

The Predominance of Real Estate in the Household Portfolio*

Laurent Barras and Sebastien Betermier[†]

August 9, 2016

Abstract

This paper investigates why household portfolios are heavily skewed toward real estate. Previous studies suggest that the large portfolio share of real estate primarily stems from non-investment-related motives as homeowners are often forced to invest heavily to buy the home they want to consume. In contrast, we show that homeowners would still invest the bulk of their wealth in real estate in a frictionless setting where they could own and consume separate amounts of housing. We provide empirical support to this argument and derive a dynamic portfolio model to study why real estate has such a strong investment appeal.

JEL Classification: G11, G12.

Keywords: Portfolio allocation, hedging, real estate.

*We thank Matthieu Bouvard, Laurent Calvet, Adolfo de Motta, Chris Hennessy, Martin Lettau, Christine Parlour, David Schumacher, Richard Stanton, Adam Szeidl, Carles Vergara-Alert, Johan Walden, Nancy Wallace, and the seminar participants at McGill 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

One striking feature of household portfolios around the world is the predominance of real estate. For example, Tracy and Schneider (2001) document that the average share of gross worth invested by U.S. households in real estate ranges between 40% and 45%, whereas the fraction invested in public stocks is only between 5% and 10% (see also Campbell, 2006; Carroll, 2002; Guiso and Sodini, 2013). The predominance of real estate is puzzling from the perspective of standard portfolio theory. Why do households invest the bulk of their wealth in their home? Shouldn't they choose a more diversified portfolio?

Despite its importance, real estate has received relatively little attention in the literature on portfolio allocation as most models focus on the allocation of financial assets only. One common view is that there is little to study because homeowners are often *forced* to invest heavily in real estate to fully own the home they want to live in. In other words, the portfolio share of real estate is bound to be large as a result of the homeowners' consumption decisions irrespective of whether it is a good investment. This constrained portfolio allocation arises because (i) real estate is a durable good that provides dual consumption and investment benefits, and (ii) there are no easily available partial ownership arrangements that would help to separate these benefits.¹

In this paper, we study how the consumption and investment properties of real estate contribute to its predominance in the portfolio. We disentangle their contributions by investigating what existing homeowners would do if they had easy access to partial ownership arrangements allowing them to own and consume separate amounts of real estate. Our main result is that they would still invest the bulk of their wealth in real estate in this frictionless environment. In other words, we show that the portfolio share of real estate due to consumption motives is relatively small. Real estate must therefore have unique investment appeal to justify the under-diversified portfolios held by homeowners.

¹Recent studies highlight the importance of the housing consumption motives and show that they affect the homeowners' risk appetite for equities (see Brueckner, 1997; Chetty and Szeidl, 2007; Cocco, 2005; Frantatoni, 1998; Damgaard et al., 2003; Flavin and Yamashita, 2002; Vestman, 2012; Yamashita, 2003; Yao and Zhang, 2005).

Teasing apart the consumption and investment motives for holding real estate is a challenging exercise. Since homeowners are typically forced to fully own their homes, their optimal portfolio share of housing, or *total housing demand*, has both an unconstrained and a constrained component.² The *unconstrained demand component* is the share of real estate that homeowners would choose in a frictionless world where partial ownership is possible. It is therefore entirely driven by investment motives. The *constrained demand component* is the additional investment that homeowners must make to fully own the home they want to consume. As such, it captures the impact of the consumption motives. Both demand components are hard to tease apart because they are not observable – standard surveys do not report what homeowners would do if partial ownership arrangements were made available to them. Moreover, both demand components are interdependent as households that desire a high real estate investment are, *ceteris paribus*, less constrained by homeownership.

The research approach that we design overcomes these issues. We first demonstrate that the unconstrained housing demand of existing homeowners is bound to be large via two simple and general economic arguments. That is to say, their portfolio share of real estate is bound to be primarily driven by investment motives. We then empirically validate this claim by designing a novel identification strategy that allows us to measure the variation in both demand components across homeowners. Finally, we show that these findings can be integrated into a unified and rational framework for portfolio choice, and provide new insights into why real estate has such a strong investment appeal.

The first part of the paper presents the two general arguments regarding why the unconstrained demand is bound to be large for existing homeowners. First, the mere fact that they have chosen to own rather than rent reveals a strong preference for housing as an investment asset. While homeowners would typically reduce their housing investment if partial ownership arrangements were available, this reduction can only be moderate. Otherwise, homeowners would not choose to own their home – they would instead rent it and just extract the consumption benefits of housing. The option to rent therefore puts a lower bound on the size of the homeowners' unconstrained demand for housing.

²While there exist solutions to relax this homeownership constraint, they are not easily available. For example, equity-sharing programs are only used in very specific contexts (Caplin, Chan, Freeman and Tracy, 1997) and hedging strategies based on REITs or the Case-Shiller indices do not correlate highly with local house prices.

The second argument is that homeowners would optimally choose a bigger home if partial ownership arrangements were available. In other words, they would own less of a bigger house which would still result in a large real estate investment. The intuition behind this argument is that the constraint of homeownership leads to under-consumption of housing as it forces households to own the full value of their home (see Brueckner, 1997).

Combining both arguments, we show that the homeowners' unconstrained and constrained demand components approximately make up $\sqrt{\phi^U}$ and $1 - \sqrt{\phi^U}$ of their portfolio share of real estate, where ϕ^U is the optimal fraction of the home that they would own in the frictionless setting. As the square-root of a fraction is larger than the fraction itself, this rule of thumb implies that the unconstrained demand dominates the constrained demand for most values of ϕ^U . For example, even if a homeowner were to only retain a 50% equity stake in a home, the unconstrained demand would account for as much as 70% of its total housing demand. Cases where the unconstrained demand is small can be ruled out because they are not consistent with the fact that homeowners have chosen to own rather than rent. Our analysis therefore demonstrates that, even if homeowners were able to sell a fraction of their home, their investment in real estate would remain high.

The second part of the paper empirically confirms this prediction. Using micro-level data from the Panel Study of Income Dynamics (PSID), we circumvent the non-observability of the unconstrained housing demand by developing an identification strategy that exploits information about landlords. Because they invest more in housing than they consume, landlords are not constrained by homeownership. We can therefore exploit the cross-sectional variation in their housing decisions to infer the variation in the unconstrained demand among the non-landlord homeowners (i.e., those constrained by homeownership). A key advantage of this cross-sectional approach is that it does not impose the level of housing demand to be the same among landlords and non-landlords.

As predicted by our theoretical arguments, the unconstrained demand explains 80% of the variation in the housing demand across non-landlord homeowners, versus 6% for the constrained demand (the 14% residual variation is captured by the covariance term). We also uncover new insights into which types of homeowners have greater unconstrained and constrained demands. The unconstrained demand is large among less wealthy and larger

households. The economic magnitude of the wealth effect is particularly strong, i.e., a one-standard deviation *decrease* in wealth increases the unconstrained demand by 11.3%. As for the constrained demand, we show that younger households with low income and large size are more constrained by homeownership. This result is consistent with the intuition that these households are more likely to be first-time homeowners operating on a tight housing budget (Cocco, 2005; Flavin and Yamashita, 2002).

The third part of the paper shows that our main findings can be integrated into a unified and rational framework for portfolio choice. We solve a stylized model based on a version of the classic Merton (1971) framework in which housing makes up part of the consumption basket. The analysis shows that the combination of multiple periods and multiple goods is sufficient to produce (i) a portfolio that is skewed toward real estate, (ii) a total housing demand that is primarily driven by the unconstrained demand component, and (iii) cross-sectional variation in the unconstrained demand that is in line with our empirical findings. Its simplicity allows us to derive closed-form expressions for the homeowners' unconstrained and constrained demands.

