be obtained either by renting or by buying a house. Let's denote by h_{t+1} the size of the house that is owned in period t and by f_t the amount that is rented.

[1]Aggregating non-durable and housing consumption by way of a Cobb-Douglass function sets the elasticity of intra-temporal substitution between the two goods at one. This value falls within the 95 percent confidence interval of the estimates in Piazzesi, Schneider and Tuzel (2007).

We assume that the two alternatives are mutually exclusive, hence the amount of housing services enjoyed by the household at time t will be:

$$s_t = D_{t+1} h_{t+1} + (1 - D_{t+1}) f_t \tag{2}$$

where D_{t+1} is an indicator function that takes a value of 1 if the household decides to own the home in which they live or a value of 0 if it rents it. The timing implied by this formulation is the following: the household enters the period with a given amount of owned housing h_t which is a state variable. It then makes a decision as to whether living in an owned house or renting it during the period and the size of housing services to acquire. If it decides to own the amount is indicated with h_{t+1} and it becomes the value of the state variable in the next period. If it decides to rent the amount of housing services is denoted with f_t. The rental price of housing services will be denoted with r_f. If the household decides to change the size of the house it incurs in a transaction cost ϕ that is proportional to its size. The transaction cost is only paid by the seller of a house.

Beside housing, there is a liquid asset in the economy. The amount owned by the household at the beginning of period t is a_t. The asset pays a return $r(a_t)$ where the dependence of the return on the amount of the asset is assumed in order to allow a wedge between lending and borrowing rates, that is, $r(a_t)|_{a_t<0} > r(a_t)|_{a_t>0}$. Borrowing is only allowed against housing collateral subject to the fact that a minimum home equity requirement must be satisfied: we denote with θ this minimum requirement. The constraint on the liquid asset then becomes $a_{t+1} \geq -(1-\theta)h_{t+1}$. For buyers, this constraint has the usual interpretation of a down payment requirement. For the remaining homeowners, we allow the possibility of extracting equity from the home they own subject to the above constraint, but only at a cost. Following Li and Yao (2007), we assume that the cost is a fraction η of the value of the house. The approach that we follow here interprets mortgages as one-period contracts that are renewed every year at no cost and has been adopted, for example, in Li, Liu, Yang and Yao (2016), Favilukis, Ludvison and Van Nieuwerburgh (2017), or more recently Vestman

(2019). It is a simplification compared to a full modeling of multi-year mortgages with fixed yearly payments and is adopted for computational reasons. A fully specified model of multi-year mortgages would in fact add both a continuous state variable and a continuous choice variable to a model where the need to use many grid points to adequately approximate the labor earnings process already imposes an important computational burden. Even with this simplification we can capture the existence of costs to tap home equity in response to shocks to income, while the fact that mortgages are renewed every year may be less important in the current framework with a constant interest rate.

In each period the household receives a certain amount of non-financial income either in the form of earnings or in the form of a social security payment. In the description of the budget constraints that follows we generically denote this amount by y_t. During working life $y_t = G(t)z_t\xi_t$ where $G(t)$ ia a deterministic component that reflects the average life-cycle profile of earnings in the population, z_t is a stochastic permanent shock following the process $lnz_{t+1} = \rho_t lnz_t + \varepsilon_t$ where $\varepsilon_t \sim N(0, \sigma^2_{\varepsilon,t})$. Finally $\xi_t \sim N(0, \sigma^2_{\xi,t})$ is a purely transitory shock. The two shocks are independently distributed over time and across agents, however their variance as well as the autocorrelation coefficient change over the working life of the agent. This formulation with age changing persistence and variance of shocks is based on Karahan and Ozkan (2013) that shows it better fits data on households earnings in the PSID compared to the traditional permanent plus transitory shock formulation. During retirement non financial income is given by a fixed social security payment so that $y_t = \overline{y}_p$ where $\overline{y}_p$ is made to depend on the realization of earnings in the last period of life.

Given the description above we can formulate the budget constraint of the households. The budget constraint takes different forms depending on the possible transitions between different home ownership statuses. There are five types of possible transitions, that is: from renting to owning and vice versa, from renting to renting and from owning to owning, the latter with the two alternative possibilities of staying in the same home or changing it. For the sake of clarity we next proceed to describe the budget constraint for each of these five alternatives separately. We will then collapse them into one by using suitable indicator functions later when setting up the dynamic optimization problem of the household. Starting with the transition from renting to renting the constraint faced by the household will be in this case:

$$c_t + r_f f_t + a_{t+1} \leq y_t + a_t(1 + r(a_t)) \tag{3}$$

plus the borrowing constraint $a_{t+1} \geq 0$. The constraint tells us that the agent in this case has resources given by non-financial income and the liquid asset plus the interest earned on it. The resources are then used to buy non-durable consumption and housing services in the rental market and for saving in the liquid asset. The latter finally can only be held in non-negative amounts given that only collateralized borrowing is assumed. Next we consider the transition from renting to owning. In this case the budget constraint is:

$$c_t + h_{t+1} + a_{t+1} \leq y_t + a_t(1 + r(a_t)) \tag{4}$$

and the borrowing constraint becomes $a_{t+1} \geq -(1-\theta)h_{t+1}$. The right hand side of the resource constraint is unchanged from the previous case. However now the agent will use the available resources for consumption, for buying a house and for saving in the liquid asset. Moreover the borrowing constraint is now relaxed so that the liquid asset carried into the next period can be negative up to the value of the house minus the down payment requirement. The opposite case of a transition from owning to renting is described by the following constraint:

$$c_t + r_f f_t + a_{t+1} \leq y_t + a_t(1 + r(a_t)) + (1-\phi)h_t \tag{5}$$

and the borrowing constraint $a_{t+1} \geq 0$. In this case the resources available to the household are given by non-financial income, the amount of the liquid asset augmented by the interest gained on it and finally the value of the house net of the selling cost. These resources are then used to purchase non durable consumption, paying the rent for housing services and to save in the liquid asset. The borrowing constraint takes again the form of a non-negativity constraint since the household sells the house so it will not have any collateral against which to borrow for the next period. The case of owning to owning transition with a change in the size of the house is described by the resource constraint:

$$c_t + h_{t+1} + a_{t+1} \leq y_t + a_t(1 + r(a_t)) + (1 - \phi)h_t \tag{6}$$

and the borrowing constraint $a_{t+1} \geq -(1 - \theta)h_{t+1}$. The right hand side of the constraint includes non-financial income, the amount of the liquid asset from the previous decision period plus the interest on it and finally the value of the previous home, the one that the household sold, net of the transaction cost. The use of these resources include non durable consumption, the amount of the liquid asset to carry into the next period and the value of the new house purchased. The borrowing constraint says that the liquid asset carried into the next period must be no less than minus the value of the newly purchased home net of the down payment requirement. Finally we describe the case of the owning to owning transition with no change in the size of the house owned. This case must be broken down into two sub cases since we allow for the extraction of home equity, that is, we allow for the possibility of increasing the amount borrowed over time but this is assumed to be costly. Suppose first that no refinancing takes places. This case occurs formally when either $a_t \geq 0$ and $a_{t+1} \geq 0$, that is, the household has no outstanding mortgage or when $a_t < 0$ and $a_{t+1} \geq a_t$ so that the mortgage is actually partially repaid. The budget constraint associated with this case is then:

$$c_t + a_{t+1} \leq y_t + a_t(1 + r(a_t)) \tag{7}$$

In this case then the households uses only the liquid resources given by non financial income and the amount of the liquid asset plus interest to purchase consumption and liquid assets to carry into the next period. The alternative case occurs when either $a_t \geq 0$ and $a_{t+1} < 0$

or when $a_t < 0$ and $a_{t+1} < a_t$. In this case the household extracts equity from the home and hence has to pay the proportional refinancing cost η so that the budget constraint will read:

$$c_t + \eta h_{t+1} + a_{t+1} \leq y_t + a_t(1 + r(a_t)) \tag{8}$$

This constraint tells us that non financial income and liquid assets plus earned interest are used to purchase consumption, assets for the next period and to pay the refinancing cost. In both cases in which the household does not change the owned house represented in equations 7 and 8 the borrowing constraint $a_{t+1} \geq -(1-\theta)h_{t+1}$ must be added to the problem. Finally we need to separately add that $h_{t+1} = h_t$ in both cases since the size of the home remains the same.

