Why Do Homeowners Invest the Bulk of their Wealth in their Home?*

Laurent Barras and Sebastien Betermier†

March 18, 2020

Abstract

Despite the well-known benefits of diversification, homeowners invest mostly in their home. A common explanation for this pattern is that homeowners are constrained to fully own the home they want to live in. We refute this explanation and show that the predominance of housing stems from its distinct investment value. We then provide clarity on the value of the housing investment. Because owning a home provides a steady stream of housing consumption, it is equivalent to purchasing a perpetual bond indexed to that home. Housing thus plays a special role in the portfolio as one of the homeowner's risk-free assets.

JEL Classification: G11, G12.

Keywords: Portfolio allocation, real estate, hedging.

*We thank Matthieu Bouvard, Laurent Calvet, João F. Cocco, Adolfo de Motta, Chris Hennessy, Martin Lettau, Stijn van Nieuwerburgh, Christine Parlour, Guillaume Rousselet, David Schumacher, Richard Stanton, Romeo Tedongap, Adam Szeidl, Carles Vergara-Alert, Johan Walden, Nancy Wallace, participants at the Eurofidai conference in Paris, and seminar participants at Laval, McGill, Sherbrooke for their useful comments. Jielei Mao and Sebastian Krolik provided excellent research assistance. Any errors are our own.

†Desautels Faculty of Management, McGill, Bronfman Bdg, 1001 Sherbrooke St West, Montreal, QC H3A 1G5. Barras: `laurent.barras@mcgill.ca`, Betermier: `sebastien.betermier@mcgill.ca`

1 Introduction

A striking feature of household portfolios around the world is the predominance of housing. In the US, Tracy and Schneider (2001) document that the average share of gross worth invested in housing ("housing share" hereafter) ranges between 40% and 45%, whereas the share invested in public stocks is only between 5% and 10% (see also Campbell (2006); Carroll (2002); Guiso and Sodini (2013)). The contrast is even greater among homeowners who represent two thirds of US households.

The fact that housing accounts for such a disproportionate share of homeowners' wealth remains a puzzle from the perspective of standard portfolio theory. Finance research is divided on why homeowners do not choose more diversified portfolios. One line of research argues that housing predominates in the portfolio because it has distinct investment characteristics. In particular, housing offers protection against changes in the rental price, brings tax benefits, serves as a savings commitment device, and represents a valuable form of collateral against which homeowners can borrow.[1]

Another line of research views a large housing share as the result of a central friction in the housing market.[2] The friction is that households have a limited choice of ownership structure: they can either own or rent their home. If they own, they get the consumption and investment benefits of housing as a combined package. If they rent, they just get the consumption benefits. What they cannot easily do is invest in a fraction of their home, via partial ownership, and rent the remaining fraction from a co-owner. Thus, being limited to owning or renting, homeowners are often required to invest a large proportion of their wealth in the home they want to live in. This restriction on the level of housing investment is commonly known as the homeownership constraint.[3]

The objective of the present paper is to evaluate both explanations for the homeowners' large housing share. This analysis is important for finance research because these expla-

[1] See Sinai and Gyourko (2004); Han (2010); Sinai and Souleles (2005); Vestman (2019); Schlafmann (2016); Yang (2009); Corradin et al. (2014); Kraft and Munk (2011); van Hemert (2010), as well as Davis and van Nieuwerburgh (2015) for a survey of the literature.

[2] See Brueckner (1997); Cocco (2005); Flavin and Yamashita (2002); Chetty and Szeidl (2007); Frantatoni (1998); Damgaard et al. (2003); Yao and Zhang (2005).

[3] Taking a mortgage loan does not eliminate the homeownership constraint. Regardless of the amount of debt chosen by the household, she still owns 100% of the home. Drawing an analogy with a firm, the homeowner (investor) is essentially constrained to be the sole shareholder of its home (the firm).

nations offer radically different views of how homeowners allocate their wealth. The first explanation implies that the large housing share is primarily driven by investment motives. It therefore calls for an in-depth analysis of the features that make housing so valuable. In contrast, the second explanation indicates that the housing share is primarily driven by consumption motives. It also implies that any technological innovation or government policy that eliminates the homeownership constraint should allow households to benefit from more diversified portfolios.[4]

Teasing apart both explanations requires us to evaluate the *incremental* impact of the homeownership constraint on the homeowner's housing share. In other words, we need to compare the homeowner's optimal housing share in her actual environment to what she would do in a hypothetical environment in which she is free to own and consume separate amounts of housing. For convenience, we refer to this hypothetical environment as "unconstrained," even though it may incorporate other frictions besides the homeownership constraint, such as transaction costs and capital requirements.

To conduct this comparative analysis, we introduce a new measure: the Unconstrained Investment Proportion (UIP) of the homeowner's housing share. The UIP is defined as the ratio of the homeowner's unconstrained housing share to her actual housing share. This measure offers a straightforward way to interpret both our theoretical and empirical results. A value of the UIP close to one indicates that the homeowner mostly invests in housing because of its distinct investment characteristics. In contrast, a value close to zero implies that her housing share is mainly driven by the homeownership constraint.

We first show that a simple model of portfolio choice produces clear theoretical insights into the homeowner's UIP. We extend the intertemporal framework of Merton (1971) to include (i) housing as a durable good that provides both consumption and investment benefits, and (ii) the housing market friction that the agent can only rent or own the home she wishes to live in. We then derive the optimal housing share both with and without the homeownership constraint. The tractability of this setup produces an intuitive closed-form solution for the homeowner's UIP.

The main prediction of the model is that the homeowner's UIP must be large. The

[4]For example, equity-sharing programs have been proposed as a way to allow homeowners to own a fractional stake in their home (Benetton, Bracke, Cocco and Garbarino, 2019; Caplin, Chan, Freeman and Tracy, 1997).

intuition for this result is simple. The homeownership constraint is costly for the homeowner because it forces her to over-invest in housing. To mitigate this cost, the homeowner makes two adjustments. The first adjustment is to decrease her level of housing consumption compared to the unconstrained environment (i.e., she buys a smaller home). The second adjustment is to switch to rental housing if the cost of the homeownership constraint ends up being too large. The model predicts that the homeowner's UIP cannot go below 50% because it is otherwise optimal for the household to rent rather than own.

To assess the robustness of this prediction, we then generalize the theoretical analysis. We apply an entirely model-free decomposition of the UIP and account for several features examined in the housing literature such as labor income, transaction costs, borrowing constraints, alternative housing investments (e.g., REITs), and other life-cycle considerations. Our analysis confirms that, regardless of the exact magnitude of the housing share, the homeowner's UIP must be large as long as she can adjust her level of housing consumption and has access to a comparable rental housing market.

Next, we empirically examine the level of the UIP in a large cross-section of US homeowners using the Panel Study of Income Dynamics (PSID). The main challenge is that homeowners only report their actual investment in housing – not the amount they would invest in a hypothetical environment where there is no homeownership constraint. Guided by the model, we develop a flexible parametric approach that overcomes this challenge. Our approach builds on the insight that there exists one subset of homeowners that are not constrained by homeownership: landlords. Unlike other homeowners, landlords are unconstrained because they have chosen to invest more in housing than they consume. Combining micro-level data on both landlords and homeowners, we are able to estimate the UIP of every homeowner as a function of its characteristics (e.g., wealth, occupation).

The empirical results significantly strengthen the predictions of the model. The average UIP is equal to 0.94, which implies that the homeownership constraint only explains 6% of the actual housing share. More strikingly, the UIP remains close to one across all individual homeowners - even in the decile of the most constrained homeowners, it is equal to 0.84 on average. We further conduct an extensive sensitivity analysis to confirm that the large UIP is a robust feature of the data.

The evidence documented in this paper reveals that the homeownership constraint has little impact on the homeowners' housing share. Therefore, it suggests that housing must have a strong investment value to justify why homeowners willingly invest the bulk of their wealth in their home. In the remaining part of the paper, we show that the simplicity of our framework allows us to shed fresh light on this topic. Since the model abstracts from the complexities of the housing market discussed above, it allows us to zero-in on the fact that housing is first and foremost a durable good. As such, housing plays a special role in the portfolio as one of the homeowner's risk-free assets.

The insight that housing is a risk-free asset stems from examining its cash flow properties. Because owning a home guarantees a steady and durable stream of housing consumption, it is equivalent to purchasing a perpetual bond indexed to that home. Early theoretical work by Fisher (1974) establishes that perpetuities indexed to individual goods like housing represent the true risk-free assets for a long-term investor. Therefore, a large housing share should not be viewed as a large position in a single risky asset. It should instead be viewed as a large position in a risk-free asset.

Perhaps paradoxically, risk-free housing also provides speculative benefits to homeowners.[5] Whereas housing is risk-free in terms of housing consumption, it remains risky in terms of non-housing consumption. For example, a short-run position in housing makes the homeowner sensitive to variations in the house price, similar to selling a long-term bond prior to its maturity (see Sinai and Souleles (2005)). Therefore, risk-free housing has dual investment value: it provides both a steady stream of housing consumption and an exposure to the future price change of the home.

There are several reasons why the risk-free benefits of housing are economically important and broader than what prior research has established. First, each home is a distinct good that differs not only in size, quality, and location, but also in the amenities it provides. Therefore, owning a home provides a risk-free claim on the right home in the right neighborhood. Second, the risk-free benefits of housing are not limited to hedging fluctuations in the rental price (Berkovec and Fullerton, 1992; Han, 2010; Sinai and Souleles, 2005). Owning a home

[5] The role of housing as a speculative asset is documented by a large number of empirical studies, including Chambers, Spaenjers and Steiner (2019); Englund, Hwang and Quigley (2002); Favilukis, Ludvigson and van Nieuwerburgh (2017); Gatzlaff (2000); Goetzmann and Ibbotson (1990); Goetzmann (1993); Iacovello and Ortalo-Magné (2003).

also protects the homeowner against the risk of having to move out and not find a comparable home. This form of "quantity risk" is common in areas with low price risk due to rent control. In comparison, the alternative strategy of renting the home and buying housing investment products (e.g., REITs) does not provide the same level of insurance.[6]

Consistent with these arguments, the empirical evidence suggests that homeowners value the risk-free benefits of housing. By estimating the unconstrained housing share, we can identify the profile of homeowners that attach a strong investment value to housing. Building on the findings of Han (2010), we find that homeowners with a licensed occupation or several children (in particular young ones) purposely choose to invest more in housing because they generally have limited flexibility to move across areas. In addition, less wealthy households choose a greater housing share. As shown by Calvet and Sodini (2014), such households have a low tolerance for risk and tend to invest in safer assets. The results also reveal that homeowners value other investment features of housing. For instance, they invest more when housing offers higher speculative benefits (i.e., higher average return and lower volatility).

Our paper contributes to several literatures. One influential literature examines the optimal level of the housing share (see Cocco (2005); van Hemert (2010); Vestman (2019); Yao and Zhang (2005)). In these models, the total amount invested in housing depends on many factors including human capital, leverage, and several housing frictions. Here, we primarily focus on the homeowner's UIP and find that the homeownership constraint contributes little to the housing share as long as the homeowner has access to homes of different sizes/qualities and a rental housing market.

Our paper also contributes to the literature that examines the welfare implications of the homeownership constraint (e.g., Cauley, Pavlov and Schwartz (2007); Flavin and Yamashita (2002)). Our model predicts that technological innovations and government policies that eliminate this constraint should have a limited impact on the housing share. The reason is that homeowners, once unconstrained, optimally increase their housing consumption. This prediction is supported by the recent findings of Benetton, Bracke, Cocco and Garbarino (2019) who document that the introduction of shared equity programs in the U.K. has led households to buy more expensive properties.

[6]This result is consistent with the empirical evidence that taking positions in REITs or housing futures contracts provide imperfect hedges against changes in the price of an individual home (Englund, Hwang and Quigley, 2002).

Finally, the paper contributes to several studies that examine the nature of the risk-free asset. In a world with a single consumption good, Campbell and Viceira (2001) show that an indexed perpetuity is the most appropriate risk-free asset for an infinitely-lived investor. With multiple consumption goods, Fisher (1974) shows that the investor has not one, but multiple risk-free assets: one indexed bond per consumption good. More recently, Cochrane (2014) shows that identifying the risk-free properties of assets becomes easier once we focus on their cash flows rather than their covariance with the state variable to be hedged. Building on these results, our cash flow analysis provides a powerful characterization of housing as a risk-free asset - a characterization that is lost when we interpret these benefits strictly in the context of the standard hedging demand of Merton (1971).

The rest of the paper is organized as follows: Section 2 develops the model, Section 3 presents the empirical analysis, and Section 4 concludes. The Internet Appendix contains all the derivations and details of the empirical estimation.

2 Theoretical Analysis

We begin the analysis by developing a simple model of the homeowner's UIP. We then relax the assumptions of the model to generalize the main insights. Finally, we discuss the implications of our analysis for the investment value of housing in the portfolio.

2.1 Setup

The model is an extension of the Merton (1971) framework that incorporates (i) housing as a durable good that provides both consumption and investment benefits, and (ii) the housing market friction that the agent can only rent or own the home she wishes to live in. We purposely keep the model simple in order to clearly characterize the channels that contribute to a large UIP. As such, we leave aside several features that are frequently examined in the housing literature, including labor income, transaction costs, borrowing constraints, taxes, and other life-cycle considerations. As explained later, all these features affect the housing share but do not change our main conclusion that the homeowner's UIP must be large.

2.1.1 Consumption Preference and Investment Assets

Consumption Preference. The agent has initial wealth W and consumes a basket of two goods over multiple periods: a perishable non-housing good (C) and a durable housing good (K). The agent's instantaneous utility is given by the Cobb-Douglas function

$$U(C,K) = \frac{1}{1-\gamma}\left(C^{\beta_C}K^{\beta_K}\right)^{1-\gamma}, \tag{1}$$

where γ is the coefficient of relative risk aversion over the entire consumption basket and β_C and β_K are the relative importance of non-housing and housing consumption inside the agent's consumption basket (with $\beta_C + \beta_K = 1$).[7] The agent has an infinite time horizon and an additively separable utility function, which yields the lifetime expected utility:

$$E\left[\int_0^\infty e^{-\delta s}\,U(C_s, K_s)ds\right], \tag{2}$$

where δ is the time discount factor. The infinite horizon setting is convenient because the agent's decisions depend on preferences and state variables but not on time.

The cost of housing services is stochastic. Using the non-housing good as numeraire, we assume a constant rent-to-price ratio ρ and write the dynamics of the unit house price P_H as

$$\frac{dP_H}{P_H} = \mu_H\,dt + \sigma_H\,dZ_H, \tag{3}$$

where μ_H and σ_H denote the instantaneous expected return and volatility and Z_H is a Wiener process. The total price of housing consumption is thus equal to $K\rho P_H$. We interpret μ_H as the expected return net of maintenance costs. We denote by α_K the share of wealth that is used for housing consumption,

$$\alpha_K = \frac{KP_H}{W}. \tag{4}$$

Investment Assets. The agent can take long or short positions in a bond and a stock fund. The bond pays out a constant risk-free rate r in terms of the non-housing good, whereas the stock price P_S is stochastic:

$$\frac{dP_S}{P_S} = \mu_S\,dt + \sigma_S\,dZ_S, \tag{5}$$

[7] This specification is consistent with the empirical evidence in Davis and Ortalo-Magné (2011) that expenditure shares of housing are remarkably constant over time and across regions.

where μ_S and σ_S are the instantaneous expected return and volatility and Z_S is a Wiener process that is uncorrelated with Z_H.[8]

As a durable good, housing can also be used for investment purposes. For each unit of owned housing, the agent receives a rental dividend ρP_H and a capital gain of dP_H. Denoting by H the number of housing units owned, we write the share of wealth invested in housing ("housing share") as

$$\alpha_H = \frac{HP_H}{W}. \tag{6}$$

The housing share measures the economic exposure of the agent to the housing market. It should not be confused with the concept of "home equity" which commonly refers to the fraction of the home financed by equity. The distinction is important because, if the agent takes a mortgage to finance the home purchase, her exposure to the housing market does not go down as home equity would suggest.

2.1.2 Structure of the Housing Market

A central friction in the housing market is that partial homeownership opportunities are unavailable to the agent (see Brueckner (1997); Cocco (2005); Flavin and Yamashita (2002); Henderson and Ioannides (1983); Yao and Zhang (2005)). We model this friction by imposing restrictions on the housing ownership ratio ϕ_H, which is defined as the fraction of the home owned by the agent:

$$\phi_H = \frac{H}{K} = \frac{\alpha_H}{\alpha_K}. \tag{7}$$

All values of ϕ_H over the interval $(0, 1)$ are assumed to be unavailable because they imply that the agent is a partial owner of her home.[9] As a result of this market incompleteness, the agent can only choose among three available options: renting ($\phi_H = 0$), full homeownership ($\phi_H = 1$), and being a landlord ($\phi_H > 1$). Panel A of Figure 1 displays these options as a function of ϕ_H.

[8] Assuming a zero stock-housing correlation is not necessary to obtain closed-form solutions but it largely simplifies the different equations of the model. This assumption is consistent with the empirical evidence documented in previous studies (e.g., Flavin and Yamashita, 2002; Goetzmann and Ibbotson, 1990)

[9] Note that the agent is free to adjust the levels of housing consumption and investment at any time (i.e., there are no trading costs). The impact of trading frictions is discussed in more detail below.

– Figure 1 here –

The agent's preference toward one of these three options depends on the ownership ratio she would optimally choose in a hypothetical world with partial ownership:

$$\phi_H^U = \frac{H^U}{K^U} = \frac{\alpha_H^U}{\alpha_K^U}, \tag{8}$$

where α_H^U and α_K^U are the unconstrained levels of housing investment and consumption relative to total wealth. As depicted in Panel B of Figure 1, an agent with ϕ_H^U above one becomes a landlord as she is free to choose ϕ_H equal to ϕ_H^U. In contrast, an agent with ϕ_H^U below one is forced to make a constrained choice between homeownership and renting. Homeownership dominates if ϕ_H^U is sufficiently close to one, whereas renting is preferable if ϕ_H^U is sufficiently close to zero.

2.1.3 The Homeowner

Our analysis centers on the investment and consumption decisions of the homeowner, i.e., an agent who has chosen homeownership over the alternative options of renting and being a landlord. For this particular agent, the homeownership constraint binds, i.e., we have:

$$\phi_H = 1 \Leftrightarrow H = K \Leftrightarrow \alpha_H = \alpha_K. \tag{9}$$

Simply put, the homeowner must consume and invest the same amount of housing. Even though she would ideally prefer to invest less in housing than she consumes ($H^U \leq K^U$), the absence of partial ownership agreements prevents her from doing so.

Unconstrained versus Constrained Housing Shares. A direct implication of the homeownership constraint is that the housing share chosen by the homeowner has both an unconstrained and a constrained component:

$$\alpha_H = \alpha_H^U + \underbrace{(\alpha_K - \alpha_H^U)}_{\alpha_H^C}. \tag{10}$$

The unconstrained component α_H^U is equal to the optimal housing share in a hypothetical environment where the homeowner is free to choose the ownership ratio ϕ_H^U (i.e., she can choose different levels of H and K). As such, α_H^U captures the investment value of housing.

By contrast, the constrained component α_H^C is equal to the incremental investment that the homeowner must make to fully own the home she wants to consume (i.e., H and K must be equal). Therefore, α_H^C captures the impact of the homeownership constraint.

The Unconstrained Investment Proportion (UIP). Building on Equation (10), we formally define the homeowner's UIP as

$$UIP = \frac{\alpha_H^U}{\alpha_H}. \tag{11}$$

The UIP allows us to tease apart the homeowner's motives for owning housing. A value close to one indicates that the housing share remains largely unchanged with or without the homeownership constraint. At the other end of the spectrum, a value close to zero implies that the homeownership constraint plays a key role in driving the investment in housing.

2.2 Theoretical Analysis of the Homeowner's UIP

We now study the drivers of the homeowner's UIP and demonstrate that it is bound to be large. The gist of the argument is the following. The homeownership constraint is costly because it leads to an over-investment in housing. To mitigate this cost, the homeowner makes multiple adjustments. First, she finds it optimal to live in a smaller home than the one she would choose in the unconstrained environment. Second, she can choose to live in rental housing if the cost of the homeownership constraint is too large. As a result of these adjustments, the homeowner's UIP must be large - a point that we formally show below.

2.2.1 Adjustment in Housing Consumption

To study the first adjustment channel, we solve for α_H^U and α_H – the two components of the UIP. We first solve the model in the unconstrained environment ($H \neq K$) to obtain the unconstrained housing share α_H^U. We then introduce the homeownership constraint ($H = K$) to obtain the housing share α_H. Our approach builds on that of Damgaard, Fuglsbjerg and Munk (2003) but provides an explicit decomposition of the homeowner's housing share into its unconstrained and constrained components.