The reason for the large unconstrained housing demand in the model is that real estate is equivalent to a perpetual indexed bond on housing consumption. As such, it provides a unique combination of speculative and hedging investment benefits (Fisher, 1974). It improves the risk-return trade-off of the overall portfolio, which is beneficial for the entire consumption basket. This speculative benefit has previously been analyzed in one-period mean-variance models and shown to be important because of the low correlation between the housing and stock markets.³ In addition, real estate represents a risk-free claim on housing consumption, which is the single largest good in the households' consumption basket. This hedging benefit is particularly valuable for households that have greater housing consumption needs and are more risk averse. These predictions are consistent with our empirical evidence that the unconstrained demand is stronger among larger households and the less wealthy that are known to be more risk-averse (Calvet and Sodini, 2014).

Of course, several other investment properties of real estate are likely to play an important role in driving the homeowners' unconstrained housing demand, such as life-cycle effects,

³See Eichholtz, Koedijk and de Roon (2000), Englund, Hwang and Quigley (2002), Gatzlaff (2000), Goetzmann and Ibbotson (1990), Goetzmann (1993), and Iacovello and Ortalo-Magné (2003).

borrowing constraints, labor income risk, taxes, or transaction costs (e.g., Corradin, Fillat and Vergara-Alert, 2014; Kraft and Munk, 2011; Van Hemert, 2010; Yang, 2009). The main takeaway of the paper is that the predominance of real estate in the portfolio *must* come from its investment properties, whatever they might be. Our portfolio choice model shows that, even in the classic Merton (1971) setting, real estate already has distinct investment value relative to traditional assets.

The paper relates to several strands of the literature on portfolio choice and real estate. Several studies highlight the crowding out effect of the homeownership constraint on the demand for stocks inside the financial portfolio (e.g., Chetty and Szeidl, 2014; Cocco, 2005; Flavin and Yamashita, 2002). Relative to these papers, we focus on the determinants of the total housing demand. Our analysis reveals that the impact of the constraint varies depending on how much the asset under consideration weighs in the portfolio. In the case of housing, the constraint only explains a small fraction of the large total demand. Many studies in real estate also examine the determinants of the decision to own or rent a home (e.g., Henderson and Ioannides, 1983) including the hedging properties of housing (Sinai and Souleles, 2005). Here, we focus on the share of housing in the portfolio of homeowners. Finally, Caplin, Chan, Freeman and Tracy (1997) and Cauley, Pavlov and Schwartz (2007) analyze the welfare gains of shared-equity programs that would allow homeowners to sell a fraction of their home. Our analysis shows that, even if homeowners were able to participate in those programs, their investment in real estate would still represent the bulk of their wealth.

The rest of the paper is organized as follows. Section 2 reviews the empirical importance of real estate in the portfolios of U.S. households. Section 3 presents the two general arguments regarding why the unconstrained demand is bound to be large for existing homeowners. Section 4 describes the empirical analysis, and Section 5 lays out the portfolio choice model. Section 6 concludes. The Internet Appendix contains all the derivations and reports additional empirical results.

2 A Review of Real Estate in the Portfolio

We first provide a brief review of the empirical importance of real estate for US households. We use data from the Panel Study of Income Dynamics (PSID), a national survey of U.S. households that is widely used in the household finance literature.⁴ This database, which is described in greater detail later, provides us with a wealth survey of a representative sample of households between 1984 and 2013.

– **Figure 1 here** –

Real estate is the cornerstone of household portfolios. In Figure 1, we plot the portfolio composition of 30 groups of households sorted by their gross worth (stocks, bonds/insurance, cash, businesses/farms, total real estate, and motor vehicles). Except for the poorest 25%, we see that real estate accounts for more than 50% of household wealth. By contrast, stocks typically represent less than 20% of wealth.

– **Figures 2 and 3 here** –

The predominance of real estate is remarkably stable both over time and across regions. Figure 2 reveals that the share in housing remains above 50% over the entire period 1984-2013. Likewise, Figure 3 shows little variation in housing shares across diverse Metropolitan Statistical Areas such as Boise City, ID or San Diego, CA. These numbers are in line with the estimates in Campbell (2006), Guiso and Sodini (2013), Poterba and Samwick (2001), and Tracy, Schneider and Chan (1999).

The large average shares in housing depicted in Figures 1, 2, and 3 are entirely driven by homeowners as renters typically invest nothing or very little in real estate. In the next two sections, we therefore focus our attention on homeowners and examine the determinants of their housing investment decisions.

⁴The PSID is conducted at the Survey Research Center, Institute for Social Research, University of Michigan.

3 Theoretical Analysis of the Housing Demand

In this Section, we determine whether the large portfolio share of real estate among existing homeowners primarily comes from investment or consumption motives. Both these motives potentially play a role because homeowners are typically forced to fully own the home they want to consume. To address this issue, we theoretically examine what homeowners would do in a frictionless environment where they can consume and own separate amounts of housing.

3.1 Notation

We begin by introducing some structure and notation to outline the key features of the real estate market. The first is that real estate is a durable good that provides a mix of consumption and investment benefits. To distinguish between them, we denote by K^T and H^T the quantities of housing consumed and owned by a household. Each unit of housing can be viewed as a one-dimensional summary of the quality of the house which accounts for size, location, and specific characteristics. The superscript T stands for *total* as we will shortly decompose these quantities into several components.

Another distinct feature of the real estate market is that households are typically limited in their ability to disentangle both benefits. This is illustrated in Figure 4 which depicts the ratio of housing investment to housing consumption $\phi^T = H^T/K^T$, also referred to as the *ownership stake* of housing. Households can either: (i) rent a home ($\phi^T = 0$), (ii) own the home ($\phi^T = 1$), (iii) or become landlords and rent out a part of the home they own ($\phi^T > 1$). However, partial ownership agreements are generally not available ($0 < \phi^T < 1$). In other words, households cannot own a fraction of the home and rent the remaining fraction from the other co-owner.

– Figure 4 here –

Existing homeowners are therefore constrained by homeownership, i.e. they must consume and own identical levels of housing ($H^T = K^T$). Their optimal decisions result from a constrained optimization and include both an unconstrained and a constrained component,

denoted by superscripts U and C :

$$K^T = K^U + K^C, \quad (1)$$

$$H^T = H^U + H^C. \quad (2)$$

The vector of unconstrained components $\{H^U, K^U\}$ represents the homeowners' optimal decisions in a frictionless or unconstrained environment where they can own and consume separate amounts of housing ($H^U \neq K^U$). The vector of constrained components $\{H^C, K^C\}$ captures the adjustments in the levels of housing investment and consumption caused by the homeownership constraint.

We can apply the same decomposition to the homeowners' portfolio share of real estate, or *total housing demand*, which is defined as $\alpha_H^T = H^T \cdot P_H / W$ where P_H is the unit price of housing owned and W is the households' wealth. Given Equation (2), the housing demand can be decomposed as

$$\alpha_H^T = \alpha_H^U + \alpha_H^C. \quad (3)$$

where α_H^U and α_H^C denote the *unconstrained* and *constrained* housing demands.

Equation (3) frames the problem for our analysis of the housing demand among homeowners. In order to determine whether their demand α_H^T is primarily driven by investment or consumption motives, we need to measure the two demand components α_H^U and α_H^C . The unconstrained demand captures the investment value of housing because it is the optimal housing share chosen by homeowners in a frictionless environment. In contrast, the constrained demand reflects the impact of consumption motives because it represents the additional investment that homeowners must make to fully own the home they want to consume.

3.2 The Predominance of the Unconstrained Demand

We now use two simple and general economic arguments to demonstrate that the unconstrained demand for homeowners α_H^U is bound to be large. We begin by presenting the intuition behind each argument and then provide a more formal analysis.

First, the choice made by homeowners to own rather than rent reveals a strong preference for housing as an investment asset. In a frictionless environment with partial ownership, existing homeowners would sell a fraction $1 - \phi^U$ of their home and decrease their investment in housing (i.e., $\alpha_H^U < \alpha_H^T$). However, this reduction can only be moderate—otherwise, homeowners would prefer to rent the house and just extract the consumption benefits of housing (i.e. they would pick $\phi^T = 0$ instead of $\phi^T = 1$). The option to rent therefore puts a lower bound on the optimal ownership stake ϕ^U and, consequently, on the size of the unconstrained demand α_H^U .

