Housing Demand During the Boom: The Role of Expectations and Credit Constraints

Tim Landvoigt*

UT Austin, McCombs School Of Business

December 2014

Abstract

Optimism about future house price appreciation and loose credit constraints are commonly considered drivers of the recent housing boom. This paper infers both mean and variance of short-run expectations of future house price growth and minimum down payment requirements from observed household choices. The expectations and credit constraints are implied by a life-cycle portfolio choice model that encompasses home ownership, housing demand, and financing choices. I estimate the parameters of this model using data from the Survey of Consumer Finances from 1995 to 2010. The main result is that both expectations of future mean price growth and minimum down payment requirements were close to their long-run averages during the boom. Subjective uncertainty about the house price growth rate, however, was increasing. Expectations and credit constraints are separately identified due to their differential effects on the intensive and extensive margins of housing demand. The increase in uncertainty about future prices helps to explain the rise in household debt. Given the option to default, greater expected volatility leads to higher optimal leverage.

*tim.landvoigt@mccombs.utexas.edu. This paper is based on a chapter of my dissertation at Stanford University. I am grateful to Martin Schneider and Monika Piazzesi for their guidance with this project. I have also benefited from comments by Manual Amador, Tim Bresnahan, Max Floetotto, Bob Hall, Nir Jaimovich, Seema Jayachandran, Ken Judd, Pete Klenow, Victor Rios-Rull, Johannes Stroebel, and Michèle Tertilt, as well as participants of the Stanford Macro Lunch Seminar and the SED 2010 Annual Meetings.

1 Introduction

During the recent housing boom, a sharp increase in house prices relative to rents coincided with low interest rates and loose lending standards in the mortgage market. These features characterized the economic environment that households were facing when they were deciding whether to rent or own, whether to upgrade or downgrade the house they own, or whether to re-finance their mortgage. A generally unobserved but important determinant of such decisions is the expected rate of house price growth that households assume in their decision-making process. This paper examines the role expectations and credit constraints played in shaping household behavior during the boom by inferring short-run expectations of future house price growth and average minimum down payment requirements from observed household choices. The expectations and credit constraints are implied by a life-cycle portfolio choice model that encompasses home ownership, housing demand, and financing choices.

The goal of this paper is not to determine the cause of the boom, but rather to test whether the choices of the majority of households during the boom can be explained by a rational model with reasonable expectations about future prices. My approach further connects the financing side of observed household choices during the boom with the extensive and intensive margins of housing demand, i.e. the decision whether to rent or own and the amount of housing services consumed. To accomplish this, I solve a life-cyle portfolio choice model with housing, and use the optimal policies to estimate expectations and credit constraints with data from the Survey of Consumer Finances (SCF) for the period 1992 to 2010[1].

We can think of the optimal policies resulting from the dynamic program as a mapping from different dimensions of household heterogeneity - age, income, wealth, homeownership status, and the house owned initially - into choices for the next period - homeownership status, house value, consumption of housing services and numéraire, and the amount saved or borrowed. In other words, given a model of optimal housing demand, for which expectations about future house prices are an input parameter, it is possible to back out implied

[1]The SCF is only conducted every three years.

expectations from observed demand.

In order to perform this inference one must assume a structure for the path of household expectations. I divide household beliefs into short-term beliefs, which dictate expectations for the next period, and long-term beliefs that are based on long-run averages and apply to all subsequent life-cycle periods. At the same time, other time-varying variables, such as house price-to-rent ratios and interest rates, are set to their observed value for each period. This way I can use the model to trace short-term variation in these observable variables and estimate the matching short-term expectations and down payment requirements, while keeping long-term household beliefs about all variables set to long-run averages.

The main finding is that estimated household expectations were relatively close to average long-run house price growth (of 2.5% annually), with slightly higher expectations at the beginning and the end of the boom. Even though the estimated mean expectations are close to the long run average, the subjective volatility during the boom years is considerably higher with an estimated standard deviation of house price growth of 25% annually. The estimation also finds an increase of the short term down payment constraint from a value of 9% in 1998 to 18.5% in 2007 (as share of the house value at the time of the purchase), with the long run constraint set to 15%.

The separate identification of the two channels affecting household choices – expectations and down payment constraints – comes from their differential impact on the intensive and extensive margins of housing demand. In particular, changes of the down payment constraint have a quantitatively larger effect on the intensive margin. Changes in expected house price growth, however, more strongly affect the extensive margin of housing demand – the decision whether to own or rent. The key feature of the model for understanding these differential effects is the transaction cost of selling houses. This transaction cost causes inertia in the choices of existing home owners who are consequently less likely to adjust their housing demand in response to short-term fluctuations in the economic environment. Hence the identification of the estimated parameters mainly relies on the choices of young households in the data who are on the margin between renting and owning. These households are financially constrained, and increasing or decreasing their optimism about future house prices does not affect the size of the house they are able to afford (the intensive margin). On the other hand, tightening or relaxing the down payment constraint directly impacts the house size for those households with a binding constraint.

The estimated upward trend in credit constraints can be understood by comparing data averages to model-predicted averages. Particularly, absent any time variation in model parameters, the model-predicted house values rise almost one-for-one with house prices since transaction costs cause most home owners to hold on to their existing houses despite the large increase in prices. The combination of low mortgage rates and high realized house price gains would then inflate model-predicted house values during boom beyond data values if down payment constraints were relaxed in addition. Hence the slight tightening of the constraints during the boom is needed to match the evolution of house value-to-income ratios in the data. Everything else equal, the tighter constraints would depress model-implied home ownership rates. The role of the slight increase in expectations towards the end of the boom is thus to offset the effect of the tightened constraints on the ownership rate.

The estimates of subjective house price growth volatility are below the long run average at the beginning and above the long run average at the height of the boom. Increased volatility (i.e., a higher standard deviation of house price growth) leads to increased leverage in the model. This is because households with defaultable mortgages are effectively holding a call option on their houses, and the value of this option increases as house price volatility rises. Households then consume part of this greater option value through higher debt today. Hence the estimates of greater house price volatility are identified from the increase in household debt during the boom period[2].

Of course, several other aspects of the data such as time-variation in rent-to-price ratios, and in the joint distribution of wealth, income, and age also affect the estimates. However, the identification arguments outlined above capture the quantitatively most important channels that drive the estimates.

Survey evidence on return expectations for houses in the US during the boom years is limited. Case, Quigley, and Shiller (2003) performed mail surveys of home *buyers* in 2002. Their point estimates suggest high capital gains expectations among buyers – between 6 and

[2]The approximately constant leverage ratios during the boom imply a substantial increase in debt due to the large increase in house values.

11 percent per year – for different regions of the US, although they are rather imprecise. Contrary, Piazzesi and Schneider (2009) report based on the 2005 Michigan Survey of Consumers that the large majority of households expressed the view that buying a house is not a good idea, and only 20% of households expected future prices to be high. The estimates in this paper are consistent with both kinds of evidence to the extent that they represent *average* expectations across potentially more optimistic buyers, and less optimistic incumbent owners and renters. Furthermore, my estimates are not outright pessimistic expectations – they merely say that household choices imply expectations of moderate growth, both at the beginning and at the peak of the boom[3].

