Leveraged Landlords: Life-Cycle Portfolio Choice With Rental Properties, Mortgages, and Margin Calls

Lucas Ramiro Szwarcberg

Presented to the Department of Applied Mathematics
in partial fulfillment of the requirements
for a Bachelor of Arts degree with Honors

Harvard College
Cambridge, Massachusetts
March 24, 2023

Abstract

Most rental housing is supplied by amateur households, rather than professionally run businesses, yet little is known about what motivates individuals to invest in this sector. I examine the hypothesis that households invest in this asset class because mortgages allow for high leverage that is not subject to margin calls. To do this, I model portfolio allocation between rental properties and stocks for investors with future expected labor income, and allow for both mortgage and stock-margin leverage. My analysis demonstrates that rental properties partially crowd out stock ownership for young investors, who are attracted to the high leverage available with mortgages. The model results show that initial wealth has a non-monotonic effect on the amount of rental housing purchased, because rental-property ownership is most attractive at medium values of wealth, which are high enough to afford the smallest possible property without taking on too much risk and low enough that future income makes leverage highly desirable. In addition, I analyze how the results change depending on the availability of leveraged stock investment, the rental yield of the property, and the riskiness of labor income.

Acknowledgements

First and foremost, I want to express my profound gratitude to my advisor, Professor John Campbell, for his unwavering kindness, encouragement, and mentorship throughout this process. His guidance during numerous meetings and email exchanges has been invaluable in shaping my thesis. Furthermore, my work builds heavily on his own research on mortgages and household finance, making him the most fitting advisor I could have hoped for.

I also owe a great deal to the faculty members who have taught and inspired me over the years, particularly Professors Greg Mankiw and Jerry Green. I am also grateful for the friendships I have had with my classmates in economics at Harvard. The intellectual exchanges I've had with both professors and friends have significantly contributed to my development as a scholar.

I would like to extend my sincere thanks to my good friend, Ramiz A. Razzak, for his valuable help during the ideation phase, as well as for reviewing multiple drafts. His keen eye for detail and thoughtful suggestions have greatly enhanced this work, reflecting the depth of our collaborative relationship and mutual respect. I am also grateful to my friends Laura Nicolae and Vitor Weiss Jung for providing valuable feedback on my thesis draft that helped to refine my work.

I have been fortunate to receive assistance from several economists in the course of my research. I am grateful for the useful guidance that Slavik Sheremirov, my section leader, provided throughout my thesis work. Special thanks go to Barry Nalebuff for kindly reading my work and offering helpful suggestions, especially regarding the role of REITs. Both he and Ian Ayres have been important inspirations, as their book initially sparked my interest in this research topic. I appreciate Marco Giacoletti's assistance in providing the ZIP-code-level risk data that was crucial for my analysis and for answering my questions about interpreting the data. I would also like to thank Adam Guren for addressing my initial inquiries about the data for my thesis.

Lastly, I want to express my heartfelt appreciation to Kristine Ma and my parents for their unwavering support and encouragement. Their continued faith in my abilities has been a driving force behind the completion of this thesis.

> [The businessman] can look forward to a high salary in the future; and with so high a present discounted value of wealth, it is only prudent for him to put more into common stocks compared to his present tangible wealth, borrowing if necessary for the purpose[.]
>
> *Samuelson (1969)*

In his seminal paper on lifetime portfolio selection, Samuelson (1969) argues that the conventional wisdom that a businessman should take greater risk in his investments than a widow or retiree because of his longer time horizon is misguided. This is because multiplying the same kind of independent risk many times increases risk rather than cancelling it out, and thus stocks are riskier when they are held longer. However, Samuelson also argues that the businessman's greater anticipated future income justifies holding more stocks, borrowing if necessary. In the limiting case with certain future income, this income can be regarded as an implicit bond holding, and thus the businessman will usually be overweight in bonds relative to his optimal portfolio choice unless he borrows, even if that optimal allocation is the same as that of the widow. What if there were a way to implement this idea and diversify risk across time, without the limitations of leveraged investment in stocks that have made it unpopular for long-term investment? In this paper, I explore the use of leveraged investment in rental properties for this purpose. I show that, because investors with a large expected income can anticipate future savings, leveraging investments in rental properties over time can allow them to bring their holding of risky assets closer to the liquidity-unconstrained optimum.

Ayres and Nalebuff (2008, 2010) have examined and recommended the use by young people of leveraged investment in stock indices, for the same reasons as Samuelson, and discussed leveraged exchange-traded funds, stocks in a margin account, futures, and deep-in-the-money call options as potential ways to do so. They suggested using wealth inclusive of expected discounted future savings as the base and investing a fixed share of that wealth, determined by risk aversion.[1] Despite their endorsement, however, leveraged investment in stocks remains a niche endeavor: about 0.25% of U.S. individuals held margin loans in 2019 (Board of Governors of the Federal Reserve System 2019). Leveraged investment in real estate, by contrast, is much more popular. 6.9% of U.S. taxpayers reported individual rental income in 2019 (Internal Revenue Service 2021), many of whom likely financed the property with a mortgage. In Canada, 7.9% of individual taxpayers reported rental income in 2020, which has increased since 2008 when it was 7.0% (Statistics Canada 2022). Furthermore, 19% of mortgaged house purchases in Canada between 2014 and 2021 were made by investors who already had another mortgage (Khan and Xu 2022), indicating that purchases of investment properties make up a significant share of the mortgage market. A central aim of this paper is to explain why leveraged real-estate investment is so much more popular among households than leveraged stock investment.

Mortgages differ from the stock-brokerage-based borrowing methods considered by Ayres and Nalebuff in several ways. Unlike these methods, mortgages take into account the borrower's income, debts, and credit history to determine the loan amount and terms. This allows qualifying investors to borrow more than they could through other methods, such as buying stock on margin. Mortgages typically require an initial down payment of 20%, as opposed to 50% for margin investing. Furthermore, mortgages can go underwater, meaning that the amount outstanding on the loan can exceed the value of the property, as long as

1. The authors provide an alternative intuition for leveraging, emphasizing improved diversification across time for a fixed level of lifetime risk. This is because one can achieve lower risk by investing \$2 in each of three periods than by investing \$1, \$2 and \$3 in each period, although the total expected return is the same. Assuming that the returns r_t in each period t are i.i.d., and that there is no compounding across periods, then the return to the first strategy is $E(r_t)(2+2+2) = 6E(r_t)$ while the return to the second strategy is $E(r_t)(1+2+3) = 6E(r_t)$. Meanwhile, the variance of the first strategy, $3\,\mathrm{Var}(2r_t) = 12\,\mathrm{Var}(r_t)$, is less than that of the second strategy, which is $\mathrm{Var}(r_t) + \mathrm{Var}(2r_t) + \mathrm{Var}(3r_t) = 14\,\mathrm{Var}(r_t)$.

While this helps illustrate the logic of time diversification, note that in a simple model of maximizing CRRA utility of final wealth, adding access to leverage will not usually result in the total risk's being held constant.

payments are being made, whereas margin-based methods never allow for a negative-equity position in the account. They instead impose a "maintenance-margin" constraint that must be satisfied at all times, liquidating the position if it is not. This means that mortgages can effectively allow greater leverage to be maintained in drawdowns compared to margin investing.

Because mortgage borrowing can only be used to purchase individual properties, whether access to leveraged investment in real estate is welfare-enhancing will naturally depend on the return and risk to holding individual real estate properties and how they compare to stock indices. One influential recent paper examines the total return to equity, housing, bonds and bills in 16 advanced economies between 1870 and 2015, and finds that residential real estate has similar average real returns but lower volatility compared to equities (Jordà et al. 2019). They show that while capital gains on residential properties tend to exceed inflation only slightly, the actual or imputed rents account for the rest of the significant risk premium. This explains why their results are not at odds with economists' conventional wisdom that real capital gains on U.S. real estate are low compared to equities: whereas the well-known Shiller U.S. house price index has geometric-average real returns of 0.56% per year, Jordà et al. (2019)'s dataset implies a similar real capital gain of 0.61% over the same period, alongside a rental yield of 5.32%, so rental yields make up the majority of the historical total return in the U.S. (IA25).

As opposed to investing in a stock index, investing in a particular rental property also carries undiversifiable risk, both between locations and between ex-ante-comparable houses in the same location. Giacoletti (2021) finds that this idiosyncratic risk affects the risk-adjusted return of housing to a lesser extent the greater the holding period is, i.e., it does not follow a random walk on an annualized basis. He establishes this using a proprietary CoreLogic dataset of property deeds of single-family residences in California between 1996 and 2018. Amaral et al. (2021) extend this analysis to a larger CoreLogic dataset covering many U.S.

metropolitan areas, and find similar results.[2] Jordà, Schularick, and Taylor (2019) find that even when accounting for this idiosyncratic risk, housing still has historical returns comparable to equities with lower risk, which they call a new risk premium puzzle. This recent literature on aggregate returns and idiosyncratic risk justifies my exploration of real estate as an attractive investment even before leverage is taken into account.

The analysis in this paper focuses on rental properties rather than owner-occupied properties. It is true that most households will be best off buying a house to live in before they buy any rental housing, for a few reasons. First, because the rent savings from living in a house one owns are usually not taxed, whereas income from a rental property is taxed, it would be tax-inefficient to rent out one property and pay rent for another. Second, renting involves moral hazard from unobservable tenant behavior, which increases the cost of renting relative to owning. Third, capital gains on owner-occupied housing are partially tax-exempt in the U.S. However, these reasons do not mean it is beneficial to buy the largest house possible: if someone buys a bigger house than they would otherwise like in order to have greater expected capital gains, they will be missing out on rental income, or on greater returns from investing the difference in the stock market. This is because a significant portion of housing returns are rents; the total unlevered returns (capital gains plus dividends or rents) to stocks and housing are comparable, but housing's capital gains alone cannot match stocks' total return. My paper therefore focuses on modeling the allocation of funds left over after potentially buying owner-occupied housing–that is, funds devoted solely to investment. This is in light of the fact that comprehensive models of the joint choice of stock ownership and owner-occupied housing ownership already exist, such as those of Yao and Zhang (2005a, 2005b) and Cocco (2005), to which my paper can be seen as complementary.[3]

In this paper, I develop a model of portfolio allocation between mortgaged rental-property

2. They additionally show the idiosyncratic component to be decreasing in metropolitan-area size (Amaral et al. 2021, 33–34). They also study the cross-section of regional average returns, showing that smaller cities have higher total returns than larger cities, but also higher risk, such that the overall Sharpe ratios are similar.

3. I discuss the separability of investment-portfolio choice from owner-occupied housing in Subsection 5.3.

investment and stock investment by investors with future labor income. I model the fact that stock investment on margin is subject to margin calls, while mortgages are not. Furthermore, while Ayres and Nalebuff's analysis assumed away risk in labor income, I generalize labor income to be stochastic. Welfare is evaluated in terms of isoelastic utility of final wealth. I calibrate the model with empirical parameter estimates, assuming conservatively that housing has the same risk and return as stocks rather than the similar return and lower risk documented by Jordà et al. (2019) and Jordà, Schularick, and Taylor (2019). I show that an investor with a typical income level must have over two years' worth of labor income saved for it to be optimal to purchase rental housing, at which point they buy a property worth 1.2 times their wealth, with a down payment of 70%. I also report the model solution for all relevant values of the wealth-to-income ratio, and analyze how the results change depending on the presence of a minimum-house-size constraint, the availability of leveraged stock investment, the rental yield of the property, and the riskiness of labor income. I also report model results for a number of changes to the calibration assumptions.

My analysis identifies the initial ratio of wealth to income as a key determinant of whether the household purchases a rental property. Since I model the fact that there is a minimum house value, rental-property ownership is not feasible for very small values of initial wealth, even with the lowest possible down payment. There is modest crowding-in from the minimum-house-size constraint: for medium values of initial wealth, households buy slightly more housing than they would if not for the constraint. However, there are also values of initial wealth for which households can afford to buy a rental property, but do not do so. Intuitively, this is because the house they would buy if there was no minimum-house-size constraint is much smaller than the house they are able to buy, such that the risk they would have to take on in order to buy a property is too high to be desirable. Hence, my model results show that rental-property ownership is most attractive at medium values of wealth, which are high enough to afford the smallest possible property without taking on too much risk and low enough that future income makes leverage highly desirable.

While households use a down payment of at least 70% in my baseline calibration, I also examine how investor heterogeneity may result in the higher levels of mortgage leverage that some real-world investors employ. I first document the interaction between household income and the minimum-house-size constraint. Whereas my baseline case sets household income to the U.S. median, higher incomes relax the minimum-house-size constraint for given values of the wealth-to-income ratio. For instance, with a household income of $375,000, households first purchase housing at the low wealth-to-income ratio of 0.16. They nearly maximize their leverage, taking out a 20% down payment on a house worth 4 times their wealth. Drawing upon empirical evidence that the top income quintile saves at higher rates, I show that higher rates of anticipated saving out of future income also lead to higher levels of housing investment and mortgage leverage. Therefore, in addition to finding a role for leveraged rental-property investing in the typical investor's portfolio with a relatively high down payment, my model predicts that higher household incomes or higher savings rates can lead to lower down-payment ratios, including some at the 20% lower bound.

My paper contributes to several literatures. First, I add to the literature on life-cycle portfolio choice, which has so far focused on leveraged investment in owner-occupied housing (Yao and Zhang 2005a, 2005b; Cocco 2005) and in stocks (Ayres and Nalebuff 2008, 2010), by introducing the option to invest in a rental property using a mortgage that isn't subject to margin calls. This has several potential consequences. First, commentators on leveraged stock investing tend to see it as idiosyncratic and unusual.[4] My model demonstrates that a sizeable number of investors may already be pursuing time diversification, only in a different form than stock investment. Indeed, this could help explain why most Americans think a rental property is a better investment than stocks (Haughwout 2021). Second, since my model's investment choices nest those of Ayres and Nalebuff (2008, 2010), while allowing for greater leverage through mortgages and for diversification through housing investment, the strategy that solves my model is welfare-improving relative to the strategies they consider.

4. For instance, Tabarrok (2022) commented that "most people won't do this because unlike a home-mortgage it's a non-standard idea."