Second, homeowners would optimally choose a bigger home if partial ownership arrangements were available. Because they would own less of a bigger house, it would still result in a large unconstrained demand α_H^U . The intuition behind this argument can be illustrated with a simple example. Consider a homeowner that holds a wealth of \$400K. In a frictionless setting, it would choose a partial ownership stake ϕ^U of 0.5 by living in a home worth \$200K and owning just \$100K of it. Its unconstrained demand α_H^U is therefore equal to 25% (100/400). When the homeownership constraint is imposed, the housing investment of \$100K is no longer optimal because it halves the targeted level of housing consumption. Likewise, buying a home worth \$200K is not optimal either because it doubles the targeted housing investment. The optimal compromise between the imbalanced consumption basket and the imbalanced portfolio is to choose a housing investment between \$100K and \$200K. Suppose, for instance, that the constrained homeowner settles for a house worth \$120K. In this case, the total demand α_H^T equals 30% (120/400), which is only 5% greater than the unconstrained demand α_H^U .⁵

Model-free decomposition. We now formally show that when both arguments are combined, they imply that the unconstrained demand α_H^U remains close to the total demand α_H^T . To make this point, we rely on a model-free decomposition of the ratio α_H^U/α_H^T into two variables: (i) the optimal ownership stake ϕ^U , and (ii) the ratio of the housing consumption levels in the unconstrained and constrained cases:

$$\frac{\alpha_H^U}{\alpha_H^T} = \frac{H^U}{W} \cdot \frac{W}{H^T} = \frac{H^U}{K^U} \cdot \frac{K^U}{K^T} = \phi^U \cdot \frac{K^U}{K^T}. \quad (4)$$

⁵The trade-off between household investment and consumption is one of the core predictions of portfolio choice models when households are constrained by homeownership (see Brueckner (1997) and the continuous-time model presented in Section 5).

where we use the restriction $H^T = K^T$ implied by the homeownership constraint. The first argument – the fact that homeowners have decided to own rather than rent their home – implies that the optimal ownership stake ϕ^U should be high. The second argument – the fact that they would increase their level of housing consumption in a frictionless setting – implies that the ratio K^U/K^T should be greater than one. The product of both terms is therefore bound to be close to one.

As a rule of thumb, suppose that existing homeowners would make the same adjustment to their housing investment and consumption levels in the frictionless environment. That is, they would decrease housing investment and increase housing consumption in the following way,

$$\frac{H^U}{H^T} = \frac{K^T}{K^U} = \sqrt{\phi^U}. \quad (5)$$

Applied to our previous example where the constrained homeowner must choose a house that is worth between \$100K and \$200K, Equation (5) implies a home value of \$140K (i.e., $100/140 = 140/200 = \sqrt{0.5}$).⁶

Under this assumption of equal adjustment, the total housing demand can be decomposed as

$$\frac{\alpha_H^U}{\alpha_H^T} = \sqrt{\phi^U}, \quad \frac{\alpha_H^C}{\alpha_H^T} = 1 - \sqrt{\phi^U}. \quad (6)$$

Figure 5 displays this decomposition for the different values of the optimal ownership stake ϕ^U chosen by homeowners. The results reveal that the unconstrained demand dominates the constrained demand for most values of ϕ^U . For example, if ϕ^U is equal to 0.70, the unconstrained demand accounts for 84% of the total housing demand. Even if ϕ^U is as low as 0.30, the unconstrained demand still represents more than 50% of the total demand. Only extremely small values of ϕ^U lead to an insignificant unconstrained demand, e.g., if ϕ^U equals 0.01, then it only represents 10% of the total demand. However, these cases can be ruled out because they cannot be reconciled with the fact that existing homeowners have chosen to own rather than rent. Therefore, our analysis demonstrates that even if homeowners were able to sell a fraction of their home, their unconstrained demand for housing would remain high.

⁶We can check that the adjustment rule in Equation (5) is consistent with the identity $\phi^U = \frac{H^U}{K^U}$, i.e., $\phi^U = \frac{H^U}{H^T} \cdot \frac{K^T}{K^U} = \sqrt{\phi^U} \cdot \sqrt{\phi^U}$.

– Figure 5 here –

4 Empirical Analysis of the Housing Demand

4.1 Data Description

We now turn to the empirical analysis to test whether the predominance of the unconstrained demand is confirmed in the data. We use micro-level data on the housing demand of homeowners from the wealth survey of the PSID over the period 1984-2013. This survey tracks a representative sample of U.S. households and contains information about their financial situation (income, wealth and its composition), their demographics (age, household size, marital status, high school or post-high school education), and their real estate holdings (value of the house, mortgage, tenure choice). It is conducted every 5 years from 1984 to 1999 and then every two years from 1999 to 2013.

We measure a homeowner’s housing share as the total investment in real estate 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). We select gross worth as the denominator 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 is either arbitrarily large or undefined.⁷ Additional details on the data requirements can be found in the Internet Appendix.

– Table I here –

Table I (Panel A) reports summary statistics for all homeowners at the end of the sample period in 2013. As defined by the PSID, the age, gender, and education variables all refer to the household head. The average homeowner is 54 years old, has a household income of \$74,000, and a gross worth of \$390,000. We classify homeowners as landlords if they own multiple homes and receive rental income in the year prior to 2013. Landlords represent 8% of the population, have similar demographic characteristics as non-landlords, and are

⁷In the Internet Appendix, we find that the empirical results still hold when replacing gross worth with net worth.

wealthier. In Panel B, we report the wealth composition of homeowners in 2013. The bulk of gross worth is invested in real estate (67%), which is in line with the statistics from Section 2. Financial assets, businesses, and autos make up the remaining 23%, 3%, and 7%.

4.2 Econometric Issues

Our empirical analysis is challenging because the unconstrained demand is not observable, i.e., we do not know whether existing homeowners would rather be partial owners of their homes. Therefore, the simple approach of measuring whether the average unconstrained demand represents a large fraction of the average housing share cannot be used.

One possible solution is to study the variation in housing shares across homeowners, which allows us to sidestep the difficult task of estimating the average level of the unconstrained demand. However, this cross-sectional approach still suffers from a lack of identification because of the joint variation in the unconstrained and constrained demands. Both demand components are interdependent as households that desire a high real estate investment are, *ceteris paribus*, less constrained by homeownership. We illustrate this point with the following example.

Consider two households that both have the same wealth of \$400K. Both ideally would like to live in a home worth \$200K but differ in their desired investment in housing. Household l has a *low* unconstrained demand for housing and would like to own just half the house ($\phi_l^U = 0.5$ and $\alpha_{H,l}^U = 25\%$). Since the partial ownership arrangement is not available, it decides to own a smaller house worth \$120K ($\alpha_{H,l}^T = 30\%$). In comparison, household h has a *high* unconstrained demand for housing and wants to fully own the house ($\phi_h^U = 1$ and $\alpha_{H,h}^U = \alpha_{H,h}^T = 50\%$).

– **Figure 6 here** –

While it may be tempting to interpret a high housing share as the outcome of a strongly binding constraint, we see that household h is both *more* keen to invest in real estate, and *less* constrained by homeownership. It is therefore the unconstrained demand that explains the cross-sectional variation in the total demand. This is illustrated in Figure 6 which plots

the differences between the total demands, $\alpha_{H,h}^T - \alpha_{H,l}^T$, and the unconstrained demands, $\alpha_{H,h}^U - \alpha_{H,l}^U$. While the unconstrained demand explains 125% of the cross-sectional variation in the total demand, the contribution of homeownership goes in the opposite direction (-25%). This example shows that we need a proper identification strategy to infer the cross-sectional variation in the unconstrained demand (i.e., the difference $\alpha_{H,h}^U - \alpha_{H,l}^U$).

4.3 Identification Strategy

We develop an identification strategy based on the investment decisions made by landlords. Unlike the other homeowners, landlords are not constrained by homeownership because they invest more in housing than they consume (see Figure 4).⁸ From the cross-sectional variation in their housing decisions, we can therefore infer the variation in the unconstrained demand across the non-landlord homeowners (i.e., those constrained by homeownership).