The finding of tighter constraints for the 2007-2010 period during the housing market bust is consistent with other evidence on tightening credit conditions during that period. However, the estimate of the 2004-07 constraint at 15% runs counter the well established notion that credit constraints were relaxed during the boom years. A lesson from the partial equilibrium exercise in this paper is that high *realized* house price gains and cheap credit in the form of low mortgage rates are sufficient to explain *average* value-to-income and loan-to-value ratios during the boom. This is consistent with the idea that the estimates represent average down payment requirements across different segments of the mortgage market, including conforming prime mortgages and subprime mortgages[4].

While quantitative survey evidence on mean expected house price gains is limited, there is even less quantitative evidence on uncertainty about future house prices during this period. The survey by Case, Quigley, and Shiller (2003) finds greater standard errors for the mean expectation of respondents in 2002 than in 1998, hinting at an increase in dispersion during the boom. Similarly, the numbers reported Piazzesi and Schneider (2009) from the 2005 Michigan Survey of Consumers are indicative of greater disagreement about future house prices during the late boom period. There was certainly a discussion among academic economists and in the media during 2004 and 2005 whether the large run-up in prices con-

[3]They also imply that households did not anticipate the large run-up in prices, but neither did they anticipate the bust.

[4]Foote, Gerardi, and Willen (2008) show that the median LTV at origination for subprime loans in Massachusetts reached 90% in 2005. Demyanyk and Hemert (2011) report an average LTV of 86% at origination for subprime loans in 2006. The 20% down payment requirement for conforming loans as defined by the GSEs remained constant throughout the boom period. Average loan-to-value ratios among all home owners as measured in the SCF are roughly constant between 36 and 38 percent throughout the boom period

stituted a bubble, see for example a special report in the Economist (2005), or studies by Case and Shiller (2005) and Himmelberg, Mayer, and Sinai (2005).

To summarize, the US housing boom of the 2000s was characterized by a large rise in house value-to-income ratios and mortgage debt, with a relatively stable aggregate leverage ratio and home ownership rate. The structural estimation exercise in this paper finds that neither overly optimistic expectations about future house prices nor extremely low down payment requirements are necessary to rationalize aggregate household choices over this period. In a frictionless model that allows to adjust housing consumption without cost, the large increase in house prices would cause many homeowners to substitute from housing consumption into other goods, and some additional force such as optimism would be needed to induce them to keep their now expensive houses. However, in a model with realistic housing transaction costs such as the one used in this paper, most homeowners simply hold on to their existing houses despite the rise in prices, and an increase in the aggregate house price translates into higher house value-to-income ratios almost one-for-one.

Further, high uncertainty about future house prices may have been a factor in contributing to the increase in household debt. In the model, greater volatility of future house price growth increases the value of the implicit call option consisting of a house and a defaultable mortgage. Hence the estimates show that even at times of growing house values and few observed defaults, the option of default may substantially affect household choices through second moments.

House prices are exogenous in this analysis, which therefore does not offer an explanation why house prices rose in the first place. It merely says that conditional on the realized path of house prices, interest rates, mortgage spreads, and rent-to-price ratios, household choices are implying expectations of moderate price growth. In any competitive equilibrium model that would generate the observed path of house prices and interest rates, the conclusions of the demand analysis in this paper would still be valid. Furthermore, the result of moderate expected growth but high uncertainty (in the sense of disagreement) is consistent with a theory of the boom that relies on a small subset of agents who are very optimistic and whose actions drive price growth during the boom, such as articulated in Geanakoplos (2010) or Piazzesi and Schneider (2009).

This paper is related to a large literature on the role of cheap credit during the housing boom. Several reduced form empirical studies including Mian and Sufi (2009), Mian and Sufi (2011) show that easier access to credit mattered for house prices at the regional level. Recent studies that embed household life-cycle models in a dynamic general equilibrium framework of the housing market to assess the importance of cheap credit include Kiyotaki, Michaelides, and Nikolov (2011) and Favilukis, Ludvigson, and Nieuwerburgh (2013). These papers have endogenous house prices and focus on the effect of relaxed credit constraints and easier access to credit more broadly on house prices[5]. Corbae and Quintin (2014) use an equilibrium model of the mortgage market to show that relaxation of payment-to-income constraints was crucial to explain the increase in mortgage debt. None of these papers estimate expectations or credit constraints to track the evolution of household choices over time. Rather, they calibrate these parameters based on evidence from other empirical studies and focus on equilibrium effects. Recent papers on the role of expectation formation in the housing market include Piazzesi and Schneider (2009), Burnside, Eichenbaum, and Rebelo (2011), and Glaeser, Gottlieb, and Gyourko (2010). These papers propose different theoretical mechanisms by which expectations of future house price gains may feed back to current house prices. Landvoigt, Piazzesi, and Schneider (2014) consider both the role of credit constraints and expectations in a equilibrium model of a local housing market. They show that loose credit constraints are important for explaining the high capital gains at the low end of the house quality distribution during the boom.

The model presented in this paper is similar to the models developed by Campbell and Cocco (2003), Cocco (2004), Yao and Zhang (2005). These papers focus on introducing housing as an additional asset in a portfolio choice setting with life-cycle labor income. They solve for optimal life-cycle positions of housing and other assets such as bonds and stocks; their emphasis is on analyzing the optimal policies for a given calibration that uses parameter values from the literature. Li, Liu, and Yao (2009) and Bajari, Chan, Krueger, and Miller (2013) perform a structural estimation of a life-cycle model with housing similar to the one in this paper, using data from the PSID. However, their focus is mainly on

[5]Other papers focusing on housing collateral constraints mainly from the perspective of risk-sharing in general equilibrium include Campbell and Hercowitz (2005), Lustig and Nieuwerburgh (2005), and Iacoviello and Pavan (2013).

using the fitted model to conduct experiments and predict future household behavior. While Li, Liu, and Yao (2009) focus on policy experiments about changes in lending conditions, Bajari, Chan, Krueger, and Miller (2013) are predicting the length and depth of the slump in the housing market. In contrast to the analysis of this paper, all of the papers listed above assume that household beliefs about future house prices are described by the same stochastic process over the life-cycle, irrespective of current economic conditions.

This paper proceeds as follows. Section 2 describes the model, discusses the assumptions, and outlines the computational solution method. Section 3 discusses the empirical strategy and its implementation, and states values of the calibrated parameters and data source. Section 4 presents the data moments entering the objective function and the estimation results. It further discusses the identification, and interprets the findings. Section 5 concludes.

2 Model

2.1 Household Problem

A household lives for years $t = 25, \dots, 100$, with a probability of survival from year $t-1$ to t of λ_t, and $\lambda_{T+1} = 0$. Every year until retirement at age $t_R = 65$, the household receives labor income Y_t that follows an exogenous stochastic process. After retirement, the household receives a constant fraction of its last labor income Y_{t_R} until death. The household chooses consumption of housing services S_t and other goods C_t (the numéraire) every year to maximize expected lifetime utility. The per-period utility function $u(C_t, S_t)$ is assumed to satisfy the usual properties of being strictly increasing and concave in its two arguments arguments. Lifetime utility is given by

$$\mathrm{E}_t \left\{ \sum_{t=0}^{T} \beta^t \left[\Lambda_t \lambda_{t+1} \, u(C_t, S_t) + \Lambda_t (1 - \lambda_{t+1}) \, B_t \right] \right\},$$

where B_t is the bequest the household leaves to its children in case it does not survive until year $t+1$, and

$$\Lambda_t = \prod_{s=0}^{t} \lambda_s$$

is the unconditional probability that the household is alive in year $t < T$.

