Homeownership and Portfolio Choice over the Generations

Gonzalo Paz-Pardo*

November 2, 2022

Abstract

Earnings are riskier and more unequal for households born in the 1960s and 1980s than for those born in the 1940s. Despite improvements in financial conditions, younger generations are less likely to be living in their own homes than older generations at the same age. By using a life-cycle model with housing and portfolio choice that includes flexible earnings risk and aggregate asset price risk, I show that changes in earnings dynamics account for a large part of the reduction in homeownership across generations. Lower-income households find it harder to buy housing, and as a result accumulate less wealth. (*JEL* D31, E21, E24, G11, J31)

The economic conditions faced by young US households are radically different from those that their parents and grandparents experienced when they were their age. There is less mobility in the labor market, career-long positions are less and less prevalent, and earnings inequality has increased. While the labor incomes of high earners have increased

*European Central Bank, DG Research, 20 Sonnemanstraße 60314 Frankfurt am Main (Germany), gonzalo.paz_pardo@ecb.europa.eu This paper is part of my PhD thesis at University College London. I'd like to thank the editor, Ayşegül Şahin, and an anonymous referee for their helpful comments. I am grateful to Mariacristina De Nardi, for her invaluable guidance and support, and to Manuel Arellano, Orazio Attanasio, Anmol Bhandari, Richard Blundell, Chris Carroll, Giulio Fella, Carlos Garriga, Michael Graber, Dirk Krueger, Jeremy Lise, Ralph Lütticke, Rory McGee, Fabrizio Perri, Morten O. Ravn, Víctor Ríos-Rull, Kamila Sommer, and David Wiczer for very helpful conversations. I would also like to thank seminar participants at UCL, UPenn, University of Minnesota, ECB, Netspar, Cintia Workshop, Minneapolis Fed, Banca d'Italia, Copenhagen Business School, Nova, University of British Columbia, Copenhagen Business School, Banco de España, Banco de Portugal, University of Copenhagen, EUI, Bundesbank Household Finance Conference, Spanish Macroeconomics Network, and Carlos III for their useful comments and suggestions. I gratefully acknowledge financial support from the ESRC for my graduate studies. See Paz-Pardo (2022) for replication material.

substantially over time in real terms, income-poorer individuals have seen their earnings stagnate or decrease.[1]

Meanwhile, homeownership for younger generations has shrunk. Within the cohort born in the 1940s, at age 35 over 70% of households were living in houses they owned. The figure was ten percentage points lower for those born in the 1960s, and more than 20 percentage points lower for the early 'Millennials' born in the 1980s. This happened in a context in which financial markets have become more developed[2] and stock market participation has been increasing for younger generations. These distributional changes cannot be perceived in aggregate homeownership rates, which have remained stable during most of this period.

This paper studies the role of these changes in household labor income dynamics and financial conditions in explaining homeownership and portfolio composition across generations. To do so, it proposes two novel contributions. First, it designs a flexible, cohort and business-cycle dependent earnings process, based on Arellano, Blundell and Bonhomme (2017), that allows shocks to household labor income to be age-varying, non-normal, non-linear, and correlated with stock market returns and house prices, as in the data. Second, it builds and calibrates a rich life-cycle model with correlated aggregate and idiosyncratic risk, in which households decide their consumption, savings, housing stocks, portfolio share of safe and risky assets, and mortgage debt. Importantly, households only need to satisfy downpayment constraints and income tests at the time of mortgage origination, which implies that the outstanding mortgage can go above the value of the house if there is a negative shock to house prices. Households can also hold liquidity whilst they have a mortgage. In the model, homeownership, in addition to being a way of accumulating wealth, provides utility and insurance against rental price risk.

I use the model to compare the life experiences of three generations, namely, those born in the 1940s, 1960s, and 1980s. I assume that an American born in the 1940s differs from younger generations in three main ways. First, they face different experiences in the labor market. I use household earnings data from the Panel Study of Income Dynamics (PSID)

[1] These facts have been established in a large literature surveyed in Acemoglu and Autor (2011) and Goldin and Katz (2009). Guvenen et al. (2017), using US administrative data, and Borella, De Nardi and Yang (2019), using survey data, find decreases in median male wages in real terms between the cohorts born in the 1940s and the cohorts born in the 1960s.

[2] Dynan, Elmendorf and Sichel (2006) describe how financial deregulation, changes in risk-assessment methods, and the expansion of secondary markets increased the fraction of households with access to credit and how much those who already had access could borrow.

to estimate the earnings process separately for all three generations, thus incorporating the changes in earnings inequality and earnings risk in a flexible, data-driven manner. I separate the persistent and transitory components of earnings, which allows me to control for potential measurement error in the survey. Second, they face different conditions in financial and housing markets. Housing has become more expensive over time with respect to average incomes, and different generations entered the labor market in different stages of the business cycle or the house price cycle. Third, the 1980s generation faced particularly looser financial constraints when they started to buy houses in the early 2000s, which I capture with a reduction in downpayment constraints.

Time, age, and cohort are explicit in the model. Average earnings, homeownership, and stock market participation at each age differ across generations as they do in the data. I do not homogeneize age profiles across cohorts and thus do not need to disentangle year and cohort effects to obtain them.[3] I adopt the actual realizations of house prices and stock market returns each year from historical data, and use the Survey of Consumer Finances (SCF), including its earlier versions dating back to 1963, to obtain information about household portfolio compositions by age and generation.

The main results are as follows. First, intergenerational changes in earnings dynamics, asset returns, and housing prices obtained from the data fully explain the differences in homeownership between the 1940s and 1960s cohorts. For the 1980s cohort, who started to buy houses in the early 2000s, looser borrowing constraints partially counteracted the effect of high house prices. I do not need to assume that preferences have changed to explain the lower homeownership rates for younger generations.

To isolate the effect of changes in labor market income dynamics, I perform a counterfactual experiment in which I attribute the earnings process of the 1940s cohort to the younger generations, whilst keeping all other elements of the model constant, including house prices. More than half of the difference in homeownership at age 30 for both generations can be accounted for by changes in earnings inequality and risk. Not all of it is due to delayed home-buying: changes in earnings dynamics still have an important effect at age 40 and afterwards. These results are robust to letting house prices adjust,

[3] Age, year, and cohort are collinear. To obtain age profiles in a sample with several cohorts and years, the usual practice is to either remove year fixed effects or cohort fixed effects, which can lead to very different implications. See Heathcote, Storesletten and Violante (2005) for a discussion on how the choice of removing year or cohort effects impacts measures of earnings and consumption inequality, and Ameriks and Zeldes (2004) for the effect on household portfolio shares.

assuming an empirically plausible level of housing supply elasticities.

The main driver of these changes is the increase in earnings inequality at labor market entry, with a more limited role for the increase in earnings risk. Households with lower initial and expected lifetime earnings find it harder or suboptimal to engage in a large expenditure like a house, which would leave them with a sizeable mortgage with respect to their current income, and thus exposed to income and house price risk.

Second, the increase in stock market participation of younger cohorts can be rationalized with a substantial reduction in stock market participation costs, which reflects easier information acquisition. Today, many workers who are starting new jobs either receive information about retirement accounts or are automatically enrolled into retirement plans like IRAs or 401(k)s.

These intergenerational changes also have implications for household wealth accumulation. In the 1940s generation, relatively poorer households bought housing. They did so to accumulate wealth, but also because they wanted to be homeowners, because it provided insurance against rental risk, and because leveraging with a mortgage allowed them to benefit from gains in the housing market. Similarly ranked households in younger generations are no longer buying houses and, because of these additional reasons to hold housing, they do not fully compensate the lack of housing wealth by saving in financial assets. Financial wealth now represents an increasing share of household portfolios, but it is more unequally distributed than housing wealth. The model predicts that lowering the cost of access to financial markets for lower and middle income households can increase their wealth holdings and reduce wealth inequality.

Overall, these findings suggest that changes in labor market income dynamics and in the housing market are having substantial effects in the life experiences of most Americans, and they can influence, in the longer term, the distribution of income and wealth, intergenerational mobility, and the effects of policies.

Related literature

This paper builds on the tradition of Bewley-Aiyagari life-cycle models that study optimal consumption and saving in the presence of realistic labor market income risk, and in particular those that have studied the role of housing on wealth accumulation (Fagereng et al., 2019) and the limited investment of households in stocks (Gomes and Michaelides

(2005), Cocco, Gomes and Maenhout (2005), Alan (2006), to name a few). Fewer contributions have explored, like I do, the interaction between housing, portfolio choice, and the life-cycle.[4] Cocco (2005) shows that younger and poorer investors have less financial wealth to invest in stocks because they start out by buying houses. My life-cycle model draws from this insight, but includes more realistic mortgages, for which downpayment constraints do not bind every period, and a much richer risk structure, both in terms of aggregate asset prices and idiosyncratic earnings risk.

I model idiosyncratic earnings risk based on a recent literature that has described its rich features (Guvenen et al. (2016)), such as age variation, non-normalities such as negative skewness, and nonlinear persistence, and their implications (De Nardi, Fella and Paz-Pardo, 2020). However, most of these processes abstract from cyclical fluctuations. I propose an extension of the econometric framework devised by Arellano, Blundell and Bonhomme (2017) that allows for business cycle variation in rich earnings dynamics in the form of a Markov-switching regime, and incorporates that the left-skewness of earnings shocks fluctuates over the cycle: during recessions, large drops in earnings become more likely (Guvenen, Ozkan and Song, 2014). This contrasts with more standard earnings process, that usually displays countercyclical variance (Storesletten, Telmer and Yaron, 2004). Another recent contribution that designs and implements an earnings process with variation in higher order moments over the business cycle is Busch and Ludwig (2017). I use a flexible nonparametric model that I estimate in panel data, while they define a rich parametric process and estimate it, à la Storesletten, Telmer and Yaron (2004), by using cross-sectional moments to identify the sequence of past shocks. Unlike theirs, my approach allows for variations in earnings dynamics over different cohorts.

Several previous studies have analyzed the implications of the changes in U.S. earnings dynamics on aggregate outcomes. Nakajima (2005) studies the impact of higher earnings inequality on portfolio allocations and asset prices with a general equilibrium model. In particular, he finds that increasing earnings inequality can spur demand for financial assets and thus decrease their return in general equilibrium, which in turn increases the demand for housing assets. This paper incorporates the increase in earnings inequality as a key force, and replicates exogenously the observed changes in asset prices.