To formalize this intuition, we model the housing share of household i in year t as a linear function of the household's characteristics (e.g., income, age):

$$\alpha_{Hi,t} = a_t + s'_U X_{i,t} + (b_t + s'_C X_{i,t}) \cdot D_{i,t}^{NL} + u_{i,t}, \quad (7)$$

where $X_{i,t}$ is a vector of K characteristics, $D_{i,t}^{NL}$ is a dummy variable that equals one if the household is a non-landlord homeowner and zero otherwise, and $u_{i,t}$ is the residual term. The terms a_t and b_t are vectors of year fixed effects that absorb any differences in the levels of housing shares between landlords and non-landlords. The vector s_U measures the slopes associated with landlords. It therefore proxies for the slopes of the unconstrained demand (i.e., the dotted line in Figure 6). The vector s_C measures the incremental slopes of the housing demand for non-landlords, so it proxies for the slopes of the constrained demand. Taking the sum of the two vectors, we obtain the vector of slopes for the total demand for housing, $s_T = s_U + s_C$ (i.e., the solid line in Figure 6).⁹

The next step is to decompose the cross-sectional variation in housing share among non-landlords into its unconstrained and constrained components. We estimate the variance

⁸In our sample, the average ratio of housing investment to housing consumption for landlords is equal to 2.1, versus 1.0 for non-landlords.

⁹The estimated vectors s_U and s_T obtained from the panel regression are numerically equivalent to computing two separate regressions for landlords and non-landlords (see Chapter 6 of Greene, 2011)

of the unconstrained demand of non-landlords, $\text{var}(\alpha_{Hi,t}^U)$, where $\alpha_{Hi,t}^U = s'_U X_{i,t}$ for each non-landlord i and year t . Likewise, we estimate the variances of the constrained and total demand of non-landlords, $\text{var}(\alpha_{Hi,t}^C)$ and $\text{var}(\alpha_{Hi,t}^T)$, where $\alpha_{Hi,t}^C = s'_C X_{i,t}$ and $\alpha_{Hi,t}^T = s'_T X_{i,t}$. To perform the variance decomposition, we simply compute the variance ratios of the two demand components as $\text{var}(\alpha_{Hi,t}^U) / \text{var}(\alpha_{Hi,t}^T)$ and $\text{var}(\alpha_{Hi,t}^C) / \text{var}(\alpha_{Hi,t}^T)$.

Our strategy builds on an early insight of Brueckner (1997), who uses landlords to analyze the average level of the housing share among non-landlords. Here, we specifically focus on its cross-sectional variation, which offers several important advantages. First, it does not require that the housing demand among landlords has the same level as the unconstrained demand among non-landlords. There may be unobservable characteristics that are specific to landlords, such as dealing with tenants, housing-specific maintenance issues, and particular tax rules.¹⁰ While these landlord-specific factors affect the intercept coefficient a_t , they leave the slope vector s_U unchanged. Second, the cross-sectional variance analysis is computed from the characteristics of non-landlords only. As such, it controls for possible differences in the dispersion of characteristics between landlords and non-landlords.¹¹

4.4 Main Results

Table II reports the estimates of Equation (7) for the period 1984-2013 and the baseline set of characteristics, which include demographic predictors (age, household size), socio-economic predictors (log income, log wealth, education), and house price properties (average growth, volatility). The theoretical motivation for using these variables as well as their ability to predict the housing share are discussed in the life-cycle models of Cocco (2005), Van Hemert (2010), and Yao and Zhang (2005). Standard errors are clustered at the household level.

– Table II here –

Variance Decomposition. Table II (Panel A) reports the variance ratios for the unconstrained and constrained components of the housing demand. Consistent with our theoretical

¹⁰Unlike regular homeowners, landlords are taxed on the rental income they earn, even though they are allowed to depreciate the value of the house they rent out.

¹¹The variance decomposition could potentially be biased if the intervals for the characteristics of landlords and non-landlords do not overlap and the relation between these characteristics and the housing share is non-linear. In the Internet Appendix we perform a careful analysis to rule out this concern.

predictions in Section 3, we find that the unconstrained demand is by far the predominant component of the housing demand among non-landlord homeowners. It explains 80% of the cross-sectional variation of the housing share, whereas the impact of homeownership only explains 6%. The remaining 14% is attributed to the covariance between the two demand components. The importance of the unconstrained demand is also reflected in the F -statistics of joint tests of the slope coefficients. While the null hypothesis $H_0: s_U = 0$ is strongly rejected in the data (p -value < 0.001), the statistical evidence against the null $H_0: s_C = 0$ is statistically weaker (p -value = 0.047).

Analysis of the Slope Coefficients. Our empirical analysis also allows us to identify which types of homeowners have greater unconstrained and constrained demands for housing, and therefore provide new insights into the drivers of the large unconstrained demand. Table II (Panel B) reports the coefficients associated with the total, unconstrained, and constrained demands. To facilitate the interpretation of the results, each coefficient is divided by the standard deviation of its associated characteristic among non-landlords.

We begin by analyzing the vector of slopes s_T for the *total* housing demand along with the associated t -statistics (first and second columns of Panel B). The model does a good job at explaining the variation in housing shares across homeowners. The R^2 is close to 30% and is higher than typical estimates in the household finance literature (e.g., Betermier, Calvet and Sodini, 2015). Among the seven household characteristics, gross worth is the most important predictor, both economically and statistically. We estimate that a one standard deviation increase in wealth reduces the housing share by 11.5% (t -statistic of -29.83). We also find that larger households tend to increase their investment in real estate (t -statistic of 13.25), while older households do the opposite (t -statistic of -6.08). Consistent with the mean-variance analysis, the demand for real estate is positively related to our measure of expected return (average growth) and negatively related to house price volatility. These results are all in line with previous studies that provide theoretical and empirical evidence for the negative effects on wealth, age, and volatility on the total housing share (e.g. Cocco, 2005; Tracy, Schneider and Chan, 1999).

We turn to the vector of slopes s_U for the *unconstrained* demand (third and fourth columns of Panel B). Quite remarkably, we find that a one standard deviation increase in

wealth reduces the unconstrained demand by 11.3% (t -statistic of -11.31), which is nearly identical to its effect on the total demand (11.5%).¹² Therefore, the negative impact of wealth on the total housing share originates from the unconstrained demand. There is also a positive relation between household size and the unconstrained demand with a t -statistic of 2.19. These results indicate that the real estate investment has special appeal to households that exhibit low wealth levels or large consumption needs.

Finally, we examine the slope coefficients s_C associated with the *constrained* demand (fifth and sixth columns of Panel B). We find that younger households with low income and large size are relatively more constrained by homeownership. These effects are consistent with the argument in Cocco (2005) and Flavin and Yamashita (2002) that that these households are more likely to be first-time home buyers operating on a tight housing budget. However, the economic magnitude of these effects is small. The maximum variation in the constrained demand following a one standard deviation change in any of the characteristics remains below 1.6%.

Altogether, our empirical analysis confirms our prediction that the unconstrained demand is the main driver of the high housing share across non-landlord homeowners. It also reveals that the unconstrained demand is particularly strong among less fortunate and larger households, while the constrained demand is higher for younger households with low income and large size.

4.5 Robustness Analysis

We now present three important robustness tests that all confirm the predominance of the unconstrained demand (additional details can be found in the Internet Appendix).