Housing has the dual role of an asset that the household can save in, and a durable consumption good that generates housing services. Households can consume housing services in two ways: they can either own or rent a house. The variable $\tau_t \in \{0, 1\}$ represents a household's decision whether to be a home owner or not in year t, with $\tau_t = 1$ indicating ownership. A house of size H_t produces housing services with the linear technology

$$S_t = \Phi(\tau_t, a_t)H_t, \tag{1}$$

where a_t is the age of the household in year t. The housing services production coefficient $\Phi(\cdot)$ generally depends on the home ownership status τ_t and age a_t. It captures age-dependent aspects of the preference for ownership that are not directly contained in this model, such as changes in household size and uncertainty about future household size. A unit of the housing asset sells for P_t units of numéraire, and can be rented for P_t^r in the rental market.

The household assumes that labor income and house price follow a Markov process with transition rule

$$[Y_t,\, P_t] = F\left([Y_{t-1},\, P_{t-1}],\, \epsilon_t\right), \tag{2}$$

where ϵ_t is a two-dimensional random vector distributed independently over time. I will specify the exact form of the transition rule below.

The rental price is pegged to the asset price through a deterministic, but potentially time-varying ratio

$$\alpha_t = \frac{P_t^r}{P_t}. \tag{3}$$

In addition to the housing asset, the household can save and borrow the amount L_t in a risk-free bond. By saving one unit of numéraire in the bond at $t-1$, the bond pays out $R_t > 1$ units at t. In order to borrow, the household has to own a house and use part of its value as collateral. In particular, when the household buys a house, it can at most borrow an amount $(1 - \delta_t)$ of the house value to finance the purchase, where δ_t is the fraction required as a down payment:

$$L_t \geq -(1 - \delta_t)P_t H_t. \tag{4}$$

Furthermore, the interest rate when borrowing is higher by a spread of ζ_t.

The budget constraint and the evolution of household wealth over time are best understood by distinguishing two cases. First, if the household did not own a house at age $t-1$, its liquid resources in period t consist of savings and interest from the previous period and current labor income. The household can use this wealth to consume the numéraire good, buy or rent units of the housing asset, and save in the risk-free asset. If the household decides to buy a house (i.e. purchase a positive amount of the housing asset), it can also borrow in the risk-free asset subject to constraint 4. Since the borrowing rate is higher than the rate for savings, the household will never optimally save and borrow at the same time. Thus it suffices to keep track of the net position L_t in the risk-free asset. This yields the following budget constraint for a household who was renting in period $t-1$

$$R_t L_{t-1} + Y_t = C_t + L_t + P_t H_t[(1-\tau_t)\alpha_t + \tau_t(1+\psi)], \tag{5}$$

subject to the down payment constraint 4, and using the fact that the rental price can be expressed in terms of the house price and the rent-to-price ratio α_t based on equation 3. The coefficient $(1+\psi)$ multiplying the expenditure on the new house in the last term accounts for a proportional maintenance cost $\psi P_t H_t$ that a homeowner must pay each period in order to offset depreciation.

The second case is that of a household who enters period t owning a house. The household may sell its current house in order to buy a new one of different size or rent instead. In this case, the sale requires payment of a transaction cost proportional to the house value, $\nu P_t H_{t-1}$. In general, the homeowner can decide to stay in the current house, and therefore not incur the transaction cost. Hence the home owner's liquid resources consist of savings and labor income as for the renter, plus the value of the house net of mortgage principal, interest, and the sales transactions cost. Denoting the decision whether to sell or keep the house by $\xi_t \in \{0,1\}$, with 0 indicating keeping the house and 1 selling, the constraint for the home owner is

$$\begin{aligned}(R_t + 1_{[L_{t-1}<0]}\zeta_t)L_{t-1} + Y_t + P_t H_{t-1} = & C_t + L_t + (1-\xi_t)P_t H_{t-1} \\ & \xi_t\{P_t H_t[(1-\tau_t)\alpha_t + \tau_t(1+\psi)] + \nu P_t H_{t-1}\}\end{aligned} \tag{6}$$

again subject to down payment constraint 4, and with τ_t indicating the ownership decision as in equation 5.

In addition to the decision whether to stay in the current house, sell and rent, or sell and buy, a home owner can also decide to default on its debt. In case of default, mortgage debt and home equity are erased, and the household incurs a cost of default κ that is proportional to its income. Hence this household's budget constraint is essentially that of a household who does not own a house coming into the period with income $(1-\kappa)Y_t$.[6] Denote the decision whether or not to default for a home owner by $d_t \in \{0,1\}$, with $d_t = 1$ indicating default.

Each household has to move with a certain probability every period, independent of all other shocks and previous periods. This shock is only relevant for home owners since it forces them to sell their house and incur the transaction cost. Renters sign period-by-period rental contracts, and thus their problem is unaffected. Let the outcome of this shock be denoted by $M_t \in \{0,1\}$, with 0 indicating that the household may keep the house and 1 that it must move.

The complete life-cycle optimization problem can be stated recursively using dynamic programming. Denote the vector of state variables at time t by $X_t = [M_t, P_t, a_t, \tau_{t-1}, H_{t-1}, Y_t, L_{t-1}]$, and the vector of choice variables $Z_t = [\tau_t, \xi_t, d_t, H_t, C_t, L_t]$. Then the value function at age $t = 0, \ldots, T-1$ is defined as

$$V_t(X_t) = \lambda_{t+1} \left\{ \max_{Z_t} u\left(C_t, \Phi(\tau_t, a_t) H_t\right) + \beta \mathrm{E}_t \left[V_{t+1}(X_{t+1}) \right] \right\} + (1-\lambda_{t+1}) B(X_t) \qquad (7)$$

subject to constraints 4, 5, and 6 and the transition equation for income and prices 2, and by

$$V_T(X_T) = B(X_T) \qquad (8)$$

for the final period.

To close the model, I still need to specify functional forms for the intra-period utility function $u(C_t, S_t)$ and the bequest function $B(X_t)$. For the utility function, I use the conventional Cobb-Douglas form for composite utility from housing services and other goods:

$$u(C_t, S_t) = \frac{\left[C_t^{1-\rho} \left(\Phi(\tau_t, a_t) H_t \right)^{\rho} \right]^{\gamma}}{1-\gamma}, \qquad (9)$$

where ρ determines the relative weight on housing services and γ is the risk-aversion parameter. The function $\Phi(\tau_t, a_t)$ that governs the age-dependent preference for renting is given by

$$\Phi(\tau_t, a_t) = 1 + (1-\tau_t) \exp(-\phi a_t).$$

[6] The constraint for the defaulting household is $(1-\kappa)Y_t = C_t + L_t + P_t H_t[(1-\tau_t)\alpha_t + \tau_t(1+\psi)]$.

with parameter ϕ. If $\phi > 0$, as will be the empirically relevant case, then the additional utility from renting is decreasing exponentially with age.