Proposition 1. *The homeowner's optimal housing share is the sum of the optimal uncon-*

strained and constrained housing shares:

$$\alpha_H = \alpha_H^U + \alpha_H^C, \tag{12}$$

where α_H^C *is the positive root to the quadratic equation:*

$$A\,\left(\alpha_H^C\right)^2 + B\,\alpha_H^C = C, \tag{13}$$

where $A = \frac{1}{2}\left[\frac{\gamma\sigma_H^2}{\rho} + \psi\right]$*,* $\psi = \frac{\sigma_H^2}{\rho}\left(\beta_K(1-\gamma)+\gamma\right)$*,* $B = 1 + \psi\alpha_H^U$*, and* $C = \alpha_K^U - \alpha_H^U$*. If the homeowner is constrained (i.e.,* $\phi_H^U < \phi_H$*), it follows from Equations* (12) *and* (13) *that:*

$$\alpha_H^U < \alpha_H < \alpha_K^U. \tag{14}$$

Proof. See the Appendix. □

Equation (14) reveals a key insight. As we introduce the homeownership constraint, the homeowner is forced to invest more in housing ($\alpha_H > \alpha_H^U$) to reach an ownership ratio ϕ_H equal to one. At the same time, it is optimal for the homeowner to adjust her level of housing consumption downward ($\alpha_K < \alpha_K^U$). The optimal housing share therefore lies between the unconstrained levels of investment and consumption.

To illustrate, we consider a homeowner endowed with a wealth of $400,000. In an unconstrained environment with partial ownership, she would choose to live in a home worth $200,000 and own $100,000 of it. Once the homeownership constraint is imposed, buying a home worth $200,000 is no longer optimal because it doubles the targeted housing investment. Likewise, the housing investment of $100,000 is not optimal either because it halves the targeted level of housing consumption. The optimal compromise between over-investment and under-consumption is to choose an intermediate investment level, say $150,000.[10]

The adjustment in housing consumption contributes to a large homeowner's UIP. The reason is that the impact of the homeownership constraint becomes muted as α_K moves close to α_H^U. Dividing Equation (14) by α_K^U, we obtain a lower bound for the UIP:

$$\phi_H^U < UIP < 1. \tag{15}$$

[10]To be clear, the housing consumption adjustment between the unconstrained and constrained settings does not trigger any trading costs. In the real world, the homeownership constraint always binds, (i.e., the unconstrained environment is hypothetical), so the homeowner chooses a home worth $150,000 on her initial purchase.

Equation (15) says that the UIP must be greater than the homeowner's unconstrained ownership ratio ϕ_H^U. For example, if a homeowner would ideally like to own 80% of her house ($\phi_H^U = 0.8$), her UIP necessarily ranges between 0.8 and 1. The lower bound ϕ_H^U has a natural interpretation: it corresponds to the UIP if the agent is not allowed to optimally adjust her level of housing consumption.[11]

2.2.2 The Option to Rent

The second adjustment channel available to the homeowner is the option to rent. If the impact of the homeownership constraint on her housing share ends up too large, the homeowner may decide to switch to rental housing. Because of this additional option, the fact that she has chosen to own her home indicates that her UIP must be large.

To extract this information, we solve the renter's problem and denote her optimal decisions using the superscript R. To do so, we maximize the agent's expected utility in Equation (2) under the constraint that no amount can be invested in the home ($H = 0$ equals zero). Then, we compare the value functions of the homeowner and the renter. This comparison is mathematically challenging but can be largely simplified via the use of a first-order Taylor expansion.[12] We obtain the following proposition that determines the condition under which owning dominates renting.

Proposition 2. *Let V and V^R correspond to the value functions of the homeowner and the renter. Using a first-order Taylor expansion of the log differences* $\log(\alpha_K) - \log(\alpha_K^R)$ *and* $\log\left(\frac{\alpha_C}{\alpha_K}\right) - \log\left(\frac{\alpha_C^R}{\alpha_K^R}\right)$*, we show that homeownership strictly dominates renting ($V > V^R$) if the following condition holds:*

$$\alpha_H^U > \alpha_H^C. \tag{16}$$

Proof. See the Appendix. □

Equation (16) provides a sufficient condition for the agent to become a homeowner. It

[11] This assumption is made by previous work that examine the impact of the homeownership constraint (e.g., Cauley, Pavlov and Schwartz (2007); Flavin and Yamashita (2002). In this case, α_H must increase all the way up to α_K^U, which implies that $UIP = \alpha_H^U/\alpha_H = \alpha_H^U/\alpha_K^U = \phi_H^U$. Using our previous example, it means that the agent must live in a home worth \$200,000. Therefore, we have: UIP=0.5 ($\frac{100{,}000}{200{,}000}$), versus UIP=0.66 ($\frac{100{,}000}{150{,}000}$) when housing consumption can be adjusted downward.

[12] Taylor expansions are commonly used in the portfolio choice literature (e.g., Campbell and Viceira (2002)).

predicts that homeownership is the preferred choice if the investment value of housing is sufficiently large or, more formally, if the unconstrained housing share is above the constrained housing share.

Combining the results in Propositions 1 and 2 yields a tighter lower bound on the values that the homeowner's UIP can take:

$$max\left(\frac{1}{2}, \phi_H^U\right) < UIP < 1. \quad (17)$$

Equation (17) shows that all homeowners must have a UIP equal or superior to 0.5, otherwise homeownership is no longer optimal. In addition, the UIP must be even larger among homeowners who are not heavily constrained by homeownership (i.e., those with a high ownership ratio ϕ_H^U) due to the optimal consumption adjustment discussed earlier.

To be clear, Equation (17) does not imply that every single homeowner invests heavily in housing. What it says is that, irrespective of the level of the housing share, the UIP conditional on being a homeowner must be large.

2.3 The Homeowner's UIP Beyond the Model

Our analysis of the homeowner's UIP summarized in Equation (17) depends on the assumptions of the model. Therefore, one might be concerned about the robustness of our conclusions. To address this issue, we now present a model-free decomposition of the UIP and include additional features of the housing market. Our results confirm that the homeowner's UIP must be large in a more general context.

2.3.1 A Model-free Decomposition

The following proposition provides a decomposition of the UIP that holds regardless of the portfolio choice model we use.

Proposition 3. *The homeowner's UIP always satisfies*

$$UIP = \left(\frac{\alpha_K^U}{\alpha_K}\right) \cdot \phi_H^U. \quad (18)$$

Proof. By definition, we have: $UIP = \frac{\alpha_H^U}{\alpha_H}$. Multiplying and dividing the right hand side by α_K^U, we obtain Equation (18). □

Equation (18) shows that the UIP is high if the housing consumption ratio α_K^U/α_K and the unconstrained ownership ratio ϕ_H^U are high. We illustrate this point in Table I by computing the UIP for various values of α_K^U/α_K and ϕ_H^U. As shown by the dotted region, the UIP is always greater than 0.5 provided that α_K^U/α_K is above 1 and ϕ_H^U is above 0.5.

– Table I here –

These conditions correspond to the two adjustment channels implied by our model. We expect the ratio α_K^U/α_K to be above one because the constrained homeowner endogenously chooses to adjust her level of housing consumption downward. Although the exact values of α_K^U and α_K depend on the model, it is always optimal to choose α_K below α_K^U.[13] Additionally, ϕ_H^U cannot be too low because homeownership would otherwise be no longer optimal. Here again, the exact value of the threshold depends on the model but its existence only requires the presence of a rental housing market.

This generalized analysis produces a valuable insight. Any model that predicts a low homeowner's UIP must be based on the premise that these two economic conditions do not apply. For instance, the housing market may not be deep enough to offer a continuum of homes of different sizes/qualities. In this case, the constrained homeowner may not be able to optimally reduce her housing consumption from α_K^U to α_K. In addition, the rental market may offer a limited supply of particular homes in sparse urban areas, thus eliminating the option to rent.[14] In both cases, the homeowner is left with few options to mitigate the impact of homeownership on her housing share, and her UIP could therefore be low.

[13]For example, Brueckner (1997) also shows that the homeownership constraint implies a lower housing consumption without the Cobb-Douglas CRRA utility specification used here.

[14]Interestingly, the recent growth in the rental market worldwide implies that the option to rent is available to a larger number of households. As noted by the Economist (2020) in its special report on housing: "Since 2010 global institutional investment in residential property has more than doubled in real terms. An expansion of corporate housing will raise standards in the rental sector. Big firms may be more professional than mom-and-pop landlords, and may benefit from economies of scale which allow them to provide better-quality accomodation at lower prices."

2.3.2 Additional Features of the Housing Market

As previously mentioned, our setup leaves aside several relevant features of the housing market. We now discuss the effect of introducing these features into the model, and show that the homeowner's UIP remains large.

Features that make housing more valuable. Our model does not include several features of housing that increase its investment value. Housing brings tax benefits via the tax shield on the mortgage interest (Sinai and Gyourko, 2004). It also serves as a savings commitment device (Schlafmann, 2016). Finally, it is a valuable form of collateral that allow households to borrow (Yang, 2009). Incorporating these features makes the homeownership constraint less binding as the household is willing to invest more in housing. Since the ownership ratio ϕ_H^U increases, the homeowner's UIP can only be larger.

Features that make housing less valuable. There are also several features of housing that decrease its investment value. Housing can be risky because it correlates with human capital wealth (Davidoff, 2006) and exposes the household to potential mortgage defaults (Elul et al., 2010). In addition, the protection that housing offers against fluctuations in the rental price may have limited value if the agent has a short time horizon (Sinai and Souleles, 2005), or if rental prices are less volatile than house prices (Campbell, Davis, Gallin and Martin, 2009) - two possibilities that are ruled out by the model. Whereas these features decrease ϕ_H^U, they have a limited impact on the homeowners' UIP. The reason is that agents with low levels of ϕ_H^U optimally choose to rent. For instance, older households commonly decide to rent after retirement because they have a shorter time horizon (Painter and Lee, 2009). Conditional on being homeowner, the UIP must therefore remain large.

Features that increase the option to rent. Our model excludes two features of the housing market that make renting appealing. First, households may obtain some investment exposure to the housing market via alternative assets, such as REITs or mortgage mutual funds. Thanks to these assets, they can remain renters and invest in housing without the constraint associated with homeownership. Second, buying and selling a house may generate significant transaction costs which can be avoided by renting. In both cases, the option to rent becomes more valuable. This implies that the lower bound on ϕ_H^U at which the agent chooses to become a renter increases. Consequently, households that have chosen to be homeowners must exhibit higher levels of UIP.[15]

Capital requirements. Finally, the model abstracts from capital requirements that limit the debt that households can take to finance their home purchase (Yao and Zhang, 2005). A key implication of capital requirements is that they limit the additional investment in housing that the household makes in response to the homeownership constraint (i.e., they impose a cap on α_H^C). For example, consider the agent described in Section 2.2.1 and suppose that her wealth of \$400,000 mostly consists of nontradeable human capital. Without sufficient cash, we assume that the maximum amount she can invest in housing is equal to \$60,000. Because this amount remains unchanged with or without the homeownership constraint, the UIP is equal to one (versus 0.66 in the case without capital requirements). Therefore, the presence of capital requirements leads to a higher UIP than our model predicts.[16]

2.4 The Homeowner's UIP and the Investment Value of Housing

A key implication of our analysis of the UIP is that housing must have distinct investment value to justify its predominance in the portfolios of homeowners. To elaborate, the average housing share among US homeowners in our sample is approximately 70% of their gross worth.[17] The lower bound of 50% from Equation (17) implies that the average homeowner invests *at least* 35% of her wealth in housing because of its investment value. This amount is around three times higher than the amount invested in stocks (10%).

The simplicity of the model allows us to shed fresh light on this topic. Since the model abstracts from the complexities of the housing market discussed above, it allows us to zero-in on the fact that housing is first and foremost a durable good. As such, it plays a distinct role in the portfolio of one of the homeowner's risk-free assets.

[15] In addition, if the household chooses to become homeowner, trading costs may cause her unconstrained and total housing shares α_H^U and α_H to drift away from their frictionless levels because the portfolio is not constantly rebalanced (Corradin, Fillat and Vergara-Alert, 2014; Grossmann and Laroque, 1990; Damgaard, Fuglsbjerg and Munk, 2003). However, the impact of this effect on the UIP is not clear because both α_H^U and α_H should drift in the same direction based on the realized return of housing.

[16] In this example, α_H^C equals zero - there is no over-investment due to the homeownership constraint because the maximum housing investment is below α_H^U. In the less extreme scenario where the maximum amount is between α_H^U and α_K^U, the constrained component α_H^C is positive, but capped. Therefore, the homeowner's UIP is still greater than that predicted by our model.

[17] This proportion is computed based on the sample of homeowners in the Panel Study of Income Dynamics (PSID) between 1984 and 2013 (see Section 3.2 for details).

2.4.1 Housing as a Risk-free Asset

By definition, a homeowner is allowed to stay in her home indefinitely. From a cash flow perspective, she therefore benefits from a steady and durable level of housing consumption. Therefore, owning a home is equivalent to purchasing a perpetual bond indexed to that particular home. This simple insight calls for a powerful interpretation of the role of housing in the portfolio. A large housing share should not be viewed as a large position in a single risky asset. It should instead be viewed as a large position in a risk-free bond indexed to housing consumption.

To formalize this concept, we build on Proposition 1 and show how the risk-free benefits contribute to the amount of wealth that the household chooses to invest in housing.

Proposition 4. *The optimal housing share that the agent chooses to invest in the unconstrained setting without the homeownership constraint is given by*

$$\alpha_H^U = \alpha_H^{MV} + \alpha_H^{RF} = \frac{1}{\gamma}\,\frac{\mu_H + \rho - r}{\sigma_H^2} + \left(1 - \frac{1}{\gamma}\right)\beta_K. \tag{19}$$

Proof. See the Appendix. □

The first term α_H^{MV} is the standard mean-variance weight that captures the risk-return trade-off of housing. The second term α_H^{RF} captures the risk-free benefits of housing. These benefits are highly valuable if the agent is risk averse (high γ) and cares a lot about housing consumption (high β_K). When the agent becomes infinitely risk-averse, the analogy with indexed bonds becomes extremely clear. In this case, the optimal portfolio consists of two perpetual indexed bonds: (i) a weight β_C is invested in a bond indexed on the non-housing good, and (ii) a weight β_K is invested in housing, which is nothing else than a bond indexed on the housing good. These weights correspond to the agent's consumption expenditures in both goods and guarantee that the agent obtains a balanced, risk-free stream of aggregate consumption.

It may seem counter-intuitive that housing provides both risky and risk-free benefits. However, this is exactly what we should expect from an index bond - a point forcefully made by Fisher (1974). Studying indexed bonds in a world with multiple consumption goods, Fisher shows that each bond indexed to a particular good provides a combination of risk-free and risky benefits. The intuition is that, in addition to being the true risk-free asset vis-a-vis the indexed good, the bond allows investors to speculate on the future price of the indexed good.[18]

Our result in Equation (19) is closely related to several studies that derive α_H^{RF} and interpret it as a standard hedging demand in the spirit of Merton (1971) - that is, housing provides a hedge against future fluctuations in the price of housing consumption.[19] Of course, this hedging-demand interpretation is fine. However, what is lost in this interpretation is the more general concept that housing is one of the agent's risk-free assets. This point becomes easier to see once we focus on the cash flows that housing provides and not on its covariance properties with the state variable. Cochrane (2014) makes a similar point in the general context of dynamic portfolio allocation: focusing on cash flow streams instead of hedging demands allows for a simpler characterization of the investor's optimal strategy.[20]

2.4.2 The Economic Significance of the Risk-free Benefits

There are several reasons why the risk-free benefits of housing are economically important and broader than the model suggests. The basic idea behind our argument is that each house is a distinct good. It differs not only in size and location, but also in the amenities it provides (e.g., schools, public transports). Therefore, a household interested in a particular house cannot easily replicate the same level and quality of housing consumption by switching to another home. In this context, being a homeowner yields a clear advantage: it provides a perpetual, risk-free claim on *the* home in *the* neighborhood the household chooses to settle in. In comparison, rental contracts are typically negotiated for a limited number of years. Therefore, renters always face the risk of having to move out and look for a different home.

The risk-free benefits of housing extend beyond the hedge it provides against future

[18]Fisher explains that "there are two sources of demand for each bond. One is a hedging demand, related to the share of that good in the consumption basket, and the other is a speculative demand, which tends to increase the demand for bonds indexed on the prices of goods expected to rise relatively rapidly."

[19]See Ben-Shahar (1998); Berkovec and Fullerton (1992); Damgaard, Fuglsbjerg and Munk (2003); Nordvik (2001); Sinai and Souleles (2005).

[20]To illustrate this point, Cochrane considers an investor with a 10-year horizon. Whereas a payoff approach immediately reveals that the right risk-free rate is a 10-year zero coupon bond, this point is hidden in the state-variable hedging optimization.

fluctuations in housing costs. In the model, we assume for simplicity that housing is a uniform consumption good that can be easily scaled up or down, which means that households are only subject to price risk (i.e., fluctuations in the price P_H). However, in many areas, rental prices are regulated and are thus not subject to large fluctuations. Households are still subject to "quantity risk" as they typically face supply shortages of housing (Glaeser and Luttmer, 2003). Owning a home protects households from this risk.

The risk-free nature of housing also helps to understand why homeowners favor owning a home to the alternative of renting the home and buying a diversified portfolio of homes via housing investment products (e.g., REITs). Although owning a diversified portfolio eliminates the idiosyncratic component of the house price, this intuition only applies to the risky component of housing, but not to its risky-free component. As Equation (19) highlights, the volatility of the house price is irrelevant for the risk-free weight because it does not affect the steady stream of housing consumption that the house provides. What instead matters is that the housing investment is "indexed" to the appropriate consumption good. In this respect, owning the home that the household chooses to settle in is the most appropriate risk-free asset.

It is natural to ask whether the risk-free benefits of housing are large enough to explain the housing share chosen by homeowners. Although a precise answer to this question requires a sophisticated life-cycle portfolio choice model, a simple back-of-the-envelope calculation suggests that they may be quite large. With a risk aversion γ of 5 and a housing preference parameter β_K of 0.2, we find that the value of α_H^{RF} in Equation (19) is equal to 16% of total wealth.[21] If total wealth includes human capital, these estimated housing shares are not far from the ones observed in the data. For example, Cocco (2005) estimates that housing accounts for 22% of total wealth among US homeowners that own more than $100,000 of assets. Since housing allow homeowners to borrow against their future income, they may therefore take on debt to reach the appropriate risk-free housing weight relative to total wealth. The resulting investment may look large when measured with respect to financial wealth, but it remains a moderate fraction of total wealth.

[21] These coefficients are consistent with those used in previous studies. For example, Yao and Zhang (2005) use a risk-aversion parameter γ equal to 5 and a housing preference parameter β_K equal to 0.2.

3 Empirical Analysis

The main prediction of our theoretical analysis is that the homeowner's UIP must be large. We now examine whether this prediction holds empirically. To this end, we develop a novel parametric approach to infer the entire distribution of the UIP for a representative sample of US homeowners.

3.1 Econometric Framework

We begin the presentation of the econometric framework with an overview of the parametric approach. We then formalize the main steps of the estimation procedure for the housing share and UIP for each homeowner.

3.1.1 Overview of the Approach

The main challenge of the empirical analysis comes from data limitation. In traditional surveys, homeowners only report their housing share α_H - not the amount they would invest if they were free to choose partial ownership arrangements. Because the unconstrained housing share α_H^U is not observable, we cannot directly compute the UIP.

To overcome this problem, we build on the insight that there exists one subset of households - the landlords - for which α_H^U is observable. Unlike homeowners, landlords are unconstrained because they have decided to invest more in housing than they consume. With this additional information, we apply a flexible parametric approach to infer the unconstrained housing share across homeowners and, ultimately, their UIP.[22]

To provide some intuition for our approach, we present a simple example based on our model. We consider a hypothetical set of agents that only differ in their observable preference for housing consumption β_K. For these agents, a higher β_K creates higher housing consumption needs and thus makes the homeownership constraint more binding (i.e., ϕ_H^U decreases with β_K). Denoting by $\bar{\beta}_K$ the cutoff value at which ϕ_H^U equals one, we classify

[22] Our estimation procedure resonates with that of Brueckner (1997) who also uses data on landlords to measure the impact of the homeownership constraint on the holdings of non-housing assets (such as equities). Here, we use a different estimation procedure and exclusively focus the homeowner's UIP.

agents as (i) landlords if their β_K is below $\bar{\beta}_K$, and (ii) homeowners if their β_K is above $\bar{\beta}_K$.