With the description of the model given above we can next move to formulate the dynamic programming problem that describes the household optimization problem. The state variables are given by the quadruple (a_t, h_t, z_t, D_t) where as anticipated in the previous description a_t and h_t are the holdings of liquid assets and the stock of owned housing at age t, z_t is the persistent component of the shock to the labor income at age t and D_t is the indicator function for the current home ownership status.[2] Setting $X \equiv (a, h, z, D)$ the value function then reads:

$$\begin{aligned} V_t(X_t) = \max_{c_t, a_{t+1}, h_{t+1}, f_t, D_{t+1}} & \{u(c_t, D_{t+1}h_{t+1} + (1 - D_{t+1})f_t) + \\ & \beta p_t \mathbf{E_t} \left[V_{t+1}(X_{t+1})\right]\} \end{aligned} \tag{9}$$

where β is the subjective discount factor. The maximization is performed under a budget constraint, a borrowing constraint and the process for labor income defined above. The borrowing constraint is simply:

$$a_{t+1} \geq -(1 - \theta)h_{t+1} \tag{10}$$

[2]In practice one could remove the indicator function as a state variable since $h_t = 0$ can stand for the agent being currently a renter and $h_t > 0$ for the household being a homeowner. We leave a separate state variable for home ownership for clarity of exposition.

In order to express compactly the budget constraint we introduce two more indicator functions. The first denoted I_m takes a value of 1 if a household that is a homeowner in both the current and next period decides to change home and a value of 0 otherwise. The second indicator function I_η takes a value of 1 if the household decides to extract home equity and 0 otherwise. Notice that I_η can take a value of 1 only if $I_m = 0$, that is the household does not change house and $a_{t+1} < 0$ and $a_{t+1} < a_t$. If $D_t = 0$, that is, the household enters the period as a renter then the budget constraint will be:

$$c_t + D_{t+1}h_{t+1} + (1 - D_{t+1})r_f f_t + a_{t+1} \leq a_t(1 + r(a_t)) + y_t \tag{11}$$

Alternatively if the household enters the period as a homeowner the budget constraint will read:

$$c_t + D_{t+1}I_m\Big\{h_{t+1} - h_t(1-\phi)\Big\} + D_{t+1}(1-I_m)I_\eta\eta h_t + (1-D_{t+1})\{r_f f_t - h_t(1-\phi)\} + a_{t+1} \leq a_t(1+r(a_t)) + y_t \tag{12}$$

Substituting the values of the decision variable D_{t+1} and the indicator functions I_m and I_η one can obtain in turn each of the budget constraints described above in equations 3 to 8. More specifically in equation 11, setting D_{t+1} to 0 gives the constraint in equation 3 while setting it equal to 1 gives equation 4. Similarly in the constraint 12 a value of D_{t+1} equal to 0 leads to constraint 5. On the other hand if in equation 12 we plug the value 1 in D_{t+1} we obtain constraint 6 if we set I_m to 1 and constraints 7 and 8 if we set I_m to 0, the latter two depending on I_η being 0 or 1.

3 Calibration

We calibrate the model to the US economy. We set some parameter values to target key macro aggregates and borrow the others from the literature.

We start by describing the parameters that are taken from external sources. The model period is 1 year. Agents enter the model at age 20 and live for a maximum of 100 years. They work the first 45 years, thus retiring at age 65 so that we set T_R and T equal to 46 and 80 respectively. We take survival probabilities from the Berkeley Mortality Database for the male population. The coefficient of relative risk aversion α is set to 2, a standard value in the literature. The interest rates are taken from Kaplan, Mitman and Violante (2020). This implies that we take a value of 0.03 for the lending rate $r(a_t)|_{a_t>0}$ and 0.04 for the borrowing rate $r(a_t)|_{a_t<0}$. The rental rate for housing services r_f is set to 0.050 based on the work of Davis, Lehnert and Martin (2008). The value of the minimum down payment θ is set at 0.2 a value commonly used in the literature and that we take from Li et al. (2016). The values that determine the process for earnings, that is, the autocorrelation coefficients ρ_t, the variances of the persistent shocks $\sigma^2_{\varepsilon,t}$ and of the temporary shock $\sigma^2_{\xi,t}$ are taken from the original estimates of Karahan and Ozkan (2013). Since these parameters are age changing there are 3 times T_R of them, hence we refer directly to the cited work for their exact values. With this exception the parameters described above are summarized in the top panel of Table 1.

The remaining parameters are calibrated internally so that certain moments generated by the simulated data match their empirical counterparts in the US data. These parameters are the subjective discount factor β, the weight of housing in the utility function σ, the minimum house size $\underline{h}$, the transaction cost for selling a house ϕ and finally the refinancing cost η. The targeted moments that are used to fix these parameters are the median wealth-to-income ratio, the median housing-to-income ratio, the home ownership rate for the working-age population, the proportion of homeowners that sell a house in each period and the proportion of homeowners that decide to refinance their housing debt in each period. We calculated the first two of these five moments using SCF data for 1989 and 1992.[3] The home ownership rate for the population below age 65 is taken from the U.S. Census Bureau, 1998. Finally, the proportion of households selling a house and refinancing a mortgage are taken from the

[3]We choose the 1989 and 1992 issues of the SCF because they overlap with the data period in the original estimation of the insurance coefficients by Blundell et al. (2008).

values reported in Kaplan et al. (2020). While all the parameters interact to determine jointly the simulated moments of interest, some have a decisively stronger impact on specific moments. In the third column of the bottom panel of Table 1 we thus report for each of the given parameters the corresponding moment that it affects more strongly.

The discount factor turns out to be equal to 0.970, a value that is in line with what is usually found in the macroeconomic literature. The weight of housing in the utility function is 0.892. The refinancing cost, needed to tap home equity and that in our model is defined over the house value turns out to be 0.006, very close to the value used in Li and Yao (2007). With respect to the minimum house size this turns out to be 4 times average earnings at age 20 or about 64000 dollars when scaled to the values in the 1992 edition of the Survey of Consumer Finance. The calibrated transaction cost for selling a house turns out to be 0.005, which is lower than what is commonly used in the literature where the value is taken from external sources. In our model the only reasons for changing home are related to changes in the value of the earnings shocks or to aging, hence in both cases changes in the optimal house size. In the data there are moves that are determined by job changes or family reasons. We conjecture that the low value of the transaction cost that we find stems from the fact that our empirical target includes all those other reasons for changing home on top of the only one allowed in the model. This conjecture is also consistent with the findings in Vestman (2019) that in a calibrated life-cycle model of housing and portfolio choice where the cost on housing transactions is set higher based on external sources very few moves occur endogenously for reasons related to consumption and savings choices.

We finally proceed to describe two further aspects of housing. First, house size is treated as a discrete variable. We employ 11 points for the grid of house sizes. The lowest point is chosen endogenously as described above. The highest point is set at a value corresponding to the 90th percentile of the house price distribution for age 20-30 households like in Attanasio, Kovacs and Moran (2024). Second, implicit in the formulation of the model is a constant price for houses. We make this assumption based on results in Goetzmann and Spiegel

(2000) who report that average yearly real house price appreciation in the US from 1980 to 2000 was only 0.5% and go on to claim that once one takes into account property taxes and maintenance cost the actual financial return on housing is even lower. Given that we do not model these housing related costs and the time framework of the estimated insurance coefficients by Blundell et al. (2008) this seems an appropriate choice.