To specify bequest utility, it is helpful to first define liquid wealth after the potential sale of the housing asset as

$$W_t = (R_t + 1_{[L_{t-1}<0]}\zeta_t)L_{t-1} + \tau_{t-1}(1-\nu)P_t H_{t-1} + Y_t. \tag{10}$$

Bequest utility depends on liquid wealth in the household's final year and the current house price

$$B(W_t, P_t) = \bar{B}\frac{(W_t/P_t^{\rho})^{1-\gamma}}{1-\gamma}, \tag{11}$$

where $\bar{B}$ is a parameter that governs the strength of the bequest motive[7]

2.2 House Price and Labor Income Processes

Since the empirical analysis will involve cross-sections of households of different age cohorts, I will use the subscript t to index the calendar year, and i to index an individual household. The age of household i in year t will be denoted by a_{it}.

A crucial step in inferring household expectations from observed decisions is the modeling of household beliefs about future house prices. This involves specifying a parametric form for the transition rule $F([Y_{it-1},\ P_{it-1}],\ \epsilon_{it})$ in equation 2 for income and house prices. First, I assume that the individual house price follows a random walk in logs, i.e. the growth rate of the house price is

$$R_{it}^H \equiv \frac{P_{it}}{P_{it-1}} = \exp(m_{t-1} + \epsilon_{it}^H), \tag{12}$$

where ϵ_t^H is a random variable with zero mean, and m_{t-1} is the deterministic drift. As is evident from the subscript, I assume that the drift parameter is common across all houses.

The labor income for household i in year t also follows a random walk in logs

$$G_{it}^Y \equiv \frac{Y_{it}}{Y_{it-1}} = \exp(f(a_{it}) + g_{t-1} + \epsilon_{it}^Y), \tag{13}$$

where $f(a_{it})$ is a deterministic life-cycle trend, g_{t-1} is aggregate income growth in year t, and ϵ_{it}^Y is a random variable with mean zero. I assume that the vector $\epsilon_{it} = (\epsilon_{it}^H, \epsilon_{it}^Y)$ is

[7]The functional form of the bequest motive ensures that the value function is homogeneous in the house price. It is also sensible since it reflects that at high house prices, a given amount of wealth buys less housing consumption.

independently distributed over time; however, the two components may have a non-zero contemporaneous covariance $\sigma_{HY,t} > 0$ that represents a potential common exposure of housing and income risks at the regional or national level

$$\mathrm{Var}(\epsilon_{it}) = \begin{bmatrix} \sigma^2_{H,t} & \sigma_{HY,t} \\ \sigma_{HY,t} & \sigma^2_{Y,t} \end{bmatrix}. \tag{14}$$

It should be noted that, from the perspective of the optimizing household, the distinction between aggregate and idiosyncratic risk is only important to the extent that aggregate risk may induce a positive correlation between income and house price growth.

I will assume that ϵ_{it} is normally distributed. For the rest of the paper, it will then be convenient to directly write the mean and standard deviation of the log-normal random variable R^H_{it} as

$$\hat{m}_{t-1} = \mathrm{E}[R^H_{it}], \text{ and}$$
$$\hat{\sigma}_{H,t-1} = \mathrm{Var}[R^H_{it}]^{1/2},$$

respectively[8].

2.3 Computational Solution

The state and the choice variables of the dynamic program given by equations 7 and 8 can be re-defined to allow for a more efficient computational solution. These transformations are also the basis for the mapping of model quantities to observables described in the next section, so I will state the important aspects here and refer the reader to appendix A for details on the transformed model and the computational approach.

First, after omitting i subscripts again for notational simplicity, we can normalize all model quantities by current income Y_t, which is equivalent to normalization by permanent income due to the i.i.d. nature of the innovations to income growth. Specifically, define $w_t = W_t/Y_t$ and $\bar{h}_{t-1} = P_t H_{t-1}/Y_t$ for the endogenous state variables and $c_t = C_t/Y_t$,

[8]In terms of the parameters m_{t-1} and $\sigma_{H,t}$, one therefore gets

$$\hat{m}_{t-1} = m_{t-1} + \frac{1}{2}\sigma^2_{H,t},$$
$$\hat{\sigma}_{H,t-1} = \left[\left(\exp(\sigma^2_{H,t}) - 1\right) \exp(2m_{t-1} + \sigma^2_{H,t})\right]^{1/2},$$

by the usual arithmetic for log-normal random variables.

$l_t = L_t/Y_t$ and $h_t = P_t H_t/Y_t$ with respect to choices. All housing related quantities are expressed in terms of expenditure since this is what we observe in the data. I reduce the choices of both owner and renter to the value of the occupied house h_t, which is possible due to the linearity of housing services production from the housing asset. Thus letting the vector of transformed state variables be given by $x_t = [M_t, \tau_{t-1}, w_t, \bar{h}_{t-1}, l_{t-1}]$ and the vector of choice variable by $z_t = [\tau_t, \xi_t, d_t, h_t, c_t, l_t]$, one can define the normalized value function $v_t(x_t) = V_t(X_t)/(Y_t P_t^{-\rho})^{1-\gamma}$ to get

$$\begin{aligned}
v_t(x_t) =& \lambda_{t+1} \left\{ \max_{z_t} \frac{\left[c_t^{1-\rho}\left(h_t \Phi(\tau_t, a_t)\right)^{\rho}\right]^{1-\gamma}}{1-\gamma} + \beta \, \mathrm{E}_t \left[v_{t+1}(x_{t+1}) \left(\frac{Y_{t+1}}{Y_t}\right)^{1-\gamma} \left(\frac{P_{t+1}}{P_t}\right)^{-\rho(1-\gamma)} \right] \right\} \\
&+ (1-\lambda_{t+1}) b(w_t) \\
=& \lambda_{t+1} \left\{ \max_{z_t} \frac{\left[c_t^{1-\rho}\left(h_t \Phi(\tau_t, a_t)\right)^{\rho}\right]^{1-\gamma}}{1-\gamma} + \beta \, \mathrm{E}_t \left[v_{t+1}(x_{t+1}) \left(G_{t+1}^{Y} (R_{t+1}^{H})^{-\rho}\right)^{1-\gamma} \right] \right\} \\
&+ (1-\lambda_{t+1}) b(w_t) \qquad (15)
\end{aligned}$$

subject to conformably rewritten budget and down payment constraints given in appendix A, and where G^Y and R^H are the growth rates of income and house price as defined in equations (13) and (12). It becomes apparent from equation (15) that the normalization of the value function eliminates two exogenous state variables for computational purposes, which are income and the house price.

In practice, the computation is best performed in terms of two different value functions (both normalized as above) and the resulting optimal policies: one for households who were renting in the previous period or those who were forced to sell and move due to the exogenous shock, and one for homeowners that have the additional option of staying in their current house. Appendix A contains details on these transformed value functions and the corresponding budget constraints and transition equations for the states. Due to the nature of the estimation procedure, the model's solution will have to be re-computed for each iteration of the estimation loop.