My paper can also be seen as a model of rental supply, and in particular of the behavior of small-scale or amateur landlords. Despite the fact that households supply the majority of rental housing (Shroder 2001, 1070), little is known about these amateur landlords, whether theoretically or empirically. Plaut and Plaut (2013, 121–123) provide a review of this small literature to which my paper contributes. To my knowledge, this is the first economic model to consider leverage opportunities as an explanation for the involvement of amateurs in the supply of rental housing.

The paper is organized as follows. Section 1 sets up the model. Section 2 describes the calibration of the model parameters to match real-world conditions. Section 3 describes my solution method. Section 4 reports the results. Section 5 discusses the implications and limitations of the model and of its results. Section 6 concludes.

1 Model

I now detail the theoretical model under which households invest. In terms of its theoretical setup, the paper most closely related to mine is Yao and Zhang (2005b), which models consumption and portfolio choice in the presence of a mortgage without maintenance margin, of future labor income and of the option to invest in stocks. My paper's setup differs from theirs in a few important ways. First, while they consider utility derived from ordinary consumption and housing services in the context of owner-occupied housing, I focus on rental properties, which has the benefit of allowing the investor to receive and save rental income rather than only being able to directly consume the housing services provided by the house. Second, I consider leverage in stocks, which they rule out. Third, whereas they model intertemporal utility from consumption, my model focuses on utility from final wealth, with a constant savings rate in the other periods. This matches the assumption from Ayres and Nalebuff's work, and has the appeal of making the consumption-savings choice exogenous, although it is a restrictive assumption (Ayres and Nalebuff 2008, 3). Because the ability to continue holding a mortgage while its value exceeds that of the home is central to my model, the setup also bears similarity to models focusing on mortgage default, such as Campbell and Cocco (2015) and Chen, Michaux, and Roussanov (2020), although those papers do not consider portfolio choice between stocks and housing.

The flow of time in the model is as follows. There are T periods, $t = 0, 1, \dots, T-1$. Each period has a beginning, middle and end. Variables without a superscript relate to the beginning of the period (e.g., S_t), whereas those with the superscripts "mid" and "end" relate to the middle and end of the period (e.g., S_t^{mid}).

At the beginning of each period, the household chooses their portfolio. In the middle of the period, the returns on the chosen risky assets are realized, and the household's liquid wealth is calculated after receiving rental income and labor income and after making the mortgage payment. At the end of the period, the household is subject to foreclosure or default if their middle-of-period liquid wealth was negative. Each period's end-of-period wealth becomes the following period's beginning-of-period wealth.

1.1 The house-purchase decision

At time 0, the household is endowed with wealth W_0 and makes a choice about the amount H_0 of housing to purchase. They can choose not to purchase housing ($H_0 = 0$), but if they purchase a nonzero amount of housing, then they must purchase a minimum of $\underline{H}_0$ (housing is lumpy).[5]

The household may finance their housing purchase by taking out a mortgage in the amount of $M_0 \geq 0$. The amount of housing the household can purchase is limited by a minimum down-payment ratio α, so $1 - \frac{M_0}{H_0} \geq \alpha$ (or equivalently, $M_0 \leq (1-\alpha)H_0$); the remainder must be financed out of existing wealth, so $H_0 - M_0 \leq W_0$.[6] Combining these two conditions, we

5. In particular, the type of housing that can be purchased with a mortgage (an individual rental property) is lumpy. I discuss real-estate investment trusts (REITs), shares of which could be purchased in very small increments, in Subsection 5.6.

6. One could also have assumed that there is a maximum mortgage payment β, such that we must have $P_0 \leq \beta$, to model loan-to-income and other credit qualifications based on the household's particular

see that the maximum house value the household can afford is given by $H_0 \leq \frac{W_0}{\alpha}$.

In the middle of each period $t \in \{0, \ldots, T-1\}$, there is a payment due on the mortgage. I model a fixed-rate mortgage (FRM) since it is the predominant mortgage type in the U.S. The real interest rate r_{FRM} is constant for all periods, and so is inflation π. The payments P_t^{mid} are calculated to have the same nominal value in all periods, as is the case in practice, and are given in real terms by a modified annuity formula: $P_t^{\text{mid}} = \frac{M_0 \times i_{\text{FRM}}}{(1-(1+i_{\text{FRM}})^{-T})(1+\pi)^t}$, where $i_{\text{FRM}} = (1 + r_{\text{FRM}})(1 + \pi) - 1$.

In the middle of each period $t \in \{0, \ldots, T-1\}$, the house yields a rent yH_t^{mid}, where y is a constant rental yield. This rental income is assumed to be net of property taxes and routine maintenance costs. This assumption of a constant rental yield is in line with the rent-or-own model of Yao and Zhang (2005a, 201), and is reasonable empirically: the standard deviation of the rental yield has historically been low, at 0.8% in the U.S. according to the database by Jordà et al. (2019). Like Giacoletti (2021), I ignore the additional property-level yield risk investigated empirically by Eichholtz et al. (2021).

The house is held until retirement at the end of period $T - 1$.[7]

1.2 The choice of positions in the brokerage account

At each time $t \in \{0, \ldots, T-1\}$, the household can take a position in cash Z_t and stocks S_t. The household may not take a short position in stocks: $S_t \geq 0$. They can take a short position in the first asset, cash, but are limited by a maintenance-margin constraint that holds in every period: $1 + \frac{Z_t}{S_t} \geq g_m$ (or equivalently, $Z_t \geq (g_m - 1)S_t$).

In practice, margin lending typically also involves an initial-margin requirement, at a higher rate.[8] As in Willen and Kubler (2006), I only consider the maintenance-margin

characteristics as of time 0. I do not develop that here, and it is not included in models such as Cocco (2005) and Yao and Zhang (2005a). Some examples of papers including a loan-to-income constraint are Campbell and Cocco (2015) and Chen, Michaux, and Roussanov (2020), which focus on mortgage default.

7. I discuss this assumption in Subsection 5.2.

8. If we defined a required initial-margin ratio g_i, we could require that when $S_t > S_{t-1}^{\text{mid}}$ for $t \geq 1$, or when $t = 0$, we must have $1 + \frac{Z_t}{S_t} \geq g_i$.

requirement.

1.3 Labor-income and risky-asset processes

I assume that the asset returns on stocks and housing, as well as the evolution of labor income, follow these respective LogNormal distributions:

$$R_{\mathrm{s},t} \sim \mathrm{LogNormal}(\mu_{\mathrm{s}} - \frac{1}{2}\sigma_{\mathrm{s}}^2, \sigma_{\mathrm{s}}^2) \tag{1}$$

$$R_{\mathrm{h},t} \sim \mathrm{LogNormal}(\mu_{\mathrm{h}} - \frac{1}{2}\sigma_{\mathrm{h}}^2, \sigma_{\mathrm{h}}^2) \tag{2}$$

$$R_{\mathrm{y},t} \sim \mathrm{LogNormal}(\mu_{\mathrm{y}} - \frac{1}{2}\sigma_{\mathrm{y}}^2, \sigma_{\mathrm{y}}^2). \tag{3}$$

for each time t, where the μ's are the means of the simple returns rather than the log returns. Furthermore, I assume that $\rho_{\mathrm{s,h}} \equiv \mathrm{Corr}(\log R_{\mathrm{s},t}, \log R_{\mathrm{h},t})$, $\rho_{\mathrm{s,y}} \equiv \mathrm{Corr}(\log R_{\mathrm{s},t}, \log R_{\mathrm{y},t})$ and $\rho_{\mathrm{h,y}} \equiv \mathrm{Corr}(\log R_{\mathrm{h},t}, \log R_{\mathrm{y},t})$ for all t, thereby parametrizing the within-period correlation of the logarithms of the different processes. Also, random variables are assumed to always be independent across periods.

The assumption of LogNormal and time-independent returns for stocks and housing is widespread in the portfolio-choice literature (Cocco 2005; Yao and Zhang 2005a). However, it is common in the literature to model the mean of the income process as age-dependent, and to allow for transitory shocks to income rather than only permanent ones (Cocco, Gomes, and Maenhout 2005). I omit those elements for simplicity, following Tan (2013) and Giacoletti (2021).

1.4 The household's flow budget constraint

The household chooses the allocation to cash, stocks, housing and mortgage at the beginning of period 0 under the following budget constraint:

$$Z_0 + S_0 + H_0 - M_0 = W_0. \tag{4}$$

In later periods, holdings of stocks and cash are chosen at the beginning of each period $t \in \{1, \ldots, T-1\}$, with a total equal to the previous period's end-of-period liquid wealth. Also, the housing at the beginning of the period will be the same as that at the end of the previous period. Therefore, the flow budget constraints at time t are:

$$Z_t + S_t = L_t = L_{t-1}^{\text{end}} \tag{5}$$

$$H_t = H_{t-1}^{\text{end}}, \tag{6}$$

where L_{t-1}^{end} and H_{t-1}^{end} are determined by the middle-of-period and end-of-period processes described below.

In the middle of each period $t \in \{0, \ldots, T-1\}$, after the household has made the choices relevant to that period, the following random variables are realized:

$$S_t^{\text{mid}} = R_{\text{s},t} S_t \tag{7}$$

$$H_t^{\text{mid}} = R_{\text{h},t} H_t \tag{8}$$

$$Y_t^{\text{mid}} = R_{\text{y},t} Y_t, \tag{9}$$

where $Y_t = Y_{t-1}^{\text{mid}}$ if $t > 0$ (so that the $R_{\text{y},t}$'s are permanent shocks to income), and where Y_0 is an exogenous parameter.

At this same time, the household receives labor income Y_t^{mid}, of which they save a constant share Γ. The household also receives rental income yH_t^{mid}, as mentioned earlier. The interest rate on the household's cash is r_{f} or r_{m} (where the "m" refers to margin), depending on whether the cash balance is positive or negative. We have that for $t \in \{0, \ldots, T-1\}$, the liquid wealth at the middle of period t is

$$L_t^{\text{mid}} = S_t^{\text{mid}} + yH_t^{\text{mid}} + (1 + r_{\text{f}}) Z_t I(Z_t > 0) + (1 + r_{\text{m}}) Z_t I(Z_t < 0) + \Gamma Y_t^{\text{mid}} - P_t^{\text{mid}}, \tag{10}$$

where I is the indicator function.

If $L_t^{\text{mid}} \geq 0$ for $t \in \{0, \ldots, T-1\}$, then $L_t^{\text{end}} = L_t^{\text{mid}}$ and $H_t^{\text{end}} = H_t^{\text{mid}}$. Otherwise, if $L_t^{\text{mid}} < 0$ for $t \in \{0, \ldots, T-1\}$, then the household cannot make their mortgage payment even if they sell all their stocks and use all cash. In this case, the house is foreclosed at value $(1-\eta)H_t$, where η represents a foreclosure penalty proportional to the price of the house. The end-of-period liquid wealth at time t is then:

$$L_t^{\text{end}} = \begin{cases} L_t^{\text{mid}} + (1-\eta)H_t - M_t & \text{if } L_t^{\text{mid}} + (1-\eta)H_t - M_t > 0 \\ \nu Y_t & \text{otherwise,} \end{cases} \tag{11}$$

where M_t is the balance of the mortgage, and in that case $H_t^{\text{end}} = 0$, and the position in housing will remain zero for all future periods. Therefore, the lender may seize the borrower's financial assets at the time of default, but does not have a claim to their future labor income.[9] These same mortgage terms are also studied in Campbell and Cocco (2015, 1506) in an extension to their main model.

The parameter ν is a fraction of current labor income that is accessible in case of default and is useful for avoiding evaluation of the utility function at zero or negative values. A similar component is used in Campbell and Cocco (2015), where it is motivated as resulting from social insurance.

1.5 The overall optimization problem

The household's optimization problem involves choosing a housing amount H_0, a mortgage amount M_0, and an initial asset allocation $\{S_0, Z_0\}$ at time 0. At each future time $t \in \{1, \ldots, T-1\}$, they will also choose asset allocation $\{S_t, Z_t\}$ under the constraint generated by the income and returns realized up to that time. The household knows the distributions of stochastic processes generating $\{R_{\text{s},t}, R_{\text{h},t}, R_{\text{y},t}\}_{t=0}^{T-1}$.

They choose these housing and other risky-asset allocations to maximize expected utility over final wealth, which we write as $E[U(L_{T-1}^{\text{end}} + H_{T-1}^{\text{end}})]$, since the house is assumed to be

9. Laws on recourse against the borrower in the case of mortgage default vary across U.S. states. Broadly, mortgages can be "recourse" or "non-recourse", as summarized in Campbell and Cocco (2015, 1505–1506). Upon default on a non-recourse mortgage, the lender has no recourse beyond seizing the house itself. The situation modeled here is equivalent to a recourse mortgage with a subsequent bankruptcy that erases the residual claim to future labor income.

sold at the end of period $T-1$. They make those choices subject to the period constraints outlined in previous sections. I assume that utility is given by the isoelastic utility function, as is standard in the literature on life-cycle portfolio choice: $U(w) = \frac{w^{1-\chi}}{1-\chi}$ for some coefficient of relative risk aversion (CRRA), χ, such that $\chi > 0$ and $\chi \neq 1$.

2 Calibration

I now turn to the calibration of the model parameters to match the practical situation studied in this paper. These are summarized in Table 1, at the end of the section.

First off, I set $T = 30$, interpreted as years, since this is both a reasonable timeline to save for retirement and is also the most common maturity for a fixed-rate mortgage. The period length of one year is useful to limit the computational complexity and is in line with the literature.[10] However, this is a simplification because in practice, a mortgage could be foreclosed after missing payments for less than a full year, and because households may be willing to rebalance stock holdings more often than once a year. A horizon of 30 years is also used in Campbell and Cocco (2003). It is worth noting that the average age of rental-property buyers is higher than this would imply, being 44 to 51 years old according to a Canadian study (Teranet 2022).

2.1 Labor income and risky-asset processes

I set the risk premium for equities, $R_{s,t} - (1 + r_f)$, to have a mean of 6.1% and a standard deviation of 16.7%, as to match the arithmetic mean and the standard deviation of the historical risk premium of the world stock index between 1900 and 2015 (Dimson, Marsh, and Staunton 2006), with dividends reinvested.