Table 1: Calibration Results: Parameter Values

Parameter	**Description**	**Source/Target Moments**	**Value**
	Panel (a): Externally calibrated		
α	Coeff. relative RA		2.000
$r(a)_{a>0}$	Lending interest rate	Kaplan et al. (2020)	0.030
$r(a)_{a<0}$	Borrowing interest rate	Kaplan et al. (2020)	0.040
θ	Downpayment	Li et al. (2016)	0.200
r_f	Rental rate	Davis et al. (2008)	0.050
	Panel (b): Internally calibrated		
β	Discount factor	Wealth-to-income (median)	0.970
σ	Non-durables weight in utility	Housing-to-income (median)	0.892
$\underline{h}$	Minimum house size	Home-ownership rate	6.235
ϕ	Transaction cost for selling	Prop of HOWs selling homes	0.005
η	Mortgage refinancing cost	Prop of HOWs ref mortgages	0.006

In table 2 we perform a check on the quality of the calibration by reporting the values of the targeted moments and their empirical counterpart. Overall the quality of the match is very good with only minor discrepancies in certain moments.

Table 2: Calibration Results: Targeted Moments

Moments	**Source**	**Data**	**Model**
Wealth-to-income (median)	SCF (1989, 1992)	1.405	1.403
Housing-to-wealth (median)	SCF (1989, 1992)	0.861	0.860
Home-ownership rate	U.S. Census Bureau, 1998	0.675	0.677
Prop of HOWs selling homes	Kaplan et al. (2020)	0.100	0.100
Prop of HOWs ref mortgages	Kaplan et al. (2020)	0.060	0.060

4 Results

In this section we report the results of the quantitative analysis of the model. We split the section in two parts. In the first subsection we present life-cycle profiles of several variables that are meant to show that beyond matching the target ratios in the calibration the model does a reasonably good job at matching wealth, consumption and home-ownership rates during working life. This validates it as a suitable tool to analyze the question of interest. The second subsection focuses on the main results of the model, that is, the ones concerning insurance coefficients against persistent and temporary shocks and how those are affected by the presence of an illiquid asset like housing. We obtain the results in the section by simulating a cohort of 50000 household over the life-cycle.

4.1 Life-cycle profiles

We start by reporting the life-cycle profiles of wealth accumulation in the model. This is done in Figure 1. Agents start life with zero wealth, then wealth start accumulating up to the retirement age, after which it declines sharply to reach again zero by the maximum allowable age. The profile for the liquid and illiquid components are similar although housing wealth

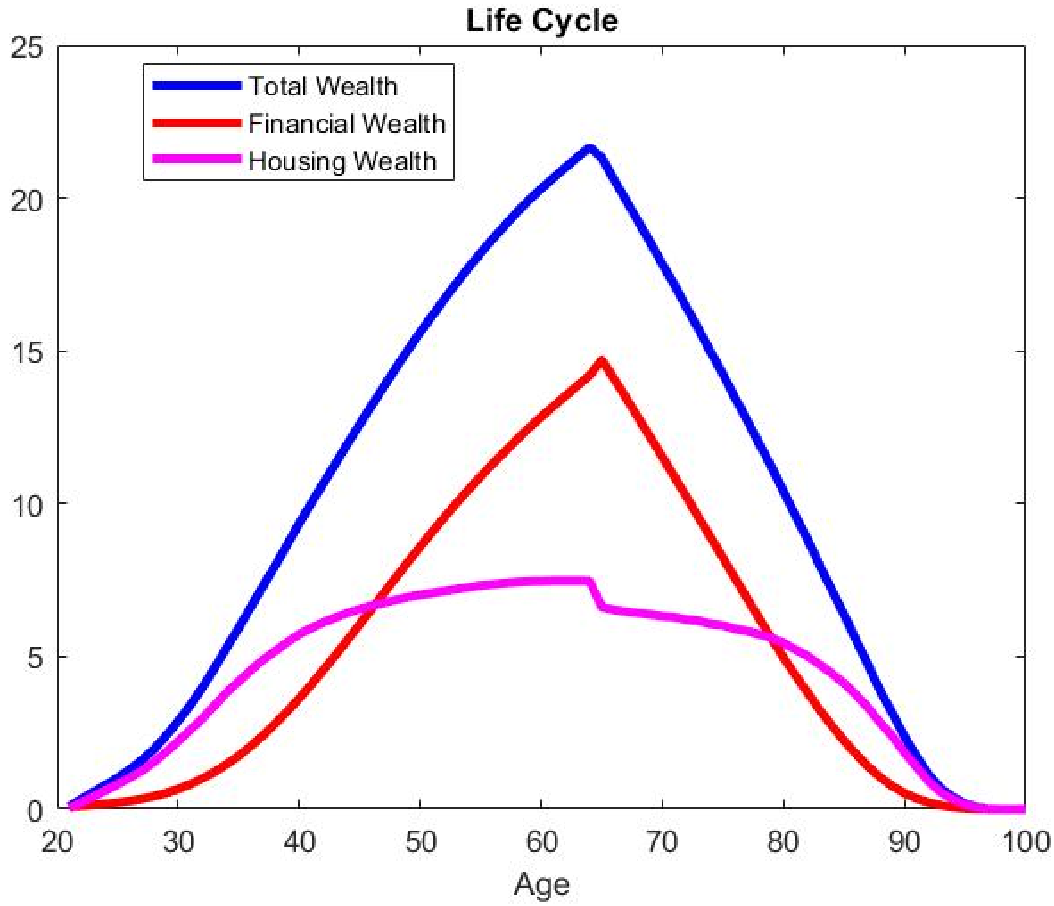

Figure 1

This figure reports the life-cycle profiles of average wealth. It does that both for total wealth and separately for financial and housing wealth.

accumulation shows a smoother variation. In the data wealth is not much decumulated after retirement like in the model. The different model behavior reflects the lack of specific features to discourage dissaving late in life like bequest motives and the risk of out-of-pocket medical expenditures. This is not a problem for our results though since our focus is on insurance against labor earnings shocks for which what matters is the life-cycle profile of wealth during working life.

To this end we next report the life-cycle profiles of net worth to earnings and of gross housing wealth to earnings in the model and in the data. The data are elaborated based on the 1989 and 1992 issues of the Survey of Consumer Finances. Figure 2 reports the median