[4]These include Flavin and Yamashita (2011), Yao and Zhang (2005), and Vestman (2012), who focuses on the role of preference heterogeneity to explain why homeowners participate more in the stock market. Becker and Shabani (2010) and Chetty, Sándor and Szeidl (2017) study the role of mortgage debt on portfolio allocations.

However, in my model there is a minimum house size, which not only generates a notion of homeownership, but also implies that households who want to hold more housing (because of precautionary reasons or because the return on financial assets is lower) might be constrained because of LTV or LTI requirements.

The increase in earnings uncertainty has also been related to the decrease in marriage rates (Santos and Weiss, 2016) and in fertility (Sommer, 2016). However, its link to homeownership rates across cohorts has received less attention. A notable exception is Fisher and Gervais (2011), who in a stationary equilibrium framework find that the increase in earnings uncertainty is a major candidate to explain the reduction in homeownership of the young between 1980 and 2000.[5] This paper builds on their contribution along several dimensions. First, I model each cohort separately, which allows me to better capture the earnings process, house prices, and cyclical histories they faced. Furthermore, I do not need to make assumptions about the convergence of future homeownership rates of younger cohorts with those we observe for the currently old. Second, in my model house prices are risky and agents can hold liquidity while they have a mortgage. Both affect the risk associated with buying a house: the former increases household exposure to risk, but the latter decreases it, because it allows them to better smooth income fluctuations. Third, I study the role of housing in the context of a richer household portfolio decision, and thus can accommodate possible substitution effects across asset classes as prices and returns change over time, which also allows me to draw conclusions about overall household wealth accumulation.

Despite the intergenerational changes, aggregate homeownership rates have been remarkably flat between the 1960s and the 1990s. They increased in the period leading up to the 2008 financial crisis: Chambers, Garriga and Schlagenhauf (2009) relate most of the increase to mortgage innovations. My model incorporates more flexible mortgages for the youngest generation and also finds that they had a positive effect in homeownership. Mankiw and Weil (1989) study the effect of generation sizes in housing demand and housing prices. They suggested that the baby bust episode that ocurred in the late 1960s would generate reduced housing demand in the 1990s. The forces described in this paper, that reduce within-cohort homeownership rates, generate effects in the same di-

[5] Fischer and Khorunzhina (2019) suggest that changes in divorce rates can also explain the reduction in homeownership rates, as increased divorce risk triggers precautionary savings for the young but reduces homeownership for older working-age households.

rection. Puzzlingly, house prices have continued to rise. This can be related with supply restrictions (Falcettoni and Schmitz, 2018), but also with the increased saving of older generations, both because it might drive down interest rates (Martin, 2005) and because they live longer and hold on to their houses for longer.

I An overview of intergenerational changes

The generations born in the 1960s and 1980s, when compared to that of the 1940s, have faced more earnings inequality, more earnings volatility, and more expensive house prices on average. They are less likely to be homeowners, but more likely to participate in the stock market. I now turn to empirical evidence from the U.S. to describe these differences in detail. Appendix A provides information about the data and sample selection.

The top panels of Figure 1 show median and mean earnings for US households by age and decade of birth. Although the earnings of the median male earner have actually decreased over these different generations, the large increase in female labor force participation acted as a counteracting force and kept median earnings constant in real terms, and higher after age 30 for the younger cohorts.[6] Comparing average and median earnings suggests that earnings have become progressively more right-skewed and the earners above the median have seen larger increases than the earners below. The middle left panel confirms this intuition and show that earnings dispersion has grown for younger cohorts. Most of the difference is already present at age 25. This large increase, together with relatively little action in the medians, implies that the earnings-poorest of more recent cohorts are relatively worse off than people in the same percentile of earnings of earlier cohorts, and conversely the earnings-richest are better off today.[7]

Apart from more inequality, for younger cohorts there is also more earnings volatility, as measured by the standard deviation of household earnings *changes* (middle right panel of Figure 1), particularly at younger ages.[8] This evolution is closely linked to the

[6] Appendix C.1.1 compares male and household earnings, and C.2.2 reports the same results using a different deflator.

[7] Part of these can reflect intergenerational changes in family composition and delayed household formation. Appendix C.2.3 shows that patterns are similar in a sample of married couples.

[8] This agrees with a plurality of studies that use different data sources, income definitions, and income processes (Gottschalk, Moffitt et al. (1994), Carr and Wiemers (2018), Braxton et al. (2021), etc.), although there is some evidence with individual male earnings pointing in the opposite direction (Bloom et al., 2017)

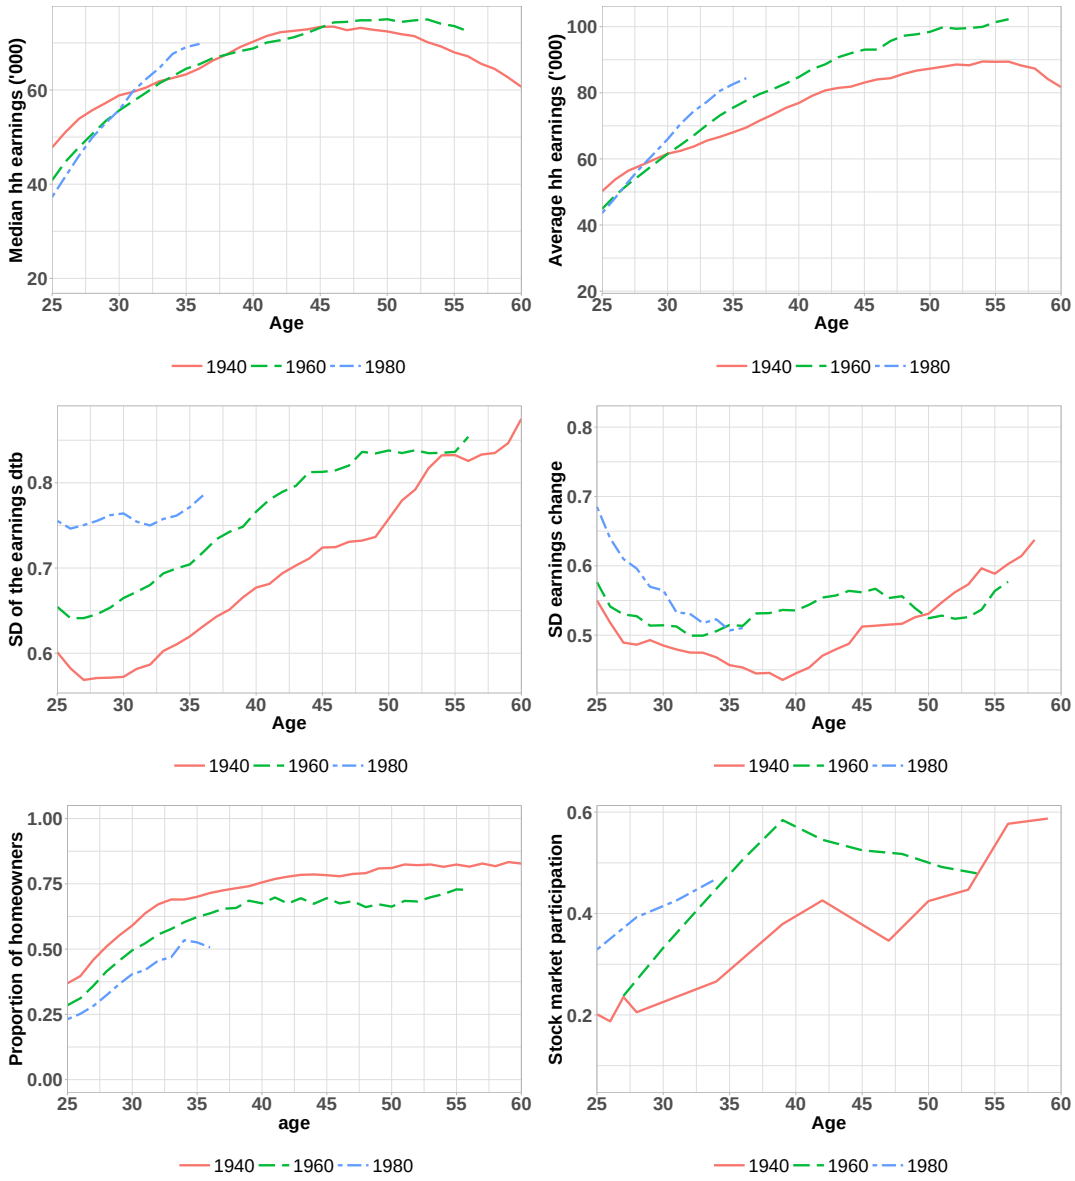

Figure 1: Intergenerational changes. Top: median (left) and average (right) household earnings. Middle: standard deviation of the log earnings distribution (left), and standard deviation of log earnings changes (right). Bottom: homeownership (left) and stock market participation (right). PSID data (Panel Study of Income Dynamics, 2017), deflated using the CPI, stock market participation from SCF (Survey of Consumer Finances (2019) and Survey of Consumer Finances (1960 to 1986)).

so-called reduction in "fluidity" in the US labor market (Davis and Haltiwanger, 2014), which has affected different workers differently. On the one hand, longer-duration jobs are an increasing proportion of the job tenure distribution and workers' median tenure has increased (Hyatt and Spletzer, 2016). On the other hand, jobs with very long tenure are becoming less likely and large disruptions to careers are more frequent than before (Molloy, Smith and Wozniak, 2021). The rate of job loss has been decreasing over time, but so have the job finding rates of the unemployed and the probability that they stay in the same sector after becoming unemployed (Fujita, 2018). In my PSID sample, the growth in earnings volatility is closely linked to the longer average duration of unemployment for younger generations, which is only partially mitigated by the reduced likelihood of job loss (see Appendix C.1.2 for details).

Parallel with the changes described earlier, homeownership rates have been falling for recent cohorts. I use the word homeownership to refer to the percentage of households that live in owner-occupied housing. In the bottom left panel of Figure 1, we observe that, at age 35, homeownership has dropped by over 10 percentage points between the cohorts born in 1940 and 1960, and by another 10 percentage points between the cohorts born in 1960 and 1980. These intergenerational changes have occurred during a time in which, remarkably, the homeownership rate has been very stable and high for historical standards, around 65%, only fluctuating noticeably during the housing boom and bust.[9]

At the same time, stock market participation has increased significantly for younger cohorts (Figure 1, bottom right), but most of it has happened through indirect stock market participation via mutual funds or retirement accounts (direct stock market participation did not change much, see Appendix C.5.2). Stock market participation also displays year effects. For instance, direct stock market participation increased significantly in the years before the 2000 stock market crash, and dropped dramatically afterwards, as it can be seen in the profile for the 1960s cohort when they were 40 years old.