While it was preferable to fix the mean stock returns according to Dimson, Marsh, and Staunton (2006) because they provide a market-cap-weighted world index return, the database

10. For instance, Campbell and Cocco (2003, 1464) use a period length of 2 years for computational tractability.

by Jordà et al. (2019) provides information about the relationship between stocks and housing returns: when weighting equally across countries, their total real returns averaged 6.66% and 6.74%, respectively, for the same 1900–2015 period.[11] This justifies setting the total risk premium for housing to have the same mean as that of stocks.

The corresponding standard deviations from Jordà et al. (2019) are 23.84% and 10.40%. However, as explained earlier, that standard deviation is not inclusive of idiosyncratic risk. Giacoletti (2021) finds idiosyncratic risk to make up 44% of the variance for the longest holding period studied. Under the assumption that the same relationship holds outside that paper's California sample, the 10.40% index standard deviation would correspond to a 13.96% total standard deviation for housing, which is still much lower than the 23.84% standard deviation of stocks from Jordà et al. (2019).

Giacoletti's definition of idiosyncratic risk pertains only to variance across similar houses within a ZIP code, but an investor buying a house for investment will care just as much about the ZIP-code-level variability not encompassed in index returns. However, my own analysis of raw data kindly provided to me by Giacoletti shows that ZIP-code-level variation only adds about 0.2% in standard deviation for his dataset, thus not closing the gap in variance with stocks.

Despite this, I assume that housing is not less risky than the similarly-performing stock market in the baseline case, as to make my model's assessment of housing as an investment more conservative, and instead report the results for the Jordà et al. (2019) parameter values as an alternative scenario in Subsection 4.6. I set the total risk premium for housing, $R_{\mathrm{h},t}(1+y)-(1+r_{\mathrm{f}})$, to have the same mean and standard deviation as equities, fixing $y=5\%$ to correspond to the mean rental yield in the U.S. between 1900 and 2015, which was 5.04% according to the database by Jordà et al. (2019). Also, the correlation between the log returns to stocks and to housing in a given period is set to $\rho_{\mathrm{s,h}}=0.2$, also from my own calculations using the Jordà et al. (2019) database. This happens to match Yao and

11. My replication package for this paper contains code for performing this calculation and other similar ones, either using the Jordà et al. (2019) database or using the data from Giacoletti (2021) mentioned below.

Zhang (2005a, 207). I set the correlation in logs between stock returns and labor-income evolution to $\rho_{\mathrm{s,y}} = 0$, as estimated by Yao and Zhang (2005a, 207), and that between housing returns and labor-income evolution to $\rho_{\mathrm{h,y}} = 0.2$, following Yao and Zhang (2005a, 207) and Campbell and Cocco (2015, 1514). I also consider alternative values for these correlations in Subsection 4.6.

Meanwhile, the labor-income process $R_{s,t}$ is parametrized to have a mean of 3% and a standard deviation of 15% in the baseline case (Tan 2013, 65). Note that the mean labor-income growth over the life-cycle reflects not only aggregate productivity growth but also the fact that older workers earn higher incomes.

2.2 Borrowing conditions

The real risk-free rate r_{f} is set to 1.25%, as to match the average U.S. Treasury bill real rate between 1900 and 2015, and the inflation rate π is set to 3% to match U.S. CPI growth over that period, where both values are calculated from the database by Jordà et al. (2019). Meanwhile, the mortgage risk premium, $r_{\mathrm{FRM}} - r_{\mathrm{f}}$, is set to 2%. As explained by Cocco (2005, 547), real-life fixed-rate mortgages involve the acquisition of an option on future inflation and on future real interest rates, which is not captured in a model like this one where inflation and real interest rates are deterministic. He sets the mortgage risk premium to 2% rather than its historical average of 3% to adjust for this option value.

I set the minimum down-payment ratio α to 20%, since this is the minimum required for a conventional mortgage not to require private mortgage insurance. Investment-property mortgages generally require higher down-payment ratios than mortgages on owner-occupied properties: for instance, Fannie Mae requires a minimum 15% down payment for a mortgage to be eligible for delivery to it in the case of a one-unit property, as opposed to a minimum 3% down payment for an owner-occupied property (Fannie Mae 2022, 2).

The penalty for foreclosure, η, is set at 15% following Yao and Zhang (2005b, 14).[12]

12. But see also Campbell, Giglio, and Pathak (2011), who find a higher average value loss of 27% upon

Meanwhile, the baseline proportion of wealth to labor income granted upon default, ν, is set to 0.02 following Campbell and Cocco (2015); this calibration aims to make defaults very unlikely under the optimal strategy in the model, rather than to match empirical default rates, since I am not focused on strategic default.

I disable margin lending in the baseline case, since it is little used by the general public. However, I consider its effect in a separate analysis in Subsection 4.3. In that case, I set the parameters to match conditions currently offered by Interactive Brokers, which offers much better terms than other U.S. brokerages, as noted by Ayres and Nalebuff (2010, 167). The margin rate r_{m} is then set to equal the risk-free rate plus 1.5%, which matches the spread currently offered by Interactive Brokers on the first $100,000 (Interactive Brokers, n.d.). Due to the Federal Reserve's Regulation T, U.S. brokers, including Interactive Brokers, usually require a maintenance-margin ratio of 25% and an initial-margin ratio of 50%. Those are the leverage ratios that must be satisfied at all times and when the funds are first borrowed, respectively. Since I only model maintenance margin, I set this ratio g_m to 50% to ensure that both constraints are satisfied, and to follow Ayres and Nalebuff (2008, 2010) who recommend rebalancing to a 50% ratio.[13]

2.3 Household preferences and resources

Having calibrated the parameters pertaining to the external environment applying to all investors, I now discuss the parameters that could vary across investors, but which I calibrate for the median investor in the baseline case.

The savings rate is set to $\Gamma = 6\%$, as to correspond to the most common rate of 401(k) contributions in the U.S. (Miller 2022).

I calibrate risk aversion by solving the model while disallowing housing, setting the savings rate to 0 and setting $T = 1$. In such a setting, $\chi = 4$ leads to about 60% of investment in

foreclosure.

13. Willen and Kubler (2006) also only model maintenance margin, yet they use a maintenance-margin ratio of 25% to account for the fact that other products such as futures allow for greater initial leverage than Regulation T margin. Using a maintenance-margin ratio of 25% makes no qualitative difference in my results.

stocks, with the rest in cash, reflecting the share of stock investment near retirement in many mutual funds targeting a certain retirement date, such as Vanguard's. Therefore, I set it to that value for the baseline case. This coefficient is also in the theoretically plausible range of 2 to 10 (Tan 2013, 65).

By the nature of the CRRA utility function, choices in the model are mostly invariant to the level of Y_0, as long as $\frac{W_0}{Y_0}$ is fixed. There is only one exception: the minimum house size $\underline{H}_0$ induces a dependence on the level of Y_0, since it is a feature of the environment that is not inherently related to the household's income. In other words, for a given fraction of labor income in initial savings, the minimum house size will be less likely to bind for a household with a higher income. Since the median household income for those aged 25 to 34 years old was around \$75,000 in 2021 (Statista Research Department 2022), I set $Y_0 = \$75{,}000$. Furthermore, I set the minimum house value to \$225,000, so it is equal to three times annual income in the baseline case. By contrast, Cocco (2005) uses a much lower minimum house value of \$20,000 (in 1992 dollars), citing survey evidence that some reported house values are that low (545). I make a more conservative assumption since it might be harder to rent out properties that are particularly low-valued compared to others in the same geographical area. Furthermore, such low-valued houses might only exist in certain locations such as cities in decline, where housing value is lower than the cost of new construction (Glaeser and Gyourko 2005), making them impractical for most investors to administer since it is easier to oversee a rental property near one's residence (Levy 2022).

I vary the ratio $\frac{W_0}{Y_0}$ of initial wealth to initial labor income in most of my results, and otherwise specify if I am fixing it for the purpose of a particular analysis. I focus on values ranging between 10^{-1} and 10^{1}, representing the range between just over a month's income and ten years' income. For comparison, a study of Swedish households showed the 10th and 90th percentiles of the wealth-to-income ratio to be 0.87 and 9.25 (Calvet et al. 2022, 25).

Variable name	Description	Value	Source
T	Number of periods	30	See text
χ	Coefficient of relative risk aversion (CRRA)	4	See text
Γ	Savings rate	6%	Miller (2022)
r_{f}	Real risk-free rate	1.25%	Jordà et al. (2019)
$E(R_{\mathrm{s},t}) - (1 + r_{\mathrm{f}})$	Mean real risk premium of stocks	6.1%, $\forall t$	Dimson, Marsh, and Staunton (2006)
$\sqrt{\mathrm{Var}(R_{\mathrm{s},t})}$	Standard deviation of real returns to stocks	16.7%, $\forall t$	Dimson, Marsh, and Staunton (2006)
$E(R_{\mathrm{h},t}(1 + y)) - (1 + r_{\mathrm{f}})$	Mean real total risk premium of housing	6.1%, $\forall t$	Set to match stocks
$\sqrt{\mathrm{Var}(R_{\mathrm{h},t}(1 + y))}$	Standard deviation of real total returns to housing	16.7%, $\forall t$	Set to match stocks
$E(R_{\mathrm{y},t}) - 1$	Mean real permanent income growth	3%, $\forall t$	Tan (2013)
$\sqrt{\mathrm{Var}(R_{\mathrm{y},t})}$	Standard deviation of real permanent-income growth	15%, $\forall t$	Tan (2013)
$\rho_{\mathrm{s,h}}$	Correlation of logs of returns to stocks and housing	20%	Jordà et al. (2019)
$\rho_{\mathrm{s,y}}$	Correlation of logs of return to stocks and permanent income growth	0%	Yao and Zhang (2005a)
$\rho_{\mathrm{h,y}}$	Correlation of logs of return to housing and permanent income growth	20%	Yao and Zhang (2005a) and Campbell and Cocco (2015)
α	Down-payment constraint	20%	See text
y	Rental yield	5%	Jordà et al. (2019)
π	Inflation	3%	Jordà et al. (2019)
$r_{\mathrm{FRM}} - r_{\mathrm{f}}$	Mortgage premium	2%	Cocco (2005)
$r_{\mathrm{m}} - r_{\mathrm{f}}$	Margin premium	1.5%	Interactive Brokers (n.d.)
g_m	Maintenance-margin ratio	50%	See text
Y_0	Initial household income	\$75,000	Statista Research Department (2022)
$\underline{H}_0$	Minimum house value	\$225,000	See text
η	Foreclosure penalty	15%	Yao and Zhang (2005b)
ν	Fraction of labor income granted upon default	2%	Campbell and Cocco (2015)

Table 1: Summary of calibrated parameter values.
See text for details.

3 Solution method

This type of model cannot be solved analytically (Campbell and Cocco 2015, 1510).[14] Yet, because the problem is finite, a solution exists and can be found through backward induction. I solve the model using numerical methods that are standard in papers on life-cycle portfolio choice such as Yao and Zhang (2005b) and Cocco (2005). I represent the end-of-period state space as $\left(\log \frac{W}{Y}, \frac{H}{W}, \frac{M}{H}\right)$ as in Yao and Zhang (2005b, 34), the period-0 choice space as $\left(\frac{H}{W}, \frac{M}{H}, \frac{S}{L}\right)$, and the choice space in other periods as $\left(\frac{S}{L}\right)$, and discretize each of those with a grid. I approximate the density functions for the income and risky-asset distributions using Gaussian quadrature.

For each admissible combination of $\left(\log \frac{W}{Y}, \frac{H}{W}, \frac{M}{H}, \frac{S}{L}\right)$, I first obtain the expected utility of end-of-period terminal wealth in period $T - 1$. If $T = 1$ (if there is only one period), then the discrete choice of $\left(\frac{H}{W}, \frac{M}{H}, \frac{S}{L}\right)$ that yields the highest terminal wealth is the optimal policy for a given initial $\frac{W}{Y}$. The value function is then the expected utility evaluated at that optimal gridpoint.

If $T > 1$, then I instead find the optimal choice of $\frac{S}{L}$ for each gridpoint in the beginning-of-period state space $\left(\log \frac{W}{Y}, \frac{H}{W}, \frac{M}{H}\right)$, and also find the corresponding period value function that will be used in computations for earlier periods. For every $t < T - 1$, iterating backwards, I again calculate the end-of-period state for each admissible combination of beginning-of-period state and choice of $\frac{S}{L}$; this end-of-period state is the next period's beginning-of-period state, and thus the utility associated with this state can be found by evaluating the following period's value function. To evaluate the following period's value function for end-of-period state values that do not lie on the state-space grid, I use multilinear interpolation of $u^{-1}(v(\cdot))$, where u is the CRRA utility function and v is the value function for period $t+1$, and evaluate

14. It is worth mentioning a simpler case that can be. In my model, position adjustment is constrained since housing is only chosen in the first period; labor income is risky; and borrowing is limited and occurs at different rates than the risk-free saving rate. If none of those three things were true in the model, then it is well known that the labor income could be treated as an implicit holding in the risk-free asset of the discounted amount, and the optimal investment in the mean-variance-optimal combination of stocks and of housing could be given as a simple formula because of the assumptions of CRRA utility and of LogNormal returns (Campbell 2018, 309).

the utility of this interpolated value, as suggested by Carroll (2022). If $t > 0$, I compute the period value function by performing a grid search over $\frac{S}{L}$ for each point in the state space; if $t = 0$, I compute the overall value function by grid search over $\left(\frac{H}{W}, \frac{M}{H}, \frac{S}{L}\right)$ for each initial $\frac{W}{Y}$.

After the model has been solved, I separately simulate the model in the forwards direction, this time using 10,000 Monte Carlo simulations per initial wealth-to-income ratio rather than discretizing the distributions. Making use of the period choice functions that were previously solved for, I obtain the mean beginning-of-period choice and state variables for all periods. This allows me to show what happens in each period on average, as affected by the decisions in previous periods, rather than only conditionally on beginning-of-period state. A similar simulation approach is used in Cocco, Gomes, and Maenhout (2005).