2.4 Discussion

Several assumptions deserve a brief discussion. First, note that the model specified above yields the optimal demands for housing conditional on age, income, wealth, home ownership status, and the price of the housing asset. I do not explicitly specify the equilibrium in the markets for the housing asset or housing services. However, the goal of this analysis is to infer implied household beliefs about future price growth, and in any competitive equilibrium households will take the house price P_t as given. Therefore, the exercise of inferring implied expectations from observed demands is well-defined without an explicit specification of equilibrium as long as the optimal demand functions are evaluated at realized equilibrium prices[9].

Transaction Cost

The most important aspect of the distinction between owning and renting arises from the transaction cost for selling houses. In the absence of the transaction cost, the recursive structure of the problem implies that in addition to the household's age, only the beginning-of-period net worth and income are relevant state variables. In other words, if there was no transaction cost, we could think of homeowners as simply purchasing the house always only for one period, and thus at the beginning of the period -after selling the house and paying back the mortgage- it is irrelevant whether a household owned or rented in the previous period. However, with the transaction cost in place, homeowners have the option of not selling their house and thus not incurring the cost. This creates inertia in homeowners' adjustments to changes in the economic environment. Hence the quantity of housing owned at the beginning of the period, H_{t-1}, becomes a state variable. To account for mobility profiles over the life-cycle observed in the data, the shock M_t is a reduced-form way of modeling that homeowners may have to move and sell their house for reasons exogenous to the model, such as job-related relocations etc. Without the mobility shocks, but with reasonably high transaction costs, the model would not be able to match the amount of sales in the data.

Leverage and Mortgage Default

Young households face a life-cycle labor income profile with a deterministic component that is increasing but not tradable. The net present value of this non-risky trend part of future labor

[9] Of course, an implicit restriction on equilibrium results from the assumed time-series properties of house prices as specified in equation 12.

income is similar to a long position in a safe asset. For realistic parametrizations of the income process and housing returns, it is optimal for the young household to offset this position by taking a short position in the actual risk-free bond. Due to the collateral constraint, going short the risk-free bond means taking out a mortgage to finance the purchase of a house, and in this way achieve the optimal portfolio composition of risky and safe assets[10]. As households in the model age, they reduce their leverage and instead hold a positive position of the safe assets. The amount of savings of old households largely depends on the strength of the bequest motive.

The possibility of default on the mortgage interacts with the optimal leverage choice. A defaultable mortgage means that households hold a call option on their house, with leverage taking on the role of the strike price. Exercising the option is equivalent to keeping the house and not defaulting on the mortgage. The net value of the option is decreasing in the cost of default κ (which is akin to the option premium): if κ was prohibitively large, the optionality would disappear and households would simply hold a long position in the housing asset. Further, the value of the call option is decreasing in the strike price (i.e. leverage), but increasing in the mean m *and* the volatility σ_H of the house price. Any increase in the value of the option makes the household wealthier today. Everything else equal, this leads to higher consumption and greater leverage today. For example, if house price volatility goes up, the option becomes more valuable ceteris paribus, and households optimally react by increasing the strike price of the option (the leverage ratio) and consuming some of this option value today. In summary, this means that leverage is increasing in the option value, so any factor that raises the option value also raises optimal leverage.

Owning versus Renting

There are two basic channels in this model that determine the household's optimal ownership decision: I will refer to the first as the "user-cost" channel and the second as the "life-cycle" channel. The "user-cost" channel is based on comparison of the contemporaneous costs and benefits of owning versus renting, such as the rent-to-price ratio α and the housing maintenance cost ψ.

[10]See Yao and Zhang (2005) for a detailed discussion of the optimal portfolio composition with labor income and housing as collateral.

The "life-cycle" channel is in part due to the upward sloping labor income profile of young households. These households want to frequently adjust the level of housing services as their incomes rise during the early part of their life-cycle. However, if they chose to become home owners, the transaction cost would punish frequent upgrades in house size, and the down payment requirement makes a house that would also be large enough later in life unaffordable to the young household. Thus, the down payment constraint in equation (4) deters young households from becoming home owners too early. Instead, the cash-poor, but human capital-rich constrained young households rent and save until they have enough cash for the down payment of a house that is large enough. As previous quantitative analyses have found, the life-cycle pattern of ownership induced by borrowing constraints is however not effective in explaining the low rate of ownership among young households who have sufficient funds for their down payment. This feature of the data necessitates the introduction of another force that keeps the ownership rate among young households low. The model specified above matches this aspect of the life-cycle profile of ownership through a preference for rental housing that is declining in household age. This preference stands in for non-financial considerations driving the home ownership decision of young household, such as uncertainty about future family size.

3 Empirical Strategy and Data

3.1 Estimation Procedure

Overview

The goal of the empirical approach is to infer changes in short-term household expectations and the magnitude of credit constraints over the period of the recent housing boom. In order to do this, I use the cross-sections from years 1992 to 2010 of the Survey of Consumer Finances (SCF), which contains detailed information on the wealth composition and income of a representative sample of U.S. households[11]. Since the data are only available in three-year increments, I set the length of a model period to three years. I estimate (i) several

[11] The Federal Reserve conducts the survey every three years. The SCF oversamples rich households who hold the majority of aggregate U.S. wealth, but also provides sampling weights that can be used to calculate statistics based on a representative U.S. sample. This paper only computes means and variances from the SCF using the sampling weights.

preference parameters, which I restrict to be identical for all periods, (ii) expected house price growth $\{\hat{m}_t\}_t$ and its volatility $\{\hat{\sigma}_{H,t}\}_t$, and (iii) average down payments requirements $\{\delta_t\}_t$, which I allow to take different values for each period.

Some assumptions about household belief formation are necessary to execute the estimation. First, I assume that households have short-term beliefs about price growth and credit constraints over the next three years, and long-term beliefs about all following life-cycle periods. The long-term beliefs are set to long-run averages based on past observed data. The short-term parameters are allowed to vary from period to period, hence being a potential source of short-term swings in household choices. In the current setup, I do not allow for heterogeneity in household beliefs. I execute the estimation of the time-varying house price process and credit constraint parameters for the years 1998, 2001, 2004, and 2007.

Secondly, to estimate the preference parameters, I use additional moments from the survey years 1992 and 1995. For these years, I assume that the economy is in a long-run "steady state", with short term beliefs about house prices and short term credit constraints being equal to their long term values. The reason for this strategy is that I do not want to overuse preference parameters such as risk aversion and discount factor to explain household choices during the extreme boom-bust episode of the years 1998-2007[12].

Preference Parameters

The four preference parameters to be estimated are

- the Cobb-Douglas weight on housing ρ,
- the discount factor β,
- the bequest strength $\bar{B}$,
- and the age preference for renting ϕ.

[12]Since all parameters are jointly estimated, preference parameters are of course partially identified from the moments of the boom years. However, including additional moments of the preceding years stabilizes the estimates to reasonable values.