These changes are closely related with the evolution of asset prices and financial conditions. The ratio of median house prices to median income has increased, on average, in

[9]Homeownership rates have also fallen for younger generations under alternative sample selection procedures (C.2.1), for married households or households with children (C.2.3) or by education groups (C.3.2). Using census data instead of the PSID reveals very similar patterns (Appendix C.5.1) and allows to show that they have also fallen across geographical areas, but that they have aggregated into a stable national average homeownership rate

the United States over the last 60 years (left panel of Figure 2, based on PSID data).[10] Younger generations, at the same age, now have to devote more years of their income to buy a home compared with their parents. The right panel of Figure 2 compares stock returns and house price growth. There are large cyclical variations in house prices, although they are not always correlated with the business cycle. These induce an additional source of variation across cohorts, as some of them may have entered the labor market in a time where house prices were cyclically low, and benefitted from the situation to make housing purchases earlier on in their lives. On the other hand, the evolution of stock returns shows larger annual fluctuations, which are more strongly correlated with the business cycle, and less autocorrelation. However, much of house price risk is idiosyncratic or local (Landvoigt, Piazzesi and Schneider, 2015), and is not captured in this comparison.

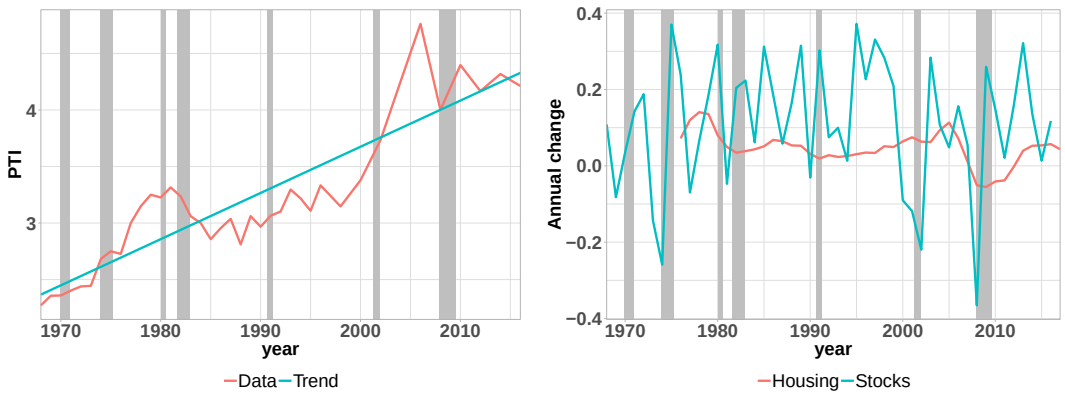

Figure 2: Evolution of house prices and asset returns. Left: median Price to Income (PTI) ratios for housing (PSID data). Right: Stock returns, S&P 500, vs. growth in Federal Housing Finance Agency (FHFA) house price index. Shaded areas correspond to NBER recessions.

The process of financial deregulation and innovation that started in the 1980s and expanded during the 1990s improved the access of households to credit, both from an extensive (more people can get credit) and intensive (the same household can borrow larger amounts) perspective. See, for instance, Gerardi, Rosen and Willen (2007) for a detailed description of the regulatory changes, the changes in the structure of the financial sector, and the new mortgage products that became available over this period. These changes were partially encouraged by policymakers, who were worried about low homeownership rates (for instance, Bill Clinton's National Homeownership Strategy).

[10] Lovenheim (2011) shows that both median and mean home price indices constructed from PSID data track Federal Housing Finance Agency repeat home sales indices very well.

Another important change was the introduction of tax-advantaged retirement accounts, such as individual retirement accounts (IRAs), which started in 1974 and became popular in the 1980s, and 401(k)s, which were introduced in 1978 and also became popular later on. Later reforms made these accounts more beneficial and less restricted, and automatic enrollment in pension plans further increased the number of stock market participants by reducing both the financial and psychological costs of enrollment.

II A business-cycle dependent earnings process

In this section I develop a flexible earnings process that can capture the differences across generations I have just described, whilst encompassing a set of elements that have been shown to be important to describe the features of household earnings risk and its implications on household consumption and self-insurance (De Nardi, Fella and Paz-Pardo, 2020). These include age-varying persistence, variance, and higher order moments, non-normalities such as high negative skewness and large kurtosis, and non-linearities such as previous-earnings-dependent persistence.

The process is based on Arellano, Blundell and Bonhomme (2017), but, on top of that, it includes three important factors: business cycle variation in earnings dynamics, including its non-normal and nonlinear features, intergenerational changes in the distribution of earnings, and intergenerational changes in earnings risk. The former is necessary because idiosyncratic risk correlates with aggregate asset price risk, which can have implications for household portfolio decisions and insurance over the business cycle. The latter two are necessary to address the questions posed in this paper.

Let $\tilde{y}_{it}$ denote the logarithm of pre-tax labor earnings, net of age effects, for household i of cohort c_i ($c_i \in \{1940, 1960, 1980\}$) living in calendar year t with age age_{it}. I assume earnings are the sum of a persistent and a transitory component:

$$\tilde{y}_{it} = \eta_{it} + \epsilon_{it} \tag{1}$$

where both have absolutely continuous distributions. The persistent component η_{ith} is assumed to follow a first-order Markov process, while the transitory component ϵ_{ith} has zero mean and is independent over time and of the persistent component.

We can introduce these assumptions by writing the processes for η and ϵ, and the initial condition for the persistent component η_1 as:

$$\eta_{it} = Q_\eta(\nu^\eta_{it}|\eta_{i,t-1}, age_{it}, c_i, \Omega^y_t), \nu^\eta_{it} \overset{iid}{\sim} U(0,1), t > 1 \tag{2}$$

$$\epsilon_{it} = Q_\epsilon(\nu^\epsilon_{it}|age_{it}, c_i), \nu^\epsilon_{it} \overset{iid}{\sim} U(0,1) \tag{3}$$

$$\eta_{i1} = Q_{\eta_1}(\nu^{\eta_1}_{it}|age_{it}, c_i, \Omega^y_t), \nu^{\eta_1}_{it} \overset{iid}{\sim} U(0,1) \tag{4}$$

Equation 2 specifies the dependence of η_{it} on its previous realization with a flexible quantile function Q_η. This function depends on the age of the household, age_{it}, its cohort, c_i, and the aggregate state of the labor market, Ω^y, which is a Markov process. Thus, the features of earnings shocks are allowed to be different in expansions and recessions. In this way, this formulation explicitly includes age, cohort, and year effects.

Q maps draws ν_{it} from the uniform distribution $U(0,1)$ into quantile draws for η. ν_{it} can be thought of as a rank: if it is 0.9, it implies that the realization of η_{it} is on the 90th percentile conditional on age and $\eta_{i,t-1}$. A similar reasoning follows for the initial realization of the persistent component, with the further simplification that it only depends on age, cohort, and the current state of the labor market; and for the transitory component, which only depends on age and cohort. I treat the transitory component as measurement error or alternatively as a fully-insurable source of earnings fluctuations.

Following Arellano, Blundell and Bonhomme (2017), to estimate the process I specify a parametric form for the quantile functions as low order Hermite polynomials. Appendix D contains details about the implementation and its comparison with a more standard canonical earnings process. While the earnings process is estimated on pre-tax rather than post-tax household earnings, most of its features regarding non-linearity and non-normality are qualitatively similar to De Nardi, Fella and Paz-Pardo (2020) and therefore I refer the interested reader to the discussion therein.

I estimate the earnings process on PSID data for all three cohorts. Given that the PSID became biennial from 1997 onwards, the period is two years for both the earnings process and the structural model. I use the full length of the PSID (1968-2017).[11] More

[11] The semiparametric implementation of the nonparametric model defined in Arellano, Blundell and Bonhomme (2017) allows to interpolate and obtain an earnings process for every state and age even if not all combinations are present in the data.

details about the data treatment, cohort definitions, and sample selection are available in Appendix A.

A Implications of the earnings process

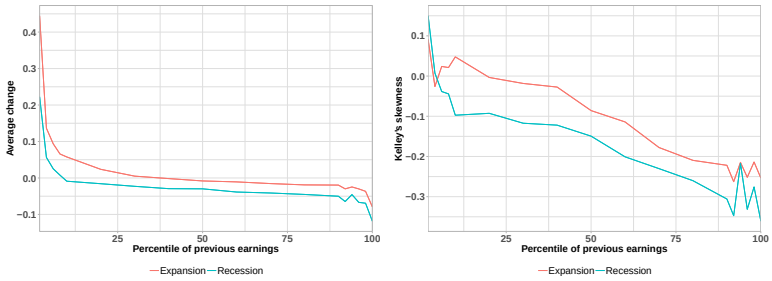

Figure 3: Average expected change in earnings (left) and Kelley's skewness of earnings changes (right). 1940s cohort.

The earnings process captures the intergenerational changes in earnings dynamics in terms of earnings inequality and uncertainty documented in Section I well (see Appendix D.2). Unlike more standard earnings processes, it also captures *countercyclical* skewness: during recessions, skewness becomes more negative, thus implying that large decreases in earnings become more likely with respect to large increases in earnings. This business cycle effect is also present in first moments: during a recession, all individuals expect lower increases in their earnings than they usually would (Figure 3).

An additional realistic feature that the Markov-switching earnings process captures is history dependence: at any point in time, the distribution of earnings for a given cohort depends on the set of expansions and recessions that the cohort has lived through. In particular, the recovery from recessions is usually sluggish. Unlike a standard process with countercyclical variance in logs, in which average earnings in levels counterfactually *increase* during recessions, my flexible earnings process generates drops in average earnings during a recession, from which agents (particularly the young) take long to recover (see Appendices D.3 and D.4 for details).

III Model

I build a life-cycle structural model to evaluate to which extent the changes in earnings and financial conditions described in Section I, modelling the former using the process

described in Section II, can account for the intergenerational differences in homeownership and portfolios I described earlier. In the model, the economy is populated by a continuum of households i that belong to cohort c. From the perspective of a cohort, age and time are equivalent and indexed by t. The model period is two years. All variables in the model are real.