I wrote the Python code to perform the above optimization and simulations from scratch, using NumPy for vectorized operations (Harris et al. 2020) and Econ-ARK for fast linear interpolation (Carroll et al. 2018).[15]

4 Results

4.1 Baseline calibration

I begin with the results for the baseline calibration discussed previously. Figure 1a shows that investors must have an initial wealth-to-income ratio of at least $\frac{W}{Y} = 10^{2/5} \approx 2.51$ for it to be optimal to buy housing. Specifically, investors with $\frac{W}{Y} \approx 2.51$ choose $\frac{H}{W} = 1.2$, $\frac{M}{W} = 0.36$ and $\frac{S}{W} = 0.16$. As seen in Figure 1b, the minimum down payment made is 70% (corresponding to a loan-to-value ratio of 30%).

Conversely, investors with low initial wealth purchase only stocks. For the lowest values of initial wealth, this is a direct result of the minimum-house-size constraint making it impossible to afford to buy a rental property even with the smallest possible down payment. However, the minimum-house-size constraint does not always have the effect of making baseline investors

15. The code is included in the replication package for this paper.

with relatively low wealth put all their money towards the minimal down payment even when this would be enough to afford a house. To see this, notice that an investor with $\frac{W}{Y} = 1$ could have purchased up to 5 times their wealth in housing with the maximum 20% down payment, whereas the minimum housing purchase was calibrated to be only 3 times annual income (and in this case income is equal to wealth). While this might seem surprising, the simple explanation for this is that the investor does not wish to take as much risk as would be necessary to buy housing at lower wealth-to-income ratios. This will be investigated more directly in the following subsection, where the minimum-house-size constraint will be relaxed.

When the investor does purchase housing, the share of their wealth invested in housing in the first period is seen to be significantly greater than that invested in stocks, by a factor of at least four. Moreover, this is the case even at high levels of wealth relative to labor income, as seen in Figure 1a. This may seem puzzling given that stocks and housing were calibrated to have identical risk and returns (with housing being slightly less attractive due to its positive correlation with income). However, it can be explained by the fact that the first period is the investor's only opportunity to purchase housing: Therefore, whereas in a one-period model without labor income and without leverage, the investor would have optimally invested the same share in both assets, in this case they are taking into account that the proportion of housing wealth will decline relative to liquid wealth over the course of the lifecycle, and accordingly purchasing more housing than stocks.

That explanation is supported by Figure 2, which shows mean beginning-of-period values of porefolio shares and of the wealth-to-income ratio, calculated using 10,000 simulations of the model solution across all periods. We see that while the share of housing ($\frac{H}{W}$) initially dominates stocks ($\frac{S}{W}$) by a large factor, the holding of stocks eventually overtakes housing after about 10 periods.

Figure 1b also shows that of the liquid wealth left over after purchasing housing with the chosen mortgage amount, the household invests 100% of it into stocks rather than cash in the first period, no matter the wealth-to-income ratio. As seen in Figure 2, the household does hold some cash in later periods, but generally cash investment remains lower than 40% even when there is little or no labor income left to earn. This is even though the risk-aversion coefficient was calibrated in Subsection 2.3 to generate only a 60% investment in stocks in the one-period model without labor income and without housing. The reason is simply that the presence of a second lightly correlated risky asset, housing, makes risky investment much more attractive than would be investment in stocks alone.

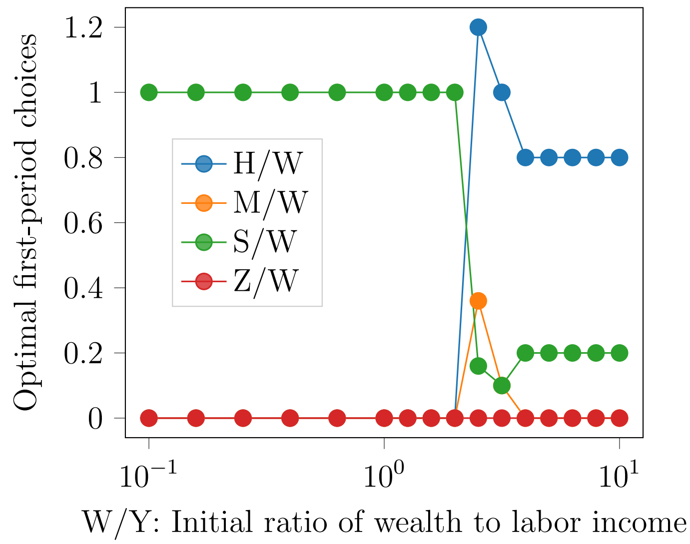

(a) Balance sheet, first period

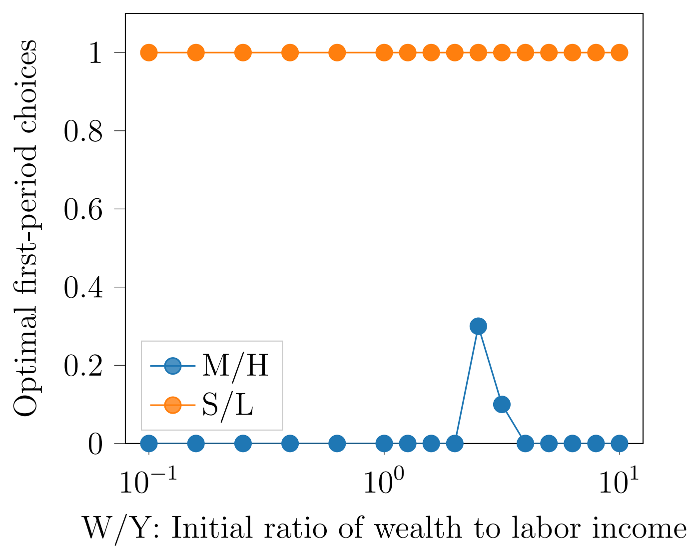

(b) LTV and stock-to-liquid-wealth ratios, first period

Figure 1: Model solution in the baseline case.
The variables in the first figure are the ratios of the value of housing, of the mortgage, of stocks and of cash with respect to wealth. The variables in the second figure are the ratios of the mortgage to the housing value (LTV ratio) and of stocks to liquid wealth, where liquid wealth is given by $L = W - H(1 - \frac{M}{H})$. I set $\frac{M}{H} = 0$ for the purposes of the graph when $H = 0$.

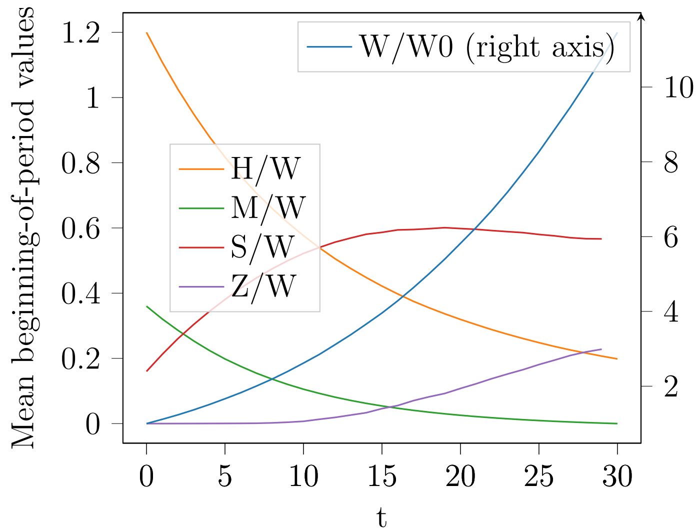

Figure 2: Dynamics of the model solution in the baseline case, with $\frac{W_0}{Y_0} = 2.51$.
The variables are the ratio of stocks to total wealth, the ratio of the value of housing to total wealth, the ratio of the mortgage to total wealth, the ratio of cash to total wealth, and the ratio of current wealth to initial wealth.
The values of $\frac{H}{W}$ and $\frac{M}{W}$ for $t = T$ are the end-of-period values for the preceding period; there are no corresponding values for $\frac{S}{W}$ and $\frac{Z}{W}$ since they are only chosen at the beginning of a period.

4.2 The effect of the minimum-house-size constraint

The first change in assumptions I consider is the removal of the minimum-house-size constraint ($\underline{H}_0 = 0$). This will clarify the extent to which the model solution in the previous subsection was affected by this constraint.

Figure 3a shows that without that constraint, households who have saved just over a month's labor income take very high leverage in housing, buying a house worth 3 times that wealth, with a large mortgage, which they were unable to do under the effect of the minimum-house-size constraint. Figure 3b shows that the household takes the maximum down payment given a wealth-to-income ratio of 0.1, and again puts all liquid wealth into stocks at all wealth levels.

Figure 4 compares the first-period solution without the minimum-house-size constraint with the previous subsection's baseline case. We see that at the level of wealth at which the investor first purchases housing in the constrained case, they actually purchase somewhat more housing than they would without a minimum-house-size constraint (the nonzero points on the $\frac{H}{W}$ curve are higher in the second graph than the first). Therefore, there is crowding-in from the minimum-house-size constraint: the household invests more in housing than they would like to in order to be able to invest in housing at all. This crowding-in from the minimum-house-size constraint happens only to a modest extent, though, since as discussed in the previous subsection, investors with $\frac{W}{Y} = 1$ could also have afforded to buy housing, but it was not optimal for them to do so.

Furthermore, as discussed in Subsection 2.3, since $\underline{H}_0$ is a feature of the external environment, unrelated to the household's income, then a high initial labor income Y_0 will make the constraint less likely to bind for a given wealth-to-income ratio $\frac{W_0}{Y_0}$. Table 2 reports the income required for the constraint not to bind for the optimal solution at each value of $\frac{W}{Y}$.

Note, however, that households with lower incomes than those may still be able to purchase another amount of housing that is not the unconstrained optimum, and some of those amounts may be preferable to no housing at all. Table 3 reports the lowest wealth-to-income ratio at which housing is purchased in the constrained optimum, for different values of household income. We see that at the high income level of $375,000, households first purchase housing at the low wealth-to-income ratio of 0.16. They nearly maximize their leverage, taking out a 20% down payment on a house worth 4 times their wealth. However, this is again due to a crowding-in effect, since a comparison with Table 2 shows that at this same ratio, the household would have purchased a house worth only 2 times their wealth if not for the minimum-house-size constraint.

To sum up, I find that the minimum-house-size constraint significantly affects the decision to purchase rental housing. When it binds, it drives the household either to not purchase housing when they would otherwise have done so, or to buy more housing than they would otherwise like. Furthermore, the greater the household's income is, the less they are affected by the constraint for given ratios of wealth to income. My results point to much higher leverage being taken by investors with high levels of household income: households with an income of \$375,000 buy housing when they have wealth-to-income ratios as low as 0.16, rather than 2.51 for the average household, and use much more leverage.

W/Y	H/W	Required Y (\$)
0.10	3.00	750,000.00
0.16	2.00	703,125.00
0.25	1.80	500,000.00
0.40	1.40	401,785.71
0.63	1.20	296,052.63
1.00	1.00	225,000.00
1.26	0.80	222,772.28
1.58	0.80	178,571.43
2.00	0.80	140,625.00
2.51	0.80	111,940.30
3.16	0.80	88,932.81
3.98	0.80	70,754.72
5.01	0.80	56,109.73
6.31	0.80	44,554.46
7.94	0.80	35,433.07
10.00	0.80	28,125.00

Table 2: Required income for minimum-house-size constraint not to bind for the optimal unconstrained first-period solution, by wealth-to-income ratio.
The first column is the same as the $\frac{H}{W}$ curve in Figure 3a. The last column is calculated by dividing the minimum house value by the product of $\frac{H}{W}$ and $\frac{W}{Y}$.

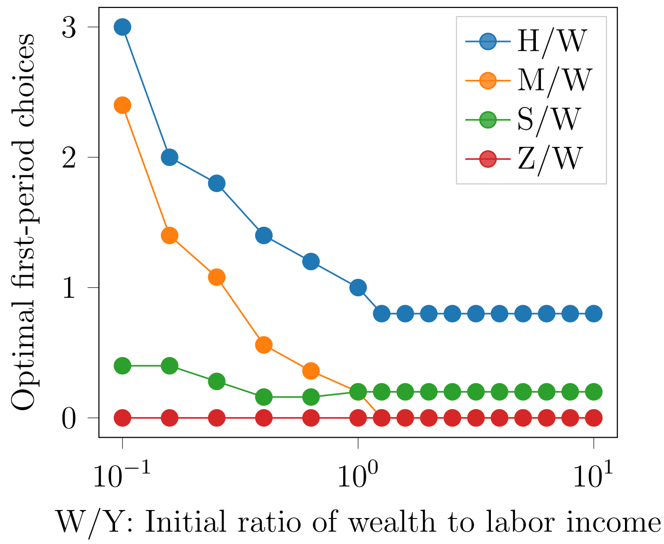

(a) Balance sheet, first period

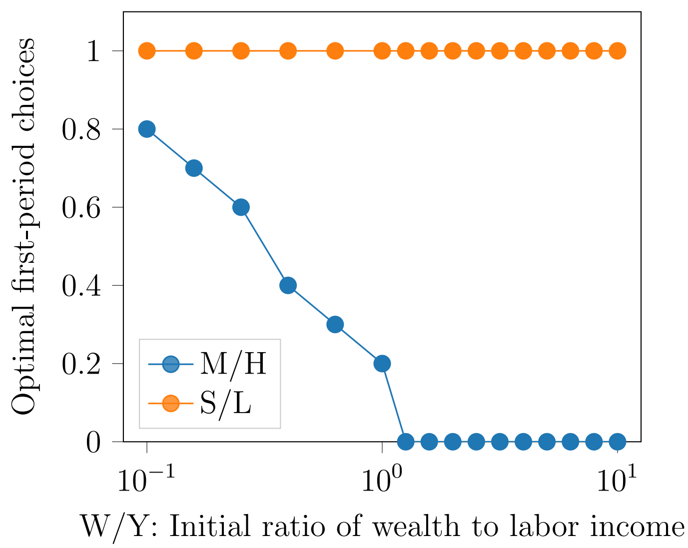

(b) LTV and stock-to-liquid-wealth ratios, first period

Figure 3: Model solution without minimum-house-size constraint.
The variables in the first figure are the ratios of the value of housing, of the mortgage, of stocks and of cash with respect to wealth. The variables in the second figure are the ratios of the mortgage to the housing value (LTV ratio) and of stocks to liquid wealth, where liquid wealth is given by $L = W - H(1 - \frac{M}{H})$.