Structure of Household Beliefs

To state the structure of household beliefs more concisely, let t denote the calendar dates in three-year increments, with $t = 1$ corresponding to 1998, $t = 2$ to 2001, and so on. At each date t, a household optimizing at this date faces an interest rate $r_t = R_t - 1$, a mortgage spread of ζ_t, a rent-to-price ratio α_t, and a minimum down payment share δ_t. Further, the household believes that mean house price growth until $t + 1$ is given by $\hat{m}_t$, with volatility $\hat{\sigma}_{H,t}$. Table 1 shows short-run parameter values for each year.

Interest rates, rent-to-price ratios and mortgage spreads are observable both to the household and the econometrician in the short-term, and I assume that at date t households sign savings, rental, or mortgage contracts until $t + 1$ subject to the rates listed in the table. To calculate the rent-to-house-price ratio, I deflate the aggregate FHFA house price index by the CPI for rental prices to obtain a series for the price-to-rent ratio. I then take the value of 5.5% as computed by Davis, Lehnert, and Martin (2008) and extrapolate this number over the sample period by scaling it with the inverse of the FHFA/CPI index growth. Mortgage spreads are computed as the difference between the 30-year fixed mortgage rate reported by Freddie Mac and yields on 20-year Treasuries[13].

The last two columns of table 1 contain realized aggregate income and house price growth for each three-year period. Real aggregate income growth is estimated from NIPA disposable household income. House price growth is calculated from the FHFA house price index (deflated by the CPI).

Mean and volatility of house price growth are latent parameters and will be inferred from observed household choices.

I further assume that households must on average at least pay for δ_t percent of the house value from their own funds when purchasing a house. Note that this parameter does not specify the average size of the down payments actually made by households. It rather determines the minimum possible down payment allowed. This parameter will be inferred from observed household choices jointly with expectations and utility parameters, under the

[13] An alternative way to isolate the mortgage spread would be to compute the difference between 1-year ARM rates and 1-year T-bill yields. However, 1-year ARM are far less common and their pricing may not be representative of the majority of mortgages. The results when using this alternative measure would be similar.

assumption that each household can borrow at the terms characterized by (r_t, ζ_t, δ_t).

Table 1 specifies beliefs for a household optimizing at date t over the next period. One still needs to specify household beliefs for all remaining life-cycle periods, i.e. for dates $t+s$, $s>0$. These long-run beliefs are constant and set to long-run averages of the data series for the variables in table 1. Table 2 shows these long-run values.

The long-run values for interest rate and rent-to-price ratio are computed from the series described above for the short-run values. I set the minimum down payment constraint to 15%. This number reflects that for the majority of borrowers over the sample period, it was possible to get a mortgage with a down payment amount below the 20% limit that the government-sponsored enterprises set for conforming loans. If we think of the parameter δ as a stand-in for the average ease of access to credit, setting it to the GSE-imposed limit for prime conforming loan is to tight, as low down payment loans were available both in the prime and subprime segments of the market from the beginning of the sample[14].

The expected long-run price growth of the housing asset is set to 2.5%. The underlying assumption is that aggregate house prices are growing at the same rate as aggregate income in the long term. The number is also consistent with average growth rates of regional and national house price indexes, such as the FHFA or the Case-Shiller S&P 500 index. The volatility of house price growth is set to 18% annually. This number reflects purely idiosyncratic house price risk, which Landvoigt, Piazzesi, and Schneider (2014) document to be between 9% and 11%. In addition, the innovation ϵ_t^Y also includes aggregate housing risk at the regional and national level, which is between 5% and 9% based on MSA house price indexes (see e.g. Flavin and Yamashita (2002)).

The way in which household beliefs are "rolling forward" through time is best illustrated by means of an example. Consider a household at date $t=1$ (i.e. in 1998). From table 1, we know that this household is facing an interest rate of 3.42%, a rent-to-price ratio of 5.40%, and a down payment requirement of δ_{1998} percent. Further, this household believes that house prices will on average grow by $\hat{m}_{1998}$ percent until 2001 with a standard deviation of $\hat{\sigma}_{H,1998}$. For all dates beyond 2001, the household believes that the values of these variables

[14]Of course a literal interpretation of δ as the minimum possible down payment for all available loan contracts in the market would imply a number around zero for most of the sample. However, this would not be representative of the typical mortgage options offered to the average borrower.

are given in table 2. In other words, the household believes that the interest rate is $r_{t+s} = 3\%$, the rent-to-price ratio $\alpha_{t+s} = 5.5\%$, etc., for all dates with $s > 0$. Once time passes and the household gets to date $t = 2$ (2001), the realizations of the variables are given by the values in table 1, and now the household believes that the long-run values from table 2 apply to all dates beyond 2004.

This structure of beliefs is consistent with mean reversion, in the sense that households believe that variables fluctuate in the short-run but always return to long-run averages. In addition, it is a computationally tractable approach.

Model-to-Data Mapping

The data do not have a panel structure, hence each year-t-sample is a cross-section of different households. Keeping this in mind, index households for each year by $i = 1, \dots, N_t$ (with i generally indexing a different household in years t and $t+1$). Then for each year t, I construct a sample $S_t \equiv \{a_{it}, \tau_{it-1}, W_{it}, P_{it}H_{it-1}, Y_{it}\}_{i=1}^{N_t}$ from the SCF, where a_{it} is the household age, τ_{it-1} indicates ownership status (rent vs. own), and the remaining variables denote net worth, house value, and labor income as defined in the previous section.

Denote the vector of short-run model parameters for year t, corresponding to table 1 in the previous section, as θ_t, and the vector of long-run parameters corresponding to table 2 and preference parameters as θ^{LR}. Given the model's optimal policy functions conditional on parameters, it is possible to calculate the optimal choices for each household in the sample, $Z(S_t, \theta_t, \theta^{LR}) = \{C_{it}, \tau_{it}, L_{it}, H_{it}\}_{i=1}^{N_t}$, with C_{it} denoting numéraire consumption, τ_{it} next period's ownership status, L_{it} the mortgage or savings amount, and H_{it} the size of the house being rented or owned in the next period. These year-t choices can in turn be mapped to year-$t+1$ state variables by simulating the house price, income, and mobility shock realizations for each household in the sample, and by applying the realized price and income growth from t to $t+1$ (the last two columns of table 1). Applying the model policies to sample S_t in this way thus leads to a simulated sample of next year's state variables $\hat{S}_{t+1}(S_t, \theta_t, \theta^{LR})$, that is a function of this year's observed state variables and the model parameters.

The estimation procedure essentially entails finding the parameter vectors $\{\hat{\theta}_t\}_{t=1998}^{2007}$ and

$\hat{\theta}^{LR}$ that minimize the distance (in a method-of-moments sense) of the simulated $t+1$-samples $\hat{S}_{t+1}$ constructed in the way outlined above, and the observed $t+1$-samples S_{t+1} for each of the years $t = 1992/1995$, 1998, 2001, 2004, and 2007. Thus it is a Simulated Method of Moments (SMM) approach applied to a dynamic model and repeated cross-sections.