A Demographics

Households are born in the model at age 25, retire at age 60 and face positive and increasing death probabilities starting at that age. They die for sure at age 86. An average demographic profile at each age is introduced in the model with a taste shifter θ_t, which represents the average OECD equivalence scale at each age, and generates age-varying marginal utility from nondurable and housing consumption.

B Preferences

Preferences are Epstein and Zin (1989) and allow to disentangle the elasticity of intertemporal substitution ψ and the risk aversion coefficient γ. Since I assume $\gamma > \frac{1}{\psi}$, they imply that agents prefer an early resolution of uncertainty, as standard in studies on the equity premium and risk-free rate puzzles, and in portfolio choice models (e.g. Cocco, Gomes and Maenhout (2005)). Thus, utility is:

$$U_{it} = [(\theta_t c_{it}^{\nu} s_{it}^{1-\nu})^{\frac{(\psi-1)}{\psi}} + \beta(\mathbb{E}_t U_{it+1}^{1-\gamma})^{\frac{1}{1-\gamma}\frac{\psi-1}{\psi}}]^{\frac{\psi}{\psi-1}} \tag{5}$$

where θ is the taste shifter described earlier, c is nondurable consumption, s is the housing service flow, β is the discount factor and ν measures the relative importance of nondurable consumption with respect to housing. This Cobb-Douglas specification assumes an elasticity of substitution of 1 between both goods, which is justified by stable shares of expenditure in housing in micro data (e.g. Davis and Ortalo-Magné, 2011). In practice, since housing in the model is discrete, this is equivalent to assuming that housing utility is a proportional scaling of the utility from nondurable consumption.

The utility value of housing s_t depends on the quality of the owned home and does not vary with its aggregate price. It is highest for owners of high-quality houses ($\bar{s}^2$),

lower for owners of low-quality houses ($\bar{s}_1$), and lowest for renters ($\bar{s}^0$).

Households value bequests left b according to De Nardi (2004). For simplicity, I assume that bequests left are not received by other generations and thus leave the economy.

$$v(b) = \phi_1(\phi_2 + b)^{\frac{(\psi-1)}{\psi}} \tag{6}$$

C Environment and technologies

C.1 Aggregate state

During each year t, the economy is in an aggregate state Ω_t that determines house prices, stock market returns, and the state of the labor market, which in turn affects the earnings process. Households know the (Markov) process governing the aggregate state, and use it to make predictions about the future, which in turn affect their decisions. Section IV.A describes each of its elements and their correlations in more detail.

C.2 Earnings

Earnings are composed of a deterministic component, which depends on age, and a stochastic persistent component η_{it}, which depends on the aggregate state of the labor market:

$$\log y_{it} = f(t) + \eta_{it}(\Omega_t) \tag{7}$$

Section II contains more details about the earnings process and its estimation. Transitory shocks may be reflecting measurement error or almost fully insurable fluctuations, so to save on computational costs I do not include them in the model.

C.3 Liquid accounts

Liquid accounts a_t are risk-free and they yield an exogenous and constant interest rate r^a. They cannot be negative: if they wish to borrow, households must apply for a specific type of financial asset, mortgages m_t, which I describe in detail in Section C.6.

$$a_{t+1} \geq 0 \tag{8}$$

C.4 Risky financial assets

Households can also hold risky financial assets or stocks f. Stock returns r_t^f depend on the aggregate state of the stock market Ω_t^f. Households cannot short financial assets:

$$f_{t+1} \geq 0. \tag{9}$$

When $f_{i,t} = 0$, households pay a fixed entry cost κ^f to start investing in stocks. This cost represents psychological, financial, and technical barriers to start investing in the stock market (opening financial accounts, acquiring information about them, etc.), and is frequently used in the portfolio choice literature (Gomes and Michaelides, 2005). Once a household participates, there are no additional costs of adjusting financial assets.

C.5 Housing

Households can either be renters, own a small house h_1 or own a big house h_2. i.e., $h_{i,t} = \{0, h_1, h_2\}$. Average house prices p_t^h depend on the aggregate state Ω_t. They are expected to grow, but fluctuate around a trend as described in Section IV.A. I abstract from housing depreciation and assume that households costlessly and fully repair their homes every period.[12] I assume that the price of the different housing sizes is a fixed fraction of average house prices.

A minimum house size h_1 is necessary so that homeownership is meaningful in the model; it is grounded both in previous literature (Cocco (2005), Attanasio et al. (2012), Sommer and Sullivan (2018)) and in empirical evidence, which shows very few households living in very small or cheap owner-occupied housing units, and few of those available for sale. The set of house sizes that households can buy is limited to 2 due to computational

[12] In the data, housing depreciates between 1 to 2 percent a year (Fraumeni, 1997). However, imposing a fixed cost of 1 or 2 percent of their housing value likely overestimates the liquid resources that households spend in home repairs (which is a median of 0.7% of family income for homeowners according to 2005-2013 PSID data) and artificially reduces housing demand. In reality, households that suffer negative shocks can postpone investments in their houses or let them depreciate.

considerations. In Appendix F.3.1 I discuss these assumptions and show that results are robust to different sizes for the small house and several specifications in which there are three different house sizes.

Housing is illiquid. Households pay a proportional transaction cost to buy or sell housing $\kappa^h p_t^{h_j}(\Omega_t)$, which depends on the price of the house which is being bought. It reflects the costs associated with selling or buying a home, which can include taxation, real estate agent fees, and other costs.

Households that do not own a home must participate in the rental market. I assume that foreign or institutional investors, who are not explicitly modelled, supply housing in the rental market, and I abstract from the equilibrium determination of house prices for tractability and simplicity[13]. The rental price $r_t^s(\Omega_t^h)$ depends on current housing prices $p_t^h(\Omega_t^h)$.[14]

$$r_t^s(\Omega_t) = \gamma^r p_t^h(\Omega_t). \tag{10}$$

During the working period, households are subject to exogenous moving shocks with probability π_{hm}. They represent events such as finding a new job in a different place or suffering a job relocation, and add to the riskiness of owner-occupied housing as an investment. In the model, when the moving shock realizes, agents sell their houses at the beginning of period, before they make their consumption and saving decisions. They must then spend that period in rental housing but can freely reoptimize afterwards. I keep π_{hm} fixed over generations based on the empirical evidence from the PSID and other studies, which I discuss in Appendix C.3.1, that the increase in earnings volatility has not been coupled with an increase in mobility for younger generations.

C.6 Mortgages

When a household wants to acquire a house of quality j, it can apply for a loan or mortgage m_t. I define mortgages so that $m_t \leq 0$. In order to get it, the household must

[13] Appendix E.4 contains an approximation to how my counterfactual results would change under endogenous determination of housing prices

[14] Davis, Lehnert and Martin (2008) show that the ratio of rents to prices stayed relatively stable over the period 1960-1995, so I assume rents are just a fixed fraction of housing prices. Rent-price ratios decreased in the early 2000s, which could be an additional channel to discourage homeownership for the young, contributing to those I explore in this paper.

fulfill two conditions: a downpayment or loan-to-value (LTV) restriction and an income test or loan-to-income (LTI) restriction. They only apply at mortgage origination.

$$m_{t+1} \geq -\lambda_h p_t^{h_j}(\Omega_t) \tag{11}$$

$$m_{t+1} \geq -\lambda_y y_{it}(\Omega_t) \tag{12}$$

where $\lambda_h < 1$. There is no uncollateralized debt.

Borrowers pay an exogenous interest rate on their debt r^b which is larger than the risk-free rate r^a. Households decide on their repayment schedule, but in every period they must at least pay the interest accrued by their debts and cannot reach their terminal age T with an unpaid mortgage balance. I do not explicitly model 30-year fixed-rate mortgages to reduce the dimensionality of the problem, but in Section V.A I show that the model generates repayment patterns that are according with the data. To reflect that households might have more difficulty acquiring a mortgage during retirement, I assume that retirees cannot upsize: if they are renters, they stay renters; if they own a home, every period they can decide to stay in the same house, downsize, or become renters.

$$m_{t+1} \geq \frac{m_t}{1 + r^b} \tag{13}$$

$$m_T = 0 \tag{14}$$

Households can extract equity from their homes in two ways. First, they can sell them and either move to rental housing or buy a new smaller or cheaper house. Second, they can delay the repayment of the mortgage principal, thus extending their mortgage duration. This assumption indirectly incorporates arrangements such as mortgage forbearance. For simplicity, I assume that they cannot increase the principal of their debt by remortgaging or accessing home equity lines of credit.

To reduce the dimensionality of the problem, due to computational considerations, I

assume that households cannot simultaneously hold a mortgage m_t, risk-free assets a_t, and risky assets f_t, but only two of the three. This assumption is weaker than modeling mortgages as negative safe assets, because it still allows households in debt to make a choice between positive safe and risky assets, as long as the choice is not interior. Therefore, mortgagors in the model are able to hold liquidity without incurring the participation cost to the stock market. Because households cannot increase the principal of their debt without selling the house and incurring transaction costs, holding some positive amount of financial assets can insure their consumption flow against relatively small negative shocks to their labor market income.

$$a_{t+1} f_{t+1} m_{t+1} = 0 \tag{15}$$

D The government

Disposable income $\lambda(y_{i,t})$ is obtained from pre-tax income $y_{i,t}$ using the tax function $\lambda(\cdot)$ (Benabou (2002)):

$$\lambda(y_{i,t}) = \lambda y_{i,t}^{1-\tau} \tag{16}$$

This specification can be negative at lower income levels and thus includes, in a parsimonious way, both progressive labor income taxation and many income-tested welfare programs, such as unemployment insurance, EITC, food stamps, etc.

The government also taxes capital income from risky and safe assets at a flat-rate τ_a and finances social security for old people $p(\cdot)$. The latter is a function of a household's last income realization.

I replicate the preferential tax treatment of owner-occupied housing: owner-occupied rents are not taxed and mortgage interest is tax-deductible. Both in the US tax code and in the model, households can choose between getting the standard deduction, which is a fixed amount, and itemization, which implies that they individually deduct qualifying expenses such as mortgage interest. Thus, only households who have a sufficiently large mortgage get the mortgage interest deduction. Furthermore, stock market losses are

deductible against asset income and labor income up to \$3,000. I also assume that the government provides housing aid to income-poor households for whom rental costs are large. In particular, the government pays all rent that is above 30% of household income.[15]

E Negative net worth and default

Both in the model and in the data, a household can have negative net worth. In the model, that can arise when a household suffers a negative housing price, income, or financial shock while holding a significant mortgage. Households can continue to hold their house as long as they are able to make interest payments to their mortgage out of their financial savings or labor income.