In Figure 2, we use the closed-form expressions of the model to plot the unconstrained and constrained housing shares for different values of β_K.[23] Both components of the housing share exhibit notable features. The unconstrained housing share α_H^U increases smoothly with β_K and is dashed beyond the cutoff $\bar{\beta}_K$ because it is not observed for homeowners. In contrast, the constrained housing share α_H^C is kinked, i.e., it is equal to zero for all landlords and only increases past the cutoff $\bar{\beta}_K$.

Our estimation approach captures these specific features by combining information on both landlords and homeowners. It uses data reported by landlords to extrapolate the unconstrained housing share past the cutoff $\bar{\beta}_K$. It also uses data reported by homeowners to capture the nonlinearity of the constrained housing share.

– Figure 2 here –

3.1.2 Estimation Procedure

We now formalize the estimation procedure. We begin by specifying the unconstrained housing share of household i at time t as a linear function of a K-vector $x_{i,t}$ of explanatory variables:

$$\alpha_{Hi,t}^U = a_t + s_U{}' \cdot x_{i,t}, \tag{20}$$

where a_t denotes the intercept and s_U is the K-vector of slope coefficients that measure the effect of $x_{i,t}$ on the unconstrained function $\alpha_{Hi,t}^U$. Intuitively, Equation (20) is a multivariate extension of Figure 2 where the difference between homeowners and landlords is not captured by the scalar β_K, but by the vector $x_{i,t}$. This vector is potentially large as it includes variables such as the household's demographic and socioeconomic characteristics, the price properties of the housing market, and power functions of the different variables. Thus, Equation (20) provides a flexible framework to capture the determinants of the unconstrained housing share.

[23] In the Appendix, we provide more information about how we calibrate the model parameters to compute the two components of the housing share (asset returns, risk aversion, and time preference).

Next, we turn to the specification of the constrained housing share α_{Hit}^C. Building on the intuition in Figure 2, we model α_{Hit}^C as a kinked function of the vector $x_{i,t}$ that takes positive values only if the homeownership constraint binds,

$$\alpha_{Hi,t}^C = \begin{cases} 0 & \text{if } d_{Hi,t} = 0, \\ g(s_C, x_{i,t}) = max(0, {s_C}' \cdot (x_{i,t} - x_{L,t})) & \text{if } d_{Hi,t} = 1, \end{cases} \tag{21}$$

where the indicator function $d_{Hi,t}$ equals one if the household is a homeowner at time t, and s_C is the K-vector of slope coefficients that measure the effect of $x_{i,t}$ on the constrained function $\alpha_{Hi,t}^C$. The k-vector $x_{L,t}$ denotes the average characteristics across landlords at time t and measures the cutoff level at which the homeownership constraint starts to bind.[24]

Equation (21) has several appealing properties. It is equal to zero for all landlords. Its magnitude depends on the k-vector $x_{i,t} - x_{L,t}$ which measures how different the homeowner is from the typical landlord in terms of underlying characteristics. Finally, it allows for the possibility that some homeowners are not constrained by homeownership (i.e., $\alpha_{Hi,t}^C$ can be null even if $d_{Hi,t}$ equals one).

Equations (20) and (21) represent the building blocks of the parametric regression. We define the housing share function as

$$\alpha_{Hi,t} = \alpha_{Hi,t}^U + \alpha_{Hi,t}^C = a_t + {s_U}' \cdot x_{i,t} + g(s_C, x_{i,t}) \cdot d_{Hi,t}, \tag{22}$$

and then jointly estimate the parameters a_t, s_U, and s_C using the surveyed housing share defined as

$$\alpha_{Hi,t}^S = \alpha_{Hi,t} + e_{i,t} = a_t + {s_U}' \cdot x_{i,t} + g(s_C, x_{i,t}) \cdot d_{Hi,t} + e_{i,t}, \tag{23}$$

where $e_{i,t}$ denotes the error term (e.g., measurement errors). Contrary to the classic linear framework, the parameters cannot be estimated using ordinary least squares because Equation (23) is nonlinear. Instead, we apply the standard Nonlinear Least Squares (NLS) methodology and discuss the details in the Appendix.

Using the estimated coefficients $\widehat{a}_t$, $\widehat{s}_U$, and $\widehat{s}_C$, we then compute the average UIP for

[24] In a multivariate setting where $x_{i,t}$ is a vector, there are many combinations of variables at which the homeownership constraint starts to bind. We capture this feature parsimoniously by interpolating around a single point $x_{L,t}$ and by imposing the theoretical restriction that the constrained demand is non-negative.

each homeowner i as

$$\widehat{UIP}_{Hi} = \frac{\widehat{\alpha}_{Hi}^{U}}{\widehat{\alpha}_{Hi}^{U} + \widehat{\alpha}_{Hi}^{C}}, \tag{24}$$

where the average unconstrained and constrained housing shares are given by:

$$\widehat{\alpha}_{Hi}^{U} = \frac{1}{T_i}\sum_{t=1}^{T_i}(\widehat{a}_t + \widehat{s}_U' \cdot x_{i,t}), \tag{25}$$

$$\widehat{\alpha}_{Hi}^{C} = \frac{1}{T_i}\sum_{t=1}^{T_i} g(\widehat{s}_C, x_{i,t}), \tag{26}$$

and $T_i = \sum_{t=1}^{T} d_{Hi,t}$ is the total number of observations available for homeowner i.

The estimation procedure requires that the housing share is correctly specified or, put differently, that $x_{i,t}$ includes all the relevant explanatory variables. In this case, the estimated coefficients $\widehat{a}_t$, $\widehat{s}_U$, and $\widehat{s}_C$ are consistent because the following orthogonality condition holds:

$$E(e_{i,t}|z_{i,t}) = 0 \quad \text{with} \quad z_{i,t} = [1, x_{i,t}', x_{i,t}^{0\,\prime}]', \tag{27}$$

where $x_{i,t}^0$ denotes the K-vector of first-order derivatives of the constrained function $\alpha_{Hi,t}^C$, i.e., $x_{i,t}^0 = \frac{\partial \alpha_{Hi,t}^C}{\partial s_C} = \frac{\partial g(s_C, x_{i,t}) d_{Hi,t}}{\partial s_C}$. Equation (27) requires that $e_{i,t}$ is orthogonal to the two vectors $x_{i,t}$ and $x_{i,t}^0$ because the parametric approach specifies both the unconstrained and constrained functions $\alpha_{Hi,t}^U$ and $\alpha_{Hi,t}^C$.[25]

In the case where the housing share is misspecified, the residual $e_{i,t}$ captures the impact of the omitted variables. The presence of omitted variable may invalidate our estimation of the UIP for two reasons. First, $\widehat{a}_t$, $\widehat{s}_U$, and $\widehat{s}_C$ are potentially biased if the orthogonality condition in Equation (27) fails. Second, we do not know whether the information contained in $e_{i,t}$ mostly affects the unconstrained or the constrained housing share.

To mitigate these concerns, we conduct an extensive sensitivity analysis which (i) examines the impact of several housing frictions (e.g., trading costs), (ii) includes a large set of

[25] A sufficient condition for Equation (27) to hold is that $e_{i,t}$ is orthogonal to $x_{i,t}$, $\frac{\partial g(s_C, x_{i,t})}{\partial s_C}$, and $d_{Hi,t}$. The last condition holds here because $d_{Hi,t}$ only depends on $x_{i,t}$. That is, being a homeowner (or a landlord) only depends on the variables $x_{i,t}$ that drive the household's housing decisions. On the contrary, the estimated coefficients would be biased if we were to use a dummy variable that depends on the error term. For instance, consider a sample of homeowners for which the housing share is only reported when it is above a threshold α_H^* (i.e., $\alpha_{Hi,t} > \alpha_H^*$). In this case, being homeowner provides information about the error term, which implies that the orthogonality condition in Equation (27) is violated (see Greene (2012, ch. 19)).

additional characteristics, and (iii) proposes a conservative estimation of the UIP in which the unmodeled residual $e_{i,t}$ is entirely attributed to the constrained housing share. The results presented in Section 3.3.3 all confirm our baseline analysis that the estimated UIP is large.

3.2 Data Description

We use micro-level data on homeowners and landlords from the wealth survey of the Panel Study of Income Dynamics (PSID) - a national survey of US households widely used in the household finance literature. This survey tracks a representative sample of US households and contains information about their real estate holdings (e.g., house value, tenure choice), their financial situation (e.g., income, wealth), and their demographics (e.g., age, household size). It is conducted every 5 years from 1984 to 1999 and then every two years from 1999 to 2013. The data requirements to be included in our sample are described in the Appendix.

We measure the housing share as the total investment in housing out of gross worth, which is defined as the sum of financial wealth (stocks, bonds, insurance, cash) and non-financial wealth (businesses, farms, real estate, and motor vehicles). Here, we favor gross worth because the measured net worth can be very small or even negative for the younger households with large educational loans. In these cases, the housing share becomes either arbitrarily large or undefined.[26]

The vector $x_{i,t}$ includes the following list of variables in our baseline specification: (i) the household's socio-economic characteristics (log income, log wealth, high-school and university education dummies, licensed occupation dummy (e.g., accountant, lawyer, pharmacist)); (ii) the household's demographic characteristics (age, marital status dummy, family size, and young children dummy); (iii) the price properties of the house (estimated expected return and volatility)[27]; (iv) the squared values of the household's continuous characteristics (log income, log wealth, and age) to capture potential non-linearities.

Our choice of explanatory variables largely builds on theoretical considerations. Variables

[26] In theory, the measure of worth should not have a significant impact on the results because worth drops out from the definition of UIP. Consistent with this prediction, we find that replacing gross worth with net worth leaves the results largely unchanged.

[27] Specifically, we compute the expected return and volatility of each individual house price (in real terms) using the time-series of home prices from the PSID database. See the Appendix for details

such as income, wealth, and age are important predictors of the housing demand in life-cycle models (e.g., Cocco (2005), van Hemert (2010), and Yao and Zhang (2005)). Variables such as the licensed occupation and young children dummies capture the household's limited flexibility in moving across different regions (e.g., Han (2010), Han (2013)). Finally, the risk-return properties of housing play a central role in standard mean-variance theory (e.g., Goetzmann (1993)).

Figures 3 and 4 show that housing represents the cornerstone asset of the portfolio in our combined sample of homeowners and landlords. The housing share remains approximately equal to 70% over the entire period 1984-2013 (Figure 3) and across diverse Metropolitan Statistical Areas such as Boise City, ID or San Diego, CA (Figure 4). The predominance of housing is therefore remarkably stable both over time and across regions.

– Figures 3 and 4 here –

Table II provides more granularity by separating homeowners and landlords at the end of the sample period (2013). We classify households as landlords if they own multiple homes and receive rental income in the year prior to 2013 or, more formally, if their ownership ratio is above one (i.e., $\phi_{Hi,t} > 1$). Panel A contains summary statistics on the composition of gross worth, debt, and financial worth. In both groups, the housing share represents, on average, more than 60% of gross worth. We see that landlords invest less in housing than homeowners (64% vs 68%). Because their ownership ratio is, by definition, higher, this result implies that landlords also consume less housing than homeowners (relative to wealth).

– Table II here –

Panel B reports the cross-sectional mean, standard deviation, and different quantiles on financial and socio-demographic characteristics. On average, landlords receive a higher income ($92,000 vs $72,000). They are also wealthier ($876,000 vs $345,000) and older (59 vs 53 years old). Finally, they have a smaller family size (0.50 vs 0.75 children) and fewer young children (19% of them have a child under 10 years old vs 28% for homeowners).

These results are in line with our econometric framework in which differences in portfolio allocations across households (Panel A) are captured by differences in their characteristics

(Panel B). Furthermore, a probit analysis confirms that these characteristics have strong power in predicting the decision to become landlord.[28]

3.3 Empirical Analysis of the Homeowner's UIP

We now present our main empirical results on the UIP, and then conduct a robustness analysis using alternative specifications.

3.3.1 Analysis of the entire Population of Homeowners

The properties of the distribution of the UIP across all homeowners are summarized in Table III. We compute the cross-sectional average and standard deviation of the UIP, and also report the 10%, 30%, 70%, and 90% quantiles.

Overall, we find that the homeowner's UIP is extremely high. The average is equal to 0.94, consistent with the fact that the unconstrained housing share is 16 times greater than the constrained housing share (67% vs. 4% of gross worth). The UIP is not only large on average - it is also close to one for the vast majority of homeowners. We see that 90% of homeowners in the population have a UIP above 0.87. In addition, the standard deviation is a mere 0.05 and thus reveals little cross-sectional variation around the average.

The statistical significance of the results is strong. The estimation is based on more than 23,000 household/year observations and achieves a good overall fit - the R^2 is close to 30%, which is higher than typical estimates in the household finance literature. The hypothesis tests further show that the unconstrained and constrained components of the housing share are statistically significant (i.e., the two null hypotheses $s_U = 0$ and $s_C = 0$ are rejected). Therefore, both components reliably contribute to the housing share, even though their economic magnitudes differ dramatically.

– Table III here –

[28] In unreported tests, we find that 75% of the pairs of observed binary responses (being a landlord or a homeowner) and their predicted values are concordant.

3.3.2 Analysis of Various Groups of Homeowners

Next, we examine the level of the UIP for different types of homeowners. In Figure 5, we plot the two components of the housing share for 10 groups of homeowners sorted on their unconstrained housing share. The model does a good job at explaining the observed housing shares among homeowners in each group (black line). Importantly, the average UIP remains close to one even among homeowners that value the investment benefits of housing the least (i.e., those with the lowest values of α_H^U).

In Figure 6, we repeat the analysis among homeowners sorted by their constrained housing share. For this analysis, we form 11 groups as the first one (group 0) includes all homeowners for which the constrained housing share is zero. Here again, we see that the UIP is always extremely large. The average UIP remains as high as 84% even among the most constrained homeowners (i.e., those with the highest values of α_H^C).

– Figures 5 and 6 here –

Altogether, the results show that the UIP is close to one for the entire cross-section of homeowners. Therefore, the empirical evidence significantly strengthens our theoretical analysis - the estimated UIP is much larger than the conservative lower bound predicted by the model in Equation (17). Our results also imply that the level of the unconstrained housing share among homeowners remains very close to the actual (observed) housing share. The summary statistics reported in Table II show that the actual homeowner's housing share in 2013 is, on average, equal to 68%. Multiplying this number by a UIP of 0.94 yields an unconstrained housing share of 64%. This is almost three times the total amount invested in financial assets (23%).

3.3.3 Alternative Specifications for the UIP

We now verify that our results are not driven by the omission of relevant variables in our specification of the housing share. For sake of brevity, we summarize our main findings below and refer the reader to the Appendix for additional details.

Housing Frictions. Households potentially face several frictions other than the homeownership constraint. Such frictions include borrowing constraints (limit on the mortgage loan), trading costs, and costs to becoming landlords (e.g., legal costs, tax treatment, managing relations with tenants). These frictions, which all contribute to the unconstrained housing share α_H^U, are not explicitly modeled in Equation (20).[29] Since they are incorporated into the residual term $e_{i,t}$, the estimator of the UIP could therefore be biased.

Theoretically, we find that these frictions are unlikely to drive our empirical results on the UIP. First, standard results in regression analysis imply that the intercept absorbs the average effect of the omitted variables. Therefore, $\widehat{\alpha}_H^U$ captures the average impact of the omitted frictions. Second, the magnitude of the costs cannot be too large, otherwise homeowners have strong incentive to rebalance their portfolios and/or make the switch to becoming landlords (Corradin, Fillat and Vergara-Alert, 2014; Grossmann and Laroque, 1990; Damgaard, Fuglsbjerg and Munk, 2003).

To validate these theoretical arguments, we re-estimate the model using only observations following a home purchase. This specification guarantees that the housing share is not driven by trading costs. To control for the costs to becoming landlords, we also re-estimate the model with quasi-landords, which are defined as homeowners that are nearly identical to landlords (see the appendix for details). The empirical results reported in Table IV reveal that the distribution of the UIP in both cases remains largely unchanged.

Additional Characteristics. In a world without housing frictions, the estimated UIP could still be biased if relevant characteristics are omitted from the vector $x_{i,t}$. To mitigate this issue, we re-estimate the two components of the housing share using a large set of additional variables. We split households into whether they reside in states where mortgage laws authorize recourse to their personal assets in the event of bankruptcy. We include dummies reflecting the ethnicity of each household, as well as characteristics that reflect the state of the housing market (e.g., the fractions of homes that are owner-occupied or vacant). Finally, we include additional neighborhood characteristics (e.g., ethnic composition, average education rate). As shown in Table IV, the cross-sectional distribution of the UIP remains largely unchanged throughout all specifications.

[29] As a reminder, α_H^U is defined as the housing share in a hypothetical setting where the agent is free to own and consume separate amounts of housing. As such, α_H^U may be affected by frictions other than the homeownership constraint.

Conservative Estimation. The inclusion of additional variables does not guarantee that the residual $e_{i,t}$ no longer contains any information about the unconstrained and constrained housing shares. To address this concern, we estimate the UIP in the most conservative manner by attributing any positive residual to the constrained housing share. In other words, we assume that any unmodeled source of variation of the observed housing share is due to the homeownership constraint. Consistent with intuition, Table IV reveals that the cross-sectional distribution of the UIP shifts to the left. However, its average level remains very large at 86%.

3.4 Additional Insights into the Investment Value of Housing

In addition to measuring the UIP, our procedure yields estimates of the coefficients s_U of the unconstrained housing share. This allows us to shed new light on the determinants of the investment value of housing - an analysis that departs from the previous literature which only focuses on the actual (observed) housing share (e.g., Cocco (2005), Han (2010)).

The leftmost columns of Table V report the marginal impact (and t-statistic) of each explanatory variable on the unconstrained housing share α_H^U. Because our baseline specification includes the squared term of each continuous variable $x_{ik,t}$ (log income, log wealth, and age), we compute the marginal impact of each variable at their average level across homeowners (i.e., when $x_{ik,t} = E(x_{ik,t})$).[30]

The empirical results give support to the view that homeowners value the risk-free benefits of housing. We find that wealth is an important determinant of the unconstrained housing share, i.e., a doubling of wealth reduces housing investment by 11.5% (t-statistic of -25.1). In addition, housing is particularly valuable for households with (i) a larger size (t-statistic of 2.3), (ii) young children (t-statistic of 3.5), and (iii) a licensed occupation (t-statistic of 2.4). These results capture the typical profile of households that strongly value a risk-free stream of housing consumption. Calvet and Sodini (2014) document that poorer households tend to be more risk averse ($1 - \gamma$ is high). Han (2010) also shows that households with a licensed occupation or several children (including young ones) have a strong preference for

[30] Formally, we compute the marginal impact as $\widehat{s}_1 + 2\widehat{s}_2\bar{x}_{ik,t}$, where $\widehat{s}_1$ and $\widehat{s}_2$ are the estimated coefficients associated with $x_{ik,t}$ and $x_{ik,t}^2$, and $\bar{x}_{ik,t}$ denotes the sample average of $x_{ik,t}$. Similarly, the variance term for the t-statistic computation is given by $var(\widehat{s}_1) + (2\bar{x}_{ik,t})^2 var(\widehat{s}_2) + (2\bar{x}_{ik,t}) cov(\widehat{s}_1, \widehat{s}_1)$, where var and cov denote the variance and covariance terms.

staying in their current homes because they have a limited flexibility in moving across areas (β_K is high). As predicted by the model in Equation (19), both effects increase the risk-free component of the unconstrained housing share.

Consistent with the standard mean-variance theory, the investment value of housing also depends on the return properties of the individual homes. A 10% increase in the expected return of the home price increases the unconstrained housing share by 5.4% (t-statistic of 4.9). In addition, a 10% increase in volatility leads to a 2%-decrease in housing investment (t-statistic of -4.4). These results highlight the dual investment value of housing discussed in Section 2.4.2. The combination of risk-free and speculative benefits contributes to the cross-sectional variation in the unconstrained housing share observed in the data.

For sake of completeness, the rightmost columns of Table V report the coefficients associated with the constrained housing share. Consistent with Table III, the magnitude of the coefficients is typically smaller - the average across all characteristics equals 4.7% versus 8.4% for the unconstrained housing share. We also see that several variables have opposite effects on the two components of the housing share (e.g., house price characteristics). This reversal arises because these variables primarily trigger changes in housing investment rather than changes in housing consumption. Therefore, an increased willingness to invest in housing also comes with a loosening of the homeownership constraint.

– Table V here –

4 Conclusion

Despite decades of financial research advocating the importance of diversifying one's portfolio, most households continue to invest the bulk of their wealth in one asset: their home. In this paper, we tease apart whether homeowners purposely choose a large housing share because they strongly value the distinct investment features of housing, or whether they are constrained to do so because of their housing consumption needs. We theoretically and empirically show that households invest in housing because they value its investment benefits. Our findings therefore imply that housing must have unique investment appeal to justify their highly skewed portfolios.