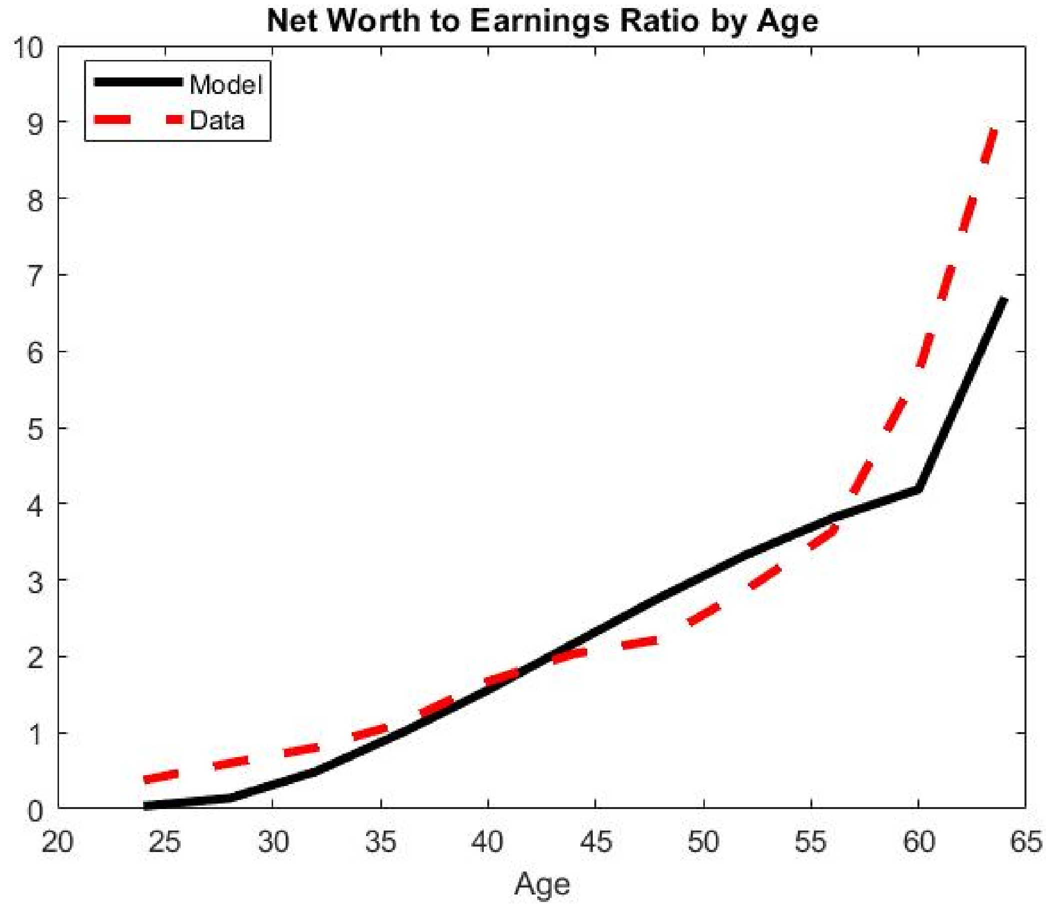

Figure 2

The figure reports the life-cycle profile of the median net worth to earnings ratio in the model (continuous line) and in the data (dashed line).

ratio of net worth to earnings for the working age population in the model and in the data. As can be seen the model simulated profile matches quite closely the profile in the data up to the age of 60 and only somewhat underestimates it between age 60 and 65.

Figure 3 reports the profile of the median gross housing wealth to earnings ratio during working life. In this case the model matches quite closely the data at all ages considered.

Next, figure 4 reports the percentage of households that are homeowners at each age in the model and compares them with the data. The model generated profile tracks the one in the data quite well again. Model home ownership rates somewhat under-predict the one in the data at younger ages while slightly over-predicting them in mid-life. This is a

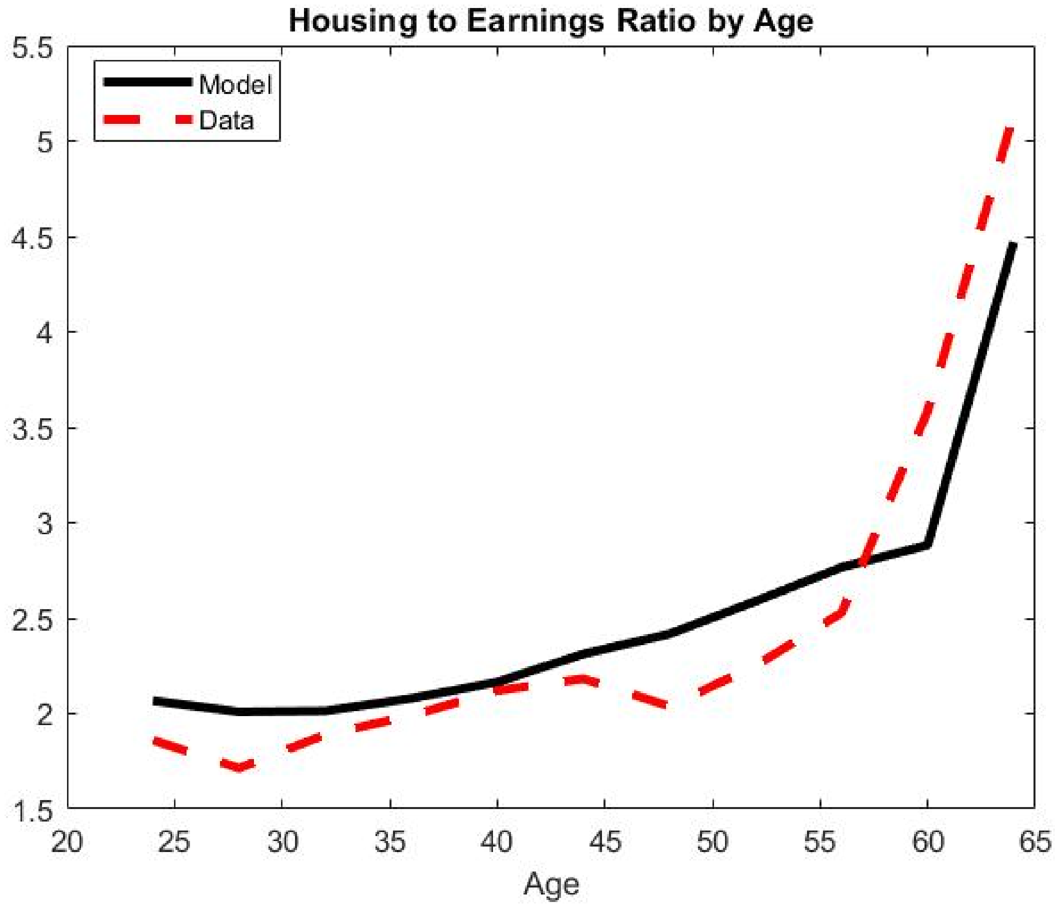

Figure 3

The figure reports the life-cycle profile of the median gross housing wealth to earnings ratio in the model (continuous line) and in the data (dashed line).

consequence of the fact that in the model every agent starts life with zero wealth, hence needs some time to accumulate the funds needed for the initial down-payment. In the data this is not true, moreover households may receive inter-vivos transfer early in life to satisfy the down-payment requirement. The fact that the model exceeds home ownership rates in mid-life stems from the under prediction mentioned above for the early part of the life-cycle and the fact that we match the aggregate average rate as a calibration target.

Finally figure 5 reports the evolution of the variance of log non durable consumption over the life-cycle in the model. This variance declines in the first few years of working life and then increases by slightly more than 10 points from about age 27 to the late part of