Households can choose to default on their mortgages. I assume they file for bankruptcy if they do so: their debt is cancelled, the creditors seize all of their housing and financial assets, they suffer a utility penalty of χ^{bk}, and they become renters with zero wealth.

F Timing

At the beginning of the period, households learn the common realization of the aggregate state Ω_t, which implies that they find out about housing prices $p_t^h(\Omega_t)$ and stock returns $r_t^f(\Omega_t)$, and their individual realization of labor income $y_t(\Omega_t)$. Jointly, those determine their net worth or cash-on-hand in period t:

$$(17)\quad coh_t = p_t^h(\Omega_t)h_t + (1+r_t^f(\Omega_t)(1-\tau_a))f_t + (1+r^a(1-\tau_a))a_t + (1+r^b)m_t + T(y_t(\Omega_t), m_t)$$

where $\lambda(\cdot)$ represents progressive taxation of labor earnings net of mortgage interest payments.

As shocks realize, mortgagors decide whether to default or continue paying their mortgages. Then, households get utility from their housing stock h_t at the beginning of the period. Then they decide on their consumption c_t and their savings for the next period, which are composed of their liquid accounts a_{t+1}, stocks f_{t+1}, and housing h_{t+1},

[15] This is a stylized representation of housing aid programs in the United States, in particular the Section 8 program (Housing Act of 1937), which provides families with low income with Housing Choice Vouchers or project based assistance. In the PSID data roughly 2.75% of working age households receive this subsidy.

minus any outstanding mortgage balance m_{t+1}.

The budget constraint and household's problem that summarize all of these elements are reported in Appendix E.1.

IV Calibration

A Aggregate state

The aggregate state of the economy in a calendar year Ω_t is the combination of the state of the labor market Ω_t^y, stock market returns r_t^f, the state of the housing market Ω_t^h and the house price p_t^h. So far, t indexed both year and age, which were equivalent from the perspective of a cohort. Naturally, calendar years and their associated states happen at different ages for different cohorts. To keep the notation in this section clear, I describe it from the perspective of a single cohort.

The aggregate state of the labor market Ω^y in $t+1$ takes two possible realizations, expansion and recession, and determines the conditional distribution of earnings shocks that agents face given their earnings in t, as described in Section II. I define a period to be recessionary if any part of it falls under an NBER-defined recession. Ω^y is persistent.

Stock market returns r^f take four possible realizations. Three of those correspond to historical averages of each tercile of the distribution of stock market returns in the S&P500 during my sample period (1963-2015). Additionally, I include a disaster state, that corresponds to the average of the lowest 5% of annual stock market realizations during this period.[16] Stock market returns do not depend on their previous realizations[17] or on housing prices, but their conditional distribution is a function of the aggregate state of the labor market $r_t^f = F(\Omega_t^y)$. Thus, stock market returns are more likely to be low when labor market income receives a bad shock, which makes these financial assets more risky from the perspective of households. The persistence in stock returns induced by the persistence in Ω^y is very small (see Appendix F.4.1).

The housing market state Ω^h denotes whether house prices are increasing or decreas-

[16] The possibility of stock market crashes is important to understand the low stock market participation and high equity premium puzzles (Bansal and Yaron (2004)), as well as the age patterns of stockholding (Fagereng, Gottlieb and Guiso, 2017).

[17] Price-dividend and earnings-price ratios are predictors of future stock returns (Campbell and Yogo, 2006), although some of the relationships between economic and financial variables and future stock performance are unstable and change over time (Pesaran and Timmermann, 1995).

ing. It is persistent, which helps the model generate procyclicality in homebuying: households buy more housing when its price is growing because they expect it to grow further. There is no idiosyncratic house price risk. Aggregate house prices evolve following:

$$p^h_{t+1} = (1 + r^h_{t+1})p^h_t + \epsilon^h_{t+1} \tag{18}$$

where $\epsilon^h_{t+1} \sim N(0, \sigma^h)$ and

$$r^h_{t+1} = \begin{cases} r^h_{low} \text{ if } \Omega^h_{t+1} = \text{decreasing} \\ r^h_{high} \text{ if } \Omega^h_{t+1} = \text{growing} \end{cases} \tag{19}$$

I estimate the persistence of Ω_y and Ω_h, the conditional distribution of stock returns $F(\Omega^y_t)$, and the two possible realizations of r^h_{t+1} directly from their empirical counterparts over the 1975-2015 period. Because the model period is two years, it may overestimate the average length of recessions; however, it also implies that the probability of exiting a recession after one model period is relatively high. I set σ^h so that, taking Equation 18 into account, the standard deviation of percentage growth rates of housing prices is 0.10, within the range reported in Landvoigt, Piazzesi and Schneider (2015). Because the correlation between Ω_h and both Ω_y and r^f is low and not statistically significant in the data, I set it to zero.

Households know the processes for the aggregate variables and form expectations about their evolution accordingly.[18]. In the simulation, the realizations of $\Omega^y_t, \Omega^h_t, r^f_t$ and p^h_t correspond to their data counterparts for each specific year (Appendix E.2 shows that the model fits the evolution of house prices and stock returns very well). For instance, when agents of the oldest cohort, born in 1942, reach 53 years of age, they face a good realization of the stock market state because 1995 was a year of high stock returns.

[18] Appendix F.4 shows that my results are robust to alternative specifications of household expectations about aggregate variables.

B Externally calibrated parameters

I set the risk aversion coefficient γ to 4, which is on the higher side of usual estimates in the macro literature, but on the lower side for the literature that rationalizes the equity premium puzzle with Epstein-Zin preferences (e.g. 10 in Bansal and Yaron (2004)). The elasticity of intertemporal substitution ψ is also disputed. I follow Kaplan and Violante (2014) for its exact quantification (see their footnote 28 for a discussion regarding this estimate) and set it to 1.5. I set the housing utility share ν to 0.2, based on NIPA data on budget shares.

I establish the risk-free rate at 1%, plus an additional 1% to account for the liquidity services of risk-free money. The mortgage interest rate is set to 4%. Both rates correspond to historical averages for the 1940s generation. I assume that the mortgage interest rate is 1% higher for retired people to reflect the more stringent credit conditions they are subject to, which is looser than assuming they cannot get a mortgage. The downpayment is 20% of the value of the house, and the income test consists in having yearly household income that is at least 1/9th of the value of the mortgage (Johnson and Li, 2010).

I set the social security replacement rate to 55% (Mitchell and Phillips, 2006). I explicitly model itemization and the standard deduction, so I already incorporate part of tax progressivity by construction. Taking this into account, I estimate the progressivity coefficient τ to be 0.085, following the procedure described in Appendix A.1.2. The parameter that controls average taxation $\lambda = 0.64$ implies an average tax rate of 35% for the average household, close to the historical level for the 1940s generation comprising federal and state taxes and FICA contributions. I set the standard deduction at a level (6% of average income) that implies that the percentage of people choosing to itemize is close to the data, which is around 30% (IRS, 2014) .[19]

The minimum housing size h_1 is chosen such that the small house is worth 50% the price of the median house in the data. This means the cheapest house in the model sells for around twice median household income in 2015 ($120,000). Housing adjustment costs are around 10% of the value of the property (Smith, Rosen and Fallis, 1988), which I distribute equally amongst seller and buyer. Rental rates are 5% of house prices (Davis, Lehnert and Martin, 2008). I set the bankruptcy penalty $\chi^{bk} = -5$, which is equivalent

[19]This is lower than its historical levels (e.g, around 10% of average income in the early 70s) because the model abstracts from itemizable expenses other than mortgage interest and local property taxes, such as out-of-pocket medical expenditure, state taxes, charitable contributions, etc.

to a loss of around 15,000 dollars for a poor, low-income renter, to the value that keeps bankruptcies around the housing boom and bust episode aligned with the data.[20]

C Internally calibrated parameters, targets, and model fit

Moment	Data	Model	Key parameter	Value
W/Y ratio	3.1	3.1	Discount factor β	0.929
Average bequest (/average income)	2.7	2.6	Bequest taste ϕ_1	12.3
Fraction of population leaving no bequests	20%	19%	Bequest taste ϕ_2	6.9
Housing ownership at age 40	76%	76%		
... of detached houses	68%	63%	Housing taste s_2	9.4
... of other housing at age 40	8.1%	12%	Housing taste s_1	1.3
Percentage buying houses at age 40	5.9%	5.5%	Moving shock π_{hm}	0.051
Stock market participation, age 40	38 %	36%	Participation cost k^f	0.39

Table 1: Targeted moments, model fit, and calibration

The model has 7 free parameters which are jointly calibrated to match 7 targets in the data. I perform the calibration for the 1940s cohort, and then keep them constant across cohorts in the experiments unless otherwise specified. Table 1 summarizes the data and the parameter which is more closely related with each of the targets. The wealth to income ratio of 3.1 corresponds to the wealth to income ratio of the bottom 95% of the wealth distribution, which I am focusing on. I obtain house ownership data from the PSID, stock market participation from the SCF, and I adjust the bequest targets for this specific cohort (see Appendix A for more details).

Matching homeownership at a particular age allows me to get an estimate of how much households enjoy owner-occupied housing, over and above its value as a financial investment and collateral. On the other hand, getting the level of stock market participation right allows me to discipline the stock market participation cost κ^f. Finally, the percentage of people buying houses after prime homebuying age is informative of the number of people who are moving for reasons that I do not model explicitly, which I summarize in the moving shock.

[20] Appendix F.3.5 shows that results are robust with respect to changes of this parameter.

As Table 1 shows, the model fits its targets very well with the associated calibrated coefficients. The discount rate is relatively low with respect to what is standard in a one-asset model. Households value housing, and the utility value of owner-occupied houses provides a further motive to hold assets beyond life-cycle and precautionary savings, which reduces the value of household patience. Besides, low levels of β are also frequent when stocks are available as an investment option with high returns.

There is a 5.1% yearly probability of receiving a shock that forces the household to move. A particularly relevant parameter is the one-off cost to start participating in the stock market k^f, which is calibrated to be 39% of average yearly earnings.