Linear Specification. One potential concern with the baseline specification is that the true relationship between the set of characteristics and the housing share is not linear. If the characteristic values for landlords and non-landlords span different intervals, the local

¹²The lower t -statistic for s_U is simply due to the fact that the number of landlords n_L is smaller than the number of non-landlords n_{NL} . The intuition can be easily illustrated with the following example. Suppose that we estimate the average housing demand for landlords ($\bar{\alpha}_U$) and for non-landlords ($\bar{\alpha}_T = \bar{\alpha}_U + \bar{\alpha}_C$). Using the standard variance formula, we find that $\bar{\alpha}_U$ and $\bar{\alpha}_C$ have respective variances proportional to $\frac{1}{n_L}$ and $\frac{1}{n_L} + \frac{1}{n_{NL}}$, whereas $\bar{\alpha}_T$ has a variance proportional to $\frac{1}{n_{NL}}$.

slope s_U estimated from landlord data could differ from the slope of the demand curve for non-landlords. This concern is most relevant for wealth because (i) it has the highest explanatory power among all characteristics, and (ii) landlords tend to be wealthier than non-landlords. To examine the potential non-linearity of the wealth effect, we perform a simple non-parametric regression of the housing share on wealth for both landlords and non-landlords. This analysis reveals two main insights. First, 80% of the wealth observations for non-landlords fall in the interval computed with landlords' wealth data. We therefore rule out the concern that the group of landlords is a non-representative sample clustered at the top of the wealth distribution. Second, we document a mild degree of non-linearity as the slope coefficient becomes flatter for low levels of wealth. If we control for this non-linearity by adding interaction terms for the less fortunate households, we find that the variance decomposition remains unchanged (78% vs 7% for the unconstrained and constrained demands).

Infrequent Rebalancing. Another potential concern is that the negative wealth effect is due to high trading costs that prevent households from rebalancing their housing shares frequently. A short-term increase in the level of financial wealth that is unrelated to housing could mechanically lead to a low housing share and vice-versa. We estimate two alternate specifications that eliminate this mechanical link. First, we replace gross wealth with its lagged value in the set of characteristics. Second, we only retain the observations following a home purchase. In both cases, the baseline results remain unchanged.

Omitted Variables. To control for potential omitted variables, we make use of an additional database in the PSID that indicates where each survey participant lives between 1968 and 2008 (at the Census Tract level).¹³ Combining this information with data from the U.S. Census Bureau, we add to equation (7) a new 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 include dummies regarding the ethnicity of each household (White/Caucasian, Asian, African-American, Hispanic), as well as the local fraction of households of each ethnic background. The introduction of these variables leaves our main results unchanged and slightly reduces the fit of the regression.

¹³This location data at the census tract level is available only by special request.

5 Housing in a Unified Portfolio Choice Framework

Our theoretical and empirical findings about the predominance of the unconstrained demand imply that real estate must have unique investment appeal for homeowners to justify their under-diversified portfolio. We now integrate these findings into a unified portfolio choice framework and provide new insights into why the unconstrained housing demand is so large. We solve a stylized model in the style of Merton (1971) where housing makes up part of the consumption basket (all the derivations can be found in the Internet Appendix).¹⁴ The model shows that the combination of multiple periods and multiple goods is sufficient to produce (i) a portfolio that is skewed toward real estate, (ii) a housing demand that is primarily driven by the unconstrained demand, and (iii) cross-sectional variation in the unconstrained demand that is in line with our empirical findings.

5.1 Setup

We consider a dynamic setting in which the household consumes a basket of two goods: a perishable non-housing good (C_t) and a durable housing good (K_t). At each time t , the utility derived by the household is given by the Cobb-Douglas function

$$U(C_t, K_t) = \frac{1}{1-\gamma} \left(C_t^\beta K_t^{1-\beta} \right)^{1-\gamma}, \quad (8)$$

where γ is the coefficient of relative risk aversion over the entire consumption basket and β is the degree to which the household values non-housing consumption relative to housing consumption. 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. If we further assume that the utility stream is additively separable and the time horizon infinite, the household's lifetime expected utility is given by

$$E \left[\int_0^\infty e^{-\delta s} U(C_s, K_s) ds \right], \quad (9)$$

¹⁴Our setup is similar to that of Damgaard, Fuglsbjerg and Munk (2003) who examine the optimization problem of a homeowner constrained by homeownership. Our setting differs in two dimensions. First, we reformulate their solution to decompose the total housing demand α_H^T into its unconstrained and constrained components α_H^U and α_H^C . Second, we also solve the optimization problem in a frictionless setting to understand the drivers of the unconstrained demand α_H^U .

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

The household 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

$$\frac{dP_{S,t}}{P_{S,t}} = \mu_S dt + \sigma_S dZ_{1,t}, \quad (10)$$

where μ_S and σ_S are constants and $Z_{1,t}$ is a Wiener process. The household can also invest H_t units in housing whose price $P_{H,t}$ is stochastic,

$$\frac{dP_{H,t}}{P_{H,t}} = \mu_H dt + \sigma_{H,1} dZ_{1,t} + \sigma_{H,2} dZ_{2,t}, \quad (11)$$

where μ_H , $\sigma_{H,1}$, and $\sigma_{H,2}$ are constants and $Z_{2,t}$ is a Wiener process that is uncorrelated with $Z_{1,t}$. We assume that, as a durable good, housing does not depreciate over time so that $dH_t = 0$ if the household does not trade. Therefore, the expected return on housing can be viewed as net of maintenance costs. Each unit of housing invested also provides a rental dividend. We assume a constant rent-price ratio ρ .

If we define $\alpha_{S,t}$, $\alpha_{H,t}$, and $\alpha_{B,t}$ as the portfolio shares invested in stocks, housing, and bonds at time t , the wealth dynamics can be written as

$$\begin{aligned} dW_t = & \left[rW_t + \begin{pmatrix} \alpha_{S,t} & \alpha_{H,t} \end{pmatrix} \begin{pmatrix} \mu_S - r \\ \mu_H - r \end{pmatrix} W_t - C_t - (K_t - H_t) \rho P_{H,t} \right] dt \\ & + W_t \begin{pmatrix} \alpha_{S,t} & \alpha_{H,t} \end{pmatrix} \begin{pmatrix} \sigma_{S,1} & 0 \\ \sigma_{H,1} & \sigma_{H,2} \end{pmatrix} \begin{pmatrix} dZ_{1,t} \\ dZ_{2,t} \end{pmatrix}. \end{aligned} \quad (12)$$

5.2 The Distinct Investment Value of Real Estate

To examine the determinants of the unconstrained demand, we focus on the frictionless setting where the household can freely choose the levels of housing consumption and investment via partial ownership agreements. Given the infinite horizon, we eliminate the time subscript from now on for notational convenience. The optimal portfolio shares are given in the following proposition.

Proposition 1. *The optimal shares in stocks, housing, and bonds in the frictionless case satisfy:*

$$\alpha_S^U = \alpha_S^{MV}, \quad (13)$$

$$\alpha_H^U = \alpha_H^{MV} + \alpha_H^H, \quad (14)$$

$$\alpha_B^U = 1 - \alpha_S^U - \alpha_H^U, \quad (15)$$

where α_S^{MV} and α_H^{MV} are the stock and housing shares in the mean-variance tangency portfolio (out of total wealth):

$$\begin{pmatrix} \alpha_S^{MV} \\ \alpha_H^{MV} \end{pmatrix} = \frac{1}{\gamma} \left[\begin{pmatrix} \sigma_{S,1} & 0 \\ \sigma_{H,1} & \sigma_{H,2} \end{pmatrix} \begin{pmatrix} \sigma_{S,1} & 0 \\ \sigma_{H,1} & \sigma_{H,2} \end{pmatrix}' \right]^{-1} \begin{pmatrix} \mu_S - r \\ \mu_H + \rho - r \end{pmatrix}, \quad (16)$$

and α_H^H is an additional hedging demand for housing:

$$\alpha_H^H = (1 - \beta) \left(1 - \frac{1}{\gamma} \right). \quad (17)$$

In this dynamic setting with multiple consumption goods, real estate has a high investment value because it brings dual benefits. First, housing provides speculative benefits by improving the risk-return trade-off of the overall portfolio. These benefits are captured the weight α_{MV}^H which determines the importance of housing in the mean-variance tangency portfolio. Second, housing provides hedging benefits regarding future housing consumption, which are captured by the weight α_H^H . It allows households to hedge against fluctuations in the future rental price.¹⁵ The combination of these investment benefits sets housing apart from traditional assets and provides a rationale for its predominance in the portfolio.