Our findings call for further research on why investment motives for owning housing are so large. Our model highlights the fact that housing, as a durable good, plays a special role as one of the homeowner's risk-free assets. However, it does not address how other features such as labor income risk, taxes, or transaction costs also contribute to the unconstrained housing share. It would be valuable to quantify these effects using a more realistic life-cycle model in the style of Cocco (2005) or Vestman (2019).

Our findings also call for further research on the broader macroeconomic implications of housing as a risk-free asset. Much research has focused on the destabilizing role that housing markets played in the financial crisis through mortgage risk-taking (Mian and Sufi, 2014). Perhaps paradoxically, it could be that the risk-free appeal of long-term housing assets contributes to the willingness of households to take on large mortgages and gain exposure to short-run fluctuations in housing prices.

We finally note that, unlike bonds, housing is a risk-free asset in positive net supply. Recent asset pricing studies have shown that economies where the risk-free asset is in positive net supply behave differently than economies where it is in zero net supply (Parlour, Stanton and Walden, 2011). The incorporation of housing in equilibrium asset pricing models, building on the work of Lustig and van Nieuwerburgh (2005) and Piazzesi, Schneider and Tuzel (2007), represents another promising avenue for future research.

References

Ben-Shahar, D. (1998), 'On the optimality of the hybrid tenure mode', *Journal of Housing Economics* **7**, 69–92.

Benetton, M., Bracke, P., Cocco, J. and Garbarino, N. (2019), Housing consumption and investment: evidence from shared equity mortgages, Technical report. Staff Working Paper No. 790, Bank of England.

Berkovec, J. and Fullerton, D. (1992), 'A general equilibrium model of housing, taxes, and portfolio choice', *Journal of Political Economy* **100**, 390–429.

Brueckner, J. (1997), 'Consumption and investment motives and the portfolio choices of homeowners', *The Journal of Real Estate Finance and Economics* **15**(2), 159–180.

Calvet, L. and Sodini, P. (2014), 'Twin picks: Disentangling the determinants of risk-taking in household portfolios', *The Journal of Finance* **69**, 867–906.

Campbell, J. (2006), 'Household finance', *Journal of Finance* **61**(4), 1553–1604.

Campbell, J. and Viceira, L. (2001), 'Who should buy long-term bonds?', *American Economic Review* **91**, 99–127.

Campbell, J. and Viceira, L. (2002), *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press.

Campbell, S., Davis, M., Gallin, J. and Martin, R. (2009), 'What moves housing markets: A variance decomposition of the rent-price ratio', *Journal of Urban Economics* **66**, 90–102.

Caplin, A., Chan, S., Freeman, C. and Tracy, J. (1997), *Housing Partnerships: A New Approach to a Market at a CrossRoads*, MIT Press, Cambridge and London.

Carroll, C. (2002), Portfolios of the rich, *in* L. Guiso, M. Haliassos and T. Jappelli, eds, 'Household Portfolios', MIT Press, Cambridge and London.

Cauley, S., Pavlov, A. and Schwartz, E. (2007), 'Homeownership as a constraint on portfolio allocation', *Journal of Real Estate Finance and Economics* **34**, 283–311.

Chambers, D., Spaenjers, C. and Steiner, E. (2019), The rate of return on real estate: long-run micro-level evidence. Working Paper.

Chetty, R. and Szeidl, A. (2007), 'Consumption commitments and risk preferences', *Quarterly Journal of Economics* **122**(2), 831–877.

Cocco, J. (2005), 'Portfolio choice in the presence of housing', *Review of Financial Studies* **18**(2), 535–567.

Cochrane, J. (2014), 'A mean-variance benchmark for intertemporal portfolio theory', *Journal of Finance* **69**(1), 1–49.

Corradin, S., Fillat, J. and Vergara-Alert, C. (2014), 'Optimal portfolio choice with predictability in house prices and transaction costs', *Review of Financial Studies* **27**(3), 823–880.

Damgaard, A., Fuglsbjerg, B. and Munk, C. (2003), 'Optimal consumption and investment strategies with a perishable and an indivisible durable good', *Journal of Economic Dynamics and Control* **28**(2), 209–253.

Davidoff, T. (2006), 'Labor income, housing prices and homeownership', *Journal of Urban Economics* **59**(2), 209–235.

Davis, M. and Ortalo-Magné, F. (2011), 'Household expenditures, wages, rents', *Review of Economic Dynamics* **14**, 248–261.

Davis, M. and van Nieuwerburgh, S. (2015), Housing, finance, and the macroeconomy, *in* J. H. Gilles Duranton and W. Strange, eds, 'Handbook of Regional and Urban Economics', chapter 12.

Elul, R., Souleles, N., Chomsisengphet, S., Glennon, D. and Hunt, R. (2010), 'What "triggers" mortgage default?', *American Economic Review: Papers & Proceedings* **100**, 490–494.

Englund, P., Hwang, M. and Quigley, J. (2002), 'Hedging housing risk', *Journal of Real Estate Finance and Economics* **24**, 167–200.

Favilukis, J., Ludvigson, S. and van Nieuwerburgh, S. (2017), 'The macroeconomic effects of housing wealth, housing finance, and limited risk-sharing in general equilibrium', *Journal of Political Economy* **125**(1), 140–223.

Fisher, S. (1974), 'The demand for index bonds', *Journal of Political Economy* **83**(3), 509–534.

Flavin, M. and Yamashita, T. (2002), 'Owner-occupied housing and the composition of the household portfolio', *The American Economic Review* **92**(1), 345–362.

Frantatoni, M. (1998), 'Homeownership and investment in risky assets', *Journal of Urban Economics* **44**, 27–42.

Gatzlaff, D. (2000), The effect of single-family housing on multi-asset portfolio allocations. Working paper.

Glaeser, E. and Luttmer, E. (2003), 'The misallocation of housing under rent control', *American Economic Review* **4**, 1027–1046.

Goetzmann, W. (1993), 'The single family home in the investment portfolio', *Journal of Real Estate Finance and Economics* **6**, 201–222.

Goetzmann, W. and Ibbotson, R. (1990), 'The performance of real estate as an asset class', *Journal of Applied Corporate Finance* **3**(1), 65–76.

Grossmann, S. and Laroque, G. (1990), 'Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods', *Econometrica* **58**, 25–51.

Guiso, L. and Sodini, P. (2013), Household finance: An emerging field, *in* G. Constantinides, M. Harris and R. Stulz, eds, 'Handbook of the Economics of Finance', Vol. 2.

Han, L. (2010), 'The effects of price risk on housing demand: Evidence from u.s. markets', *The Review of Financial Studies* **23**(11), 3889–3928.

Han, L. (2013), 'Understanding the puzzling risk-return relationship for housing', *The Review of Financial Studies* **26**(4), 877–928.

Henderson, J. and Ioannides, Y. (1983), 'A model of housing tenure choice', *The American Economic Review* .

Iacovello, M. and Ortalo-Magné, F. (2003), 'Hedging housing risk in london', *Journal of Real Estate Finance and Economics* **27**(2), 191–209.

Kraft, H. and Munk, C. (2011), 'Optimal housing, consumption, and investment decisions over the life-cycle', *Management Science* **57**(6), 1025–1041.

Lustig, H. and van Nieuwerburgh, S. (2005), 'Housing collateral, consumption insurance, and risk premia: An empirical perspective', *The Journal of Finance* **60**(3), 1167–1219.

Merton, R. (1971), 'Optimum consumption and portfolio rules in a continuous-time model', *Journal of Economic Theory* **3**, 373–413.

Mian, A. and Sufi, A. (2014), *House of Debt: How they (and you) caused the great recession, and how we can prevent it from happening again*, The University of Chicago Press.

Nordvik, V. (2001), 'A housing career perspective on risk?', *Journal of Housing Economics* **10**, 456–471.

Painter, G. and Lee, K. (2009), 'Housing tenure transitions of older households: life cycle, demographic, and familial factors', *Regional Science and Urban Economics* **39**, 749–760.

Parlour, C., Stanton, R. and Walden, J. (2011), 'Revisiting asset pricing puzzles in an exchange economy', *Review of Financial Studies* **24**, 629–674.

Piazzesi, M., Schneider, M. and Tuzel, S. (2007), 'Housing, consumption, and asset pricing', *Journal of Financial Economics* **83**, 531–569.

Schlafmann, K. (2016), Housing, mortgages, and self control, Technical report. CEPR Discussion Paper 11589.

Sinai, T. and Gyourko, J. (2004), 'the (un)changing geographical distribution of housing tax benefits: 1980-2000', *Tax Policy and the Economy* **18**, 175–208.

Sinai, T. and Souleles, N. (2005), 'Owner-occupied housing as a hedge against rent risk', *The Quarterly Journal of Economics* **120**(2), 763–789.

Tracy, J. and Schneider, H. (2001), 'Stocks in the household portfolio: A look back at the 1990s', *Federal Reserve Bank of New York Current Issues in Economics and Finance* **7**(4), 1–6.

van Hemert, O. (2010), 'Household interest rate risk management', *Real Estate Economics* **38**(3), 467–505.

Vestman, R. (2019), 'Limited stock market participation among renters and homeowners', *Review of Financial Studies* **32**(4), 1494–1535.

Yang, F. (2009), 'Consumption over the life cycle: How different is housing?', *Review of Economic Dynamics* **12**(3), 423–443.

Yao, R. and Zhang, H. (2005), 'Optimal consumption and portfolio choices with risky housing and borrowing constraints', *Review of Financial Studies* **18**(1), 197–239.

Table I
Model-Free Decomposition of the UIP

The table reports possible values that the homeowner's UIP can take as a function of 1) the optimal housing ownership ratio ϕ_H^U in the unconstrained environment, and 2) the ratio $\frac{\alpha_K^U}{\alpha_K}$, where α_K^U and α_K denote the levels of housing consumption in the unconstrained and constrained environments (relative to wealth):

$$UIP = \phi_H^U \times \frac{\alpha_K^U}{\alpha_K}$$

The dotted region in the lower right corner indicates the set of values of the UIP where the minimum ownership ratio is 50% and the introduction of the homeownership constraint leads to a decrease in the level of housing consumption (i.e., $\alpha_K < \alpha_K^U$).

ϕ_H^U	α_K^U/α_K								
	0	0.25	0.5	0.75	1	1.25	1.5	1.75	2
0	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.1	0	0.03	0.05	0.08	0.10	0.13	0.15	0.18	0.20
0.2	0	0.05	0.10	0.15	0.20	0.25	0.30	0.35	0.40
0.3	0	0.08	0.15	0.23	0.30	0.38	0.45	0.53	0.60
0.4	0	0.10	0.20	0.30	0.40	0.50	0.60	0.70	0.80
0.5	0	0.13	0.25	0.38	0.50	0.63	0.75	0.88	1.00
0.6	0	0.15	0.30	0.45	0.60	0.75	0.90	1.00	1.00
0.7	0	0.18	0.35	0.53	0.70	0.88	1.00	1.00	1.00
0.8	0	0.20	0.40	0.60	0.80	1.00	1.00	1.00	1.00
0.9	0	0.23	0.45	0.68	0.90	1.00	1.00	1.00	1.00
1	0	0.25	0.50	0.75	1.00	1.00	1.00	1.00	1.00

Table II
Summary Statistics for Households

The table reports summary statistics on wealth characteristics (Panel A) and financial and demographic characteristics (Panel B) among homeowners (first set of columns) and landlords (second set of columns) from the 2013 wave of the PSID. Landlords are defined as households that received rental income on housing properties during the year prior to the survey (i.e., they own more housing than they consume).

Panel A: Wealth Characteristics

	Homeowners					Landlords				
	Mean	Standard deviation	Percentiles 25th	50th	75th	Mean	Standard deviation	Percentiles 25th	50th	75th
Gross worth										
Financial worth	0.23	0.24	0.03	0.12	0.39	0.24	0.22	0.07	0.17	0.40
Housing	0.68	0.25	0.49	0.75	0.89	0.64	0.25	0.47	0.70	0.85
Business	0.02	0.10	0.00	0.00	0.00	0.08	0.18	0.00	0.00	0.00
Autos	0.07	0.09	0.02	0.04	0.09	0.04	0.04	0.01	0.03	0.05
Debt (relative to Gross Worth)										
Mortgage	0.36	0.66	0.00	0.26	0.59	0.19	0.25	0.00	0.09	0.30
Other debts	0.07	0.34	0.00	0.00	0.04	0.01	0.03	0.00	0.00	0.01
Financial worth										
Stocks	0.10	0.23	0.00	0.00	0.00	0.18	0.28	0.00	0.00	0.34
IRA	0.25	0.36	0.00	0.00	0.53	0.35	0.37	0.00	0.21	0.71
Bonds	0.06	0.19	0.00	0.00	0.00	0.07	0.19	0.00	0.00	0.00
Cash	0.60	0.42	0.13	0.77	1.00	0.40	0.39	0.05	0.25	0.87
Number of observations	2,315	2,315	2,315	2,315	2,315	207	207	207	207	207

Table II
Summary Statistics for Homeowners - *Continued*

Panel B: Financial and Demographic Characteristics

	Homeowners					Landlords				
	Mean	Standard deviation	Percentiles 25th	Percentiles 50th	Percentiles 75th	Mean	Standard deviation	Percentiles 25th	Percentiles 50th	Percentiles 75th
Financial characteristics										
Household income ($)	71,832	62,626	37,315	58,474	88,283	92,455	57,818	51,241	81,164	118,127
Gross worth ($)	345,449	499,058	110,274	202,739	393,834	876,574	1,183,893	271,917	507,532	969,175
Licensed occupation dummy	0.07	0.26	0.00	0.00	0.00	0.10	0.30	0.00	0.00	0.00
Self-employed dummy	0.11	0.32	0.00	0.00	0.00	0.25	0.43	0.00	0.00	0.00
Demographic Characteristics										
Age	53.19	14.86	41.00	53.00	64.00	58.84	14.62	49.00	61.00	68.00
Married dummy	0.80	0.40	1.00	1.00	1.00	0.82	0.39	1.00	1.00	1.00
Number of children	0.75	1.10	0.00	0.00	2.00	0.52	1.02	0.00	0.00	1.00
Young child dummy	0.28	0.45	0.00	0.00	1.00	0.19	0.39	0.00	0.00	0.00
Retired dummy	0.24	0.43	0.00	0.00	0.00	0.42	0.49	0.00	0.00	1.00
High school dummy	0.28	0.45	0.00	0.00	1.00	0.26	0.44	0.00	0.00	1.00
Post-high school dummy	0.24	0.43	0.00	0.00	0.00	0.21	0.41	0.00	0.00	0.00
Number of observations	2,315	2,315	2,315	2,315	2,315	207	207	207	207	207

Table III
Cross-Sectional Distribution of the Homeowner's UIP

The table reports the estimation results from the panel regression of the surveyed housing share $\alpha^S_{Hi,t}$ for each household i (homeowner or landlord) on a vector of financial, demographic, and house price characteristics $x_{i,t}$ estimated between 1984 and 2013 using micro-level data from the PSID,

$$\alpha^S_{Hi,t} = a_t + s'_U x_{i,t} + \max(0, s'_C (x_{i,t} - x_{L,t}))d_{Hi,t} + e_{i,t},$$

where $d_{Hi,t}$ is a dummy variable that equals 1 if the household is a homeowner and 0 otherwise, and $x_{L,t}$ denotes the average characteristics across landlords. We refer to s_U and s_C as the slope vectors of the unconstrained and constrained components of the housing share, which are defined as $\alpha^U_{Hi,t} = a_t + s'_U x_{i,t}$ and $\alpha^C_{Hi,t} = \max(0, s'_C(x_{i,t} - x_{L,t}))d_{Hi,}$. The UIP for each homeowner i is computed as $UIP_i = \alpha^U_{Hi} / (\alpha^U_{Hi} + \alpha^C_{Hi})$, where α^U_{Hi} and α^C_{Hi} are the components of the housing share averaged across all observations associated with homeowner *i.* The table reports the moments and quantiles of the cross-sectional distribution of the homeowner's UIP_i, the total housing share $\alpha_{Hi} = \alpha^U_{Hi} + \alpha^C_{Hi}$, and its two components α^U_{Hi} and α^C_{Hi}. The table also reports the statistical significance of the hypothesis tests that the slope coefficients are jointly equal to zero, i.e., H_o: $s_U = 0$ and H_o: $s_C = 0$.

	Mean	Std. Dev	Percentiles					Joint Tests	
			10%	30%	50%	70%	90%	F-stat	p-value
UIP	0.94	0.05	0.87	0.92	0.95	1.00	1.00		
Total housing share	0.72	0.12	0.55	0.68	0.74	0.79	0.84	83.5	<.001
Unconstrained housing share	0.67	0.11	0.53	0.63	0.69	0.73	0.80	99.5	<.001
Constrained housing share	0.04	0.04	0.00	0.00	0.03	0.06	0.10	5.8	<.001
R-square	30.21%								
Number of observations	23,566								

Table IV
Alternative Specifications of the Homeowner's UIP

Each row of the table reports the moments and quantiles of the cross-sectional distribution of the homeowner's UIP using alternative specifications for the model of the housing share. Specification (i) is based on the sample of new buyers only. Specification (ii) replaces landlords with quasi-landlords defined as homeowners with characteristics that are nearly identical to those of landlords. Specification (iii) is based on the sample of homeowners residing in mortgage-recourse states. Specification (iv) includes ethnicity dummies. Specification (v) adds neighborhood characteristics on the local housing market and homeowners from US census data. Specification (vi) considers a conservative estimation of the UIP that attributes any unmodeled variation of the observed housing share to the constrained housing share. Additional details on these alternative specifications can be found in the Appendix.

	Mean	Std. Dev	Percentiles				
			10%	30%	50%	70%	90%
(i) Based on new buyers	0.94	0.05	0.87	0.92	0.95	0.98	1.00
(ii) Based on quasi-landlords	0.87	0.05	0.81	0.85	0.87	0.89	0.92
(iii) Recourse state	0.95	0.05	0.88	0.92	0.95	0.99	1.00
(iv) Ethnicity	0.93	0.05	0.86	0.90	0.94	0.98	1.00
(v) Neighborhood characteristics							
Housing market	0.94	0.05	0.87	0.90	0.94	0.98	1.00
Homeowners	0.91	0.06	0.83	0.87	0.90	0.95	1.00
(vi) Conservative estimation	0.86	0.14	0.71	0.81	0.89	0.94	1.00

Table V
Panel Regression of Housing Shares

The table reports the estimation results from the panel regression of the observed housing share $\alpha^S_{Hi,t}$ for each household *i* (homeowner or landlord) on a vector of financial, demographic, and house price characteristics $x_{i,t}$ estimated between 1984 and 2013 using micro-level data from the PSID,

$$\alpha^S_{Hi,t} = a_t + s'_U x_{i,t} + \max(0, s'_C (x_{i,t} - x_{L,t}))d_{Hi,t} + e_{i,t},$$

where $d_{Hi,t}$ is a dummy variable that equals 1 if the household is a homeowner and 0 otherwise, and $x_{L,t}$ denotes the average characteristics across landlords. We refer to s_U and s_C as the slope vectors of the unconstrained and constrained components of the housing share, which are defined as $\alpha^U_{Hi,t} = a_t + s'_U x_{i,t}$ and $\alpha^C_{Hi,t} = \max(0, s'_C(x_{i,t} - x_{L,t}))d_{Hi,t}$. The table reports the individual slope coefficients for s_U, and s_C and their t-statistics in separate columns. We cluster observations for each household to compute the variance of the coefficients.

	s_U		s_C	
	Estimate	t-stat	Estimate	t-stat
Financial Characteristics				
Gross worth (log)	-0.112	-25.06	0.002	0.35
Household income (log)	0.008	1.27	0.006	0.78
Licensed Occupation	0.035	2.42	-0.031	-1.84
Demographic Characteristics				
Age	0.002	3.34	-0.004	-5.70
Number of children	0.014	2.35	0.006	0.90
Young child dummy	0.049	3.48	-0.043	-2.82
Married dummy	-0.017	-1.56	0.001	0.11
High school dummy	-0.020	-2.07	0.000	-0.03
Post-high school dummy	-0.011	-0.80	0.009	0.61
House Price Characteristics				
Average growth	0.541	4.88	-0.367	-2.83
Volatility	-0.197	-4.41	0.089	1.62

R-square	30.21%
Number of observations	23,566

Figure 1
Feasible Opportunity Set in the Housing Market

In Panel A, we illustrate the feasible options in the housing market in terms of the agent's housing ownership ratio ϕ_H. The dotted line represents the set of unavailable partial ownership opportunities (i.e., ϕ_H cannot take values between 0 and 1). In Panel B, we illustrate the agent's tenure decision. The agent chooses to either become a renter, a homeowner, or a landlord depending on the optimal homeownership ratio ϕ_H^U that she would choose in an unconstrained environment.