There is scarce data about the initial wealth of the 1940s cohort at labor market entry. I set the initial condition of the model to the most conservative possibility that is consistent with the observed homeownership (20% equity on the house for the initial homeowners) and stock market participation (1$ in stocks for the initial stockholders).

Appendix B briefly describes the solution method of the model.

V Results

A Untargeted moments, 1940s cohort

The model replicates life-cycle homeownership profiles, the patterns of house buying by age, and stock market participation for the 1940s cohort very well (Figure 4), particularly during the working age period. Most households become homeowners between ages 25-35, and then the share of households that live in their own home stabilizes around 80%. In the model, young households do not participate in the stock market because they are concentrating their resources in saving for a downpayment and starting to pay their mortgages, rather than spending time and resources in acquiring information and access to the stock market.

Many households hold mortgages at the same time as they start investing in stocks. Figure 4, bottom right panel, shows that in the model households pay back their mortgages slowly, a feature which is not targeted in the calibration. Thus, the model suggests that the horizon of available mortgage products closely resembles what households would choose if they were to freely decide on their repayment schedule.

The model implies relatively low homeownership rates for this generation during re-

tirement. The reason is that, in the model, many households release equity from their homes and become renters as they approach retirement, taking advantage of the high house prices they face around this period. Although retirees' housing decisions are responsive to this type of financial incentives (McGee, 2021), the effect is stronger in the model than in the data. Additionally, there are a set of reasons to hold housing during retirement which are not explicitly modelled and would bring homeownership for retirees up, including a desire to age in one's home and self-insurance against uncertain medical expenditures (Nakajima and Telyukova, 2020).

In the model many households make housing adjustments right before the retirement period, during which they are no longer allowed to upsize their homes.[21] It then closely tracks the low proportion of retired home buyers. Given the fixed entry cost structure, households that enter the stock market find it optimal not to exit: as a result, stock market participation stays high during retirement.

The model is successful in replicating portfolio patterns by wealth (Figure 5). A standard portfolio choice model would yield stock holding patterns which are mildly *decreasing* rather than increasing in wealth (Gomes and Michaelides, 2005). In this model, the role of housing and the correlation of labor income and stock returns reduce the incentive of the income-poorer to participate in the stock market. Richer individuals, on the other hand, have sufficient resources available even after buying their homes, and they invest them in the stock market, in which they reap higher returns that in turn make them wealthier.

B Explaining intergenerational differences in homeownership

Keeping constant all preference parameters, I now turn to studying which are the key intergenerational changes that explain the reduction in homeownership for younger cohorts. In this experiment, cohorts differ in four ways. First, younger cohorts face more unequal and riskier earnings processes. Second, the exogenous house prices and stock returns correspond to those that each generation actually faced, so that, for younger generations, the median earner needs to spend more years of income to buy a house. Third, there have been changes in financial conditions. On the one hand, different mortgage products

[21] In a version of the model where this restriction is not imposed, this spike does not occur and all of the main conclusions still hold, but the model overestimates how many housing sales there are in retirement and how many retired households still hold mortgages (see Appendix F.3.2).

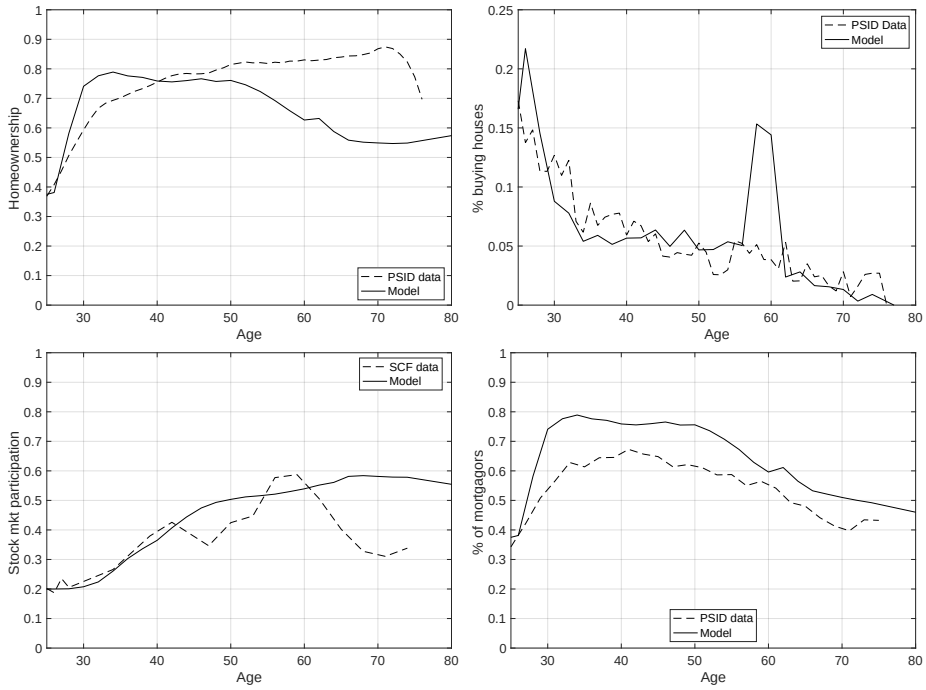

Figure 4: Life-cycle profiles for the 1940s cohort. Top left: homeownership by age; top right: proportion of households buying a house by age; bottom left: stock market participation; bottom right: percentage of all households with a mortgage by age.

were available to the 1980s generation during their homebuying years, which I replicate as a reduction in downpayment requirements. Namely, I assume that the maximum LTV ratios of mortgages increased from their baseline level of 80% to 97.5% between 2000 and 2010, after which they unexpectedly went back to normal.[22] On the other hand, I reduce stock market participation costs to match the stock market participation profile (see Section V.C). Fourth, I input to each generation their specific average demographic profile by age, which I report in C.4, which captures the effect on consumption needs of differential timings in marriage and childbearing. For a cleaner comparison, I keep initial wealth at age 25 constant across generations.[23]

Figure 6 shows the homeownership rates for each of the three cohorts in the data, compared with the profile implied by the model. Notably, keeping preference parameters

[22] Glaeser, Gottlieb and Gyourko (2012) use housing industry data and show that for most of the 1998-2008 period the 75th percentile of LTV ratios at origination was above 95%, with the 90th percentile consistently around 100%. Duca, Muellbauer and Murphy (2011), using American Housing Survey data, show that average LTV ratios for first time buyers, which were stable around 0.80-0.85 in the 1980s and early 1990s, jumped up to 0.90-0.95 during the 2000s.

[23] This assumption is conservative, as it is likely that younger generations are entering the labor market later and with less wealth. In Appendix F.1 I provide results for the case in which all agents start at zero wealth. All conclusions are unchanged, but the model with initial zero wealth underestimates homeownership at earlier ages.

Figure 5: Bottom: portfolio shares of assets by wealth decile at retirement age (left: PSID data, right: model).

constant, the model very closely replicates the decrease in homeownership that occurred between the 1940s and the two latter cohorts.[24]

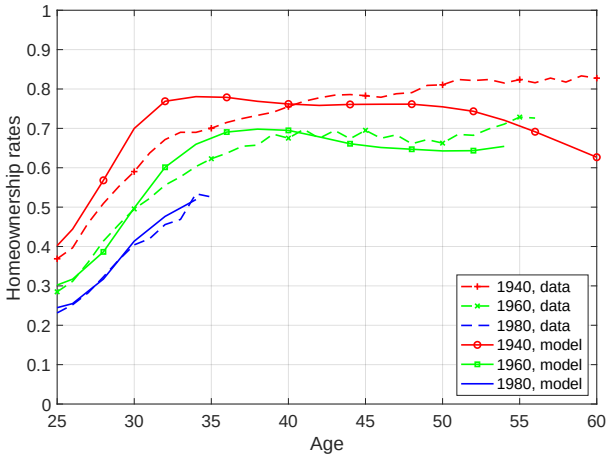

Figure 6: Homeownership by cohorts, data vs. model

B.1 Decomposing the decrease in homeownership

I now turn to evaluating, using the model, which are the key factors that drove the decrease in homeownership. Table 2 shows the results of a Shapley-Owen decomposition in which I evaluate the relative contributions of six key elements in explaining the reduction in homeownership at different ages: initial earnings inequality, earnings risk thereafter, changes in average housing price-to-income ratios, histories of aggregate shocks, average

[24]Figure 6 only represents the 1960s and 1980s cohorts up to the ages in which I can fully observe them in the data. In Appendix F.5 I show the model-implied homeownership rates for the future under different simulated realizations of the aggregate state.

demographic structure at each age, costs of participation in the stock market and, for the 1980s generation, changes in financial conditions. These are the only differences across cohorts in the model. Thus, by counterfactually changing them one by one I can quantify their relative contribution to the difference between the observed profile for a given generation and that of the 1940s.[25]

	1960s generation			1980s generation	
Age	30	40	50	30	35
Total	-10	-8	-14	-19	-18
Earnings	84	90	7	87	56
initial inequality	76	18	-37	70	27
risk	8	72	44	17	30
Aggregates	16	-2	92	119	126
house price trend	35	112	43	76	84
histories	-19	-114	49	43	42
Financial conditions	0	0	-2	-108	-83
stock participation costs	0	0	-2	-2	-1
borrowing conditions	0	0	0	-107	-81
Demographics	0	13	3	2	1

Table 2: Contribution of each factor in the change in homeownership with respect to the 1940s generation (% of the change), by age

At age 30, changes in labor market outcomes explain 80 percent of the homeownership gap of the 1960s generation with respect to the 1940s, mostly due to initial earnings inequality. With a more unequal earnings distribution, and little average increases in earnings, households in low ranks of the income distribution have lower initial and expected lifetime earnings than their counterparts in older generations. These households face two issues when they consider buying a house. First, they are financially constrained, as they need to save for a downpayment and pass an income test. Second, they are aware that having a large mortgage with respect to their incomes is risky, as negative shocks

[25] All elements in the decomposition have potential interaction effects, which means that shutting them on and off alternatively would not sum to 100% of the change. The Shapley-Owen decomposition allows to obtain the total contribution of each element to the change by considering its contribution to every possible permutation of the other factors being on and off, and averaging over all of these.

could take them to a situation in which they must reduce a lot their nondurable consumption to make mortgage payments. Thus, they choose to be renters. For some this is a delay in the decision to buy houses, but for some this state is relatively persistent. At age 40, earnings dynamics still explain 90% of the homeownership gap.