The reason for why real estate provides dual investment benefits is that it is equivalent to a perpetual indexed bond, i.e., it yields a perpetual and fixed stream of housing consumption H_t . In his analysis of the properties of indexed bonds, Fisher (1974) provides an intuitive explanation for their dual investment value: “Where there are many goods and many indexed bonds, 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

¹⁵Ortalo-Magné and Rady (2002) and Sinai and Souleles (2005) emphasize the hedging value of housing as a key driver of the decision to own or rent a home.

rise relatively rapidly, given equal real rates on all bonds. The presence of the speculative demands means that the indexed bonds are not held in the portfolio in the same proportion as their expenditure shares.”

The nature of real estate as an indexed bond on housing consumption explains why the hedging demand only depends on the two preference parameters β and γ . The household has a strong hedging demand if it cares a lot about housing consumption ($1-\beta$ is high) and is highly risk averse (γ is high).¹⁶ Properties of the house price are not relevant because they do not affect the steady stream of housing consumption that real estate provides. For example, an infinitely risk-averse investor should simply invest a fraction β in bonds and a fraction $1 - \beta$ in real estate, regardless of their risk-return profiles.

5.3 Predictions for Homeowners’ Portfolios

We now solve the homeowner’s problem by imposing the homeownership constraint ($H_t = K_t$), and show that the model can reproduce our previous findings regarding the predominance of the unconstrained housing demand and its cross-sectional variation among homeowners. The solution to this constrained optimization decomposes the total housing demand into its unconstrained and constrained components α_H^U and α_H^C . The homeowner’s optimal portfolio shares are given in the following proposition.

Proposition 2. *The homeowner’s optimal shares in stocks, housing, and bonds satisfy:*

$$\alpha_S^T = \alpha_S^U - \frac{\sigma_{H,1}}{\sigma_{S,1}} \alpha_H^C, \quad (18)$$

$$\alpha_H^T = \alpha_H^U + \alpha_H^C, \quad (19)$$

$$\alpha_B^T = \alpha_B^U - \left(1 - \frac{\sigma_{H,1}}{\sigma_{S,1}}\right) \alpha_H^C, \quad (20)$$

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

$$A_q (\alpha_H^C)^2 + B_q \alpha_H^C + C_q = 0, \quad (21)$$

and the parameters A_q , B_q , and C_q are defined as:

$$A_q = \frac{\gamma \sigma_{H_2}^2}{\rho} \left(1 + \frac{(1-\beta)(1-\gamma)}{2\gamma}\right), \quad B_q = 1 + \frac{\alpha_H^U \sigma_{H_2}^2}{\rho} (1 - \beta(1-\gamma)), \quad C_q = \alpha_H^U \left(1 - \frac{1}{\phi^U}\right).$$

¹⁶If β equals one, there is no hedging demand and housing enters the portfolio like a typical stock, representing a constant share of the tangency portfolio.

To examine the predictions of the model, we compute the portfolio shares with calibrated input data summarized in Table III (additional details on the data are reported in the Internet Appendix). The asset parameters are based on annual log real return data on stocks, short-term bonds, and real estate. We select 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. For each asset, we compute real returns from 1968 to 2013 when the PSID dataset is available using as deflator an inflation index that excludes the price of the housing good (as in the model). 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 real price growth of housing is equal to 2.6% per year and is adjusted for annual depreciation costs of 1.4% based on the estimates of Leigh (1980). Its volatility and stock correlation are respectively equal to 16.3% and 0.8%.¹⁷ We set the rent-price ratio equal to 3.5% based on the recent estimates of Campbell, Davis, Gallin and Martin (2009).

– Table III here –

For the utility parameters, we set the baseline risk aversion coefficient γ equal to 5, similar to Cocco (2005) and Yao and Zhang (2005). Under the Cobb-Douglas specification, the housing preference parameter $1 - \beta$ is equal to the fraction of consumer expenditures spent on housing consumption. We set this fraction equal to 0.25, which corresponds to the estimate in Davis and Ortalo-Magné (2011). Finally, we set the time discount factor δ equal to 0.28 to match the average ownership stake of 0.86 measured across all renters, homeowners, and landlords in the PSID, which implies that the homeownership constraint is binding ($\phi^U = 0.86 < 1$).¹⁸

Predominance of the Housing Demand The results show that real estate predominates in the homeowner’s portfolio, which is consistent with the empirical evidence documented

¹⁷These moments are computed as the median of the mean, volatility, and stock correlation of the price growth across individual homes in the PSID. This database departs from the Case-Shiller indices which underestimate the price volatility of individual homes.

¹⁸While convenient, the infinite horizon assumption in the model typically generates low housing consumption, which makes the homeownership constraint non-binding ($\phi^U > 1$ and $\alpha_H^C = 0$). Choosing a higher value of δ than in previous model calibrations allows us to decrease the household’s effective horizon (see Chapter 2.2.3 of Campbell and Viceira, 2002) and match the consumption-investment ratio in the data.

in Section 2. In Table IV (first row), we find that the total share invested in housing in the baseline case represents 50% of total wealth, versus 18.5% and 31.5% for stocks and short-term bonds. The tilt towards real estate is preserved when we consider homeowners with different levels of risk aversion γ and housing preference $1 - \beta$. As Table V shows, the total housing share ranges from 39% to 70% of the portfolio for γ between 3 and 7 (Panel A) and for $1 - \beta$ between 0.15 and 0.35 (Panel B).

– **Tables IV and Table V here** –

Predominance of the Unconstrained Demand The model also confirms that the unconstrained housing demand represents the bulk of the total housing demand. In the baseline case, Table IV (second row) shows that the unconstrained demand is equal to 47% of the portfolio, which represents 94% of the total demand (i.e., $47/50$) in the baseline case. This result comes from the dual investment value of real estate. The household invests 27% of its wealth to harness the speculative benefits of housing, which is consistent with the arguments of Goetzmann and Ibbotson (1990) and Goetzmann (1993) that housing provides substantial diversification gains because of its low correlation with the stock market. The homeowner also invests another 20% to harness the hedging benefits of real estate. The combination of both benefits therefore leads to an unconstrained housing demand that exceeds the demand for stocks.

In comparison, the constrained demand only represents 6% of the total housing demand. This result is close to our rule-of-thumb approximation in Equation (6) which predicts that the constrained demand should represent 7% of the total housing demand ($= 1 - \sqrt{0.86}$). As discussed in Section 3.2, the constrained demand is small because (i) the homeowners' optimal ownership stake in housing is high ($\phi^U = 0.86$ in the baseline case), and (ii) the homeownership constraint forces homeowners to reduce their housing consumption (i.e., move to a smaller home). This second point is shown in Table IV (fourth column) where the optimal level of housing consumption drops from 55% to 50% of total wealth when the homeownership constraint is imposed.

The predominance of the unconstrained demand holds across different types of homeowners, even those that are substantially more constrained by homeownership. For instance,

consider the household with a large housing preference parameter $1 - \beta$ equal to 0.35. Table V (Panel B) shows that this household would like to own only 70% of its home ($\phi^U = 0.7$) because its housing consumption is significantly larger than its housing investment. Yet its constrained demand represents less than 15% of the total housing demand – a result that is again predicted by our rule-of-thumb approximation ($1 - \sqrt{0.70} = 16\%$).

Cross-Sectional Variation in the Unconstrained Demand. Finally, the model is consistent with our empirical findings showing that the unconstrained demand is larger among less fortunate and larger households. As documented by Calvet and Sodini (2014), households that are less fortunate tend to be more risk averse. Table V (Panel A) shows that an increase in risk aversion produce two opposite changes: (i) an increase in the hedging demand; (ii) a decrease in the speculative demand. Empirically, the variation in the speculative demand is likely to be muted because investors account for estimation risk on the asset parameters (Jorion, 1986). In contrast, the hedging demand is independent of these parameters, which means that the unconstrained demand is likely to be greater for the more risk-averse households.