Panel A. Set of available housing options

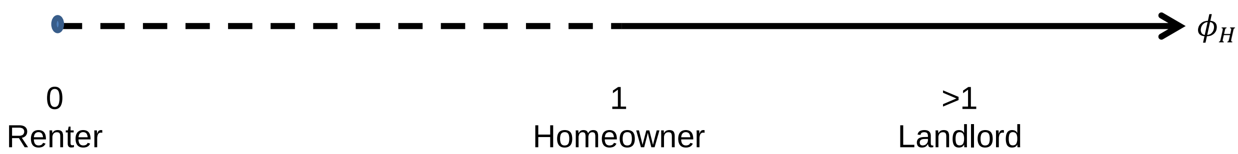

Panel B. Optimal tenure choice

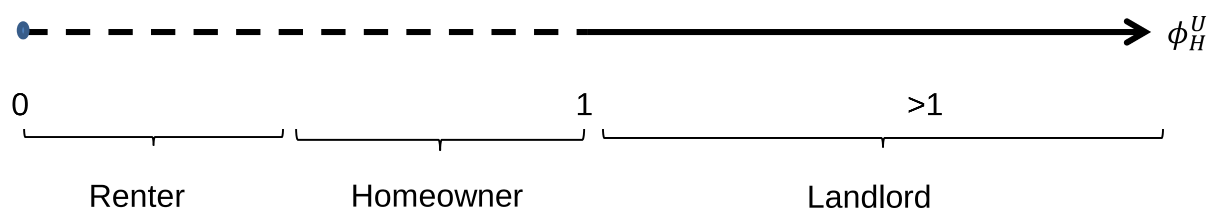

Figure 2
Illustration of the Estimation Procedure

The figure illustrates the parametric approach that we use to empirically estimate the homeowner's UIP. We consider an example based on the model where households only differ in terms of their preference for housing consumption β_K (the other parameters are provided in the Appendix). The sample includes both homeowners and landlords. The solid line represents the household's housing share, α_H. The dotted line represents the unobserved housing share α_H^U that the homeowner would choose in an unconstrained environment.

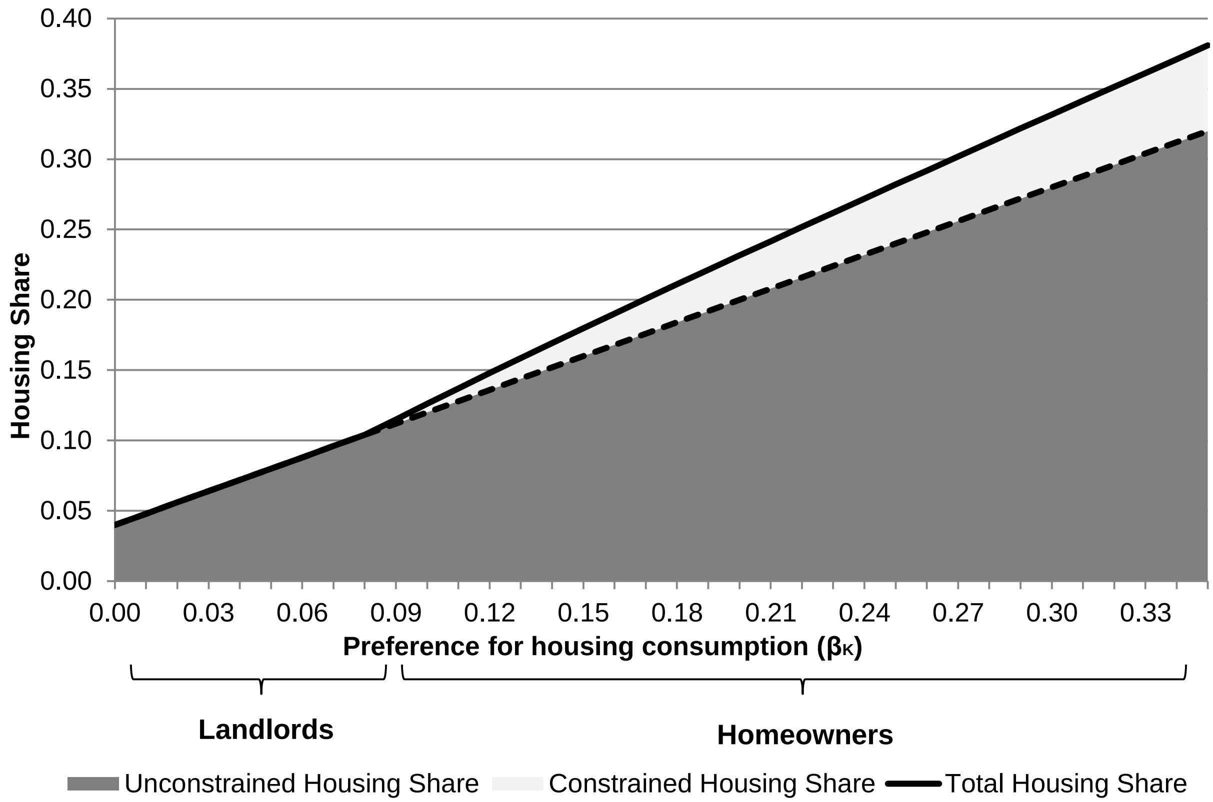

Figure 3
Portfolio Shares Across Time

The figure reports the average composition of gross worth of a representative sample of homeowners and landlords from the PSID between 1984 to 2013. For each year in the survey, we compute and report the (equally-weighted) average share of gross worth invested in housing, stocks, bonds, autos, and business.

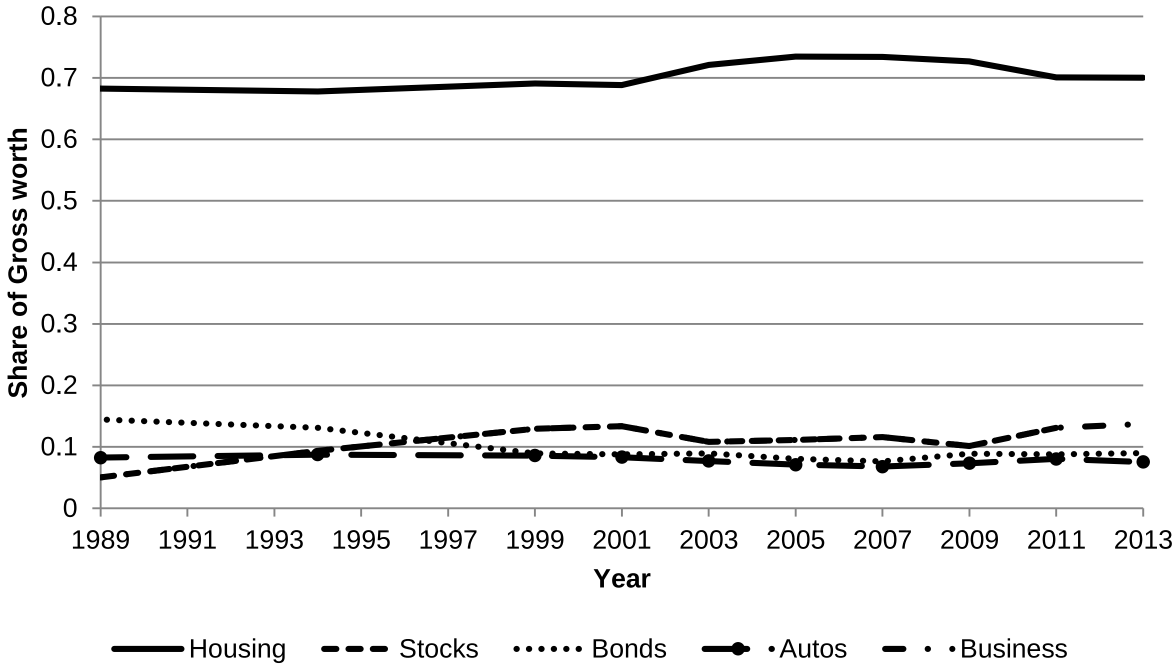

Figure 4
Portfolio Shares Across Areas

The figure reports the average composition of gross worth of a representative sample of households in the PSID for various Metropolitan Statistical Areas (MSA). For each MSA, we compute the (equally-weighted) average shares of gross worth invested in housing, stocks, bonds, autos, and business. We then compute the time-series average of each share over the years when the survey is conducted between 1984 and 2013.

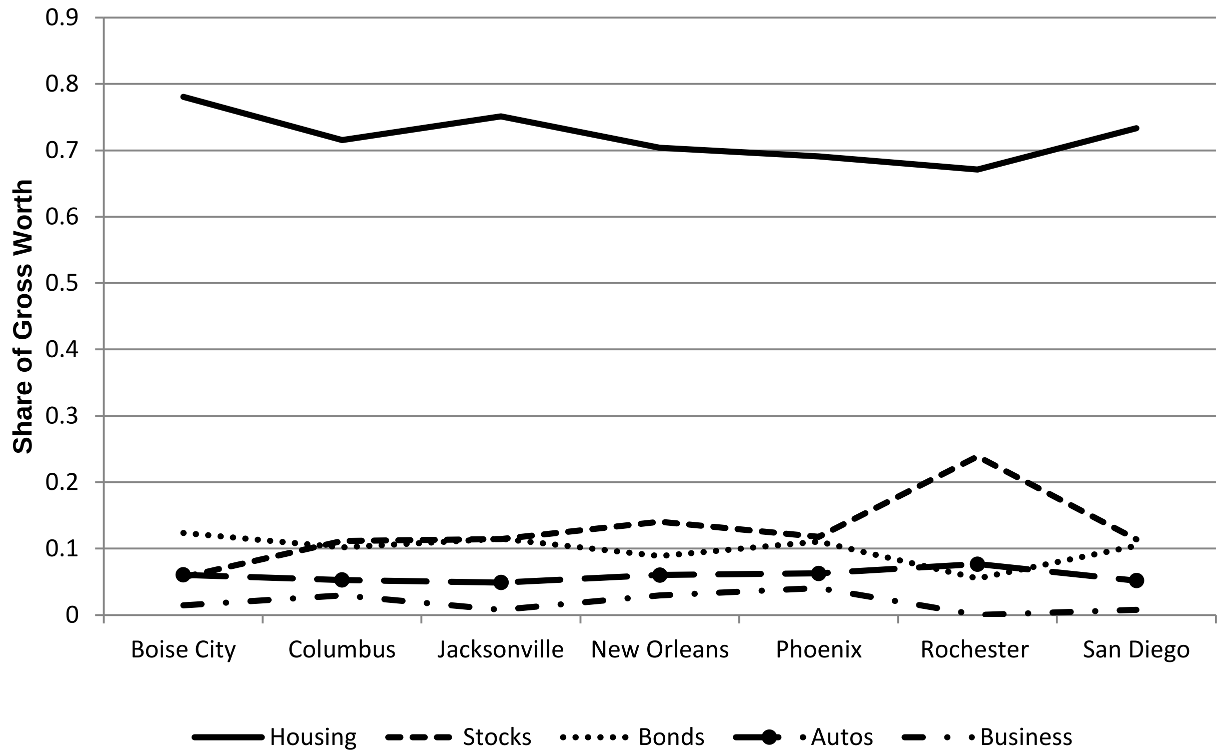

Figure 5
Housing Share Decomposition among Homeowners Sorted on their Unconstrained Housing Share

The figure illustrates the predicted unconstrained housing shares (solid dark grey), constrained housing shares (solid light grey), and observed housing shares (solid black line) across the subsample of homeowners. Homeowners are sorted into 10 equally-formed groups.

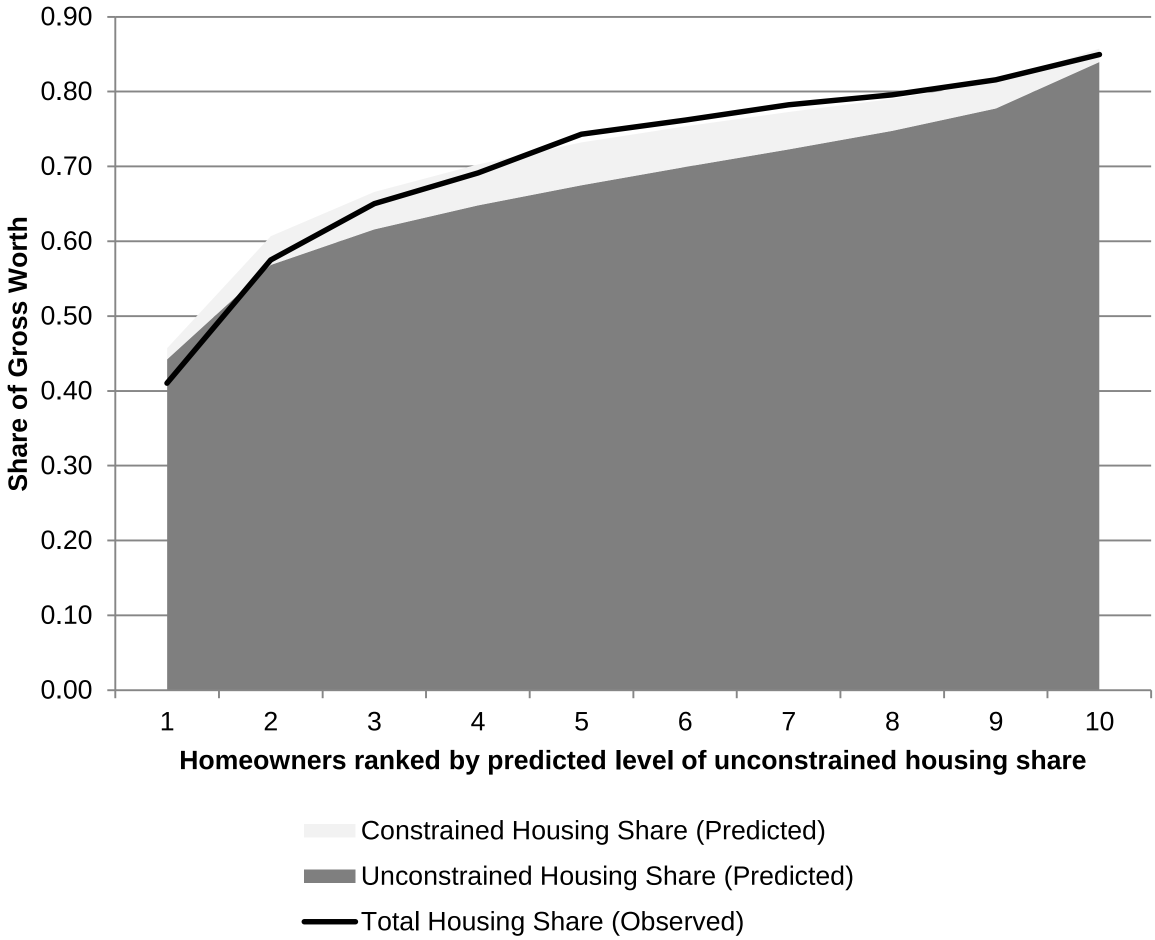

Figure 6
Housing Share Decomposition among Homeowners
Sorted on their Constrained Housing Share

The figure illustrates the predicted unconstrained housing shares (solid dark grey), constrained housing shares (solid light grey), and observed housing shares (solid black line) across the subsample of homeowners. Homeowners are sorted into 11 groups: the first group (#0) for which the constrained housing share is zero, and ten equally-formed groups for which the constrained housing share is positive.

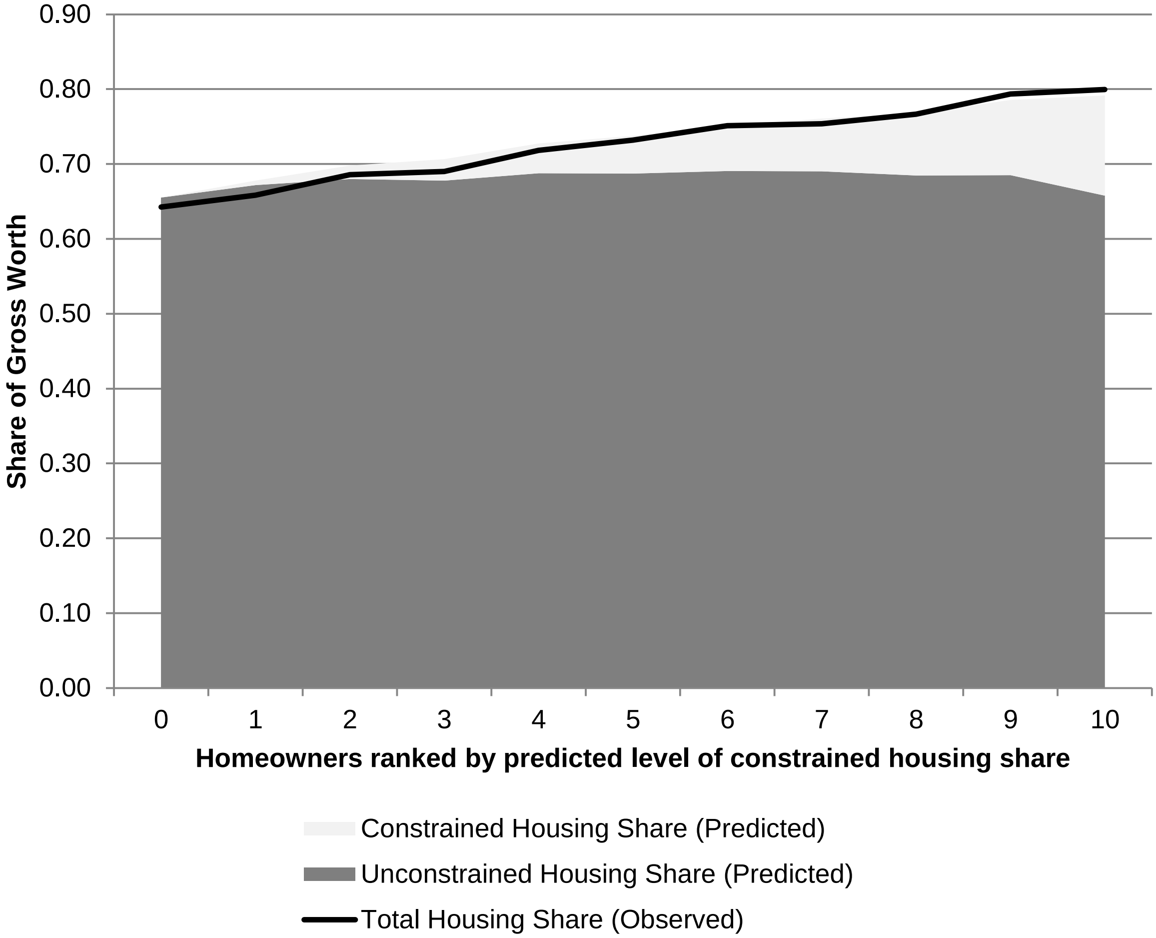

Internet Appendix for "Why do Homeowners Invest the Bulk of their Wealth in their Home?"

Laurent Barras and Sebastien Betermier*

March 18, 2020

This Internet Appendix presents a full derivation of the model, describes the econometric framework, and provides additional information on the empirical results.

*Desautels Faculty of Management, McGill, Bronfman Bdg, 1001 Sherbrooke St West, Montreal, QC H3A 1G5, Barras: laurent.barras@mcgill.ca, Betermier: sebastien.betermier@mcgill.ca

Contents

1 Derivation of the Portfolio Choice Model

1.1 Setup

We consider a dynamic model in which the agent consumes a basket of two goods: a perishable non-housing good (C_t) and a durable housing good (K_t). Here, a unit of housing is a one-dimensional summary of the quality of the house which accounts for size, location, and specific characteristics. At each time t, the utility derived by the agent is the Cobb-Douglas function

$$U(C_t, K_t) = \frac{1}{1-\gamma}\left(C_t^{\beta_C} K_t^{\beta_K}\right)^{1-\gamma}, \tag{IA-1}$$

where γ is the coefficient of relative risk aversion over the entire consumption basket and β_C and β_K are the relative importance of non-housing and housing consumption inside the agent's consumption basket (with $\beta_C + \beta_K = 1$). If we further assume that the utility stream is additively separable and the time horizon infinite, the agent's lifetime expected utility is given by

$$E\left[\int_0^\infty e^{-\delta s} U(C_s, K_s) ds\right], \tag{IA-2}$$

where δ is the time discount factor.