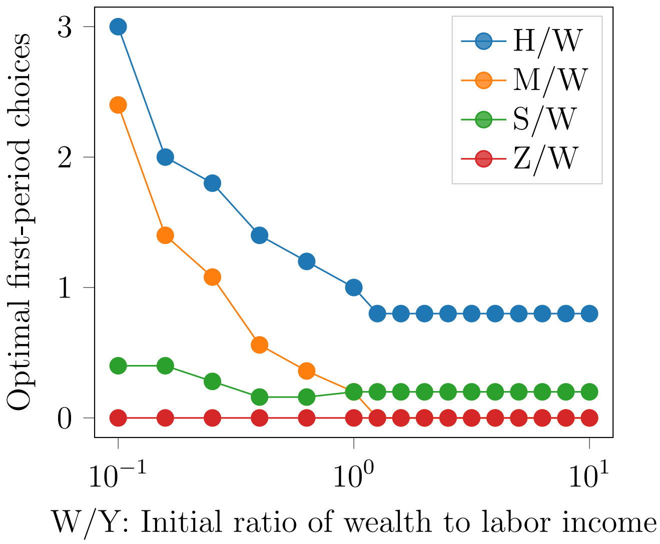

(a) Without minimum-house-size constraint

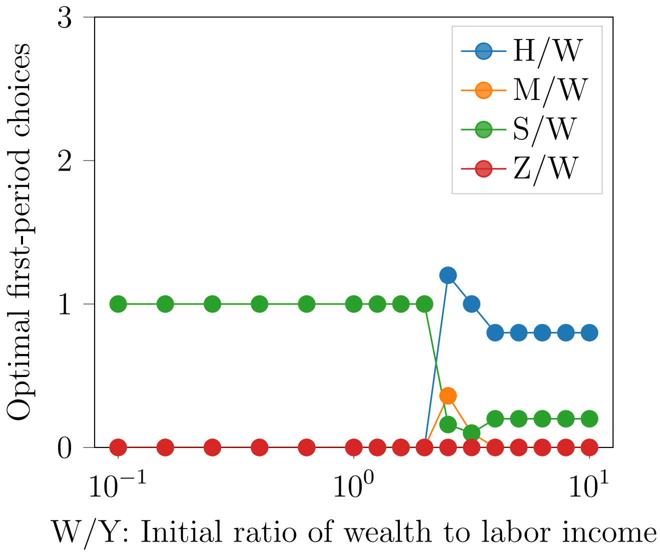

(b) With minimum-house-size constraint

Figure 4: Comparison of the model solutions for the first period without and with the minimum-house-size constraint.
Figure (a) is identical to that in Figure 3a, while the data in figure (b) is the same as that in Figure 1a. The solutions are graphed on the same scale to facilitate comparison.
The variables are the ratios of the value of housing, of the mortgage, of stocks and of cash with respect to wealth.

Household income	W/Y	H/W	M/W	S/W	Z/W	M/H
(1) \$37,500	5.01	1.20	0.24	0.04	0.00	0.20
(2) \$75,000	2.51	1.20	0.36	0.16	0.00	0.30
(3) \$150,000	1.00	1.60	0.64	0.04	0.00	0.40
(4) \$225,000	0.63	1.60	0.64	0.04	0.00	0.40
(5) \$375,000	0.16	4.00	3.20	0.20	0.00	0.80
(6) \$750,000	0.10	3.00	2.40	0.40	0.00	0.80

Table 3: Lowest initial wealth-to-income ratio for which housing is purchased, and first-period portfolio choice at that ratio, according to the level of household income.
The minimum-house-size constraint, $\underline{H}_0 = \$225{,}000$, is applied in computing these results. The case with a household income of \$75,000 corresponds to the baseline case, and thus the model solution matches Figure 1a in that case.

4.3 The effect of allowing margin investing

I now consider the effect of allowing the use of margin to invest in stocks, with a maintenance-margin requirement of 50%, as discussed in Subsection 2.2. This allows my model to nest that of Ayres and Nalebuff (2008, 2010).

I report the results in Figure 5a, this time restoring the minimum-house-size constraint. Low levels of initial wealth result in large leverage in stocks, with the maintenance-margin requirement binding under the optimal solution when $\frac{W_0}{Y_0} = 0.1$.

A comparison with Figure 5b shows that access to stock margin does not affect the choice of housing investment: it can be verified that the $\frac{H}{W}$ and $\frac{M}{W}$ curves are identical between the two graphs. This may be surprising, since we might have expected that the impossibility of leveraging stocks would lead to more investment in housing, which allows for mortgage leverage, but this turns out not to be the case. Put differently, the other differences between housing and stocks in the model, such as the fact that housing allows for leverage without a maintenance margin, and the fact that housing can only be purchased in the first period, have a greater influence on the choice of housing in the baseline calibration than whether margin leverage in stocks is available.

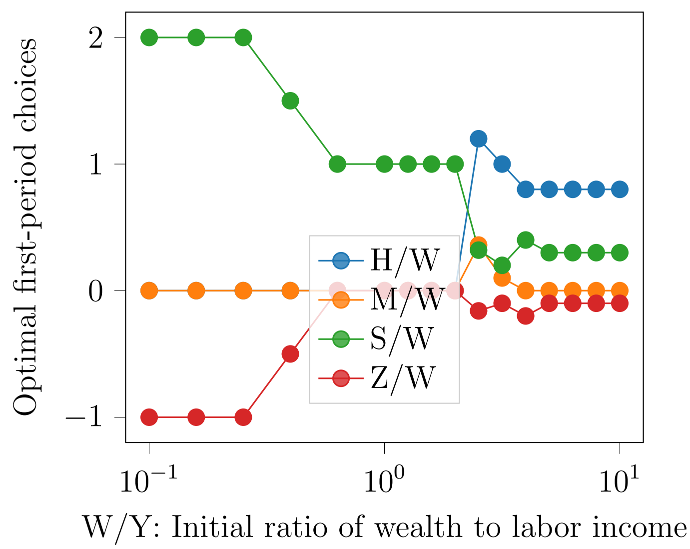

(a) With margin investing

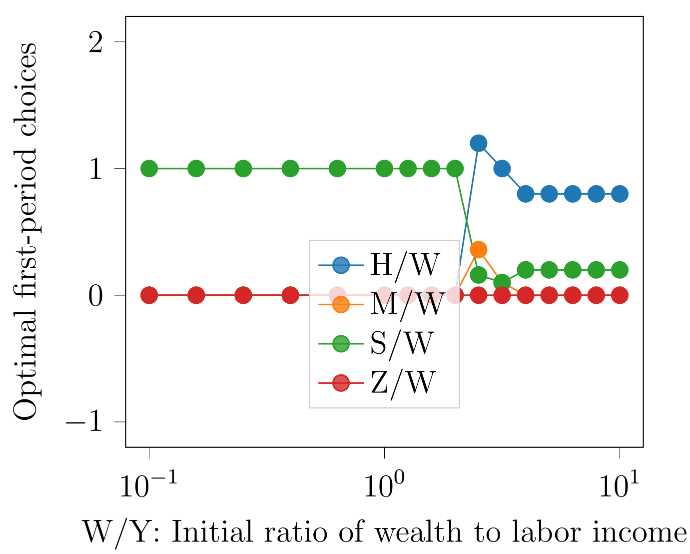

(b) Without margin investing

Figure 5: Model solution for the first period with and without margin investing.
The data in figure (b) is the same as that in Figure 1a, but is graphed on the same scale as the solution with margin investing to facilitate comparison.
The variables are the ratios of the value of housing, of the mortgage, of stocks and of cash with respect to wealth.

4.4 The effect of risk in labor income

A key motivation behind the present analysis was to explore future labor income as a basis for leveraged investing, and accordingly it is of great interest to analyze the effect of varying the riskiness of this future labor income. For this exercise, I return to the baseline case, with no margin investing and with the minimum-house-size constraint, and fix $\frac{W_0}{Y_0} = 2.51$, which was the level of wealth at which investors first purchased housing in Subsection 4.1.

The results are reported in Figure 6; recall that the baseline calibration set the standard deviation of labor income to 0.15. We see that use of mortgage leverage is somewhat sensitive to the riskiness of labor income, and more so than housing investment is. If the standard deviation is set to zero, then the household takes out a larger mortgage, and uses the greater remaining cash to buy stocks. If the standard deviation from the baseline calibration is increased to 0.45, then the household no longer purchases housing, and instead takes risk through a greater holding of stocks. One implication of this finding is that households with more-stable labor income, such as government employees, may tend to use more leverage in purchasing housing, and conversely, those with very risky future income may not purchase any at all.

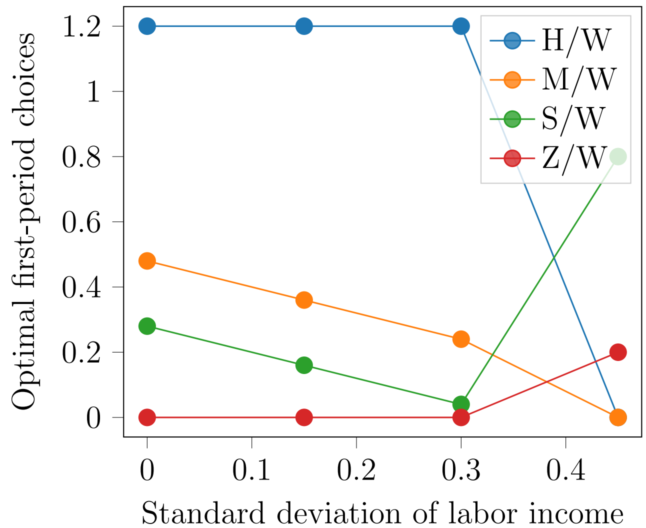

Figure 6: Model solution for the first period, with $\frac{W_0}{Y_0} = 2.51$ and varying the standard deviation of labor income.
The variables are the ratios of the value of housing, of the mortgage, of stocks and of cash with respect to wealth.

4.5 The effect of rental yield

Rental-property investors often strongly emphasize the importance of properties' cash flow, which in my model is equivalent to the rental income minus the mortgage payment. One book states that: "Here's a simple rule without exception: you never invest in cash-flow-negative properties" (Ross and Giannini 2017). Another book states that: "Break even on cash flow or lose money on cash flow, and you are on a path to financial ruin" (Turner 2015). From the economist's perspective, negative cash flow does not necessarily make a rental-property investment bad in expectation, both because the property is expected to appreciate and because the mortgage payment includes a principal payment towards paying off the mortgage. However, if the rental income is lower than the mortgage payment, such that cash flow is negative, then the household must use some of its labor income or other liquid savings to cover the mortgage payment. In this section, I examine how my model's results change as I vary rental yield, while holding the total risk and return fixed by varying capital gains accordingly. This allows me to assess whether the books' warnings about negative cash flow can be replicated in my model. Examining the effect of rental yield is also useful because it may differ across available properties: for instance, the share of returns that comes from rental yields rather than capital gains is larger for smaller cities in the U.S. relative to larger ones (Demers and Eisfeldt 2022).

Before performing any analyses, I note that even though a greater rental yield reduces the reliance on labor income or liquid wealth to cover the mortgage payment, higher rental yields are not unambiguously preferable in my model. This is because positive cash flow can be reinvested into stocks but not housing, so a higher rental yield decreases the household's ability to diversify between stocks and housing.

I report the optimal first-period choices, as well as the value function, for a range of rental yields in Figure 7. We see that the empirically calibrated yield used in the baseline calibration, which was $y = 5\%$, is in fact near what is optimal. Furthermore, there are substantial utility losses from low yields: In the case of zero yield, the household optimally chooses to forego purchasing housing entirely. To compare this to the books' recommendation, note that the break-even yield to have zero initial cash flow will depend on the size of the mortgage, but is otherwise pinned down by the mortgage-payment formula given in Subsection 1.1 once we have specified the risk-free rate, the mortgage risk premium and the inflation rate. For the baseline calibration, it can be computed that 5.97% is the yield that corresponds to zero initial cash flow with a 20% down payment, which is close to the 6% point that is optimal according to Figure 7b. Therefore, the books' recommendation to seek out non-negative cash flow is roughly accurate in my model.

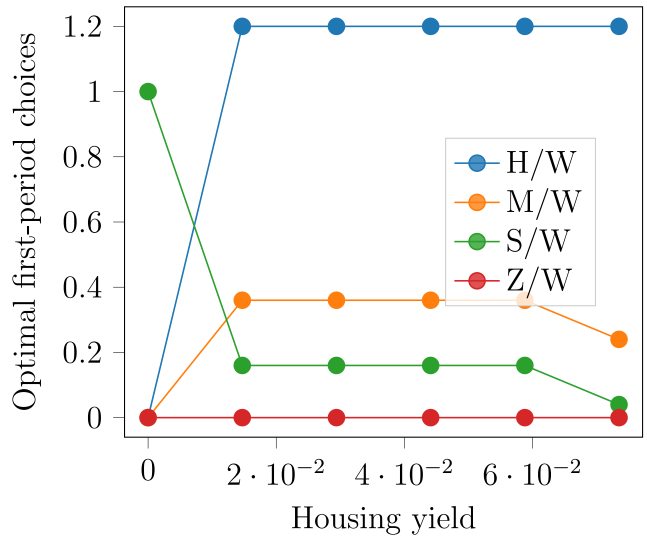

(a) Balance sheet, first period

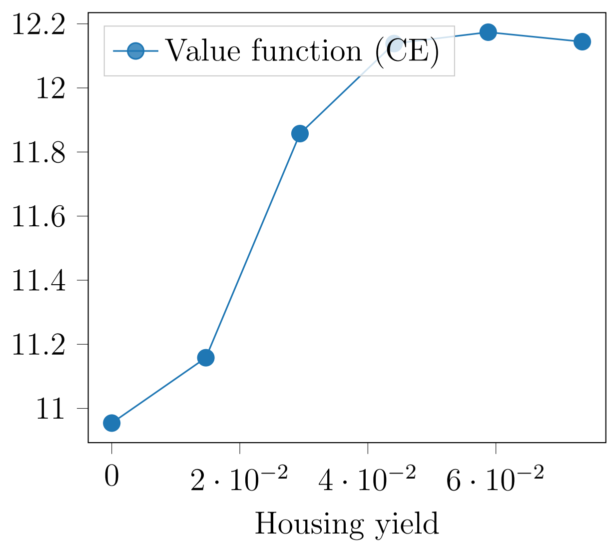

(b) Value function, first period

Figure 7: Model solution for the first period, with $\frac{W_0}{Y_0} = 2.51$ and varying the rental yield, for a fixed total risk premium.
The variables in the first graph are the ratios of the value of housing, of the mortgage, of stocks and of cash with respect to wealth. The variable in the second graph is the certainty-equivalent of the value function, and thus of final wealth, as a factor of the initial wealth.