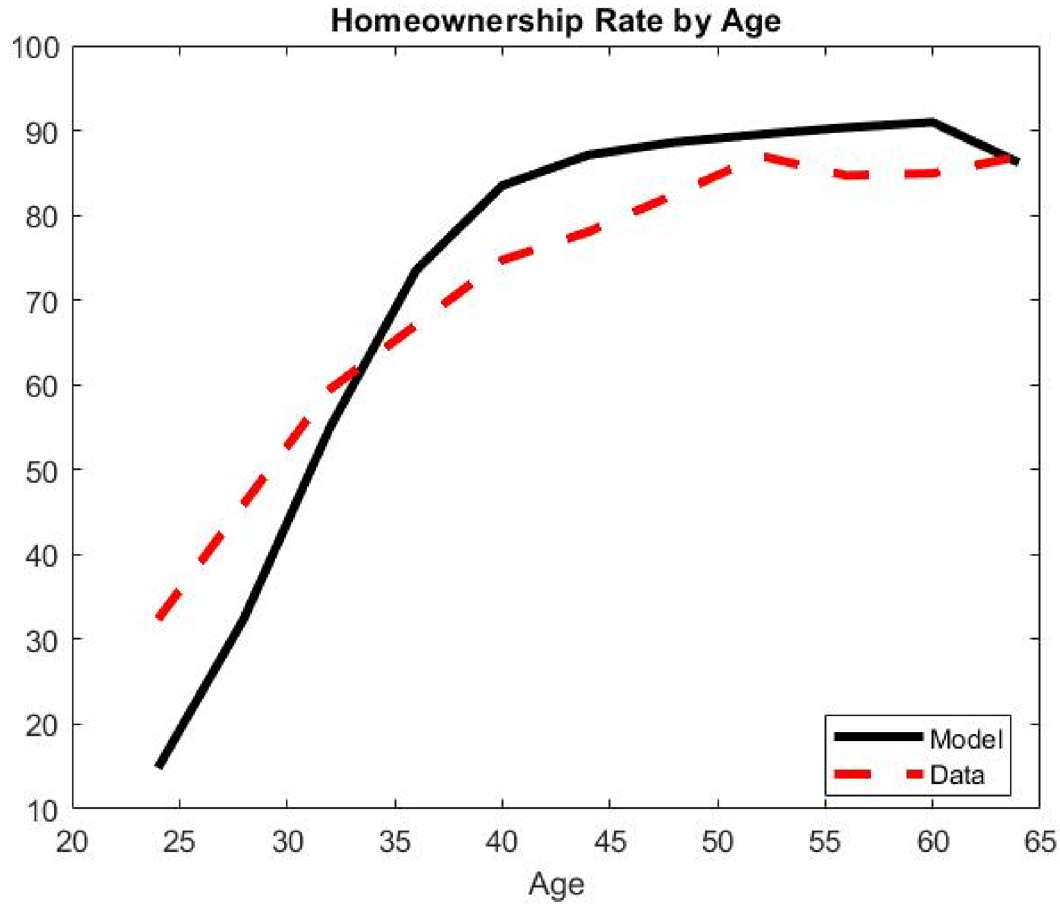

Figure 4

The figure reports the life-cycle profiles of the home ownership ratio in the model (continuous line and in the data (dashed line).

the life-cycle. The monotonically increasing path starting at about age 30 is in line with the findings reported in Heathcote, Perri and Violante (2010). At the quantitative level the growth in consumption inequality is also in line with the estimates in the cited paper when cohort effects are taken into account and slightly overestimates those that the authors obtained when controlling for year effects.[4]

Overall the comparison performed above suggests that the model developed here represents with reasonable accuracy the main patterns of wealth accumulation, home ownership

[4]Where our model differs from the estimates in Heathcote et al. (2010) is in the reduction in the variance of log consumption early in life, which is not observed in the data at least for the age range 25 to 30 for which the estimates are available.

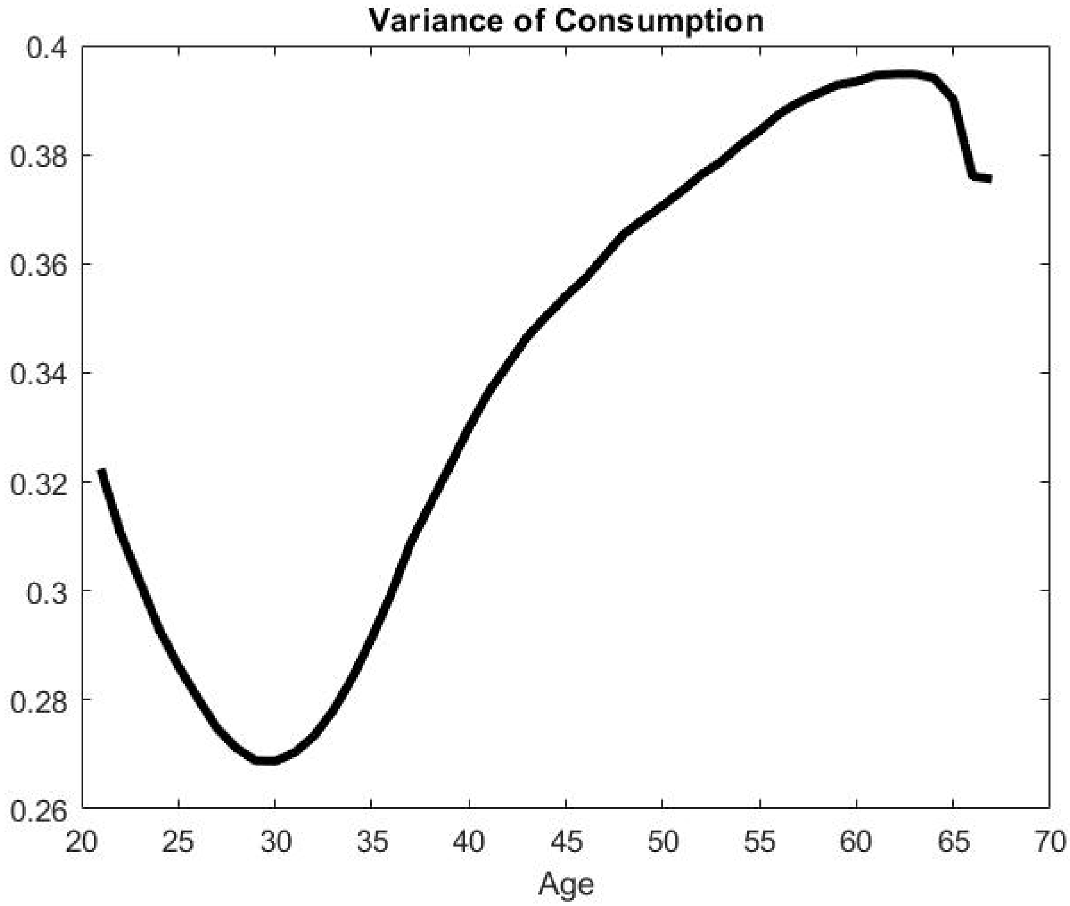

Figure 5

The figure reports the life-cycle profile of the variance of the log of non durable consumption in the population.

rates and volatility of consumption observed in the data.

4.2 Insurance coefficients

The main question we address in this paper is how households' consumption responds to shocks to the persistent and transitory components of the earnings process. This is done by computing so called insurance coefficients. Following **?** we define insurance coefficients as:

$$\phi^x = 1 - \frac{cov(\Delta c_{i,t}, x_{i,t})}{var(x_{i,t})} \tag{13}$$

where $c_{i,t}$ is the log deviation of consumption for household i at age t from its deterministic life-cycle trend and $x_{i,t}$ is a shock to the labor income process against which we measure insurance, in our model the persistent/permanent and temporary shock. In the model shocks are observable hence one can directly apply the formula in equation 13. Shocks to the earnings process are not directly observable in the data. However BPP showed that under suitable identifying restrictions the variances and covariances needed to evaluate the coefficients in equation 13 can be estimated from observable variables alone, that is, current and lagged income and consumption. We report estimates from the model using the directly observed shocks and equation 13 which we call "Model True"and compare them with BPP estimates obtained from the data. The latter will be called "Data BPP"coefficients.