As for household size, larger households are likely to have larger housing consumption needs. Table V (Panel B) reveals that an increase in the housing preference parameter $1 - \beta$ has a positive impact on the hedging demand but leaves the speculative demand unchanged. The model therefore predicts that the unconstrained demand should be large for households with large consumption needs, which is in line with the empirical evidence.

6 Conclusion

Despite the decades of financial research advocating the importance of diversifying one’s portfolio, most households continue to invest the bulk of their wealth in just one asset: their home. Nevertheless, real estate has received relatively little attention in the literature on asset allocation because its large portfolio share is often viewed as the consequence of non-investment-related motives. In this paper, we tease apart the relative importance of the consumption and investment benefits that real estate provides. We show that existing homeowners would still invest the bulk of their wealth in real estate in a frictionless setting

where partial ownership arrangements are easily available. In other words, real estate *must* have unique investment appeal to justify their under-diversified portfolios.

Our findings call for further research on why the unconstrained housing demand is so large. It is critical to better understand how one asset can possibly account for the bulk of household wealth. While our model highlights the distinct role of real estate as a perpetual indexed bond on housing consumption, it does not incorporate other features such as labor income risk, taxes, or transaction costs that likely contribute to the large unconstrained demand. It would be valuable to quantify these effects using a more realistic life-cycle model in the style of Cocco (2005) and Yao and Zhang (2005).

The distinct nature of housing as a risk-free asset also has implications for research on asset pricing anomalies (e.g. Mehra and Prescott, 1985). Typical consumption-based asset pricing models assume a risk-free asset in zero-net supply. However, housing is an asset in positive supply that provides risk-free benefits. As Parlour, Stanton and Walden (2011) demonstrate, introducing a small risk-free consumption stream to an endowment economy has a substantial impact on the levels of the market risk premium, the risk-free rate, and the asset volatilities. The development of ‘multiple-tree’ models that incorporate real estate – building on the recent work of Lustig and Van Nieuwerburgh (2005), Martin (2013), and Piazzesi, Schneider and Tuzel (2007) – represents another promising avenue of future research.

References

- Betermier, S., Calvet, L. and Sodini, P. (2015), Who are the value and growth investors?, Working paper. forthcoming *Journal of Finance*.
- 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. (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.
- Chetty, R. and Szeidl, A. (2007), ‘Consumption commitments and risk preferences’, *Quarterly Journal of Economics* **122**(2), 831–877.
- Chetty, R. and Szeidl, A. (2014), The effect of housing on portfolio choice. Working Paper.
- 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.
- Davis, M. and Ortalo-Magné, F. (2011), ‘Household expenditures, wages, rents’, *Review of Economic Dynamics* **14**, 248–261.
- Eichholtz, P., Koedijk, K. and de Roon, F. (2000), The portfolio implications of home ownership. Working paper.
- Englund, P., Hwang, M. and Quigley, J. (2002), ‘Hedging housing risk’, *Journal of Real Estate Finance and Economics* **24**, 167–200.
- 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.
- 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.
- Greene, W. (2011), *Econometric Analysis*, Pearson; 7th Edition.
- 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.
- 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.
- Jorion, P. (1986), ‘Bayes-stein estimation for portfolio analysis’, *Journal of Financial and Quantitative Analysis* **21**, 279–292.
- Kraft, H. and Munk, C. (2011), ‘Optimal housing, consumption, and investment decisions over the life-cycle’, *Management Science* **57**(6), 1025–1041.
- 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.
- 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.
- Martin, I. (2013), ‘The Lucas Orchard’, *Econometrica* **81**, 55–111.
- Mehra, R. and Prescott, E. (1985), ‘The equity premium: A puzzle’, *The Journal of Monetary Economics* **15**, 145–161.
- Merton, R. (1971), ‘Optimum consumption and portfolio rules in a continuous-time model’, *Journal of Economic Theory* **3**, 373–413.
- Ortalo-Magné, F. and Rady, S. (2002), ‘Tenure choice and the riskiness of non-housing consumption’, *Journal of Housing Economics* **11**, 266–279.
- 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.
- Poterba, J. and Samwick, A. (2001), Household portfolio allocation over the life cycle, *in* S. Ogura, T. Tachibanaki and D. Wise, eds, ‘Aging Issues in the United States and Japan’, Chicago: University of Chicago Press, pp. 65–103.
- 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.
- Tracy, J., Schneider, H. and Chan, S. (1999), 'Are stocks overtaking real estate in household portfolios?', *Federal Reserve Bank of New York Current Issues in Economics and Finance* **5**(5), 1–6.
- Van Hemert, O. (2010), 'Household interest rate risk management', *Real Estate Economics* **38**(3), 467–505.
- Vestman, R. (2012), Limited stock market participation among renters and homeowners. Working Paper, Stockholm University.
- Yamashita, T. (2003), 'Owner-occupied housing and investment in stocks: An empirical test', *Journal of Urban Economics* **53**, 220–237.
- 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
Summary Statistics for Homeowners

The table reports summary statistics on the financial and demographic characteristics (Panel A) and wealth characteristics (Panel B) of homeowners from the 2013 wave of the PSID. We consider all homeowners (first set of columns), the subset of landlords (second set of columns), and the subset of non-landlord homeowners (third set of column). Landlords are defined as homeowners that received rental income on real estate properties during the year prior to the survey. Additional details on the reported variables can be found in the Internet Appendix.

Panel A: Financial and Demographic Characteristics						
	All Homeowners		Landlords		Non-Landlords	
	Mean	Standard deviation	Mean	Standard deviation	Mean	Standard deviation
<i>Financial characteristics</i>						
Household income (\$)	73,525	62,492	92,455	57,818	71,832	62,626
Gross worth (\$)	389,042	603,655	876,574	1,183,893	345,449	499,058
<i>Demographic Characteristics</i>						
Age	53.65	14.92	58.84	14.62	53.19	14.86
Household size	2.75	1.32	2.46	1.23	2.78	1.33
Married dummy	0.80	0.40	0.82	0.39	0.80	0.40
High school dummy	0.28	0.45	0.26	0.44	0.28	0.45
Post-high school dummy	0.24	0.43	0.21	0.41	0.24	0.43
Number of observations	2,522	2,522	207	207	2,315	2,315

Table I
Summary Statistics for Homeowners - *Continued*

Panel B: Wealth Characteristics						
	All Homeowners		Landlords		Non-Landlords	
	Mean	Standard deviation	Mean	Standard deviation	Mean	Standard deviation
<i>Gross worth</i>						
Financial worth	0.23	0.24	0.24	0.22	0.23	0.24
Real estate	0.67	0.25	0.64	0.25	0.68	0.25
Business	0.03	0.11	0.08	0.18	0.02	0.10
Autos	0.07	0.09	0.04	0.04	0.07	0.09
<i>Debt (relative to Gross Worth)</i>						
Mortgage	0.34	0.63	0.19	0.25	0.36	0.66
Other debts	0.06	0.33	0.01	0.03	0.07	0.34
<i>Financial worth</i>						
Stocks	0.10	0.24	0.18	0.28	0.10	0.23
IRA	0.25	0.36	0.35	0.37	0.25	0.36
Bonds	0.06	0.19	0.07	0.19	0.06	0.19
Cash	0.58	0.42	0.40	0.39	0.60	0.42
Number of observations	2,522	2,522	207	207	2,315	2,315

Table II
Panel Regression of Housing Shares

The table reports the panel regression of the housing share $\alpha_{H,i,t}$ 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,

$$\alpha_{H,i,t} = a_t + s'_U X_{i,t} + (b_t + s'_C X_{i,t}) D_{i,t}^{NL} + \varepsilon_{i,t},$$

where $D_{i,t}^{NL}$ is a dummy variable that equals 1 if the household is a non-landlord homeowner and 0 otherwise. We refer to s_U and s_C as the slope vectors of the unconstrained and constrained components of the housing demand. We refer to $s_T = s_U + s_C$ as the slope vector of the total housing demand. Panel A reports, for each demand component, (i) the variance ratio of its predicted values relative to those of the total demand (e.g., $\text{var}(s'_U X_{i,t}) / \text{var}(s'_T X_{i,t})$), (ii) the standard deviation of its predicted values (e.g., $\text{stdev}(s'_U X_{i,t})$), and (iii) the statistical significance of the hypothesis test that its slope coefficients are jointly equal to zero, (e.g., $H_0: s_U = 0$). Panel B reports the individual slope coefficients for s_T , 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, and standardize each coefficient by the standard deviation of its associated characteristic among non-landlords.