To implement the belief structure in practice, I need to solve the dynamic program once for each estimation period t. For the estimation periods 1998 to 2007, short-run and long-run parameters differ. This difference requires to first compute a long-run value function using the long-run parameters, and then a separate short-run value function for each period using the long-run value function as continuation value. Fortunately, expected house price growth and the credit constraint are fixed at their long-run values for the initial period (the combined 1992/1995 sample). Hence in this initial period the distinction between short- and long-run does not apply, and the initial value function can be used as the long-run value function in all subsequent periods. Thus computing the model for one parameter combination and for all estimation periods requires solving the life-cycle dynamic program five times.

Estimation

As objective function for the estimation step I use a weighted sum of squared deviations of a set of data averages from averages of the simulated sample. Since the data are repeated cross-sections and the model is dynamic in nature, a pseudo-panel approach is needed to apply the SMM approach described above.

The basic methodology follows Browning, Deaton, and Irish (1985). Using the same notation as above, let $\hat{S}_t$ and S_t denote the simulated and the data samples for year t, respectively. Since the sample $\hat{S}_t$ was generated by applying the model solution to the year-t data sample S_{t-1}, these samples generally consist of different individual households, so it is not possible to state moment conditions at the level of an individual observation. However, one can divide each sample into Q cells based on observed characteristics that are stable between times t and $t+1$, which here is a three-year period between two consecutive SCF samples. Index cells by $q = 1, \ldots, Q$, and let $g_{qt} = g(q, S_t)$ denote a K-vector of sample means for cell q in year t, where in the application the elements of g_{qt} are the average homeownership rate, the value-to-income ratio and the loan-to-value ratio (i.e. $K = 3$).

In practice, I use seven birth cohorts and three education groups to get a total of $Q = 21$ cells. Let the vector $\hat{g}_{qt}(\theta) \equiv g(q, \hat{S}_t; \theta^{LR}, \theta_t)$ denote the vector of sample means for the same variables, but computed from the simulated sample. By treating each cell q as an observation with variables taking on the values of cell means, I can hence create a pseudo-panel with Q observations.

Let g_q and $\hat{g}_q(\theta)$ denote the TK-vectors of the stacked cell means for all T years. Then the TK sample moment conditions are

$$\frac{1}{Q}\sum_{q=1}^{Q} g_q - \hat{g}_q(\theta) \equiv G_Q - \hat{G}_Q(\theta) = 0, \tag{16}$$

and for the case of fewer than TK parameters in θ, the Generalized Method of Moments (GMM) objective function to be minimized in θ is

$$(G_Q - \hat{G}_Q(\theta))' D (G_Q - \hat{G}_Q(\theta)), \tag{17}$$

where D is a positive definite weighting matrix. I use the inverse variance-covariance matrix of the data moments for D, i.e. $D = \hat{\text{Cov}}(g_q)^{-1}$. Note that G_Q and $\hat{G}_Q$ are simply the aggregate sample means in the real and simulated data, for all K variables and T years. However, for the computation of the estimated variance-covariance matrix of the moment conditions, it is necessary to have the pseudo-panel structure and a well-defined notion of an observation.

Equation 17 is a conventional GMM objective function with a constant weighting matrix, and the asymptotic standard errors can generally be obtained in the well-known way (see e.g. Wooldridge (2002)). Since this is a simulation estimator, the estimated covariance matrix of the moment conditions needs to be adjusted by a factor taking into account the number of simulations. Appendix B contains details on how the standard errors are calculated, drawing on the econometric results of Pakes and Pollard (1989) and Hall and Rust (2002).

3.2 Other Parameters

Table 3 shows those parameters of the model that I do not estimate and that do not vary over the time period included in the estimation.

All parameters are annual. The sales transaction cost and the maintenance share are in line with the values used by other studies of the housing market. The transaction cost reflects the actual cost of selling such as realtor's fees and the cost of moving for homeowners (over renters). The maintenance share is the fraction of the house value that homeowners have to spend to offset depreciation.

The annual standard deviation of the shock to permanent income growth is set to 13% based on the results of Cocco, Gomes, and Maenhout (2005). The correlation of both shocks is set to 0%, based on the low estimate by Flavin and Yamashita (2002). Other studies have found similarly low correlations.

Finally, I take three sets of parameters from the literature that enter the houshold problem due to its life-cycle character.

- The deterministic part of labor income growth ($f(a)$ in equation 13) follows a third-degree polynomial whose coefficients are taken from Cocco, Gomes, and Maenhout (2005), and thus are consistent with the shock to income growth. Specifically, I use coefficients decsribing the income profile of high-school graduates estimated by Cocco, Gomes, and Maenhout (2005) using data from the PSID. The life-cyle profile has the common hump-shape.

- The survival probabilities λ_a are computed from the mortality rates reported by the National Center of Health Statistics.

- For the life-cycle profile of mobility (i.e. the probailities of moving) I use the estimates by Li, Liu, and Yao (2009). The basic shape of the mobility rate function over the life-cycle is convex and declining in age.

3.3 Data

For each of the SCF surveys from 1992 to 2010, I use the prepared extract sample of the SCF[15]. I remove all observations with the household head being younger than 25 years of age, which is the starting age of the life-cycle labor income profile I use. I take labor income

[15] These samples already contain some pre-generated variables, and some observations with unlikely answers have been removed.

to be broadly definined as the sum of wage income, income from social security and other retirement funds, income from own businesses, and government transfers. As defintion of networth, I use the pre-generated variable "networth" from the SCF, which is the balance of all household assets and liabilities. For the house value of homeowners, I use the SCF variable "houses", which is the value of the primary residence[16]. As the mortgage principal of homeowners, I use the SCF variable "mrthel", which includes home equity loans and other types of loans that use the primary residence as collateral.

Further, I remove all households with more than 5 million dollars of networth (in year 2000 dollars) from the sample. The life-cycle income process of these very wealthy households is usually not well described by the one assumed in equation 13, since a large fraction of their income is from dividends and capital gains. The problem is aggravated by the fact that these households tend be older, with traditional sources of retirement income only being a very small fraction of their overall income. The removal of these households has the additional advantage of being able to economize on grid points during the estimation. The disadvantage is a loss of about 15% of raw observations for each year, but due to the strong oversampling of wealthy households in the SCF this only amounts to about 1.5% of effective observations after applying the SCF-provided sampling weights.

4 Results

4.1 Target Moments and Estimation Results

As moments in the objective function, I use the average homeownership rate, the value-to-income ratio and the loan-to-value ratio for each of the years 2001, 2004, 2007, and 2010, and for the combined sample 1992/1995. In addition, I include the average loan-to-value ratio among older households (age 58 or older) for the initial 1992/1995 sample. This gives 16 moments and 12 parameters when only the utility parameters, the means of the house price growth process and the minimum down payment shares are estimated. Four parameters are added when the volatilities of house price growth are estimated in addition. Table 4 displays the targeted moments. All moments are sample means computed

[16] This implies that other real estate investments of the household will be included in networth and hence are counted as savings in the sense of the model.

using SCF sampling weights. House values and loan-to-value ratios are reported only for homeowners (and are zero for renters). The choice of these moments rests mainly on their natural connection to model quantities. The homeownership rate is calculated as the sample average of households' discrete own-versus-rent decisions. Similarly, the house value-to-income ratio is the sample average of a state variable of the model, and the loan-to-value ratio is the ratio of two choice variables, mortgage principal and house value. The model is designed to capture several important features of homeownership, house size, and general life-cycle mortgage dynamics; hence these moments represent the set of quantities that the model is best suited to match. The model also makes predictions about the household net worth-to-income ratio, but due to the lack of other, higher-yielding assets such as stocks, it is impossible for the model to match general wealth dynamics in a period like the late 1990s, and thus I do not include this moment in the set of targets.