The estimator used by BPP though requires that certain assumptions be satisfied otherwise the estimates will be biased. Among them is the fact that solutions are interior, that is, they are away from the borrowing constraint. Another assumption that is required is that the persistent shock be truly permanent, something that several studies have questioned and that we do not assume in our model. For these reasons in a separate section we will also report insurance coefficients computed by applying the BPP estimator to our simulated data. We call those coefficients "Model Based BPP". Comparison of true and estimated coefficients in the simulated data will give us an idea of the potential bias in the data based estimates, hence on the validity of comparing data and model coefficients.[5]

Table 3 reports the aggregate results, that is, the results that are computed over all the population of working age households. In the top row we report the insurance coefficients in the data, based on the estimates in BPP. The values are 0.36 for the permanent shock and 0.95 for the temporary shock. The meaning of these figures is that, according to the data, given a shock to earnings only 5 percent of it is passed to consumption if the shock is temporary and 64 percent is passed when the shock is permanent. The next panel reports the values of the true insurance coefficients computed from simulated data in the model, first

[5]The procedure outlined in this and the previous paragraph is based on Wu and Krueger (2021)

Table 3: Insurance coefficients: Model true and data BPP

	Permanent shock	Temporary shock
Data	0.36	0.95
No housing	0.480	0.890
Baseline	0.440	0.880
$\phi = 0$	0.451	0.888
$\eta = 0$	0.445	0.886
$\phi = 0, \eta = 0$	0.453	0.889
$\phi = 0.07, \eta =$ baseline	0.426	0.872
No selling, No refinancing	0.416	0.866

for a model with only one liquid asset and then for the baseline model with housing. The two models are otherwise equally calibrated, in particular the median wealth-to-earnings ratio for total wealth is kept constant. In the model where all wealth is considered fully liquid the insurance coefficient for the temporary shock is 0.89 and the insurance coefficients against the persistent shock is 0.48. This is 0.12 points above the empirical one. Earnings shocks in the model have limited persistence hence they are easily insurable. When part of wealth is made by illiquid housing the insurance coefficient against the temporary shock is 0.88, basically unchanged. The insurance coefficient against the permanent shock instead goes down to 0.44 reducing by a third the gap with the data counterpart. The intuition is simple: housing is illiquid, hence part of the household wealth cannot be used to smooth consumption unless a cost is paid to tap equity, or worse to sell the house and move to one of different size or to renting. As a consequence the insurance coefficients against permanent shock suffers a sizable decrease.

In the third panel of table 3 we check the impact on the insurance coefficients of eliminating in turn the frictions that make housing illiquid. This is obtained by changing the selling cost ϕ and the refinancing cost η. These are the only parameters that we touch in each of the experiment, that is, we do not recalibrate the other parameters to match all the targets in the baseline model. Setting ϕ to 0 increases the insurance coefficients against temporary shocks from 0.880 to 0.888, almost 1 percentage point. The effect of setting η to 0 is somewhat smaller: in this case the coefficients raises to 0.886. Finally when both parameters are set to 0 the insurance coefficient against the temporary shock raises to 0.889. Moving to the insurance coefficients against the persistent shock, when the selling cost is set to 0, the insurance coefficient raises from 0.440 to 0.451, when the refinancing cost is set to 0 the insurance coefficient becomes 0.446, finally when both are set to 0 the insurance coefficient further increases to 0.453. In the fourth panel we report results for two experiments that on the contrary raise the degree to which housing is illiquid. First we increase the selling cost to 0.07, that is, a value that is in the range of values considered by papers that take it from external sources rather than calibrating it internally as we did here. Under this higher selling cost the first line of the fourth panel of table 3 shows that the insurance coefficient against permanent shocks decreases to 0.426, while the coefficient for temporary shocks decreases to 0.872. This result strongly suggests that under an alternative modeling strategy with exogenous moving shocks and transaction costs on housing based on external sources, the impact of introducing illiquid housing on the insurance coefficients would not change quantitatively in a significant way compared to the case considered here. Finally we performed a further experiment where we raised both ϕ and η to levels such that no refinancing and no selling occurs.[6] In this extreme case the insurance coefficient against temporary shocks further drops to 0.866 and that against persistent shocks drops to 0.416 down from 0.44 in the baseline. Overall, then the effects of changing the degree to which housing wealth is illiquid does not generate big effects on the degree of insurance against the shocks.

We then move to the life-cycle profiles of those coefficients to understand the reasons behind the muted response of the observed coefficients to changes in the degree of liquidity of housing. Figure 6 reports the true insurance coefficient against persistent shocks calculated

[6] While we could easily raise η to a level such that no refinancing occurs, we bumped into instabilities in the execution of the code past certain values of ϕ. Strictly speaking then house sales are not driven exactly to zero.

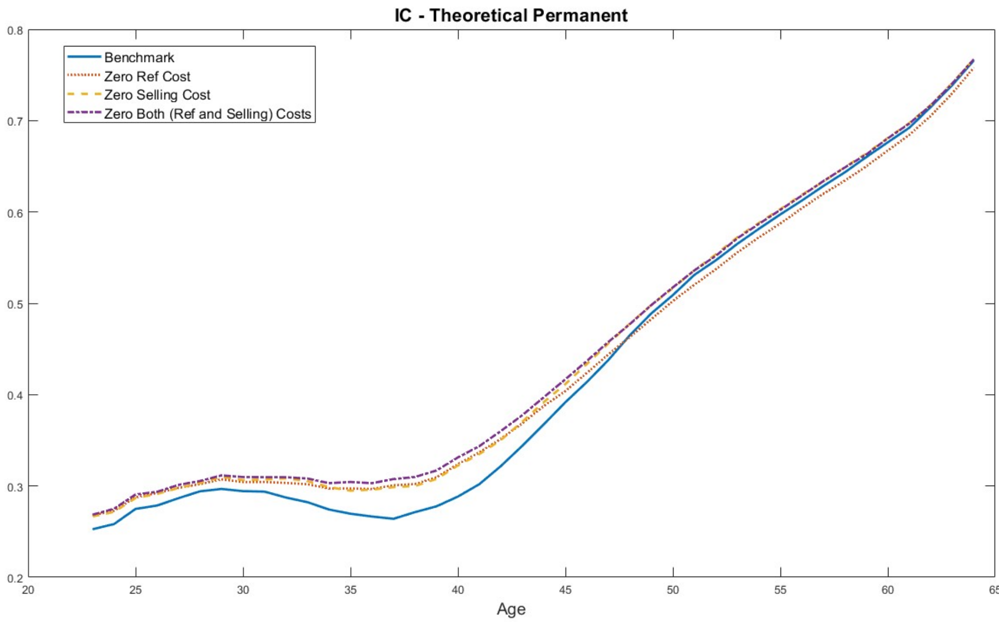

Figure 6

This figure reports the life-cycle profile of the insurance coefficients against the persistent shock in the benchmark model and in alternative specifications of the model where the selling cost and/or the refinancing cost have been turned off.

on model data by age. The life-cycle profile of the coefficient starts flat at a value of 0.24 at the beginning of working life and with some fluctuations remains roughly speaking constant up to age 40. After that it takes a decisive turn and increases up to values of almost 0.8 by retirement age 65. Turning to the comparative analysis of the different lines we see that they are close to each other from entry into the labor market to age 30, then the line for the benchmark case significantly falls relative to the one corresponding to the cases where frictions in selling the house and/or refinancing debt are eliminated. Finally the four lines get closer again and become virtually overlapping from age 50 on. At the peak the difference between the insurance coefficient in the benchmark case and in the zero cost case is a bit higher than 0.04, a sizable one sixth of the benchmark model value at the same age. The interpretation of this result can be understood by looking jointly at figures 4, and 7. In general it must be observed that the frictions in selling the house or refinancing housing debt are relevant essentially for those homeowners that have recently bought a home, hence