The agent can invest in a short-term risk-free bond and a risky stock fund without any short-sales restrictions. Using the non-housing good as numeraire, we denote by r the constant risk-free rate of the bond and write the dynamics of the stock price $P_{S,t}$ as

$$\frac{dP_{S,t}}{P_{S,t}} = \mu_S\, dt + \sigma_S\, dZ_{S,t}, \tag{IA-3}$$

where μ_S and σ_S are constants, and $Z_{S,t}$ is a Wiener process. The agent can also invest in H_t units of housing whose price $P_{H,t}$ is stochastic,

$$\frac{dP_{H,t}}{P_{H,t}} = \mu_H\, dt + \sigma_H\, dZ_{H,t}, \tag{IA-4}$$

where μ_H, and σ_H are constants, and $Z_{H,t}$ is a Wiener process that is uncorrelated with $Z_{S,t}$. We assume that, as a durable good, housing does not depreciate over time so that $dH_t = 0$ if the agent does not trade. According to this specification, the expected return on housing can be viewed as net of maintenance costs.

The agent can bridge the difference between H_t and K_t via the rental market. We assume a constant rent-to-price ratio ρ, which means that the price of renting one unit of housing is $\rho\, P_{H,t}$. Consequently, the net amount spent on the rent is $(K_t - H_t)\,\rho\, P_{H,t}$. We define the ownership ratio $\phi_{H,t}$ as the fraction of the home owned by the agent, i.e., $\phi_{H,t} = H_t/K_t$. The agent is (i) a renter if $\phi_{H,t} = 0$, (ii) a partial homeowner if $0 < \phi_{H,t} < 1$, (iii) a homeowner if $\phi_{H,t} = 1$, and (iv) a landlord if $\phi_{H,t} > 1$.

Combining this information, we can write the dynamics of the wealth process as

$$\begin{aligned} dW_t = &\Big[rW_t + \Theta_{S,t}(\mu_S - r) + H_t P_{H,t}(\mu_H - r) - C_t - (K_t - H_t)\,\rho P_{H,t}\Big]\,dt \\ &+ \Theta_{S,t}\,\sigma_S\, dZ_{S,t} + H_t P_{H,t}\,\sigma_H\, dZ_{H,t} \end{aligned} \tag{IA-5}$$

where $\Theta_{S,t}$ is the total amount invested in the stock. In terms of the Sharpe ratios $\lambda_S = \frac{\mu_S - r}{\sigma_S}$ and $\lambda_H = \frac{\mu_H - r}{\sigma_H}$, Equation (IA-5) becomes

$$\begin{aligned} dW_t = &\Big[rW_t + \Theta_{S,t}\lambda_{S,t}\sigma_{S,t} + H_t P_{H,t}\lambda_{H,t}\sigma_{H,t} - C_t - (K_t - H_t)\,\rho P_{H,t}\Big]\,dt \\ &+ \Theta_{S,t}\,\sigma_S\, dZ_{S,t} + H_t P_{H,t}\,\sigma_H\, dZ_{H,t}. \end{aligned} \tag{IA-6}$$

1.2 Format of the Solution

The infinite-horizon setting is convenient because the portfolio decisions depend on preferences and state variables but not on time. For notational convenience, we therefore eliminate the time subscript from now on. At time t, the agent's value function V is given as

$$V(W, P_H) = \sup_{\Gamma} E\left[\int_0^\infty e^{-\delta s} U(C_\tau, K_\tau) ds, \right] \tag{IA-7}$$

where $\Gamma = \{C, K, \Theta_S, H\}$ is the set of admissible controls. The Hamilton-Jacobi-Bellman (HJB) equation of this problem can be written as

$$\delta V(W, P_H) = \sup_{\Gamma} \Big[U(C, K) + E\left[dV(W, P_H)\right]\Big]. \tag{IA-8}$$

We begin by expanding (IA-8) as

$$\begin{aligned} \delta V = \sup_{\Gamma} &\Big[U(C, K) + V_W E(dW) + \frac{1}{2} V_{WW} E(dW)^2 \\ &+ V_{P_H} E(dP_H) + \frac{1}{2} V_{P_H P_H} E(dP_H)^2 + V_{W P_H} E(dW\, dP_H)\Big]. \end{aligned} \tag{IA-9}$$

We then plug in the processes for dW and dP_H:

$$\delta V = \sup_{\Gamma} \Big[U(C,K) + V_W \left(rW + \Theta_S \lambda_S \sigma_S + H P_H \lambda_H \sigma_H - C - K\rho P_H \right) + \frac{1}{2} V_{WW} \left(\Theta_S^2 \sigma_S^2 + H^2 P_H^2 \sigma_H^2 \right) + V_{P_H} \mu_H P_H + \frac{1}{2} V_{P_H P_H} \sigma_H^2 P_H^2 + V_{W P_H} \sigma_H^2 H P_H^2 \Big]. \quad \text{(IA-10)}$$

Building on the work of Damgaard, Fuglsbjerg and Munk (2003), we guess that the value function has the form

$$V(W, P_H) = \frac{1}{1-\gamma} \kappa P_H^{-\beta_K(1-\gamma)} W^{1-\gamma}. \quad \text{(IA-11)}$$

where κ is a constant. Since V is homogeneous, we can write $V(W, P_H) = P_H^{\beta_C(1-\gamma)} v(\tilde{W})$, where $v(\tilde{W}) = \frac{1}{1-\gamma} \kappa \tilde{W}^{1-\gamma}$, and $\tilde{W} = W/P_H$. Therefore, all the derivatives of $V(W, P_H)$ can be re-expressed in terms of $v(\tilde{W}) = v$, P_H, and $\tilde{W}$:

$$V_W = P_H^{\beta_C(1-\gamma)-1} v_{\tilde{W}}, \qquad V_{WW} = P_H^{\beta_C(1-\gamma)-2} v_{\tilde{W}\tilde{W}}, \quad \text{(IA-12)}$$

$$V_{P_H} = P_H^{\beta_C(1-\gamma)-1} v_{P_H}, \qquad V_{P_H P_H} = P_H^{\beta_C(1-\gamma)-2} v_{P_H P_H}, \quad \text{(IA-13)}$$

$$V_{W P_H} = P_H^{\beta_C(1-\gamma)-1} v_{\tilde{W} P_H}, \quad \text{(IA-14)}$$

where $v_{\tilde{W}}$, $v_{\tilde{W}\tilde{W}}$, v_{P_H}, $v_{P_H P_H}$, and $v_{\tilde{W} P_H}$ are given as

$$v_{\tilde{W}} = \kappa \tilde{W}^{-\gamma}, \quad \text{(IA-15)}$$

$$v_{\tilde{W}\tilde{W}} = -\gamma \kappa \tilde{W}^{-\gamma-1}, \quad \text{(IA-16)}$$

$$v_{P_H} = \beta_C(1-\gamma) v - \tilde{W} v_{\tilde{W}}, \quad \text{(IA-17)}$$

$$v_{P_H P_H} = (\beta_C(1-\gamma) - 1)\beta_C(1-\gamma) v - 2(\beta_C(1-\gamma) - 1)\tilde{W} v_{\tilde{W}} + \tilde{W}^2 v_{\tilde{W}\tilde{W}}, \quad \text{(IA-18)}$$

$$v_{\tilde{W} P_H} = (\beta_C(1-\gamma) - 1) v_{\tilde{W}} - \tilde{W} v_{\tilde{W}\tilde{W}}. \quad \text{(IA-19)}$$

The HJB Equation (IA-10) becomes easier to work with once we apply a couple of changes of variables. First, we define $\tilde{C} = C/P_H$ and $\tilde{\Theta}_S = \Theta_S/P_H$. Once we apply this change of variables to (IA-10), we can eliminate the term $P_H^{\beta_C(1-\gamma)}$ to get a modified HJB Equation that is independent of P_H:

$$\delta v = \sup_{\Gamma} \Bigg[U(\tilde{C}, K) + v_{\tilde{W}} \left(r\tilde{W} + \tilde{\Theta}_S \lambda_S \sigma_S + H \lambda_H \sigma_H - \tilde{C} - K\rho \right) + \frac{1}{2} v_{\tilde{W}\tilde{W}} \left(\tilde{\Theta}_S^2 \sigma_S^2 + H^2 \sigma_H^2 \right) + v_{P_H} \mu_H + \frac{1}{2} v_{P_H P_H} \sigma_H^2 + v_{\tilde{W} P_H} H \sigma_H^2 \Bigg]. \tag{IA-20}$$

We then apply a second change of variables in order to work with portfolio shares. Let $\alpha_C = C/W$, $\alpha_K = KP_H/W$ denote the values of non-housing and housing consumption relative to W, and $\alpha_S = \Theta_S/W$ and $\alpha_H = HP_H/W$ denote the shares invested in stocks and housing. It follows that $\alpha_C = \tilde{C}/\tilde{W}$, $\alpha_K = K/\tilde{W}$, $\alpha_S = \tilde{\Theta}_S/\tilde{W}$, and $\alpha_H = H/\tilde{W}$. The HJB Equation (IA-20) thus becomes

$$\delta v = \sup_{\Gamma} \Bigg[\tilde{W}^{1-\gamma} U(\alpha_C, \alpha_K) + v_{\tilde{W}} \tilde{W} \left(r + \alpha_S \lambda_S \sigma_S + \alpha_H \lambda_H \sigma_H - \alpha_C - \alpha_K \rho \right) + \frac{1}{2} v_{\tilde{W}\tilde{W}} \tilde{W}^2 \left(\alpha_S^2 \sigma_S^2 + \alpha_H^2 \sigma_H^2 \right) + v_{P_H} \mu_H + \frac{1}{2} v_{P_H P_H} \sigma_H^2 + v_{\tilde{W} P_H} \tilde{W} \alpha_H \sigma_H^2 \Bigg]. \tag{IA-21}$$

1.3 The Unconstrained Problem

We begin our analysis by solving the simplest portfolio problem in which the agent is free to be a partial owner of her home (i.e., ϕ_H can take any value). This allows us to prove Proposition 4 in the main text. To this end, we derive the optimal decisions of the agent using the full choice set $\Gamma = \{C, K, \Theta_S, H\}$. The optimal shares and the value function v are denoted by the superscript U (for unconstrained):

$$v^U = \frac{1}{1-\gamma} \kappa^U \tilde{W}^{1-\gamma}. \tag{IA-22}$$

The first order conditions with respect to α_C, α_K, α_S, and α_H are

$$U_{\alpha_C^U} = \beta_C \left(\alpha_C^U\right)^{\beta_C(1-\gamma)-1} \left(\alpha_K^U\right)^{\beta_K(1-\gamma)} = v_{\tilde{W}}^U \tilde{W}^\gamma = \kappa^U, \tag{IA-23}$$

$$U_{\alpha_K^U} = \beta_K \left(\alpha_C^U\right)^{\beta_C(1-\gamma)} \left(\alpha_K^U\right)^{\beta_K(1-\gamma)-1} = \rho v_{\tilde{W}}^U \tilde{W}^\gamma = \rho\kappa^U, \tag{IA-24}$$

$$\alpha_S^U = -\frac{v_{\tilde{W}}^U}{v_{\tilde{W}\tilde{W}}^U \tilde{W}} \frac{\lambda_S}{\sigma_S} = \frac{\lambda_S}{\gamma\sigma_S}, \tag{IA-25}$$

$$\begin{aligned} \alpha_H^U &= -\frac{v_{\tilde{W}}^U}{v_{\tilde{W}\tilde{W}}^U \tilde{W}} \frac{\lambda_H}{\sigma_H} - \frac{v_{\tilde{W}P_H}^U}{v_{\tilde{W}\tilde{W}}^U \tilde{W}} \\ &= \frac{\lambda_H}{\gamma\sigma_H} + \beta_K\left(1 - \frac{1}{\gamma}\right) = \frac{\bar{\lambda}_H}{\gamma\sigma_H}, \end{aligned} \tag{IA-26}$$

where $\bar{\lambda}_H = \lambda_H + \beta_K(\gamma - 1)\sigma_H$ is a modified Sharpe ratio that takes into account the additional hedging value of housing. Equation (IA-26) corresponds to Equation (19) in the main text.

We can also solve for the optimal levels of consumption for the non-housing and housing goods. After merging Equations (IA-23) and (IA-24), we obtain

$$\frac{\alpha_C^U}{\alpha_K^U} = \frac{\rho\beta_C}{\beta_K}, \tag{IA-27}$$

which implies that the ratio of non-housing consumption to housing consumption only depends on the rental price of housing. From Equation (IA-27), we can then re-express α_C^U and α_K^U as

$$\alpha_C^U = \beta_C \left(v_{\tilde{W}}^U\right)^{-\frac{1}{\gamma}} \tilde{W}^{-1}\epsilon = \beta_C(\kappa^U)^{-\frac{1}{\gamma}}\epsilon, \tag{IA-28}$$

$$\alpha_K^U = \frac{\beta_K}{\rho} \left(v_{\tilde{W}}^U\right)^{-\frac{1}{\gamma}} \tilde{W}^{-1}\epsilon = \frac{\beta_K}{\rho}(\kappa^U)^{-\frac{1}{\gamma}}\epsilon, \tag{IA-29}$$

where $\epsilon = \beta_C^{\frac{1-\gamma}{\gamma}} \left(\frac{\rho\beta_C}{\beta_K}\right)^{-\frac{\beta_K(1-\gamma)}{\gamma}}$. The value of κ^U can be found by inserting the optimal controls into the HJB Equation (IA-21),

$$\begin{aligned} 0 = &-v^U \left(\delta - r - \mu_H\beta_C(1-\gamma)\bar{\beta} - \frac{1}{2}\sigma_H^2\beta_C(1-\gamma)\left(\beta_C(1-\gamma) - 1\right)\right) \\ &+ \frac{\gamma}{1-\gamma}\left(v_{\tilde{W}}^U\right)^{\frac{\gamma-1}{\gamma}}\epsilon + v_{\tilde{W}}^U \tilde{W}(r - \mu_H) - \frac{1}{2}\frac{\left(v_{\tilde{W}}^U\right)^2}{v_{\tilde{W}\tilde{W}}^U}(\lambda_S^2 + \bar{\lambda}_H^2). \end{aligned} \tag{IA-30}$$

If we divide Equation (IA-30) by $\tilde{W}^{1-\gamma}$ and $\frac{\gamma}{1-\gamma}$, the $\tilde{W}$ term cancels out, and the value of κ^U can be written as

$$\kappa^U = \left(\frac{\epsilon}{\frac{\delta}{\gamma} - \frac{1-\gamma}{\gamma}\bar{r} - \frac{1-\gamma}{2\gamma^2}(\lambda_S^2 + \bar{\lambda}_H^2)}\right)^\gamma, \tag{IA-31}$$

where $\bar{r} = r - \mu_H \beta_K + \frac{1}{2}\sigma_H^2 \left((\beta_C(1-\gamma)-1)(\beta_C-2)-\gamma\right)$. Given (IA-31), the consumption levels of consumption for the non-housing and housing goods reduce to

$$\alpha_C^U = \beta_C \left(\frac{\delta}{\gamma} - \frac{1-\gamma}{\gamma}\bar{r} - \frac{1-\gamma}{2\gamma^2}(\lambda_S^2 + \bar{\lambda}_H^2)\right), \tag{IA-32}$$

$$\alpha_K^U = \frac{\beta_K}{\rho} \left(\frac{\delta}{\gamma} - \frac{1-\gamma}{\gamma}\bar{r} - \frac{1-\gamma}{2\gamma^2}(\lambda_S^2 + \bar{\lambda}_H^2)\right). \tag{IA-33}$$

The optimal ownership ratio is therefore equal to

$$\phi_H^U = \frac{\alpha_H^U}{\alpha_K^U} = \frac{\frac{\bar{\lambda}_H}{\gamma \sigma_H}}{\frac{\beta_K}{\rho}\left(\frac{\delta}{\gamma} - \frac{1-\gamma}{\gamma}\bar{r} - \frac{1-\gamma}{2\gamma^2}(\lambda_S^2 + \bar{\lambda}_H^2)\right)}. \tag{IA-34}$$

1.4 The Homeowner's Problem

1.4.1 Solution

Next, we solve for the homeowner's problem by imposing the constraint that the agent must fully own the home she lives in (i.e., ϕ_H must be equal to one). This analysis allows us to prove Proposition 1 in the main text. To this end, we assume that the agent only has access to the reduced choice set $\Gamma = \{C, K, \Theta_S\}$, and that H must be equal to K. To simplify notation, we do not use any superscript to denote the homeowner's optimal decisions and value function

$$v = \frac{1}{1-\gamma}\kappa \tilde{W}^{1-\gamma}. \tag{IA-35}$$

We denote by ξ the constant that links the value functions of the unconstrained agent and the homeowner, i.e. $\xi = \frac{\kappa}{\kappa^U}$. In other words, v is equal to ξv^U.

We derive a formulation of the solution that allows us to identify the additional impact of homeownership constraint ($\phi_H = 1$) on the agent's optimal decisions.[1] Each control variable

[1] Our derivation and solution to the homeowner's optimization problem differ from the approach proposed by Damgaard, Fuglsbjerg and Munk (2003) which does not distinguish between the unconstrained and constrained components of the homeowner's housing share.

can be written as:

$$\alpha_C = \alpha_C^U + \alpha_C^C, \tag{IA-36}$$

$$\alpha_K = \alpha_K^U + \alpha_K^C, \tag{IA-37}$$

$$\alpha_S = \alpha_S^U + \alpha_S^C, \tag{IA-38}$$

$$\alpha_H = \alpha_K = \alpha_H^U + \alpha_H^C, \tag{IA-39}$$

where the term with superscript U corresponds to the decision that the homeowner would take in an unconstrained setting (where ϕ_H can be freely chosen), and the term with superscript C corresponds to the additional impact of the homeownership constraint on the homeowner's decision.

From the previous section, we already know the solutions for the unconstrained decisions. Therefore, solving for the homeowner's optimal decisions is equivalent to solving for the constrained components α_C^C, α_K^C, α_S^C, and α_H^C. In fact, as we will see below, the full problem boils down to solving for α_H^C.

From the HJB Equation (IA-21), we first derive the first order conditions with respect to α_C, α_K and α_S as

$$U_{\alpha_C} = \beta_C \left(\alpha_C\right)^{\beta_C(1-\gamma)-1} \left(\alpha_K\right)^{\beta_K(1-\gamma)} = v_{\tilde{W}} \tilde{W}^{\gamma} = \kappa, \tag{IA-40}$$

$$\alpha_K = \alpha_H = \alpha_H^U - \frac{1}{v_{\tilde{W}\tilde{W}} \tilde{W}^2 \sigma_H^2} \left(\tilde{W}^{1-\gamma} U_{\alpha_K} - \rho v_{\tilde{W}} \tilde{W} \right), \tag{IA-41}$$

$$\alpha_S = -\frac{v_{\tilde{W}}}{v_{\tilde{W}\tilde{W}} \tilde{W}} \frac{\lambda_S}{\sigma_S} = \alpha_S^U. \tag{IA-42}$$

The last equation implies that, in our setting, the homeownership constraint has no impact on the optimal level of stock investment (e.g., $\alpha_S^C = 0$).[2]

Merging Equations (IA-40) and (IA-41), we obtain the consumption ratio of non-housing consumption to housing consumption as

$$\frac{\alpha_C}{\alpha_K} = \frac{\alpha_C^U}{\alpha_K^U} - \frac{\beta_C}{\beta_K} \frac{v_{\tilde{W}\tilde{W}} \tilde{W}}{v_{\tilde{W}}} \sigma_H^2 \alpha_H^C, \tag{IA-43}$$

$$= \frac{\beta_C}{\beta_K} \left(\rho + \gamma \sigma_H^2 \alpha_H^C \right), \tag{IA-44}$$

[2]This result comes from our assumption that the stock and housing prices are uncorrelated. If the correlation is positive, then the homeownership constraint leads to a crowding out effect on the investment in stocks.