4.6 The effect of other parameter assumptions

In the baseline case evaluated earlier, there was a role for a mortgaged rental property in the representative investor's portfolio, but this was a somewhat limited one since the household uses large down payments of at least 70% as discussed in Subsection 4.1. In practice, many investors use the least down payment required, i.e., 20%.[16] Given that my calibration was for an average investor, this observed behavior could be explained by the heterogeneity of various parameters among investors, whether they are labor-income riskiness and household income, as already discussed, or others such as risk aversion and savings rate. In this section, I explore the effects of varying those other parameters to assess for what attributes additional heterogeneity is most significant. I also test alternate values of parameters of the investment environment, such as returns and correlations between returns.

Table 4 reports model results under a range of different parameter assumptions, with Row 1 showing the baseline case. For conciseness, I report the lowest wealth-to-income ratio for which housing is purchased, and the optimal portfolio choice at that ratio. As we have seen, that ratio is usually also the ratio for which the chosen allocation to housing is the largest.

Rows 2 to 6 report results for some of the other scenarios considered in previous subsections. A noteworthy case is that of riskless labor income, which I had considered in Subsection 4.4 for a particular wealth-to-income ratio, whereas I now fix the wealth-to-income ratio in the way described above. Row 5 shows that with riskless labor income, households only need one year's labor income in savings in order to buy housing, and they take out a mortgage with the minimum 20% down payment at that wealth level. Therefore, assuming riskless labor income as in Ayres and Nalebuff (2008, 2010) is sufficient to generate the maximum level of leverage.

The first new situation I examine is that of a lower coefficient of relative risk aversion, which was set to 4 in the baseline case. A coefficient of relative risk aversion of 2 has frequently been used by other researchers, such as Campbell and Cocco (2015, 1511) and Ayres and

16. One book about rental-property investing mentions that: "Many of the examples in this book will use conventional 20% down payments as the norm for planning and strategizing for the future, as this approach is the most common" (Turner 2015).

Nalebuff (2008, 2010). Row 7 shows that this leads to more than double the maximum housing purchase ($\frac{H}{W} = 3$) than in the baseline case ($\frac{H}{W} = 1.2$), with a 20% down payment. Therefore, either a different assumption about the median investor's risk aversion or sufficient heterogeneity in risk aversion can generate much higher levels of housing purchasing and of mortgage leverage than the baseline case.

Another important parameter is the mortgage premium, which measures the spread between the risk-free rate at which cash can be deposited and the rate at which mortgage funds are borrowed, and which is set to 2% in the baseline case. In the special case of zero mortgage premium, it is always better to take out the maximum mortgage for a given choice of housing, since the additional funds will earn the same borrowing rate when deposited, and since a higher mortgage increases the option value of defaulting on the mortgage. This logic is confirmed by Row 8, which shows that the household takes out the maximum mortgage with a 20% down payment. While a lender is unlikely to ever offer a mortgage with zero risk premium, since the default risk must be priced, this case is conceptually important and is also the baseline case in Yao and Zhang (2005b).

More broadly, the mortgage premium might vary between investors depending on creditworthiness, or across time depending on market conditions. Rows 9 and 10 show that mortgage premia of 1% and 3% modestly affect housing and mortgage choice relative to the baseline case. While these are both one point away from our calibrated average, a decrease in the mortgage premium affects portfolio choice more than an increase, so there is an asymmetry in the impact of changes in the mortgage premium.

Next, I vary the correlation between stock returns and labor income, which was calibrated as 0 in the baseline case. Row 11 reports the results when the correlation between stocks and labor income is 0.2, which is the same as the correlation between housing and labor income in the baseline case. The model turns out not to be sensitive to this assumption: the results are exactly the same as in the baseline case.

Another parameter that is worthwhile to vary is the correlation between housing and labor income, which was set to 0.2 in the baseline case. A positive correlation between housing and labor income decreases the attractiveness of housing as an investment namely because it makes low rental income more likely to coincide with low labor income, thus increasing the odds of missing mortgage payments. Furthermore, it makes it such that housing is less attractive than stocks even if we disregard their leverage and liquidity differences, despite their having the same expected return and risk, because stocks, unlike housing, were calibrated as being uncorrelated with labor income. Row 11 shows that eliminating the correlation between housing and labor income modestly raises the maximum housing choice to 1.6 times wealth rather than 1.2 times wealth.

Another way in which housing is closely coupled with labor income in this model is through the rate at which labor income is saved. Indeed, the original motivation for this paper was to examine whether the presence of future labor income generates a demand for the leveraged purchase of rental properties. Row 13 confirms this intuition by showing that if there is no future labor income, the investor still buys housing for some levels of initial wealth, but requires more wealth to want to do so, and does not take out a mortgage at all. Furthermore, at the lowest wealth-to-income ratio for which housing is purchased, the investor without future savings does not invest in stocks in the first period. While this may seem counter-intuitive since there are no future savings to be invested into stocks, it is less surprising if we recall that the rental income can still be invested into stocks in later periods.

There is evidence that savings rates vary significantly both within (Venti and Wise 1998) and across (Dynan, Skinner, and Zeldes 2004) income brackets, thus making heterogeneity in the savings rate particularly important. For instance, Dynan, Skinner, and Zeldes estimate that households in the top income quintile save 12% of permanent income (420, 422), which is double our baseline value of 6%.[17] Row 14 shows that with a savings rate of 12%, households buy housing if they have a wealth-to-income ratio over 2.00, in contrast to 2.51 in the baseline

17. This is the figure for what Dynan, Skinner, and Zeldes call "active" savings, which are savings exclusive of capital gains. The rate of active savings is most analogous to my model's savings rate, which is with respect to labor income and does not encompass returns on the investment portfolio such as capital gains.

case. Furthermore, the minimum down-payment ratio drops to 60%. I also report results in Rows 15 and 16 for the higher savings rates of 20% and 30%. Row 16 that the very high anticipated savings rate of 30% can justify the maximal leverage observed among some real-estate investors: this assumption leads to total crowding out of stock investment in the first period, and a purchase of the biggest possible house with the smallest possible down payment.

Lastly, I consider the impact of using parameter estimates from Jordà et al. (2019), which exhibit a higher risk-adjusted return for housing than for stocks, and which I did not use in the baseline case in order to make the case for investing in rental properties more conservative. With the adjustments for idiosyncratic and ZIP-code-level housing risk that I suggested in Subsection 2.1, these correspond to a mean of 6.6% and a standard deviation of 23.84% for real stock returns, and a mean of 6.74% and a standard deviation of 14.16% for real housing returns. This is in contrast to a mean of 7.35% and a standard deviation of 16.7% for both assets in my baseline case. The results in Row 17 show that applying the results from Jordà et al. (2019) raises peak housing investment to 1.6 times wealth, with a 60% down payment.

	W/Y	H/W	M/W	S/W	Z/W	D.p. ratio
(1) Baseline case	2.51	1.20	0.36	0.16	0.00	0.70
(2) No minimum-house-size constraint	0.10	3.00	2.40	0.40	0.00	0.20
(3) Household income = $375,000	0.16	4.00	3.20	0.20	0.00	0.20
(4) Allowing stock margin	2.51	1.20	0.36	0.32	-0.16	0.70
(5) Riskless labor income	1.00	3.00	2.40	0.40	0.00	0.20
(6) Rental yield = expected total return	3.16	1.00	0.10	0.10	0.00	0.90
(7) CRRA = 2	1.00	3.00	2.40	0.40	0.00	0.20
(8) Mortgage premium = 0%	1.58	2.00	1.60	0.42	0.18	0.20
(9) Mortgage premium = 1%	2.00	1.60	0.80	0.20	0.00	0.50
(10) Mortgage premium = 3%	2.51	1.20	0.24	0.04	0.00	0.80
(11) Correlation between stocks and labor income = 0.2	2.51	1.20	0.36	0.16	0.00	0.70
(12) No correlation between housing and labor income	2.00	1.60	0.64	0.04	0.00	0.60
(13) Savings rate = 0% (no future labor income)	3.16	1.00	0.00	0.00	0.00	1.00
(14) Savings rate = 12%	2.00	1.60	0.64	0.04	0.00	0.60
(15) Savings rate = 20%	1.58	2.00	1.20	0.20	0.00	0.40
(16) Savings rate = 30%	0.63	5.00	4.00	0.00	0.00	0.20
(17) Jordà et al. (2019) return parameters	2.00	1.60	0.64	0.04	0.00	0.60

Table 4: Lowest initial wealth-to-income ratio for which housing is purchased, and first-period portfolio choice at that ratio, for different parameter settings.
Scenarios 1 through 6 were studied in previous subsections, and the results are also reported here for ease of comparison.
The last column indicates the down-payment ratio, which equals $1 - \frac{M/W}{H/W}$.

5 Discussion

In this section, I first discuss the implications of my parameter choices for the generalizability of my model. I also discuss the relationship between my model assumptions and the business reality of buying and operating a rental property. I then discuss the relationship between my model of a household's portfolio allocation and whether they separately buy a property to live in. Also, I discuss the possibility of taking out mortgages with down payments of less than 20%. I also discuss the implications of not modeling taxes. I then discuss the relevance of liquid real-estate investment funds to my problem. Finally, I discuss the broader implications of my findings in general equilibrium.

5.1 Robustness to parameter values and distributional assumptions

My model assumes that the returns to each asset are independent across time, as is common in the literature. If I had instead assumed mean reversion in returns, as in Campbell et al. (2001), this may have strengthened the incentive to use mortgage leverage, which does not have a maintenance-margin requirement, over stock leverage, which does. This is because under mean reversion, a liquidation during a downturn would result in missing out on the highest expected returns. One way this could be flexibly studied would be to repeat my analysis using historical returns over sequences of years for different cohorts (e.g., for the 30 years starting in 1950, in 1951, etc.), as in Ayres and Nalebuff (2008, 2010). However, this would come at the expense of a very low number of independent observations: with a 30-year investment horizon, we would only have about five non-overlapping observations across the entire 1870–2015 dataset from Jordà et al. (2019).

Furthermore, the attractiveness of rental investments might differ according to the particular locale. Amaral et al. (2021) suggest that smaller cities have higher total returns and higher risk than larger cities. While they report that the overall Sharpe ratios are similar, the ability to take on more risk for a given down payment means that investing in a smaller city may effectively relax the down-payment constraint if it is binding.

Buying and selling housing also involves transaction costs, such as appraisal, notary fees and broker fees, which I do not explicitly model. These would have been unlikely to significantly affect my results, since the horizon is long and since those transactions only occur once and cost less than a year's worth of returns (usually around 4%). Transaction costs are, in a sense, captured through the assumption that housing is only bought and sold once.

5.2 Modeling of rental properties

I begin this subsection on the limitations imposed by my model assumptions by talking about how some of them relate to the practical reality of how much skill and time investment is required to buy an individual property as an investment and of renting it out. The idiosyncratic risk of a house's return, which is the difference between its return and that of comparable houses in the same location, can be thought of as potentially originating from frictions at the time of house purchase and sale, or from uninsurable events that occur during the holding period. Alternately, what appears to be idiosyncratic risk from transaction frictions could in fact be attributable to unobserved skill in buying and selling houses. Other research has shown that there is variability in this regard: for example, real-estate agents in Singapore bought their own houses for 2.54% less on average than non-agents (Agarwal et al. 2019). My model implicitly assumes that the investor has average ex-ante skill in buying and selling houses. Therefore, an unsophisticated real-estate investor may achieve worse returns than the model calibration predicts, whereas they could easily achieve average returns in the stock market by holding an index fund.

The same consideration applies to rents: the investor is implicitly assumed to be of average skill in renting out and managing the property. Furthermore, in the sources I have consulted, the calibration of the ratio of net rents to gross rents factors in the labor cost of maintenance through depreciation, but not the labor cost or opportunity cost of time of managing the property (finding tenants, bookkeeping, hiring contractors, etc.), which is surely greater than that for an index-fund investment, and which is not explicitly modeled.

Another assumption of my model is that the investor may purchase only a single house, only in the first period, and sell it only in the last period. The structure of the model makes this unlikely that selling the house earlier would increase utility of final wealth for the median investor, unless the household also had the opportunity to buy another property.[18] In particular, the model could be extended by allowing houses to be bought and sold in any period, subject to transaction costs and to newly meeting the down-payment constraint, while maintaining for simplicity the assumption that the household owns only one house at once.[19] The model should then also directly account for the fact that annualized idiosyncratic risk is larger for shorter holding periods (Giacoletti 2021). The option to trade houses would also naturally pair with the option to perform a cash-out refinance of existing properties to convert either paid-down principal or capital gains into liquid wealth, as in Yao and Zhang (2005b, 11).[20] I omitted these elements in my model to limit computational complexity, since these changes would make the previous period's values of $\frac{H}{W}$ and $\frac{M}{H}$ into state variables for every period, as this would be needed to compute the transaction costs and to ensure the down-payment constraint isn't applied to existing non-refinanced mortgages.

This type of model also implicitly treats housing as limited by the availability of fixed factors such as land, since investment and depreciation are exogenous and deemed to be included in the net yields. Alternative models make investment and depreciation endogenous, and thus treat the housing stock as adjustable (Chen, Michaux, and Roussanov 2020, 333).