Panel A: Variance Decomposition and Joint Tests

	Ratio of Variances	Std. Dev.	Joint Test	
			F-stat	p-value
Total demand	100%	0.123	104.030	<.0001
Unconstrained demand	80%	0.110	24.160	<.0001
Constrained demand	6%	0.031	1.900	0.047

Panel B: Regression Slope Coefficients of the Housing Share

	s_T		s_U		s_C	
	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat
Financial Characteristics						
Household income (log)	-0.001	-0.45	0.009	0.88	-0.011	-1.08
Gross worth (log)	-0.115	-29.83	-0.113	-11.31	-0.003	-0.24
Demographic Characteristics						
Age	-0.019	-6.08	-0.008	-0.73	-0.012	-1.05
Household size	0.035	13.25	0.020	2.19	0.016	1.76
Married dummy	-0.015	-4.76	-0.019	-1.89	0.004	0.38
High school dummy	-0.011	-4.10	-0.024	-2.63	0.013	1.47
Post-high school dummy	0.000	-0.17	-0.008	-1.05	0.008	1.05
House Price Characteristics						
Average growth	0.024	5.92	0.011	1.13	0.013	1.19
Volatility	-0.032	-7.72	-0.018	-1.88	-0.015	-1.41
Adjusted R-square	27.43%					
Number of observations	23,566					

Table III
Moments of Asset Returns and Baseline Parameters

The table reports the moments of asset returns between 1968 and 2013 (Panel A) and the corresponding parameters for the baseline calibration (Panel B). In Panel A, the return on short-term bonds is the 1-month Treasury Bill rate from CRSP, the portfolio of stocks is the value-weighted NYSE/AMEX/NASDAQ index from CRSP, and housing represents a typical home from the Panel Study of Income Dynamics. The moments are computed using log real annual returns, where the deflator is an inflation index of the non-housing good. For housing, we track the annual price evolution of all houses reported by US households that have at least ten consecutive annual price observations. We then take the median of the mean, volatility, and stock correlation of the real price growth across all individual homes. The parameters in Panel B are chosen to match the moments in Panel A.

Panel A: Moments of Asset Returns, 1968-2013		
Asset	Estimate	
Bonds		
Average	0.011	
Stocks		
Average return	0.045	
Volatility	0.191	
Housing		
Average price growth	0.026	
Volatility	0.163	
Correlation with stock market	0.008	
Panel B: Baseline Parameters		
	Symbol	Parameter
<i>Asset parameters</i>		
Risk-free rate (B)	r	0.011
Stocks (S)		
Average return	μ_S	0.045
Exposure to shock 1	$\sigma_{S,1}$	0.191
Exposure to shock 2	$\sigma_{S,2}$	0.000
Housing (H)		
Average price growth net of depreciation	μ_H	0.012
Exposure to shock 1	$\sigma_{H,1}$	0.001
Exposure to shock 2	$\sigma_{H,2}$	0.163
Rent-price ratio	ρ	0.035
<i>Utility parameters</i>		
Risk aversion	γ	5.000
Time discount factor	δ	0.279
Preference for housing	$1-\beta$	0.250

Table IV
Baseline Portfolio Allocation

The table reports the homeowner's optimal portfolio and housing consumption decisions according to the model. In the first row (total demand), we calculate the portfolio shares in short-term bonds, stocks, and housing chosen by the homeowner. We also report its housing consumption decisions as measured by the ratio of housing consumption to wealth and the ratio of housing investment to housing consumption. In the second row (unconstrained demand), we re-calculate the same statistics assuming that the household is unconstrained by homeownership. The third and fourth rows decompose the unconstrained demand for stocks and housing into a speculative demand and a hedging demand. The last row reports the constrained demand due to homeownership. The input parameters for the calibration are provided in Table III.

	Portfolio Share			Housing Consumption Ratios	
	Bonds	Stocks	Housing	Consumption / Wealth	Investment / Consumption
Total demand	0.32	0.18	0.50	0.50	1.00
Unconstrained demand	0.35	0.18	0.47	0.55	0.86
Speculative demand		0.18	0.27		
Hedging demand		0.00	0.20		
Constrained demand	-0.03	0.00	0.03	-0.05	0.14

Table V
Cross-Sectional Variation in the Housing Demand

The table reports the different components of the demand for housing for five values of the risk aversion parameter (Panel A) and the preference for housing consumption parameter (Panel B) centered around the baseline values. All the other parameters are set equal to the baseline values from Table III. In the first row of each Panel (total demand), we report the housing share of the homeowner. In the second row (unconstrained demand), we re-calculate the housing share assuming that the household is unconstrained by homeownership. The third and fourth rows decompose the unconstrained demand for housing into a speculative demand and a hedging demand. The fifth row represents the constrained demand for housing due to homeownership. The final row reports the optimal ratio of housing investment to housing consumption in the unconstrained case (Φ^U).

Panel A: Risk Aversion					
	Risk aversion (γ)				
	3	4	5	6	7
Total demand	0.70	0.57	0.50	0.45	0.42
Unconstrained demand	0.62	0.52	0.47	0.43	0.41
Speculative demand	0.45	0.34	0.27	0.22	0.19
Hedging demand	0.17	0.19	0.20	0.21	0.21
Constrained demand	0.08	0.05	0.03	0.02	0.01
Investment / Consumption					
Unconstrained Ratio	0.76	0.81	0.86	0.90	0.94
Panel B: Preference for Housing Consumption					
	Preference for housing ($1-\beta$)				
	0.15	0.20	0.25	0.30	0.35
Total demand	0.39	0.43	0.50	0.57	0.63
Unconstrained demand	0.39	0.43	0.47	0.51	0.55
Speculative demand	0.27	0.27	0.27	0.27	0.27
Hedging demand	0.12	0.16	0.20	0.24	0.28
Constrained demand	0.00	0.00	0.03	0.06	0.08
Investment / Consumption					
Unconstrained Ratio	1.22	1.00	0.86	0.77	0.70

Figure 1
Portfolio Shares by Wealth Group

The figure reports the average composition of gross worth for different groups of households sorted by their gross worth. The results are based on a representative sample of households from the PSID between 1984 to 2013. For each year in the survey, we sort households into 30 groups based on their gross worth and we compute the (equally-weighted) average shares in bonds, stocks, real estate, business, and autos. We then compute the time-series average of each share over the years when the survey is conducted.




Figure 2 Portfolio Shares Across Time

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




Figure 3
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 in bonds, stocks, real estate, business, and autos. We then compute the time-series average of each share over the years when the survey is conducted between 1984 and 2009.




Figure 4 Feasible Opportunity Set in the Real Estate Market

This Figure illustrates the feasible options in the housing market, as a function of the investment-consumption ratio Φ . The dotted line represents the set of unavailable partial ownership opportunities.




Figure 5
Decomposition of the Housing Demand

The figure reports the ratio of the unconstrained demand to the total demand for housing as a function of the ratio of housing investment to housing consumption (Φ^U) chosen by the household in the frictionless case. We assume that the switch from full homeownership to partial ownership occurs via an equal adjustment in the levels of housing investment and consumption.




Figure 6

Illustration of the Cross-Sectional Variation in the Housing Demand

This figure illustrates the cross-sectional variation in the housing demand between two households that are constrained to be homeowners. Household *l* has a low unconstrained housing demand and household *h* has a high unconstrained housing demand. We plot their unconstrained demand (dotted line) relative to the demand of household *l*. We also plot their total housing demand (solid line), which accounts for the impact of the homeownership constraint.