We then plug Equation (IA-43) into the HJB Equation (IA-21) to obtain

$$\begin{aligned}\delta v = & \tilde{W}^{1-\gamma} U(\alpha_C, \alpha_K) + v_{\tilde{W}} \tilde{W} r - v_{\tilde{W}} \tilde{W} \left(\alpha_C + \alpha_K\right) \\ & - \frac{1}{2} v_{\tilde{W}\tilde{W}} \tilde{W}^2 \left[(\alpha_S^U)^2 \sigma_S^2 + (\alpha_H^U)^2 \sigma_H^2\right] + \frac{1}{2} v_{\tilde{W}\tilde{W}} \tilde{W}^2 \left[(\alpha_S^C)^2 \sigma_S^2 + (\alpha_H^C)^2 \sigma_H^2\right] \\ & + v_{P_H} \mu_H + \frac{1}{2} v_{P_H P_H} \sigma_H^2. \end{aligned} \tag{IA-45}$$

We can split the right-hand side of Equation (IA-45) into two terms:

$$\delta v = T_1 + T_2, \tag{IA-46}$$

where T_1 only includes the optimal controls of the unconstrained agent, and T_2 includes all the remaining elements. Both terms are given by

$$\begin{aligned} T_1 = & \tilde{W}^{1-\gamma} \xi U\left(\alpha_C^U, \alpha_K^U\right) + \xi v_{\tilde{W}}^U \tilde{W} r - \xi v_{\tilde{W}}^U \tilde{W}\left(\alpha_C^U + \alpha_K^U \rho\right) \\ & - \frac{1}{2} \xi v_{\tilde{W}\tilde{W}}^U \tilde{W}^2 \left[(\alpha_S^U)^2 \sigma_S^2 + (\alpha_H^U)^2 \sigma_H^2\right] + \xi v_{P_H}^U \mu_H + \frac{1}{2} \xi v_{P_H P_H}^U \sigma_H^2, \end{aligned} \tag{IA-47}$$

and

$$\begin{aligned} T_2 = & \tilde{W}^{1-\gamma}\left[U\left(\alpha_C, \alpha_K\right) - \xi U\left(\alpha_C^U, \alpha_K^U\right)\right] \\ & - v_{\tilde{W}} \tilde{W}\left[\alpha_C^C + \alpha_K^C \rho\right] + \frac{1}{2} v_{\tilde{W}\tilde{W}} \tilde{W}^2 \left[(\alpha_S^C)^2 \sigma_S^2 + (\alpha_H^C)^2 \sigma_H^2\right]. \end{aligned} \tag{IA-48}$$

We know that (i) $v = \xi v^U$ and (ii) $T_1/\xi = \delta v^U$ as per the HJB Equation (IA-21) in the unconstrained setting. This implies that

$$\delta v = T_1, \tag{IA-49}$$

$$T_2 = 0. \tag{IA-50}$$

We can now solve for $T_2 = 0$ to obtain the solution for α_H^C. Inserting the first-order conditions in Equations (IA-40)-(IA-42) into Equation (IA-50), and then dividing all the terms by $\tilde{W}^{1-\gamma}$ and κ, we obtain the following quadratic equation:

$$A\left(\alpha_H^C\right)^2 + B\,\alpha_H^C + C = 0, \tag{IA-51}$$

where the parameters A, B, and C are defined as

$$A = \frac{1}{2}\left[\frac{\gamma \sigma_H^2}{\rho} + \psi\right], \quad B = 1 + \psi \alpha_H^U, \quad C = \alpha_K^U - \alpha_H^U, \tag{IA-52}$$

and $\psi = \frac{\sigma_H^2}{\rho}(1-\beta_C(1-\gamma)) = \frac{\sigma_H^2}{\rho}(\beta_K(1-\gamma)+\gamma)$. Equation (IA-51) corresponds to Equation (13) in the main text.

The existence of a solution requires that $B^2 - 4A\,C \geq 0$. Only one root leads to positive values for α_C and α_K:

$$\alpha_H^C = \frac{-2C}{B + \sqrt{B^2 - 4A\,C}}. \tag{IA-53}$$

Together with the set of unconstrained decisions derived in the previous section, the solution to α_H^C determines: (i) the optimal level of housing investment α_H (from Equation (IA-39)), (ii) housing consumption α_K (from Equations (IA-37) and (IA-39)), and (iii) non-housing consumption α_C (from Equation (IA-43)).[3] Finally, the proportionality constant ξ can be found as the ratio of the marginal utilities over non-housing consumption in the unconstrained and constrained cases,

$$\xi = \frac{v_{\tilde{W}}}{v_{\tilde{W}}^U} = \left(\frac{\alpha_C}{\alpha_C^U}\right)^{\beta_C(1-\gamma)-1} \left(\frac{\alpha_K}{\alpha_K^U}\right)^{\beta_K(1-\gamma)}. \tag{IA-54}$$

1.4.2 Implications for the homeowner's UIP_H

Building on the solution presented above, we now examine its implications for the homeowners' UIP_H. We show that the impact of the homeownership constraint on the homeowner's housing share is muted because of the downward adjustment in housing consumption. This analysis allows us to complete the proof of Proposition 1 in the main text.

We consider the optimal value of the constrained component α_H^C as given by Equation (IA-53). The term A is strictly positive and the term C is strictly negative when the homeownership constraint binds. Consequently, the term $B^2 - 4A\,C$ inside the square root must exceed B^2, which yields the following upper bound for α_H^C:

$$\alpha_H^C < -\frac{C}{B} = -\frac{\alpha_K^U - \alpha_H^U}{B} \tag{IA-55}$$

From Equation (IA-55) we then obtain an upper bound for the housing share α_H:

$$\alpha_H = \alpha_H^U + \alpha_H^C \leq \alpha_H^U + \frac{\alpha_K^U - \alpha_H^U}{B}, \tag{IA-56}$$

[3] If the unconstrained homeownership ratio ϕ_H^U is exactly equal to one (i.e., homeownership is optimal), then the term C is equal to zero and we have $\alpha_C = 0$. In this special case, the homeowner's decisions therefore all correspond to those of the unconstrained agent.

or equivalently,

$$\alpha_H^U \leq \alpha_H \leq \alpha_K^U. \tag{IA-57}$$

Equation (IA-56) corresponds to Equation (14) in the main text.

1.5 The Renter's Problem

1.5.1 Solution

In our final analysis, we solve for the renter's problem by imposing the constraint that the agent cannot invest in housing (i.e., $\phi_H = 0$ must be equal to zero). This analysis allows us to prove Proposition 2 in the main text. To this end, we assume that the agent only has access to the reduced choice set $\Gamma = \{C, K, \Theta_S\}$, and that H must be equal to 0. The optimal shares and the value function v are denoted by the superscript R (for renter):

$$v^R = \frac{1}{1-\gamma}\kappa^R \tilde{W}^{1-\gamma}. \tag{IA-58}$$

The first order conditions with respect to α_C, α_K and α_S are

$$U_{\alpha_C^R} = \beta_C \left(\alpha_C^R\right)^{\beta_C(1-\gamma)-1} \left(\alpha_K^R\right)^{\beta_K(1-\gamma)} = v_{\tilde{W}}^R \tilde{W}^{\gamma} = \kappa^R, \tag{IA-59}$$

$$U_{\alpha_K^R} = \beta_K \left(\alpha_C^R\right)^{\beta_C(1-\gamma)} \left(\alpha_K^R\right)^{\beta_K(1-\gamma)-1} = \rho v_{\tilde{W}}^R \tilde{W}^{\gamma} = \rho\kappa^R, \tag{IA-60}$$

$$\alpha_S^R = -\frac{v_{\tilde{W}}^R}{v_{\tilde{W}\tilde{W}}^R \tilde{W}}\frac{\lambda_S}{\sigma_S} = \alpha_S^U. \tag{IA-61}$$

Similar to the unconstrained case, we can combine Equations (IA-59) and (IA-60) to obtain

$$\frac{\alpha_C^R}{\alpha_K^R} = \frac{\rho\beta_C}{\beta_K}, \tag{IA-62}$$

where α_C^R and α_K^R can be written as

$$\alpha_C^R = \beta_C \left(v_{\tilde{W}}^R\right)^{-\frac{1}{\gamma}} \tilde{W}^{-1}\epsilon = \beta_C(\kappa^R)^{-\frac{1}{\gamma}}\epsilon, \tag{IA-63}$$

$$\alpha_K^R = \frac{1}{\rho}(\beta_K) \left(v_{\tilde{W}}^R\right)^{-\frac{1}{\gamma}} \tilde{W}^{-1}\epsilon = \frac{\beta_K}{\rho}(\kappa^R)^{-\frac{1}{\gamma}}\epsilon. \tag{IA-64}$$

Inserting these optimal controls into the HJB Equation (IA-21), we obtain

$$\begin{aligned} 0 = &-v^R \left(\delta - r - \mu_H\beta_C(1-\gamma) - \frac{1}{2}\sigma_H^2\beta_C(1-\gamma)\left(\beta_C(1-\gamma)-1\right)\right) \\ &+ \frac{\gamma}{1-\gamma}\left(v_{\tilde{W}}^R\right)^{\frac{\gamma-1}{\gamma}}\epsilon + v_{\tilde{W}}^R \tilde{W}(r-\mu_H) - \frac{1}{2}\frac{\left(v_{\tilde{W}}^R\right)^2}{v_{\tilde{W}\tilde{W}}^R}\lambda_S^2. \end{aligned} \tag{IA-65}$$

It follows that κ^R satisfies

$$\kappa^R = \left(\frac{\epsilon}{\frac{\delta}{\gamma} - \frac{1-\gamma}{\gamma}\bar{r} - \frac{1-\gamma}{2\gamma^2}\lambda_S^2} \right)^{\gamma}. \tag{IA-66}$$

The key difference between the renter and the unconstrained agent is that the renter does not have access to the housing investment asset. Therefore, the augmented Sharpe ratio of housing $\bar{\lambda}_H$ does not appear in Equation (IA-66). The absence of housing leads to the following adjustment in the levels of housing and non-housing consumption (relative to the unconstrained case),

$$\alpha_C^R = \alpha_C^U + \beta_C \frac{(1-\gamma)}{2\gamma^2}\bar{\lambda}_H^2, \tag{IA-67}$$

$$\alpha_K^R = \alpha_K^U + \frac{\beta_K}{\rho} \frac{(1-\gamma)}{2\gamma^2}\bar{\lambda}_H^2. \tag{IA-68}$$

1.5.2 Implications for the homeowner's UIP_H

We know examine the implications of the solution to the renter's problem for the homeowner's UIP_H. As explained in Section 2.2 of the main text, homeownership is the outcome of a constrained choice which includes renting as an alternative option. Therefore, the agent chooses homeownership only if the value function she derives from owning dominates the value function she derives from renting. Building on this intuition, we explore how the option to rent imposes a lower bound on the homeowner's UIP_H.

Indifference between owning and renting To begin, we consider the threshold case where the agent is indifferent between owning and renting (i.e., $v^R = v$). This condition holds when $\kappa^R = \kappa$ or, alternatively, when $U_{\alpha_C^R} = U_{\alpha_C}$. Expanding the formulations of the marginal utilities yields the following expression:

$$\beta_C \left(\alpha_K^R\right)^{-\gamma} \left(\frac{\alpha_C^R}{\alpha_K^R}\right)^{\beta_C(1-\gamma)-1} = \beta_C \left(\alpha_K\right)^{-\gamma} \left(\frac{\alpha_C}{\alpha_K}\right)^{\beta_C(1-\gamma)-1}, \tag{IA-69}$$

or in log forms:

$$-\gamma \left[\log(\alpha_K^R) - \log(\alpha_K)\right] = \left(\beta_C\left(1-\gamma\right)-1\right)\left[\log\left(\frac{\alpha_C}{\alpha_K}\right) - \log\left(\frac{\alpha_C^R}{\alpha_K^R}\right)\right]. \tag{IA-70}$$

To simplify the algebra, we use a first-order Taylor expansion of the log function around the points $\log(\alpha_K^R)$ and $\log(\alpha_C^R/\alpha_K^R)$ to obtain

$$\log(\alpha_K^R) - \log(\alpha_K) = \frac{\alpha_K^R - \alpha_K}{\alpha_K} + R_1, \tag{IA-71}$$

$$\log\left(\frac{\alpha_C}{\alpha_K}\right) - \log\left(\frac{\alpha_C^R}{\alpha_K^R}\right) = \frac{\frac{\alpha_C}{\alpha_K} - \frac{\alpha_C^R}{\alpha_K^R}}{\frac{\alpha_C^R}{\alpha_K^R}} + R_2. \tag{IA-72}$$

Both R_1 and R_2 are the remaining second-order terms that vanish as α_K and α_C get close to α_K^R and α_C^R, i.e., $R_1 = O(||n||^2)$ where $||n|| = \left((\alpha_K - \alpha_K^R)^2 + \left(\frac{\alpha_C}{\alpha_K} - \frac{\alpha_C^R}{\alpha_K^R}\right)^2\right)^{\frac{1}{2}}$.

Using Equations (IA-71) and (IA-72) along with the general notation R for the sum of second-order terms, we can rewrite Equation IA-70 as

$$-\gamma\left[\frac{\alpha_K^R - \alpha_K}{\alpha_K}\right] = (\beta_C(1-\gamma) - 1)\left(\frac{\frac{\alpha_C}{\alpha_K} - \frac{\alpha_C^R}{\alpha_K^R}}{\frac{\alpha_C^R}{\alpha_K^R}}\right) + R. \tag{IA-73}$$

We then plug Equations (IA-43) and (IA-62) into the RHS of Equation (IA-73), make use of the homeownership constraint that $\alpha_H = \alpha_K$, and apply the decomposition $\alpha_K = \alpha_K^U + \alpha_K^C$ to obtain

$$(\alpha_K^R - \alpha_K^U) - \alpha_K^C = \psi\,\alpha_H^C\,\alpha_H + R. \tag{IA-74}$$

Finally, we insert Equation (IA-68) into the LHS of Equation (IA-77) to obtain

$$\frac{\beta_K}{\rho}\frac{(1-\gamma)}{2\gamma^2}\bar{\lambda}_H^2 = \alpha_K^C + \psi\alpha_H^C\alpha_H^U + \psi(\alpha_H^C)^2 + R \tag{IA-75}$$

Going back to the homeowner's problem, we know that the quadratic equation (IA-51) for α_H^C must hold. Equation (IA-51) can be rewritten as

$$\alpha_K^C + \psi\alpha_H^C\alpha_H^U + \psi(\alpha_H^C)^2 = \frac{\sigma_H^2}{2\rho}\beta_K(1-\gamma)(\alpha_H^C)^2 \tag{IA-76}$$

This last formulation is particularly convenient because it has a similar form to the condition in Equation (IA-73) for which renting and owning yield the same value function. Comparing Equations (IA-76) and (IA-73), we obtain the following relation:

$$\frac{\beta_K}{\rho}\frac{(1-\gamma)}{2\gamma^2}\bar{\lambda}_H^2 = \frac{\sigma_H^2}{2\rho}\beta_K(1-\gamma)(\alpha_H^C)^2 + R. \tag{IA-77}$$

Inserting the optimal value of α_H^U from Equation (IA-26) into Equation (IA-77), we obtain

$$\frac{\beta_K\,(1-\gamma)}{2\rho}(\alpha_H^U)^2 = \frac{\beta_K(1-\gamma)}{2\rho}(\alpha_H^C)^2 + R. \tag{IA-78}$$

Ignoring the second-order term R, Equation (IA-77) requires that $\alpha_H^C = \alpha_H^U$ or, equivalently, that

$$UIP_H = \frac{\alpha_H^U}{\alpha_H} = \frac{1}{2}. \tag{IA-79}$$

Although we cannot determine the sign of the approximation error R, there are two reasons why its magnitude should be small. First, the two error terms associated with R_1 to R_2 in Equations (IA-71) and (IA-72) have offsetting effects because they have the same signs. To this point, we note that R_1 and R_2 are negative because the log function is concave. In addition, both terms are multiplied by negative numbers equal to $-\gamma$ and $\beta_C\,(1-\gamma)-1$. Second, the approximation is small because the different terms are close to one another - that is, α_K^R, α_K, $\frac{\alpha_C^R}{\alpha_K^R}$ and $\frac{\alpha_C}{\alpha_K}$ are all between zero and one (α_K^R is typically greater than α_C^R because it captures the value of the house instead of the rent).

Conditions under which homeownership dominates renting We now examine the conditions under which owning dominates renting (i.e., $v > v^R$). In the case where $\gamma > 1$, the homeowner's value function exceeds that of the renter when $\kappa > \kappa^R$. The first order Taylor expansion in Equation (IA-78) becomes

$$\frac{\beta_K\,(1-\gamma)}{2\rho}(\alpha_H^U)^2 < \frac{\beta_K(1-\gamma)}{2\rho}(\alpha_H^C)^2, \tag{IA-80}$$

which holds only if $\alpha_H^U > \alpha_H^C$.

In the other case where $\gamma < 1$, the homeowner's value function exceeds that of the renter when $\kappa < \kappa^R$. Equation (IA-78) becomes

$$\frac{\beta_K\,(1-\gamma)}{2\rho}(\alpha_H^U)^2 > \frac{\beta_K(1-\gamma)}{2\rho}(\alpha_H^C)^2. \tag{IA-81}$$

which holds only if $\alpha_H^U > \alpha_H^C$. Equation (IA-81) corresponds to Equation (16) in the main text and thus completes the proof of Proposition 2.

2 Econometric Framework

2.1 Estimation Procedure

This section provides additional information on the procedure we use to estimate the coefficients of the unconstrained and constrained components of the housing share. The total sample contains information on a set of n households (homeowners and landlords) over T periods. For each household i ($i = 1, ..., n$) and each time period t ($t = 1, ..., T$), we write the housing share reported in the survey as

$$\alpha^S_{Hi,t} = \alpha^U_{Hi,t} + \alpha^C_{Hi,t} + e_{i,t} \tag{IA-82}$$

$$= a'_U y_t + s'_U x_{i,t} + max(0, s'_C (x_{i,t} - x_{L,t}))d_{Hi,t} + e_{i,t}, \tag{IA-83}$$

where y_t is a T-vector whose t^{th} element is one and the others are zero, $x_{i,t}$ is a K-vector of explanatory variables, $x_{L,t}$ is the K-vector of average characteristics across landlords, $d_{Hi,t}$ is an indicator function equal to one if household i is a homeowner (and zero otherwise), and $e_{i,t}$ is an error term. The T-vector $a_U = [a_{U,1}, ..., a_{U,T}]'$ contains the time-fixed effects, and s_U, s_C are the K-vectors of coefficients that measure the effect of $x_{i,t}$ on the unconstrained and constrained housing shares, respectively.

Equation (IA-83) can be written more compactly as $\alpha^S_{Hi,t} = a'_U y_t + s'_U x_{i,t} + \alpha^C_{Hi,t} + e_{i,t}$, where $\alpha^C_{Hi,t} = g(s_C, x_{i,t})d_{Hi,t}$ is the constrained housing share and h is a nonlinear function that depends on the vector $x_{i,t}$ and the coefficient vector s_C. Stacking together all the observations for household i, we obtain

$$\alpha^S_{Hi} = Y_i a_U + X_i s_U + \alpha^C_{Hi} + e_i, \tag{IA-84}$$

where $\alpha^S_{Hi} = [\alpha^S_{Hi,1}, ..., \alpha^S_{Hi,T_i}]'$ is the T_i-vector of reported housing shares, $Y_i = [y_1, ..., y_{T_i}]'$ is a $T_i \times T$ matrix that capture the time-fixed effects, $X_i = [x_{i,1}, ..., x_{i,T_i}]'$ is a $T_i \times K$ matrix of variables, $\alpha^C_{Hi} = [\alpha^C_{Hi,1}, ..., \alpha^C_{Hi,1}]'$ is the T_i-vector of the constrained housing share, and $e_i = [e_{i,1}, ..., e_{i,T_i}]'$ is the T_i-vector of error terms. The number of observations T_i is specific to each household and can be lower than the total number of observations T (i.e., the panel regression is potentially unbalanced).

To estimate the $(T + 2K)$-vector $\beta = (a'_U, s'_U, s'_C)'$, we apply the Nonlinear Least Squares

(NLS) method which minimizes the sum of squared residuals:

$$S(\beta) = \frac{1}{2}\sum_{i=1}^{n}(\alpha_{Hi}^{S} - Y_i a_U - X_i s_U - \alpha_{Hi}^{C})'(\alpha_{Hi}^{S} - Y_i a_U - X_i s_U - \alpha_{Hi}^{C}). \quad \text{(IA-85)}$$

The first-order condition for this minimization problem is given by

$$\sum_{i=1}^{n} Z_i' e_i = 0, \quad \text{(IA-86)}$$

where the $T_i \times (T+2K)$ matrix Z_i is defined as $[Y_i, X_i, X_i^0]$. The $T_i \times K$ matrix X_i^0 is defined as $[x_{i,1}^0, ..., x_{i,T_i}^0]'$, where $x_{i,t}^0$ is the K-vector of first order derivatives of the constrained housing share function $\alpha_{Hi,t}^C$ with respect to s_C, i.e., $x_{i,t}^0 = \frac{\partial \alpha_{Hi,t}^C}{\partial s_C}$. Equation (IA-86) shows that the orthogonality condition is similar to the one obtained in the classic Ordinary Least Squares (OLS) approach, except that the $T_i \times (T+2K)$ matrix Z_i includes the first order derivatives X_i^0 as regressors.

To smooth out the kink associated with the function $g_{i,t}$ and thus guarantee that the derivatives can be computed for all values of s_C, we follow Zang (1980) and approximate the function h around zero with a smooth function $g_s = g_s(s_C, x_{i,t})$:

$$g_s = \frac{1}{4}\frac{(s_C'(x_{i,t} - x_{L,t}))^2}{\lambda} + \frac{1}{2}(s_C'(x_{i,t} - x_{L,t})) + \frac{1}{4}\lambda \quad \text{if } -\lambda \leq s_C'(x_{i,t} - x_{L,t}) \leq \lambda, \quad \text{(IA-87)}$$

where $\lim_{\lambda \to 0} f_s = f$.