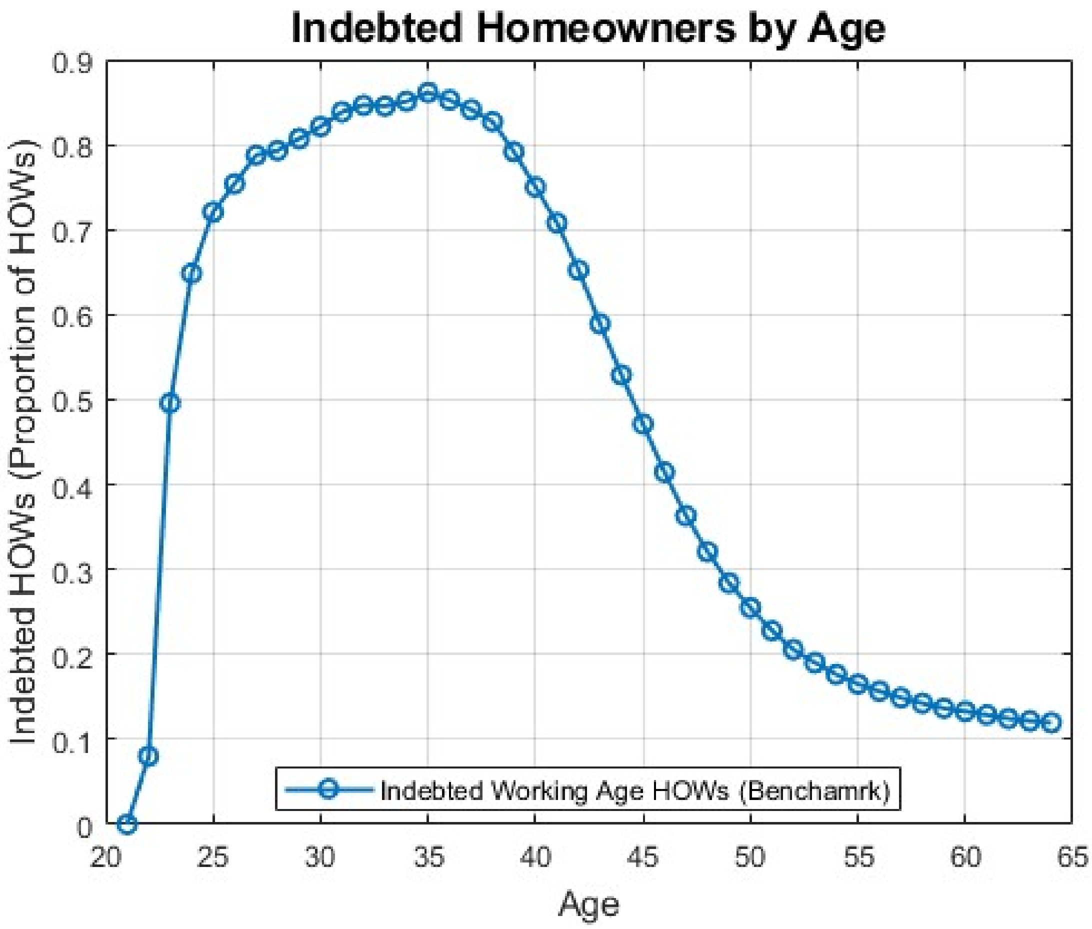

Figure 7

This figure reports the fraction of households that own a house and that have negative liquid wealth over the working life.

they have not yet fully repaid their debt so that in case of negative earnings shocks they would need to tap home equity or sell the house to smooth non durable consumption.

As figure 4 shows, before age 30 the fraction of homeowners is still relatively small, on average around 30 percent, so the fraction of the total population of this age group that could benefit from eliminating the transaction costs is small too. Starting around age 30 the fraction of homeowners picks up very quickly. At the same time for about fifteen years a sizable proportion of those homeowners are still repaying their debt. As figure 7 shows about 80 percent of homeowners are indebted at age 30 and while on a steep decline this percentage is around 40 at age 45. As a consequence a somewhat higher fraction of this age group is in a position to benefit from the removal of the transaction costs. Approximately past age 45 though, while the fraction of homeowners has reached the plateau of 80 percent of the total population, the fraction of those households that still hold debt declines to small values. The fraction of households of those age groups that potentially benefit from the elimination of the selling and refinancing cost again becomes small, so the insurance coefficient curves of the benchmark case and those obtained when removing the frictions in the housing market get again very close to each other.

Summarizing insurance coefficients in the working-age population can be thought of as an average of age-specific coefficients. Since for most age groups, a small fraction of the population is affected by the lack of liquidity of housing, removing transaction costs has an impact on insurance coefficients only for households in a narrow age range, and hence a very small impact on the overall coefficient.

4.3 Assessing the insurance coefficients

In the previous section we showed that under the assumption that shocks that are not purely temporary still have only a small degree of persistence, like in recent estimates, the insurance coefficients against these shocks exhibited by a standard incomplete market model (SIM) with one fully liquid asset exceeds the ones estimated in the work by Blundell et al. (2008). When the model is extended to take into account the fact that a substantial fraction of households' wealth is tied in the form of illiquid housing, the gap between the insurance coefficients against persistent shocks generated by the SIM model and the one in BPP estimates is reduced by about a third but still remains above the empirical counterpart. When it comes to the temporary shock the insurance coefficient generated by the one asset model is slightly below the one estimated in the data, about 6 percent less. In this case adding illiquid housing to the model further reduces the insurance coefficients but by a quantitatively negligible amount.

The validity of these results rests on the assumption that the BPP methodology to estimate insurance coefficients does not suffer from major biases, and hence it correctly measures insurance in the data. These biases may arise for two reasons. First the BPP estimator is unbiased under certain assumptions that are violated at the borrowing constraint. Second the strategy to separately identify the two shocks from observations on earnings alone is valid as long as the more persistent shock is actually permanent but it is not in the case of persistent but not permanent shocks.[7] In order to assess the impact of the above mentioned violations of the identifying assumptions, in this section we apply the BPP methodology to model simulated data and compare the resulting insurance coefficients, that we label "model based BPP"coefficients with the true ones. The results from applying the BPP methodology to model simulated data are reported in table 4. Looking at the first column in the table we see that the model based coefficient for the persistent shock is 0.465, that is, 0.015 units below the true coefficients in the case of the model with no housing and it is 0.428 in the model with housing, that is 0.012 points below the true one. The second column in table 4 shows that the "model-based BPP"insurance coefficient against the temporary shock is 0.846 in the model with one asset and 0.829 in the model with housing. The difference with the true coefficients is now 0.045 and 0.051 respectively. Overall in the case of the permanent coefficients the size of the bias is very small. Moreover, if the bias in the data estimates is similar in sign and size to the one in the model, the conclusion that adding housing brings the model coefficients closer to the estimated one would be further strengthened. In fact, in

[7] See Kaplan and Violante (2010).

Table 4: Insurance coefficients: Model based BPP

	Permanent shock	Temporary shock
No housing	0.465	0.846
Baseline	0.428	0.829

this case the true coefficient in the data would also be slightly higher than BPP estimates, reducing the gap between model and data coefficients and doing so proportionally more for the model with housing where this gap is already smaller. The difference between estimated and true model-based coefficients is somewhat larger in the case of the temporary coefficients but it would still not change the main conclusion that the SIM model under the stated assumptions in this paper leads to insurance coefficients against temporary shocks somewhat smaller than in the data, both in the model with and without illiquid housing and that the difference between the two models is minor in this respect.