Earnings inequality and risk are closely linked. Even if everyone faced the same distribution of earnings shocks, their impact would depend on their earnings at labor market entry. However, to separate both, I compute the contribution of changes in earnings dynamics over and above initial realizations. At age 40, riskier earnings explain 70% of the drop in homeownership rates. The higher volatility of earnings discourages households from engaging in a large, risky expenditure like a house. At later ages, the dependence on initial earnings realizations progressively dies out and it is harder to disentangle the effects of initial inequality and risk.

The intuition about earnings inequality and earnings risk is supported by the empirical evidence shown in Figure 7. The gap in homeownership rates between the 1940s and 1960s generations is larger for the lowest earners, which is consistent with the contribution of earnings inequality, but there are also differences all across the earnings distribution, which is consistent with the role of earnings risk.

The 1960s generation entered the labor market in a period of cylically low house prices, which explains the low or negative contribution of aggregate conditions to the change in homeownership until the 2000s boom in house prices (when this generation was around 45-50 years of age).

Despite later household formation and a lower number of children for younger generations, the change in the average number of people in a given household at each age (θ_t in the model), which affects consumption needs, has a small effect on homeownership rates.

The 1980s generation entered the labor market in a radically different period. House prices were high both from a secular and cyclical perspective, but financial constraints were laxer. For the 1980s generation, prices alone would have explained the drop in homeownership, but the lower downpayment requirements counteracted most of the potential decrease. This result suggests that changes in financial conditions were key to prevent homeownership rates of younger cohorts to plummet in a context of unstable, unequal earnings and high house prices. The remainder of the difference, over 80 percent at age 30, is also accounted for by earnings dynamics.

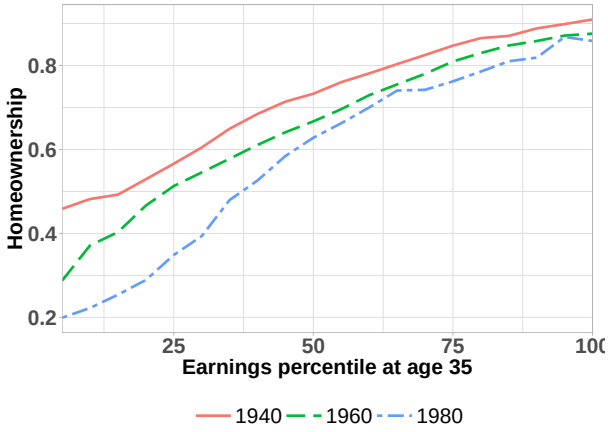

Figure 7: Homeownership by cohorts, by percentile of the earnings distribution at age 35

In all of these experiments, earnings dynamics are computed on household income, so they embed other factors that changed over the generations such as the timing of family formation. However, as shown in Appendix F.2.2, these results are robust to focusing on married couples alone. Additionally, these counterfactual experiments assume that housing supply is perfectly elastic and so house prices would not react to the increase in housing demand induced by the change in the earnings process. In Appendix E.4 I relax this assumption and show that a reduction in earnings inequality and risk would imply a significant increase in homeownership for younger cohorts even if we assume that the increase in demand would drive prices up.

C Explaining the changes in stock market participation

Understanding the increase in stock market participation documented in Section I requires taking into account not only the changes in earnings dynamics and asset returns, but also the progressive reduction in the cost of access to financial markets over time, which is partially related with the introduction of tax-advantaged, employer-sponsored retirement plans.

Figure 8 shows the implications of the model in terms of stock market participation when these changes are taken into account. More specifically, it assumes that stock market participation costs are 50% lower for the 1960s and 85% lower for the 1980s generation than they were for the 1940s generation, and additionally that the initial share of people with positive participation in the stock market has increased over the generations from

20% (1940s and 1960s) to 30% (1980s). Both of these changes capture the reduction in information costs and the effect of auto-enrolment.[26] If the reduction of participation costs is not taken into account, the profiles generated by the model are counterfactual (right panel of Figure 8). The fiscal incentives of IRAs and 401(k)s do not explain the increase in stock market participation either (see Appendix F.3.4).

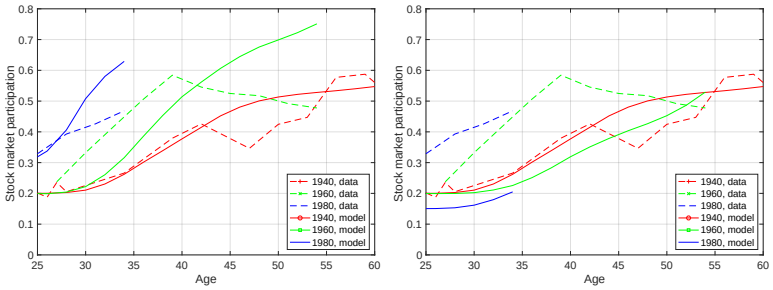

Figure 8: Stock market participation by age and cohort, data vs model. Left: main model; right: constant participation costs

D Implications for wealth accumulation

As a result of these changes in earnings dynamics and asset returns, many younger households are accumulating less wealth than similarly-ranked households in earlier generations (left panel of Figure 9). The right hand side panel of Figure 9 shows that housing played a key role in this change. In the model where households do not enjoy owner-occupied housing (solid lines), most households in the 1980s generation save more, which is consistent with the precautionary savings motive induced by their higher earnings volatility. The complete model (dotted lines) can instead replicate a decrease in wealth accumulation for the bottom 70% of the distribution. Households in the 1940s and 1960s generations used to save more because they bought houses, partially because they enjoyed owning them, which made them act as an indirect way of forced savings. Additionally, homeowners could use their mortgages as leverage. When, because of changes in earnings dynamics, house prices, and financial conditions, households cannot access houses or buy them later, these forces are not in play and households save less. Consistently with this channel, the wealth holdings of the median household at each age have dropped over the generations

[26] Appendix F.3.3 shows that changing the fixed cost of participation for per-period participation costs can generate similar patterns.

(Figure 10), although average wealth holdings have remained relatively stable because of higher wealth accumulation by the richest.

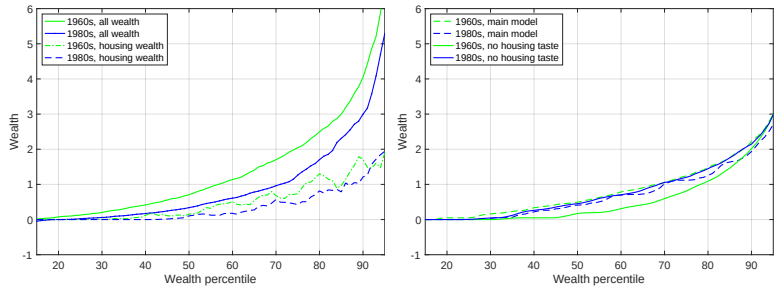

Figure 9: Net worth by wealth percentile, ages 30-40, by generation. Left, SCF data (not available for 1940s with sufficient granularity); right: model implied. Units are multiples of average income. For clarity, the top 5% and bottom 15% of the wealth distribution are not reported.

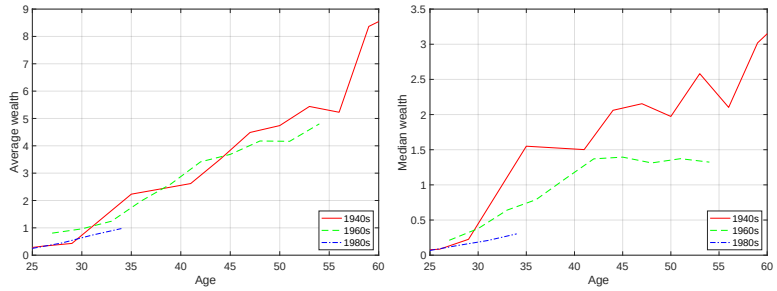

Figure 10: Median (left) and average (right) net worth by age and generation, SCF data. Units are multiples of average income.

Given that 35-year-olds frequently do not have fully equity on their homes, the gap in wealth accumulation between generations is likely to grow as households age and miss out on house price appreciation. In Appendix F.5 I conduct a simulation exercise in which I predict, under a set of assumptions, the evolution of homeownership rates for the 1980s generation beyond 2015-2020.[27] The median prediction shows that the 1980s generation homeownership rates will stay below those of the 1960s as they age.

However, financial wealth is not a perfect substitute for housing wealth. Because fewer households participate in the stock market than become homeowners, and because relatively poorer households do not benefit from leverage when they invest in stocks, it is likely that these changes in household portfolios will lead to increased wealth inequality.

[27]This analysis does not incorporate, due to the delay in availability of survey data, the Covid-19 recession and the subsequent changes in labor market conditions.

In particular, the model predicts that wealth inequality will continue to grow, unless stock market participation costs are reduced such that almost all of the population accesses the stock market by age 60. The larger the share of households that participate in the stock market, the stronger the negative effect on wealth inequality.

E Alternative specifications

Appendix F shows a set of robustness checks. They show that the main messages in the paper are robust to starting households at zero wealth (F.1), different assumptions about the earnings process (F.2), including considering changes in marital dynamics and family formation (F.2.2), different specifications of the asset structure, including alternative versions of the discretization of houses (F.3.1) or per-period stock market participation costs (F.3.3), and different assumptions about the dynamics of aggregate variables (F.4), including local correlation of income shocks and house prices (F.4.3). Appendix F.2.1 shows that a canonical earnings process would overestimate the intergenerational decrease in homeownership by overweighting the role of large initial inequality.

VI Conclusion

In this paper, I study how changes in earnings dynamics over different cohorts have affected their homeownership and portfolio choice decisions. First, I provide empirical evidence, extracted from PSID and SCF data dating back to the 1960s, that there has been a secular increase in household earnings inequality and risk, together with substantial reductions in homeownership and an increase in stock market participation.

Second, I design a flexible earnings process that accomodates rich features of earnings risk, which can be correlated with the aggregate performance of the economy and asset returns. This process replicates features of earnings data by age, over the earnings distribution, and over the business cycle, including the sluggish recovery after a recession.

Third, I develop a rich life-cycle model of housing and portfolio choice with a relatively parsimonious parametrization. Key elements are a taste for owner-occupied housing, a minimum size for houses, transaction costs, and stock market entry participation costs. I use the model to explain the intergenerational changes I observe in the data without assuming preference changes across generations. Differences in earnings dynamics account

for more than half of the reduction in homeownership at ages 30-35.