Contrary to standard OLS, there is no closed-form solution to Equation (IA-86). Therefore, we minimize the sum of squared residuals numerically using the standard Gauss-Newton algorithm (see Davidson and MacKinnon (2004) (ch. 6)) in which the parameter values are updated as follows:

$$\widehat{\beta}(s+1) = \widehat{\beta}(s) + \left(\sum_{i=1}^{n} Z_i'(s) Z_i(s)\right)^{-1}\left(\sum_{i=1}^{n} Z_i'(s) e_i(s)\right), \quad \text{(IA-88)}$$

where $\widehat{\beta}(s+1)$ and $\widehat{\beta}(s)$ are the estimated parameters at iterations $s+1$ and s, and $Z(s)$, $e(s)$ are computed based on $\widehat{\beta}(s)$. The iterating process ends when the orthogonality condition $\sum_{i=1}^{n} Z_i'(s) e_i(s)$ is sufficiently close to zero. To implement this algorithm, we need to determine the value of the initial vector $\widehat{\beta}(0)$. For the T-vector $\widehat{a}_U(0)$ and the K-vector $\widehat{s}_U(0)$, we fit a linear regression based on the data reported by the landlords. Similarly, the K-vector $\widehat{s}_C(0)$ is obtained from a linear regression of the estimated constrained housing share $\widehat{\alpha}^C_{Hi,t}(0)$ on the vector $\left(x_{i,t} - x^L_t\right))d_{Hi,t}$, where $\widehat{\alpha}^C_{Hi,t}(0)$ is defined as $\alpha^S_{Hi,t} - \widehat{a}_U(0)'y_t + \widehat{s}_U(0)'x_{i,t}$.[4]

2.2 Properties of the Estimators

To derive the large sample properties of the NLS estimator $\widehat{\beta}$, we assume that the regressors are well behaved in the sense that $\frac{1}{n}\sum_{i=1}^{n} Z_i'Z_i$ converges in probability to a positive definite matrix Q. Under certain regularity conditions on the function $S(\beta)$ discussed by Davidson and MacKinnon (2004) (ch. 6)) and Greene (2012) (ch. 7), we can show that the $(T+2K)$-vector of estimated parameters $\widehat{\beta}$ is consistent. This property essentially requires that the set of orthogonality conditions converges in probability towards zero, i.e.,

$$plim \frac{1}{n}\sum_{i=1}^{n} Z_i'e_i = 0. \tag{IA-89}$$

We can further show that $\widehat{\beta}$ is asymptotically normally distributed as the number of observations n grows large:

$$\sqrt{n}\left(\widehat{\beta} - \beta\right) \rightarrow N\left(0, Q^{-1}VQ^{-1}\right). \tag{IA-90}$$

The $(T+2K)\times(T+2K)$ matrix V allows for potential cross-dependence between the error terms associated with each household (i.e., observations are clustered at the household level):

$$V = \sum_{i=1}^{n} V_i = \sum_{i=1}^{n} Z_i'\Omega_i Z_i, \tag{IA-91}$$

where Ω_i denotes the covariance matrix of the error vector e_i. We estimate V using the cluster-robust estimator

$$\widehat{V} = \sum_{i=1}^{n} Z_i'\widehat{e}_i\widehat{e}_i'Z_i, \tag{IA-92}$$

where $\widehat{e}_i$ is the estimated error vector for household i. As shown by Liang and Zeger (1986), we can then plug $\widehat{V}$ in Equation (IA-92) to obtain a consistent estimator of the covariance matrix of $\widehat{\beta}$, i.e.,

$$plim \frac{1}{n}\left(\widehat{Q}^{-1}\widehat{V}\widehat{Q}^{-1}\right) = \frac{1}{n}\left(Q^{-1}VQ^{-1}\right), \tag{IA-93}$$

[4] We have also considered alternative sets of starting values in which $\widehat{s}_C(0)$ is equal to zero or equal to $\widehat{s}_U(0)$, and find that the estimated parameters converge to the same values.

where $\widehat{Q}$ is equal to $\frac{1}{n}\sum_{i=1}^{n} Z_i' Z_i$.

2.3 Hypothesis Tests

Given the asymptotic results in Equation (IA-90), we can conduct a hypothesis test to determine whether the K-vector of coefficients associated with the unconstrained or constrained housing share is equal to zero. To this end, we denote by M_K the $K \times (T+2K)$ matrix that impose the K restrictions implied on β (either $H_0 : s_U = 0$ or $H_0 : s_C = 0$). We can then compute the F-test statistic as

$$F = \frac{1}{K}(M_K\widehat{\beta})' \left(\frac{1}{n} M_K \widehat{Q}^{-1} \widehat{V} \widehat{Q}^{-1} M_K' \right) (M_K\widehat{\beta}), \tag{IA-94}$$

where the large-sample distribution of F is that of $\frac{1}{K}$ times a chi-squared variable with K degrees of freedom.

2.4 Alternative Specifications

2.4.1 Housing Frictions

Several housing frictions have an impact on the housing share, including borrowing constraints, trading costs, and costs to becoming landlords. Whereas these frictions have an impact on the unconstrained housing share, they are not explicitly modeled in our baseline specification in Equation (IA-83). Therefore, the omission of these frictions could potentially bias the estimated UIP.

To elaborate, suppose that the vector $x_{i,t}$ is a scalar and that the time-fixed effects are constant (i.e., $a_{U,t} = a_U$). In this case, the unconstrained and constrained housing share functions α_H^U and α_H^C can be written as

$$\alpha_{Hi,t}^U = a_U + s_U x_{i,t}, \tag{IA-95}$$

$$\alpha_{Hi,t}^C = max(0, s_C(x_{i,t} - x_{L,t}))d_{Hi,t}. \tag{IA-96}$$

Housing frictions affect the observed housing share $\alpha_{Hi,t}^S$ through the residual component of Equation (IA-83). The residual can be written,

$$e_{i,t} = F_{i,t} + \epsilon_{i,t}, \tag{IA-97}$$

where $F_{i,t}$ captures the impact of housing frictions.

Suppose that a homeowner wishes to rebalance its portfolio because the characteristic $x_{i,t}$ has changed. In a world without frictions, the observed housing share should be equal to the sum of the unconstrained and constrained components in Equation (IA-96) (i.e., $F_{i,t} = 0$). However, in a world with frictions, the homeowner may prefer a buy and hold strategy, in which case the housing share is equal to the optimal level determined at the last rebalancing date q. In this case, the friction term $F_{i,t}$ is equal to:

$$F_{i,t} = \alpha_{Hi,t} + d_{i,t}^{NO}(\alpha_{Hi,t-q} - \alpha_{Hi,t}), \tag{IA-98}$$

where $\alpha_{Hi,t} = \alpha_{Hi,t}^{U} + \alpha_{Hi,t}^{C}$, and $d_{i,t}^{NO}$ is a dummy variable equal to one if there is no rebalancing at time t, and zero otherwise.[5] Equation (IA-97) reveals that the orthogonality condition fails if $F_{i,t}$ is correlated with $z_{i,t} = [1, x_{i,t}, x_{i,t}^{0}]$. In this case, the estimated coefficients $\widehat{a}_U$, $\widehat{s}_U$, and $\widehat{s}_C$ are biased. This in turn distorts the estimated components of the housing shares and the level of the UIP_H.

There are several reasons why housing frictions are unlikely to drive the level of the estimated UIP. One reason is that the average impact of the omitted frictions $E(F_{i,t})$ is absorbed by the estimated intercept $\widehat{a}_U$ (e.g., see Greene (2012) (ch. 6)). Therefore, this standard result in regression analysis leaves the estimated unconstrained housing share largely unchanged (i.e., the average level of $\widehat{\alpha}_H^U$ should be similar with and without the housing frictions). Another reason is that the impact of the different costs (trading costs, costs to becoming landlords) cannot be too large. Portfolio choice models with transaction costs show that the probability of rebalancing increases as the housing share deviates from its target (Corradin, Fillat and Vergara-Alert, 2014; Grossmann and Laroque, 1990; Damgaard, Fuglsbjerg and Munk, 2003). From an econometric perspective, this means that $F_{i,t}$ is bounded as a large value for $\alpha_{Hi,t-q} - \alpha_{Hi,t}$ typically comes with a low value for $d_{i,t}^{NO}$.

To confirm these theoretical arguments, we re-restimate the model on a subset of homeowners for which the frictions do not apply (i.e., we condition the sample on $d_{i,t}^{NO} = 0$). To control for trading costs, we re-estimate our model only using the observations reported by households following a home purchase. For these specific observations, the reported housing

[5] For simplicity, we ignore the fact that the housing share in a buy and hold strategy is not exactly equal to $\alpha_{Hi,t-q}^{NO}$ because it also depends on the cumulative q-period returns of the different assets.

shares are consistent with our model because the portfolio has just been rebalanced. To control for the costs to becoming landlords, we re-estimate the model after replacing landlords with quasi-landlords. We define quasi-landlords as homeowners with similar characteristics as landlords based on a probit analysis.[6] The rationale behind this approach is that quasi-landlords remain unconstrained (like landlords), but have not paid the costs to becoming landlords. As discussed in the paper, the empirical results obtained with these alternative specifications remain largely unchanged.

2.4.2 Conservative Estimation

Apart from housing frictions, the estimated UIP could still be biased if we omit relevant characteristics in the vector $x_{i,t}$. In this case, the residual term $e_{i,t}$ contains information about the unconstrained and constrained housing shares, which may invalidate our empirical analysis of the UIP.

To address this issue, we propose a conservative estimation of the UIP, where we make the strong assumption that any unmodeled source of variation of the observed housing share is due to the homeownership constraint. Formally, we estimate the UIP of each homeowner as

$$\widehat{UIP}_{Hi} = \frac{\widehat{\alpha}_{Hi}^{U}}{\widehat{\alpha}_{Hi}^{U} + \widehat{\alpha}_{Hi}^{C,*}}. \tag{IA-99}$$

The new estimator of the constrained housing share $\widehat{\alpha}_{Hi}^{C,*}$ is defined as the maximum value between $\widehat{\alpha}_{Hi}^{C}$ and $\widehat{\alpha}_{Hi}^{C} + \widehat{e}_i$, where $\widehat{e}_i = \frac{1}{T_i}\sum_{t=1}^{T_i} \widehat{e}_{i,t}$. As discussed in the paper, we find that even in this worst-case scenario, the estimated UIP remains large.

3 Empirical Analysis

3.1 Datasets

This section provides additional information about the data used in the paper and the construction of variables. We use panel data from the Panel Study of Income Dynamics

[6] Specifically, we estimate a probit model for the decision to be landlord based on the vector $x_{i,t}$. We then form the group of quasi-landlords by taking the 20% of homeowners with the highest predicted score.

(PSID), a national survey of U.S. households that is widely used in the household finance literature. The survey, which now tracks more than 7,000 U.S. households, is conducted on an annual basis from 1968 to 1997 and on a bi-annual basis from 1997 to 2013. It contains information on the households' income, demographics (age, household size, marital status, high school or post-high school education), and real estate holdings (value of the house, mortgage, tenure choice). The PSID also provides a wealth survey that contains additional information on the households' wealth (e.g., holdings in stocks, bonds, and businesses). Prior to 1999, the wealth survey was only conducted in 1984, 1989, 1994. Starting in 1999, the wealth survey is part of the regular survey conducted by the PSID. Unless stated otherwise, our main empirical analysis is based on the years when the wealth survey is conducted.

To conduct our sensitivity analysis, we use another database in the PSID that indicates the location of each survey participant at the Census Tract level between 1968 and 2009. This location data is available only by special request. Combining this information with data from the U.S. Census Bureau, we obtain an additional set of neighborhood characteristics that includes the fractions of homes that are owner-occupied, vacant, recreational, and for sale, the local unemployment rate, the average education rate, and the average number of years spent by the household in the neighborhood. We also obtain dummies regarding the ethnicity of each household (White/Caucasian, Asian, African-American, Hispanic), as well as the local fraction of households for each ethnic background. Finally, we calculate for each census tract the proportions of active professionals that are farmers, executives, technicians, and workers, as defined by the U.S. Census.

When information about some characteristics is missing, we interpolate using the available date. For example, values for total taxable family income are unavailable between 1994 and 1996 and in 2001, so we use a linear interpolation from the first available surrounding years. Similarly, civil status is not available between 1994 and 1997. Similarly, the census variables covers the years 1980, 1990, and 2000. Therefore, we use a linear interpolation to infer values for all the other years. For the years after 2000, we interpolate the data using the census years 1990 and 2000. For the rural areas that were not followed in the 1980s, we interpolate the data using the 1990 and 2000 census values.

We impose a series of filters to guarantee the reliability of the data for the empirical analysis. We exclude households that belong to the Survey of Economic Opportunities, which is a special subsample of the PSID drawn from lower income levels that allows researchers to study poverty. Households must also (i) appear in the survey more than once, (ii) have a gross worth and real total taxable family income greater than \$1,000, (iii) a strictly positive financial wealth (cash, stocks, and bonds), (iv) nonmissing observations for the market values of the different assets. To eliminate outliers, we finally remove the top 0.1% values of gross worth and real estate wealth and the observations for which the values of rental income earned by landlords are in the top 0.1%. When we include the additional census variables, we exclude census areas where 100% of the active professionals are in the same category (farmers, executives, technicians, workers).

3.2 Definition of the Household

Following the convention of the PSID, we define the household as a head member. Each year, the PSID automatically selects one member of an interviewed family unit to be the head and ranks all the other members in terms of their relationship to that person (e.g., partner, child, sibling). For multiple-member units, the PSID convention is that the head must be older than 16 years old and have the greatest financial responsibility. If that person is female and she has a partner (husband, boyfriend living in the same unit, or civil partner), then her partner is designated as the head, unless he is incapacitated. We add the restriction that the household head member must be older than 18 years and younger than 100 years. As defined by the PSID, the age, gender, and education variables all refer to the household head, while the income and wealth variables are all aggregated at the household level.

To make sure that a household can be properly compared over time, we impose that any serious change to the household composition results in the disappearance and/or creation of a new household. A household no longer exists when there is a change in the marital status of the head couple. A change can come from either divorce, separation, death, or a new partnership in the case of a head member who used to be single. In the event any of the other household members keeps being interviewed by the PSID afterwards, we consider him/her as a member of a new household. In addition, a new household is created (i) when a member of an existing household who is not the head or his partner (e.g. child, sibling) leaves and creates his or her own household, and (ii) when the PSID extends the sample to new families and interviews them for the first time.

3.3 Asset Values and Wealth

In the wealth surveys, households are asked to report the market value of their assets via the following type of question: *if you sold all [the amount in asset x that you or anyone in your family own] and paid off anything you owed on it, how much would you have?* From their responses we compute the market value invested in each asset class. To compute the total value invested in housing, we take the sum of (i) the value of the household's primary house and (ii) the net value of his other real estate properties.

We use gross worth to compute the household's housing share and its wealth. It includes (i) the net value invested in stocks (including mutual funds and retirement accounts), (ii) the net value invested in bonds (including the cash value in life insurance policies), (iii) the amount of cash (including checking and savings accounts, certificates of deposits, government savings bonds, and Treasury bills), (iv) the total value invested in housing (primary house and other real estate properties), (v) the net value invested in farms and private businesses, and (vi) the net value invested in cars.

We also conduct our analysis using net worth, which is defined as gross worth minus (i) the amount remaining on non-collateralized debt (such as credit card charges, student loans, medical bills, or loans from relatives) and (ii) the mortgage on the household's primary house. Before 1999, the value invested in stocks included both retirement and non-retirement accounts. Afterwards, households were asked to report the values of both accounts separately. To be consistent with the pre-1999 years, we take the sum of these two accounts.

3.4 Return Properties of the Homes

We compute the average and volatility of the real return of home prices as follows. We first build the time-series of home prices using the entire PSID data available from 1968 to 2013. We then deflate the home prices using an inflation index that excludes the price of the housing good (as in the model). Following Piazzesi et al. (2007), we form this index from the NIPA consumption tables that excludes housing services, durable goods, cloth and shoes (see Table 2.3.5 Personal Consumption Expenditure by Major Type of Product).[7] Finally, we compute the expected return and volatility of each individual house price (in real terms) over all the available years.[8] We eliminate observations for which the average value and/or the volatility of the house's price growth is greater than 100% per year. We also require at least four return observations for each house.

3.5 Calibration of the Model Parameters

In this section, we explain in more detail the illustrative example presented in Figure 3. This example examines how the unconstrained and constrained housing shares α_H^U and α_H^C change for different types of households. To compute α_H^U and α_H^C, we use the closed form expressions of the model based on the parameter values described below.

The parameters of the asset returns are based on the annual real return data on stocks, short-term bonds, and real estate. We use the value-weighted NYSE/AMEX/NASDAQ index and the 1-month Treasury Bill from CRSP as proxies for stocks and bonds, and rely on the PSID database to track the price evolution of individual houses. This database departs from the Case-Shiller indices which underestimate the price volatility of individual homes. For each asset, we compute real returns using as deflator an inflation index that excludes the price of the housing good (as in the model).

The model parameters are based on the estimated first and second return moments between 1968 and 2013 when the PSID data is available. The average real return and volatility of the stock index are equal to 4.5% and 19.1% per year, while the average real risk-free rate equals 1.1% per year. The parameters for housing are obtained by taking the median across all individual homes. The real price growth of housing is equal to 2.6% per year after adjusting for the annual depreciation costs of 1.4% (based on the estimates of Leigh (1980)), and its volatility is equal to 16.3%. We further set the rent-to-price ratio equal to 3.5% based on the recent estimates of Campbell, Davis, Gallin and Martin (2009).

[7] This index has a correlation of 84% with the standard CPI obtained from Robert Shiller's website over the 1968-2013 period.

[8] For the years after 1997 where the PSID survey is only available every other year, we simply annualize the percentage changes in house prices.

We consider a set of households that have different preferences for housing consumption β_K, but identical levels of risk aversion and time preference. For the risk aversion coefficient, we set γ equal to 5, similar to Cocco (2005) and Yao and Zhang (2005). For the time preference coefficient, we choose a low value of δ equal to 0.25. As β_K increases, both housing investment and consumption rise (i.e., both α_H^U and α_H^K increase β_K). By choosing a low δ, we have that α_H^K rises faster than α_H^U because households are impatient in their consumption decisions. Therefore, β_K is negatively related to ϕ_U, and positively related to α_H^C (i.e., the household becomes more constrained). Alternatively, we could choose a high value for δ. In this case, the shapes of α_H^U and α_H^C in Figure 3 remain largely unchanged. The only difference is a change of sign - β_K becomes negatively related to α_H^C, which implies that α_H^C rises below the cutoff value $\bar{\beta}_K$, and is equal to zero otherwise.

REFERENCES

Campbell, S., Davis, M., Gallin, J. and Martin, R. (2009), 'What moves housing markets: A variance decomposition of the rent-price ratio', *Journal of Urban Economics* **66**, 90–102.

Cocco, J. (2005), 'Portfolio choice in the presence of housing', *Review of Financial Studies* **18**(2), 535–567.

Corradin, S., Fillat, J. and Vergara-Alert, C. (2014), 'Optimal portfolio choice with predictability in house prices and transaction costs', *Review of Financial Studies* **27**(3), 823–880.

Damgaard, A., Fuglsbjerg, B. and Munk, C. (2003), 'Optimal consumption and investment strategies with a perishable and an indivisible durable good', *Journal of Economic Dynamics and Control* **28**(2), 209–253.

Davidson, R. and MacKinnon, J. (2004), *Econometric Theory and Methods*, Oxford University Press.

Greene, W. (2012), *Econometric Analysis*, Pearson; 7th Edition.

Grossmann, S. and Laroque, G. (1990), 'Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods', *Econometrica* **58**, 25–51.

Leigh, W. (1980), 'Economic depreciation of the residential housing stock of the united states: 1950-1970', *The Review of Economic Studies and Statistics* **62**(2), 200–206.

Liang, K. and Zeger, S. (1986), 'Longitudinal data analysis for generalized linear models', *Biometrika* **73**(1), 13–22.

Piazzesi, M., Schneider, M. and Tuzel, S. (2007), 'Housing, consumption, and asset pricing', *Journal of Financial Economics* **83**, 531–569.

Yao, R. and Zhang, H. (2005), 'Optimal consumption and portfolio choices with risky housing and borrowing constraints', *Review of Financial Studies* **18**(1), 197–239.

Zang, I. (1980), 'A smoothing-out technique for min-max optimization', *Mathematical Programming* **19**, 61–77.