4.4 The model with permanent plus temporary earnings shocks

In this section we report results for a version of the model where the labor earnings process is characterized by a fully permanent and a purely temporary shock. Even though recent evidence points to the fact that this is too simple a characterization of the earnings process to match the joint data on income and consumption, this was the original specification used in Kaplan and Violante (2010) and we report it here since the comparison with the benchmark case will help us gauge the relative importance of persistence in shocks versus illiquid housing wealth in determining the insurance coefficients. The parameters of the earnings process in this case are taken directly from the study of Kaplan and Violante (2010), that is, we set the variance of the permanent shock to 0.01 and the variance of the temporary one to 0.05. The remaining parameters are calibrated exactly as in the benchmark economy. We consider the case corresponding to the baseline parametrization of the benchmark economy and we report both the model true and the model BPP coefficient. This is done in table 5. Starting with the first row in the top panel of the table we can see that in the one asset model the true insurance coefficient against the permanent shock is 0.267 and the corresponding coefficient for the temporary shock is 0.877. The former is about 0.2 below the corresponding coefficient in the model using the earnings process in Karahan and Ozkan (2013), the latter is only about 0.01 below. The effects of introducing housing in the model parallel those obtained under the Karahan and Ozkan (2013) process: the insurance coefficient against the permanent shock falls by a sizable, but not big amount, that is, about 0.04; the coefficient of the temporary shock falls by a smaller amount, that is, 0.02 points. The intuition for the results presented above is that shocks in the Karahan and Ozkan (2013) process show a limited degree of persistence, especially at young ages, hence they are easier to insure. In turn, this promotes precautionary savings allowing households to better insure not only the persistent shock but also shocks that are purely temporary.

The bottom panel of 5 reports the model-based BPP coefficients. A comparison with the model true coefficients immediately reveals that in the case of the permanent shock these are substantially lower than the true ones: in the one asset model the drop is by 0.09, from 0.267 to 0.176 and in the model that adds illiquid housing the drop is about 0.10, from 0.232 to 0.129. This result follows from the fact that the conditions under which the BPP coefficients are unbiased are valid for interior solutions only, thus fail at the borrowing constraint. Permanent socks are much less insurable than shocks that only show some degree of persistence. For this reason households have a weaker incentive to save hence they more often hit the borrowing constraint. This result also points to the fact that if the true data generating process for earnings is one that consists of a permanent plus temporary shock then the estimates in Blundell et al. (2008) are downward biased, hence a direct comparison of the true coefficients calculated on the model simulated data with the one estimated on actual data would substantially increase the gap between the SIM model implied insurance and the true one.

We can summarize the results from this section and its comparison with those for the

Table 5: Insurance coefficients: Model with permanent plus temporary earnings shocks

	Permanent shock	Temporary shock
Data	0.36	0.95
Model true		
No housing	0.267	0.877
Housing	0.232	0.856
Model based BPP		
No housing	0.176	0.874
Housing	0.129	0.854

baseline model in the following way. Introducing a second asset that is illiquid and interpreting it as housing reduces the insurance coefficients. This reduction is about 0.04 for both persistent and permanent shocks. This number is non-negligible but not big either. Even if housing was much less liquid as in some of the sensitivity analysis experiments performed in section 4.2 the reduction in the insurance coefficients would not exceed 8 percentage points. On the other hand, the effects can be more dramatic if one changes the characteristics of the earnings process: moving from the traditional permanent plus temporary and identically distributed shocks to a more carefully specified process that allows age changing and less than unitary persistence and also age changing variance of shocks like the one estimated by Karahan and Ozkan (2013) increases the insurance coefficients against permanent shocks by about 20 basis points. This is strong enough to overturn the qualitative findings associated with considering illiquid wealth: in the latter case, this would move the insurance coefficients closer to the one in the data, whereas in the former it would move them further away.

5 Conclusions

In this research we have constructed a version of the standard incomplete market model that features a richer earnings dynamics modeled after the work of Karahan and Ozkan (2013) and two goods, non durable consumption and durable housing services. Housing services can be purchased through renting or owning, in which case housing is treated as an illiquid asset that provides a side collateral value. We used the model to revisit the question first explored in Kaplan and Violante (2010) of checking if the SIM model can match the insurance coefficient for non durable consumption proposed in Blundell et al. (2008). Our analysis shows that moving from the traditional permanent plus temporary process for earnings to a more realistic one that exhibits more limited and age changing persistence has a dramatic effect on insurance coefficients leading the model to show substantially more insurance against persistent shocks than the data. Introducing illiquid housing has instead a sizable but more limited impact on consumption smoothing. Since housing indeed reduces the extent of insurance it then moves the model somewhat closer to the data.

References

Aguiar, Mark, Mark Bils, and Corina Boar, "Who are the hand-to-mouth," *Review of Economic Studies*, 2024.

Attanasio, Orazio, Agnes Kovacs, and Patrick Moran, "Temptation and commitment: A model hand-to-mouth behavior," *Journal of the European Economic Association*, 2024, *22* (4), 2025–2073.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption inequality and partial insurance," *American Economic Review*, 2008, *98* (5), 1887–1921.

—, —, and Itay Saporta-Eksten, "Consumption inequality and family labor supply," *American Economic Review*, 2016, *106* (2), 387–435.

Campanale, Claudio and Marcello Sartarelli, "Life-cycle wealth accumulation and consumption insurance," *Journal of Macroeconomics*, 2024, *79*, 1–14.

Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White, "The distribution of wealth and the marginal propensity to consume," *American Economic Review*, 2023, *113* (12), 3173–3212.

Cerletti, Enzo and Josep Pijoan-Mas, "Durable goods, borrowing constraints and consumption insurance," *Center for Economic Policy Research, Discussion Paper*, 2012, (DP9035).

Davis, Morris A, Andreas Lehnert, and Robert F Martin, "The Rent-price ratio for the aggregate stock of owner-occupied housing," *Review of Income and Wealth*, 2008, *54* (2), 279–284.

Favilukis, Jack, Sydney C. Ludvison, and Stijn Van Nieuwerburgh, "The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium," *Journal of Political Economy*, 2017, *125* (1), 140–225.

Goetzmann, William N. and Matthew Spiegel, "The policy implications of portfolio choice in underserved mortgage markets," *Yale international center for finance working paper*, 2000, (00-18).

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante, "Unequal we stand: An empirical analysis of economic inequality in the UNited States, 1967-2006," *Review of Economic Dynamics*, 2010, *13*, 15–51.

Kaplan, Greg and Giovanni L Violante, "How much insurance beyond self-insurance?," *American Economic Journal: Macroeconomics*, 2010, *2* (1), 53–87.

___ **and** ___ , "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 2014, *82* (4), 1199–1239.

___ **, Kurt Mitman, and Giovanni L Violante**, "The housing boom and bust: Model meets evidence," *Journal of Political Economy*, 2020, *128* (9), 3285–3345.

Karahan, Fatih and Serdar Ozkan, "On the persistence of income shocks over the life cycle: Evidence, theory, and implications," *Review of Economic Dynamics*, 2013, *16* (3), 452–476.

Li, Wenli and Rui Yao, "The life-cycle effects of house price changes," *Journal of Money, Credit and banking*, 2007, *39* (6), 1375–1409.

___ **, Haiyong Liu, Fang Yang, and Rui Yao**, "Housing over time and over the life cycle: a structural estimation," *International Economic Review*, 2016, *57* (4), 1237–1260.

Nardi, Mariacristina De, Giulio Fella, and Gonzalo Paz-Pardo, "Nonlinear household earnings dynamics, self-insurance and wealfare," *Journal of the European Economic Association*, 2020, *18* (2), 890–926.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel, "Housing, consumption and asset prices," *Journal of Financial Economics*, 2007, *83* (3), 531–569.

Sodini, Paolo, Stijn Van Nieuwerburgh, Roine Vestman, and Ulf von Lilienfeld-Toal, "Identifying the Benefits from Homeownership: A Swedish Experiment," *American Economic Review*, 2023, *113* (12), 3173–3212.

Vestman, Roine, "Limited stock market participation among rentners and homeowners," *Review of Financial Studies*, 2019, *32* (4), 1494–1535.

Wu, Chunzan and Dirk Krueger, "Consumption insurance against wage risk: familiy labor supply and optimal progressive income taxation.," *American Economic Journal: Macroeconomics*, 2021, (13-1), 79–113.