18. However, the option to sell early, even without the option to buy another property, could still be an improvement for some investors even if not for the median investor, since it would allow households with particularly low realizations of income to lower their holding of risky assets.

19. The assumption that the investor owns only a single house at once seems reasonable as a suggested departure from investing only in the stock market, since administrative burden would likely scale with the number of houses. However, it is worth noting that investing a given amount of capital in several properties by purchasing them in partnership with other people could be a way to decrease idiosyncratic risk and thus increase risk-adjusted returns.

20. This should not be confused with refinancing that is for the purpose of obtaining a lower interest rate, as in Campbell and Cocco (2003). Cash-out refinancing could be added to the model while maintaining my assumption of a time-invariant real risk-free rate.

5.3 Interaction with owner-occupied housing

As mentioned in the introduction, my model is framed independently of the household's decision of whether to purchase owner-occupied housing. If they were to purchase owner-occupied housing, this would be assumed to be financed with the portion of labor income that is not encompassed in the savings rate, or with funds excluded from the initial investable wealth. Therefore, my model centers on the choice of the "investment portfolio", that is to say, of assets that do not have any consumption features.

An immediate question is whether the two problems are in fact separable, even conditionally on an exogenous savings rate. In other words, does owning one's primary residence affect the optimal proportions of stocks, rental housing, rental-property mortgages and cash? There is a particular context in which the answer is likely to be "no", though it may depend on the specific utility formulation. This is when an infinitely lived household's ownership of their primary residence perfectly hedges their future need for housing services, and the residence can thus be seen as a risk-free asset providing those particular services (Barras and Betermier 2020). This is sometimes summarized as saying that "everyone is born short a house". Once that short is covered with owner-occupied housing, investing in a rental property, even one whose value is correlated with the primary residence, need not be riskier than investing in stocks.[21]

A few factors may make that hedging explanation more or less applicable. In practice, households are not infinitely lived, and if the primary residence were to be sold at the same retirement age T as the rental housing in the model, then this would surely affect portfolio choice. In this case, it would be preferable for the rental property's value to be less correlated with that of the primary residence. However, the assumption of the household's being infinitely lived can be rationalized in a few ways. One justification is that households will also need housing services in retirement, as will their bequest beneficiaries. For instance,

21. For an alternative perspective, see Lockwood et al. (2022), who argue that renting may be preferable to owning for many households from a hedging perspective, because rents positively covary with local labor-income fluctuations and thus hedge them.

Yao and Zhang (2005b), whose agents optimize over period consumption and bequests rather than the terminal wealth used in my model, incorporate housing-services consumption in the bequest beneficiaries' utility and thus extend the hedging motive for a primary residence beyond the household's lifetime, though the extent of this depends on the strength of the bequest motive.

Furthermore, a future extension of this model could endogenize consumption of both perishable goods and housing services during the working years, and thereby endogenize the choice of both the savings rate and the choice to buy owner-occupied housing. Even if primary-residence ownership were separable from investment-portfolio choice, as discussed above, the choice to buy owner-occupied housing would likely still interact with the savings rate. For instance, if the hedging motive were very strong, the household might not save much until they are able to purchase a primary residence. This would be best examined by combining my model with one like that of Yao and Zhang (2005b), who incorporate the buy-or-rent decision to study such effects.

5.4 Mortgages with low or no down payment

In my analysis, I assume that the minimum down payment is 20%, as is usually required for conventional mortgages. However, there are U.S. mortgage products allowing for even greater leverage, which I briefly discuss in this section. These seem to be popular among real-estate investors, as judged by the large number of sales of books such as *The Book on Investing in Real Estate with No (and Low) Money Down* (Turner 2020).

A prime example of such a mortgage product is that offered by the Federal Housing Administration (FHA), which allows for down payments of only 3%. These mortgages can be taken out for properties of up to four units, with the caveat that the owner must live in one of the units. Projected net rental income may be included in the borrower's income to qualify for the mortgage (U.S. Department of Housing and Urban Development 2011, 2.B.8). These loans typically come at a higher cost, namely because of the cost of mortgage insurance which cannot be cancelled for a FHA mortgage with a down payment of less than 10% unless the property is refinanced (Dehan 2023).

The existence of such mortgages further bolsters the case for modeling rental-property investment as part of life-cycle portfolio choice, as I do in this paper. This is because a down payment of 3% brings the leverage achievable with rental-property investment even further above that of regular margin investing, with its 25% maintenance-margin and 50% initial-margin requirements.

My model could be extended to consider the trade-off between higher interest rates or origination costs and lower down payments. This would allow investors with low initial wealth to invest in housing, whereas in my model results the investors with the least savings only have access to the stock market. Future research could also investigate whether multi-unit FHA mortgage holders are more likely to be young and have low savings than holders of conventional rental-property mortgages, as the logic of my model would suggest.

5.5 Taxes

Another limitation of my analysis is that I ignored the impact of taxes, as in Cocco (2005, 564). Capital-gain taxes in particular would require keeping track of the basis of each investment as an additional state variable, adding computational complexity. I now discuss how modeling taxes might affect the results.

One way that taxes could affect portfolio choice in this setting is by reducing both the return and risk of risky assets, thus inducing the investor to use more leverage in both stocks and housing than predicted by the model to compensate for this dampening of risk (Sandmo 1985, 295).

Differential tax treatment of stocks and housing could also affect investment decisions. For instance, net rental income is taxed as ordinary income at the federal level in the U.S., whereas long-term capital gains on either stocks and housing, as well as qualified dividends on stocks, are taxed at a lower rate, thus favoring stocks over housing. By contrast, interest deductibility can be more favorable for housing than for stocks: Mortgage interest is deductible as a business expense even when taking the standard deduction, while margin interest is not deductible when taking the standard deduction. Furthermore, so-called 1031 exchanges allow for the deferral of capital gains tax when selling a real-estate property and purchasing another, which is not possible with stocks. However, this is not directly relevant in this model since only one property is ever purchased. Finally, stocks and other liquid securities can be held in tax-advantaged accounts such as Individual Retirement Accounts or 401(k) retirement plans, whereas individual rental properties usually cannot be.

5.6 REITs

I now briefly discuss real-estate investment trusts (REITs). REITs are, in most cases, companies that own and manage real estate, and many have shares available for purchase on stock exchanges. They thus avoid the illiquidity and minimum purchase size involved in owning individual properties. Furthermore, they may have higher returns through economies of scale, specialization in property management, and market power in renting. They also minimize idiosyncratic risk through the holding of many properties across different locales. Some studies suggest that REITs have similar return profiles to individual properties (Hoesli and Oikarinen 2021).

In this paper, I have assumed that the only type of housing that can be purchased is an individual property. As mentioned previously, this is because my analysis centers on the fact that purchasing an individual property is necessary to obtain a mortgage, which allows for high leverage that is not subject to maintenance margin. While many REITs are internally leveraged, I am not aware of any for which the leveraged returns are comparable to those of a rental property with a 20% down payment.

I will however mention a possible extension to this paper. REITs could be incorporated as an additional asset that could be traded in each period and be purchased with margin leverage, similarly to stocks, but that would have a very high correlation with the rental property and a low correlation with stocks. In the cases with relatively low initial wealth that are of interest in the paper, it is likely that the addition of REITs would make little difference in the first period, since the similar individual-property asset would allow for more leverage. Importantly, though, a scenario in which the only allowed assets are stocks and REITs would be useful for comparison to the scenario where all assets are investable. Such a comparison would isolate the benefit provided to the investor by mortgage leverage, whereas in the current setup the benefits of a mortgage can only be evaluated in combination with the diversification benefit of allowing housing investment.

5.7 General-equilibrium implications

I now briefly discuss the broader implications of my findings for financial and real-estate markets. I have shown that young investors allocate a greater share of their portfolio to individual rental properties than they would if it weren't for the particular advantages of mortgages. Households are strictly worse off than if they could purchase REITs with a mortgage, assuming that REITs can replicate the performance of rental properties while diversifying away their idiosyncratic risk.[22] Furthermore, this tendency for investors to directly hold properties to obtain higher leverage likely translates into a higher proportion of properties being held by households than if mortgages could be used to purchase REITs. This helps explain the high proportion of rental housing that is supplied by households. There are many policy implications: for instance, one study showed that non-professional landlords select tenants less fairly than professional ones do, but are not as quick to evict those who fall behind on rent (Decker 2023). Another study showed that individual landlords respond with much less avoidance to the introduction of rent control than do corporate landlords (Diamond, McQuade, and Qian 2019).

22. Public REITs' returns may in fact be closer to those of stocks than to housing, in which case they could be considered to be part of the stock index in my model and would not warrant particular analysis. Here, I am considering how an idealized REIT would relate to my model, without delving into the empirical literature on REIT returns in detail. If there are no products available that provide even unleveraged exposure to the real-estate market without idiosyncratic risk, then there is even more reason for households to invest in individual properties, without needing to invoke leverage as a justification.

Furthermore, investors are worse off than if they could also use mortgages to buy stocks, rather than only housing. The inability to purchase stocks with a mortgage (i.e., a collateralized loan without maintenance margin) could cause real estate to tend to have a lower required risk-adjusted return than stocks in general equilibrium. Real estate may then be less affordable for prospective owner-occupiers than if mortgaged stock investment were available.

Mortgaged stock investment is in fact available in Australia: National Australia Bank offers an "Equity Builder" loan that seems to function exactly like a mortgage, with down-payment ratios as low as 25% (National Australia Bank 2021). It is secured by an investment in a stock index fund, and is not subject to margin calls, which is advertised as an advantage of the product, in line with the premise of this paper that a lack of maintenance margin is a sought-after feature.

It is worth reasoning about why such products are not more widely available. Even if stocks are riskier than real estate in absolute terms, the lender can simply require a higher down-payment ratio to make the probability of negative equity the same for both asset classes. Furthermore, the lender could require (and directly enforce through integration with a brokerage account) that the borrowed funds only be invested in approved index funds rather than riskier individual stocks.

Campbell (2006, 1586–1587) also noted a lack of innovation in product offerings, as well as variability in the uptake of product types across countries, in the market for mortgages on owner-occupied houses. He argued that the presence of unsophisticated households may dull the incentive to offer a new product to sophisticated households if the existing product involves a cross-subsidy from unsophisticated to sophisticated households. The argument, which concerns adjustable-rate mortgages as the new (or lesser-known) product and fixed-rate mortgages as the existing product, is as follows. In the United States, the predominant mortgage type for owner-occupied houses is a fixed-rate mortgage, which households benefit from canceling and refinancing at a lower rate if the prevailing interest rate is lower than the initially agreed rate at some point during the mortgage term. However, because a substantial share of households are not sophisticated or attentive enough to optimally refinance, the cancellation risk that is priced into the cost of the fixed-rate mortgage will be lower than the actuarially fair price (as construed under the assumption that all households were acting optimally) when mortgage supply is competitive. Therefore, sophisticated consumers of a fixed-rate mortgage receive a cross-subsidy from unsophisticated ones when they refinance optimally. At the same time, we can assume that the unsophisticated households who do not optimally refinance are inherently less likely than sophisticated households to use a new product because of status-quo bias. Campbell argues that there may then not be a market for the new product, since the sophisticated households who might be willing to purchase it would have to pay the actuarially fair price for it, rather than benefiting from a cross-subsidy from unsophisticated households as they did in the market for the existing product.

I believe that the same logic could be extended to the case where a mortgage on stocks is the new product and a mortgage on a rental property is the existing product. Sophisticated households might benefit from a cross-subsidy in mortgages for rental properties because of more-optimal refinancing or strategic default behavior than unsophisticated households. Under these assumptions, a mortgaged-stock-investment product may not attract enough customers to be commercially viable if the sophisticated households benefit more from being pooled with the unsophisticated real-estate purchasers than from purchasing the new product at the actuarially fair price.

Government intervention is another factor that could possibly favor the use of mortgage borrowing to buy real estate over a new mortgage product that allowed buying stock indices. In the language of my model, the mortgage risk premium may be lower than it would be if not for implicit government subsidies, for instance through government-sponsored enterprises (GSEs) such as Fannie Mae.

In short, my paper's partial-equilibrium finding that investors may purchase rental properties because of the greater leverage provided by mortgages naturally raises questions about why comparable leverage isn't usually available for purchasing diversified real-estate portfolios or stock indices. Future research might explore the consequences of this for owner-occupier purchasers, as well as the market dynamics leading such financing to only be supplied for individual properties.

6 Conclusion

In this paper, I developed a model of portfolio choice with uncertain labor income, investment in rental properties and in stocks, and borrowing through mortgages and stock margin. My model captures the fact that margin investment is subject to a maintenance-margin constraint and thus to margin calls, while mortgages are not. I calibrated the model based on the empirical evidence, carefully treating issues such as the idiosyncratic risk of individual properties, and solved the model numerically. My model showed that households with low initial wealth don't purchase any housing, whereas those with medium values of wealth relative to income purchase the most housing and use leverage, with their down payment crowding out investment in stocks.

My model highlights that the constraint of a minimum house size is central to the portfolio-choice decision, and is more likely to bind for investors with a lower income for a given wealth-to-income ratio. When it binds, the constraint either makes the household purchase no housing, or it leads them to buy more housing than they otherwise would, which I refer to as crowding-in. Furthermore, higher savings rates, which have previously been shown to be associated with higher household incomes, also make nonzero housing investment optimal for lower wealth-to-income ratios. Therefore, while my model generates a modest demand for housing investment and mortgage leverage for the median investor with future labor income, it predicts that better-off households use higher leverage for given ratios of wealth to income, mainly because of the minimum-house-size channel and also because of the savings-rate channel. This is in line with the literature that suggests that the rich earn a higher return on their investments (Bach, Calvet, and Sodini 2020). By the logic of time diversification, this likely means that rich households would earn a higher risk-adjusted return on their investments as well, thanks to the mortgage leverage available with rental properties.

Whereas previous models of life-cycle portfolio choice studied young investors' use of leveraged stock investment to diversify across time, my analysis shows that mortgaged rental-property investment can achieve time diversification even more effectively, which aligns both with its much greater popularity over margin investing, and with the popular perception that rental properties are a better investment than stocks. Of course, one could make a simpler behavioral argument that people are more familiar with real estate than with stocks, and would rather invest in what they fully understand, even if investing in stock indices requires little specific knowledge. Future research could aim to disentangle the factors that contribute to the differential appeal of rental properties and stocks, including access to leverage, psychological biases, macroeconomic trends, and life-cycle stage.