My findings suggest that intergenerational changes are important for studies of household earnings, consumption, and wealth accumulation. At any point in time, the cross-sectional distribution of the economy is formed by many different households who have lived through different histories of shocks at different points in their lives. Acknowledging this fact matters to understand the economic decisions that have led them to be where they are today, and thus to infer parameters to study the effects of policies or the evolution of the economy. These results are of interest to policymakers who care about homeownership, intergenerational redistribution, and the evolution of inequality.

Finally, this paper also adds to a burgeoning literature that points out that considering household portfolio compositions is important for many macroeconomic questions, such as consumption responses to shocks or wealth accumulation over the life cycle.

References

Acemoglu, Daron, and David Autor. 2011. "Skills, tasks and technologies: Implications for employment and earnings." In *Handbook of labor economics*. Vol. 4, 1043–1171. Elsevier.

Alan, Sule. 2006. "Entry costs and stock market participation over the life cycle." *Review of Economic Dynamics*, 9(4): 588–611.

Ameriks, John, and Stephen P Zeldes. 2004. "How do household portfolio shares vary with age." *Working paper*.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme. 2017. "Earnings and consumption dynamics: a nonlinear panel data framework." *Econometrica*, 85(3): 693–734.

Attanasio, Orazio P, Renata Bottazzi, Hamish W Low, Lars Nesheim, and Matthew Wakefield. 2012. "Modelling the demand for housing over the life cycle." *Review of Economic Dynamics*, 15(1): 1–18.

Bansal, Ravi, and Amir Yaron. 2004. "Risks for the long run: A potential resolution of asset pricing puzzles." *The journal of Finance*, 59(4): 1481–1509.

Becker, Thomas A, and Reza Shabani. 2010. "Outstanding debt and the household portfolio." *Review of Financial Studies*, 23(7): 2900–2934.

Benabou, Roland. 2002. "Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency?" *Econometrica*, 70(2): 481–517.

Bloom, Nicholas, Fatih Guvenen, Luigi Pistaferri, John Sabelhaus, Sergio Salgado, and Jae Song. 2017. "The great micro moderation."

Borella, Margherita, Mariacristina De Nardi, and Fang Yang. 2019. "The lost ones: the opportunities and outcomes of non-college educated americans born in the 1960s." NBER Working Paper.

Braxton, J Carter, Kyle Herkenhoff, Jonathan Rothbaum, and Lawrence Schmidt. 2021. "Changing income risk across the US skill distribution: Evidence from a generalized Kalman filter." *working paper*.

Busch, Christopher, and Alexander Ludwig. 2017. "Higher-order income risk over the business cycle: A parametric approach." mimeo.

Campbell, John Y, and Motohiro Yogo. 2006. "Efficient tests of stock return predictability." *Journal of financial economics*, 81(1): 27–60.

Carr, Michael D, and Emily E Wiemers. 2018. "New evidence on earnings volatility in survey and administrative data." Vol. 108, 287–91.

Chambers, Matthew, Carlos Garriga, and Don E Schlagenhauf. 2009. "Accounting for changes in the homeownership rate." *International Economic Review*, 50(3): 677–726.

Chetty, Raj, László Sándor, and Adam Szeidl. 2017. "The effect of housing on portfolio choice." *The Journal of Finance*, 72(3): 1171–1212.

Cocco, Joao F. 2005. "Portfolio choice in the presence of housing." *Review of Financial studies*, 18(2): 535–567.

Cocco, Joao F, Francisco J Gomes, and Pascal J Maenhout. 2005. "Consumption and portfolio choice over the life cycle." *Review of financial Studies*, 18(2): 491–533.

Davis, Morris A, and François Ortalo-Magné. 2011. "Household expenditures, wages, rents." *Review of Economic Dynamics*, 14(2): 248–261.

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

Davis, Steven J, and John Haltiwanger. 2014. "Labor market fluidity and economic performance." National Bureau of Economic Research.

De Nardi, Mariacristina. 2004. "Wealth Inequality and Intergenerational Links." *Review of Economic Studies*, 71: 743–768.

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

Duca, John V, John Muellbauer, and Anthony Murphy. 2011. "House prices and credit constraints: Making sense of the US experience." *The Economic Journal*, 121(552): 533–551.

Dynan, Karen E, Douglas W Elmendorf, and Daniel E Sichel. 2006. "Can financial innovation help to explain the reduced volatility of economic activity?" *Journal of monetary Economics*, 53(1): 123–150.

Epstein, Larry G, and Stanley E Zin. 1989. "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework." *Econometrica: Journal of the Econometric Society*, 937–969.

Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso. 2017. "Asset Market Participation and Portfolio Choice over the Life-Cycle." *The Journal of Finance*, 72(2): 705–750.

Fagereng, Andreas, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik. 2019. "Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains." *Working paper*.

Falcettoni, Elena, and James Schmitz. 2018. "Public Policies to Alleviate the Housing Crisis." *mimeo*.

Fischer, Marcel, and Natalia Khorunzhina. 2019. "Housing Decision with Divorce Risk." *International Economic Review*.

Fisher, Jonas DM, and Martin Gervais. 2011. "Why has home ownership fallen among the young?" *International Economic Review*, 52(3): 883–912.

Flavin, Marjorie, and Takashi Yamashita. 2011. "Owner-occupied housing: life-cycle implications for the household portfolio." *The American Economic Review*, 101(3): 609–614.

Fraumeni, Barbara. 1997. "The measurement of depreciation in the US national income and product accounts." *Survey of Current Business - US Department of Commerce*, 77: 7–23.

Fujita, Shigeru. 2018. "Declining labor turnover and turbulence." *Journal of Monetary Economics*, 99: 1–19.

Gerardi, Kristopher, Harvey S Rosen, and Paul Willen. 2007. "Do households benefit from financial deregulation and innovation? The case of the mortgage market." National Bureau of Economic Research.

Glaeser, Edward L, Joshua D Gottlieb, and Joseph Gyourko. 2012. "Can cheap credit explain the housing boom?" In *Housing and the financial crisis*. 301–359. University of Chicago Press.

Goldin, Claudia Dale, and Lawrence F Katz. 2009. *The race between education and technology*. Harvard University Press.

Gomes, Francisco, and Alexander Michaelides. 2005. "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence." *The Journal of Finance*, 60(2): 869–904.

Gottschalk, Peter, Robert Moffitt, et al. 1994. "The Growth of Earnings Instability in the US Labor Market." *Brookings Papers on Economic Activity*, 25(2): 217–272.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2016. "What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk?" Working paper, University of Minnesota.

Guvenen, Fatih, Greg Kaplan, Jae Song, and Justin Weidner. 2017. "Lifetime incomes in the United States over six decades."

Guvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy*, 122(3): pp. 621–660.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante. 2005. "Two views of inequality over the life cycle." *Journal of the European Economic Association*, 3(2-3): 765–775.

Hyatt, Henry R, and James R Spletzer. 2016. "The shifting job tenure distribution." *Labour Economics*, 41: 363–377.

IRS. 2014. "Individual Income Tax Returns Complete Report." *Internal Revenue Service, Publication 1304*.

Johnson, Kathleen W, and Geng Li. 2010. "The Debt-Payment-to-Income Ratio as an Indicator of Borrowing Constraints: Evidence from Two Household Surveys." *Journal of Money, Credit and Banking*, 42(7): 1373–1390.

Kaplan, Greg, and Giovanni L Violante. 2014. "A model of the consumption response to fiscal stimulus payments." *Econometrica*, 82(4): 1199–1239.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider. 2015. "The housing market (s) of San Diego." *American Economic Review*, 105(4): 1371–1407.

Lovenheim, Michael F. 2011. "The effect of liquid housing wealth on college enrollment." *Journal of Labor Economics*, 29(4): 741–771.

Mankiw, N Gregory, and David N Weil. 1989. "The baby boom, the baby bust, and the housing market." *Regional science and urban economics*, 19(2): 235–258.

Martin, Robert F. 2005. "The baby boom: predictability in house prices and interest rates." *FRB International Finance Discussion Paper*, , (847).

McGee, Rory. 2021. "Old Age Savings and House Price Shocks." *Working paper*.

Mitchell, Olivia S, and John WR Phillips. 2006. "Social security replacement rates for alternative earnings benchmarks."

Molloy, Raven, Christopher Smith, and Abigail K Wozniak. 2021. "Changing stability in us employment relationships: A tale of two tails." *Journal of Human Resources*, forthcoming.

Nakajima, Makoto. 2005. "Rising earnings instability, portfolio choice, and housing prices." *Unpublished Manuscript*.

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

Panel Study of Income Dynamics. 2017. "Public use dataset." *Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI*.

Paz-Pardo, Gonzalo. 2022. "Data and Code for: Homeownership and Portfolio Choice over the Generations." *American Economic Association*, [publisher], Inter–university Consortium for Political and Social Research [distributor], doi https://doi.org/10.3886/E182663V1.

Pesaran, M Hashem, and Allan Timmermann. 1995. "Predictability of stock returns: Robustness and economic significance." *The Journal of Finance*, 50(4): 1201–1228.

Santos, Cezar, and David Weiss. 2016. ""Why not settle down already?" a quantitative analysis of the delay in marriage." *International Economic Review*, 57(2): 425–452.

Smith, Lawrence B, Kenneth T Rosen, and George Fallis. 1988. "Recent developments in economic models of housing markets." *Journal of economic literature*, 26(1): 29–64.

Sommer, Kamila. 2016. "Fertility choice in a life cycle model with idiosyncratic uninsurable earnings risk." *Journal of Monetary Economics*, 83: 27–38.

Sommer, Kamila, and Paul Sullivan. 2018. "Implications of US tax policy for house prices, rents, and homeownership." *American Economic Review*, 108(2): 241–74.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 2004. "Cyclical Dynamics in Idiosyncratic Labor-Market Risk." *Journal of Political Economy*, 112: 695–717.

Survey of Consumer Finances. 1960 to 1986. *University of Michigan. Survey Research Center. Economic Behavior Program.*

Survey of Consumer Finances. 2019. *Federal Reserve Board, U.S. Treasury Department, and National Opinion Research Center.*

Vestman, Roine. 2012. "Limited stock market participation among renters and home owners."

Yao, Rui, and Harold H Zhang. 2005. "Optimal consumption and portfolio choices with risky housing and borrowing constraints." *Review of Financial studies*, 18(1): 197–239.