Bibliography

Note that citations with a page number preceded by "IA" indicate a page number within the Internet Appendix or Online Appendix.

Agarwal, Sumit, Jia He, Tien Foo Sing, and Changcheng Song. 2019. "Do Real Estate Agents Have Information Advantages in Housing Markets?" *Journal of Financial Economics* 134, no. 3 (December 1, 2019): 715–735. https://doi.org/10.1016/j.jfineco.2019.05.008.

Amaral, Francisco, Martin Dohmen, Sebastian Kohl, and Moritz Schularick. 2021. "Superstar Returns." SSRN Scholarly Paper 3992670. December 27, 2021. https://papers.ssrn.com/abstract=3992670.

Ayres, Ian, and Barry Nalebuff. 2008. "Life-Cycle Investing and Leverage: Buying Stock on Margin Can Reduce Retirement Risk." SSRN Scholarly Paper 1149340. June 1, 2008. https://papers.ssrn.com/abstract=1149340.

Ayres, Ian, and Barry Nalebuff. 2010. *Lifecycle Investing: A New, Safe, and Audacious Way to Improve the Performance of Your Retirement Portfolio.* New York, NY: Basic Books. ISBN: 978-0-465-01829-1.

Bach, Laurent, Laurent E. Calvet, and Paolo Sodini. 2020. "Rich Pickings? Risk, Return, and Skill in Household Wealth." *American Economic Review* 110, no. 9 (September): 2703–2747. https://doi.org/10.1257/aer.20170666.

Barras, Laurent, and Sebastien Betermier. 2020. "Why Do Homeowners Invest the Bulk of Their Wealth in Their Home?" SSRN Scholarly Paper 2776350. March 18, 2020. https://papers.ssrn.com/abstract=2776350.

Board of Governors of the Federal Reserve System. 2019. "Survey of Consumer Finances." Accessed February 5, 2023. https://www.federalreserve.gov/econres/aboutscf.htm.

Calvet, Laurent E., John Y. Campbell, Francisco Gomes, and Paolo Sodini. 2022. "The Cross-Section of Household Preferences." SSRN Scholarly Paper 4262670. October 27, 2022. https://papers.ssrn.com/abstract=4262670.

Campbell, John Y, Stefano Giglio, and Parag Pathak. 2011. "Forced Sales and House Prices." *American Economic Review* 101, no. 5 (August 1, 2011): 2108–2131. https://doi.org/10.1257/aer.101.5.2108.

Campbell, John Y. 2006. "Household Finance." *The Journal of Finance* 61, no. 4 (August): 1553–1604. https://doi.org/10.1111/j.1540-6261.2006.00883.x.

Campbell, John Y. 2018. *Financial Decisions and Markets: A Course in Asset Pricing.* Princeton: Princeton University Press. ISBN: 978-0-691-16080-1.

Campbell, John Y., and João F. Cocco. 2003. "Household Risk Management and Optimal Mortgage Choice." *The Quarterly Journal of Economics* 118, no. 4 (November 1, 2003): 1449–1494. https://doi.org/10.1162/003355303322552847.

Campbell, John Y., and João F. Cocco. 2015. "A Model of Mortgage Default." *The Journal of Finance* 70, no. 4 (August): 1495–1554. https://doi.org/10.1111/jofi.12252.

Campbell, John Y., João F. Cocco, Francisco Gomes, Pascal J. Maenhout, and Luis M. Viceira. 2001. "Stock Market Mean Reversion and the Optimal Equity Allocation of a Long-Lived Investor." *Review of Finance* 5, no. 3 (December 1, 2001): 269–292. https://doi.org/10.1023/A:1013860504885.

Carroll, Christopher D. 2022. "Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems," December 9, 2022. Accessed December 20, 2022. https://llorracc.github.io/SolvingMicroDSOPs/.

Carroll, Christopher D., Alexander M. Kaufman, Jacqueline L. Kazil, Nathan M. Palmer, and Matthew N. White. 2018. "The Econ-ARK and HARK: Open Source Tools for Computational Economics." In *Proceedings of the 17th Python in Science Conference,* edited by Fatih Akici, David Lippa, Dillon Niederhut, and M Pacer, 25–30. https://doi.org/10.25080/Majora-4af1f417-004.

Chen, Hui, Michael Michaux, and Nikolai Roussanov. 2020. "Houses as ATMs: Mortgage Refinancing and Macroeconomic Uncertainty." *The Journal of Finance* 75, no. 1 (February): 323–375. https://doi.org/10.1111/jofi.12842.

Cocco, João F. 2005. "Portfolio Choice in the Presence of Housing." *Review of Financial Studies* 18 (2): 535–567. https://doi.org/10.1093/rfs/hhi006.

Cocco, João 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. https://doi.org/10.1093/rfs/hhi017.

Decker, Nathaniel. 2023. "How Landlords of Small Rental Properties Decide Who Gets Housed and Who Gets Evicted." *Urban Affairs Review* 59, no. 1 (January): 170–199. https://doi.org/10.1177/10780874211041513.

Dehan, Andrew. 2023. "FHA Mortgage Insurance Removal: A Guide." RocketMortgage, January 10, 2023. Accessed February 22, 2023. https://www.rocketmortgage.com/learn/fha-mortgage-insurance-removal.

Demers, Andrew, and Andrea L. Eisfeldt. 2022. "Total Returns to Single-Family Rentals." *Real Estate Economics* 50, no. 1 (March): 7–32. https://doi.org/10.1111/1540-6229.12353.

Diamond, Rebecca, Tim McQuade, and Franklin Qian. 2019. "Who Pays for Rent Control? Heterogeneous Landlord Response to San Francisco's Rent Control Expansion." *AEA Papers and Proceedings* 109 (May 1, 2019): 377–380. https://doi.org/10.1257/pandp.20191021.

Dimson, Elroy, Paul Marsh, and Mike Staunton. 2006. "The Worldwide Equity Premium: A Smaller Puzzle." SSRN Scholarly Paper 891620. April 7, 2006. https://papers.ssrn.com/abstract=891620.

Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes. 2004. "Do the Rich Save More?" *Journal of Political Economy* 112, no. 2 (April): 397–444. https://doi.org/10.1086/381475.

Eichholtz, Piet, Matthijs Korevaar, Thies Lindenthal, and Ronan Tallec. 2021. "The Total Return and Risk to Residential Real Estate." *The Review of Financial Studies* 34, no. 8 (August 1, 2021): 3608–3646. https://doi.org/10.1093/rfs/hhab042.

Fannie Mae. 2022. "Eligibility Matrix," October 5, 2022. https://singlefamily.fanniemae.com/media/20786/display.

Giacoletti, Marco. 2021. "Idiosyncratic Risk in Housing Markets." *The Review of Financial Studies* 34, no. 8 (July 16, 2021): 3695–3741. https://doi.org/10.1093/rfs/hhab033.

Glaeser, Edward L., and Joseph Gyourko. 2005. "Urban Decline and Durable Housing." *Journal of Political Economy* 113, no. 2 (April): 345–375. https://doi.org/10.1086/427465.

Harris, Charles R., K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, et al. 2020. "Array Programming with NumPy." *Nature* 585, no. 7825 (September 17, 2020): 357–362. https://doi.org/10.1038/s41586-020-2649-2.

Haughwout, Andrew. 2021. "Do People View Housing as a Good Investment and Why?" Liberty Street Economics, April 5, 2021. Accessed November 20, 2022. https://libertystreeteconomics.newyorkfed.org/2021/04/do-people-view-housing-as-a-good-investment-and-why/.

Hoesli, Martin, and Elias Oikarinen. 2021. "Does Listed Real Estate Behave like Direct Real Estate? Updated and Broader Evidence." *Applied Economics* 53, no. 26 (June 3, 2021): 3023–3042. https://doi.org/10.1080/00036846.2020.1870921.

Interactive Brokers. n.d. "Interactive Brokers Margin Rates." Accessed February 6, 2023. https://www.interactivebrokers.com/en/?f=46376.

Internal Revenue Service. 2021. "Statistics of Income—2019 Individual Income Tax Returns." Publication 1304. December. https://www.irs.gov/pub/irs-pdf/p1304.pdf.

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

Jordà, Òscar, Moritz Schularick, and Alan Taylor. 2019. "The Total Risk Premium Puzzle." Working Paper 25653. Cambridge, MA: National Bureau of Economic Research, March. https://doi.org/10.3386/w25653.

Khan, Mikael, and Yang Xu. 2022. "Housing Demand in Canada: A Novel Approach to Classifying Mortgaged Homebuyers." Staff Analytical Note 2022-1. Bank of Canada, January 13, 2022. https://doi.org/10.34989/san-2022-1.

Levy, Antoine. 2022. "Housing Policy with Home-Biased Landlords: Evidence from French Rental Markets." Mimeo. June 30, 2022. Accessed February 23, 2023. https://sites.google.com/view/antoine-levy/home/research.

Lockwood, Lee M., Lorenz Kueng, Pinchuan Ong, and Scott R. Baker. 2022. "Renting Hedges Wage Risk." Mimeo. August 12, 2022. Accessed February 21, 2023. https://bit.ly/36IPr0s.

Miller, Stephen. 2022. "Most Common 401(k) Savings Default Rate Rises to 6% of Pay." SHRM, January 19, 2022. Accessed January 28, 2023. https://www.shrm.org/resourcesandtools/hr-topics/benefits/pages/most-common-401k-savings-default-rate-rises-to-6-percent-of-pay.aspx.

National Australia Bank. 2021. "NAB Equity Builder." Product Brochure. November. Accessed February 22, 2023. https://www.nabmarginlending.com.au/content/dam/nel/pdf-forms/Equity_Builder_app-kit/NAB_Equity_Builder_Product_Brochure.pdf.

Plaut, Pnina O., and Steven E. Plaut. 2013. "Who Wants to Be a Landlord? Factors That Affect the Inclination of Israeli Households to Rent out Property." *International Real Estate Review* 16 (1): 119–133. https://www.gssinst.org/irer/2020/04/28/who-wants-to-be-a-landlord-factors-that-affect-the-inclination-of-israeli-households-to-rent-out-property/.

Ross, Calum, and Simon Giannini. 2017. *The Real Estate Retirement Plan: An Investment and Lifestyle Solution for Canadians.* Dundurn Press, March 21, 2017. ISBN: 978-1-4597-3841-6.

Samuelson, Paul A. 1969. "Lifetime Portfolio Selection by Dynamic Stochastic Programming." *The Review of Economics and Statistics* 51 (3): 239–246. https://doi.org/10.2307/1926559.

Sandmo, Agnar. 1985. "The Effects of Taxation on Savings and Risk Taking." In *Handbook of Public Economics,* 1:265–311. Elsevier. ISBN: 978-0-444-87612-6.

Shroder, Mark. 2001. "What Makes a Landlord? Ownership of Real Estate by US Households." *Urban Studies* 38, no. 7 (June 1, 2001): 1069–1081. https://doi.org/10.1080/00420980120051657.

Statista Research Department. 2022. "Median Household Income by Age U.S. 2021." Statista, October 11, 2022. Accessed February 3, 2023. https://www.statista.com/statistics/233184/median-household-income-in-the-united-states-by-age/.

Statistics Canada. 2022. "The Daily — a Portrait of Landlords and Their Rental Income: A Tax Data Perspective," November 2, 2022. Accessed November 21, 2022. https://www150.statcan.gc.ca/n1/daily-quotidien/221102/dq221102b-eng.htm.

Tabarrok, Alex. 2022. "Lifecycle Investing." Marginal Revolution, October 13, 2022. Accessed February 6, 2023. https://marginalrevolution.com/marginalrevolution/2022/10/lifecycle-investing-2.html.

Tan, Sinan. 2013. "The Role of Options in Long Horizon Portfolio Choice." *The Journal of Derivatives* 20, no. 4 (May 31, 2013): 60–77. https://doi.org/10.3905/jod.2013.20.4.060.

Teranet. 2022. "Market Insight Report," June. Accessed February 5, 2023. https://www.teranet.ca/wp-content/uploads/2022/06/Teranet-Market-Insight-Quarterly-Report-Q2-2022.pdf.

Turner, Brandon. 2015. *The Book on Rental Property Investing: How to Create Wealth With Intelligent Buy and Hold Real Estate Investing.* Illustrated edition. Denver: BiggerPockets, October 28, 2015. ISBN: 978-0-9907117-9-7.

Turner, Brandon. 2020. *The Book on Investing in Real Estate with No (and Low) Money Down.* 2nd edition. Denver: BiggerPockets, June 30, 2020. ISBN: 978-1-947200-97-5.

U.S. Department of Housing and Urban Development. 2011. "Mortgage Credit Analysis for Mortgage Insurance on One- to Four-Unit Mortgage Loans." HUD 4155.1. March 24, 2011. Accessed February 22, 2023. https://www.hud.gov/program_offices/administration/hudclips/handbooks/hsgh/4155.1.

Venti, Steven F., and David A. Wise. 1998. "The Cause of Wealth Dispersion at Retirement: Choice or Chance?" *Papers and Proceedings of the American Economic Association* 88, no. 2 (May): 185–191. JSTOR: 116916. https://www.jstor.org/stable/116916.

Willen, Paul, and Felix Kubler. 2006. "Collateralized Borrowing and Life-Cycle Portfolio Choice." Federal Reserve Bank of Boston Public Policy Discussion Paper 06-4. March. https://www.bostonfed.org/publications/public-policy-discussion-paper/2006/collateralized-borrowing-and-life-cycle-portfolio-choice.aspx.

Yao, Rui, and Harold Zhang. 2005a. "Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints." *Review of Financial Studies* 18 (1): 197–239. https://doi.org/10.1093/rfs/hhh007.

Yao, Rui, and Harold Zhang. 2005b. "Optimal Life-Cycle Asset Allocation with Housing as Collateral." Working Paper. Accessed September 5, 2022. https://personal.utdallas.edu/~hxz054000/Yao%26Zhang2005_pub.pdf.