I use aggregate moments since all parameters are assumed to be identical across age and income groups. However, because the model's key mechanisms rest on its life-cycle character, I will examine the fit across household age and wealth in the next section to see whether the general life-cycle shape of model-implied moments lines up with their data counterparts.

Table 5 shows results of the estimation step. The asymptotic standard errors in parenthesis are calculated using the GMM formula for the case with a constant weighting matrix. Appendix B contains details on how the standard errors were computed. Specification (1) keeps the short-run volatility of house price expectations fixed at the long-run value of 0.18, while specification (2) also estimates these parameters.

The point estimates of the preference parameters in specification (1) and (2) are very close, suggesting that they are pinned down by average levels of the different moments across all periods. In general, the addition of the volatilities as free parameters does not significantly affect the estimates of other parameters, and the standard errors of the volatility estimates in specification (2) imply that the long-run volatility of 0.18 is within the 95% confidence interval for each year.

The point estimates of the minimum down payment shares exhibit an increasing trend in specification (2), even though they are not statistically different from the long-run value of 0.15 at the 5% level. The point estimates of the mean growth rate are all close to the long-run mean of 2.5%. The estimates for 1998 and 2007 are somewhat higher, while the estimate in 2001 is clearly lower. The estimated subjective volatility of house price growth is below the long-run mean in 1998, but above the mean during the boom years of 2001 and 2004.

Specification (1) is overidentified by four moment conditions. The J test for overidentifying restrictions yields a test statistic of 8.76, implying that the model cannot be rejected at the 5% level (the $\chi^2_{4,0.95}$ threshold value is 9.45).

4.2 Identification

In the following, I will explore the sources of identification for the results[17]. First, I will analyze the features of the data that compel the estimates of the preference parameters. Then I will argue that most of the identifying variation comes from young households who are on the margin of becoming home owners. The transaction cost creates an inactivity region that prevents older existing homeowners from adjusting their house size or selling their house in response to small changes in expectations or credit constraints. Finally, I will discuss how expected house price growth, borrowing constraints and the volatility of house price growth are separately identified.

Preference Parameters

The chief source of identification for the preference parameters are the four moments in the base year (1992/1995) of the estimation, for which beliefs and credit constraint are set to their long-run values. The Cobb-Douglas weight ρ is the most important determinant of value-to-income ratios in the model, and thus is identified from the mean of this ratio in the data. The discount factor β determines model leverage and is hence identified from LTV ratios in the data. The estimated value of 0.85 is for one model period of three years, implying an annual discount factor of 0.94. The discount factor interacts with the bequest motive to determine the effective age-dependent discount factor in the model. The parameter $\bar{B}$ that governs the strength of the bequest motive is therefore identified from the leverage ratio of older households. The combined estimates of β and $\bar{B}$ imply a reasonable life-cycle profile of

[17] I am using the term identification loosely throughout this section as referring to those aspects of model and data interaction that allow me to tell apart the different effects.

effective discount factors when combined with the survival probabilities. The parameter of the rental preference factor, ϕ, regulates the home ownership rate among young households. It is identified from the average home ownership rate in the data.

Transaction Cost and Persistence in Choices

Before discussing of the estimation results for down payments and expectations, it is helpful to examine the model's cross-sectional fit along a specific dimension. Particularly, table 6 shows, both for model and data, the fraction of home owners who did not purchase their home in the last three years (at the parameter estimates from table 5).

The model matches this aspect of the data rather well. The numbers demonstrate that, both in model and data, the majority of young home owners have recently purchased their home, while the majority of older owners live in the same house that they bought more than three years ago. This is the case for the base year 1992/1995 as well as for the subsequent years of the sample that experienced substantial changes to interest rates, rent-to-price ratios, and the level of house prices. Despite these changes in the economic environment that existing home owners were facing, the transaction costs of selling their house prevented these owners from adjusting their housing choices.

The flip side of this inertia for existing home owners is that most of the reaction in response to short term changes in the environment, both for the intensive and extensive margin, will come from young household who are first time home buyers.

Expectations, Down Payment Constraints, and Volatility

The three sets of time-varying model parameters – expectations, down payments constraints, and house price volatilities, are identified from three sets of moments that represent different choice margins – tenure, house value-to-income ratios and loan-to-value ratios.

Generally each parameter simultaneously affects all three choice margins for a given year. Therefore describing the identification amounts to understanding which moment is quantitatively most important for each type of parameter.

The estimates for down payment constraints are mainly identified from the intensive margin of housing demand, i.e., house value-to-income ratios. This is because the two other parameters governing expected house price gains (mean and volatility) are less powerful in determining model-predicted house values. There are two reasons for the relatively weak effect of short-term expectations on optimal house values. First, given the large transaction cost of selling a house, existing home owners will not adjust their house size in response to a moderate short term change in expected house prices. Secondly, new home *buyers* (previously renters) are mostly at their leverage constraint. Hence changing these buyers' degree of optimism has little effect on their optimal house values.

However, changing the minimum down payment for these constrained buyers directly affects their optimal house value choice. Loosely speaking, the much stronger impact of the collateral constraints on model-implied house values relative to the other parameters forces the estimation to "use" variation in the constraints to match data house values. This means in turn that the estimates of the down payment constraints are effectively pinned down by the house value-to-income ratios in the data.

Even though expectations have a limited effect on the intensive margin, they do have a large effect on the extensive margin – they decision whether to own or rent. Again home buyers who enter the period as renters are the source of identification. Their short term benefit of owning versus renting is directly affected by expected house price gains. Thus households on the margin of buying will decide to advance (delay) their purchase of a house in response to a positive (negative) change in expected price growth.

It follows logically that the third set of parameters to be estimated – price growth volatilities – are mainly identified from the third set of moments, which are the loan-to-value ratios. The main effect of an increase in subjective house price risk is greater optimal leverage, through the call option channel discussed in section 2.4 above. Since debt is frictionlessly adjustable in the model, households consume the additional future wealth from the increased value of the call option by borrowing more[18]. The effects on the intensive and extensive margins of housing demand are smaller, and they are ambiguous. A rise in the option value makes owning a house more attractive, but from a portfolio perspective, higher house price

[18]It turns out that household optimization keeps the option value roughly constant. When volatility increases and the option value rises holding everything else equal, households choose higher leverage which is equivalent to choosing a higher strike price of the call option. This reduces the option value and increases current consumption.

risk makes housing less attractive as an investment.

It is instructive to break the identification argument in several steps using graphical representations of the objective function. Each step involves comparing data moments to model-implied moments for a different set of model inputs, while the utility function parameters are set to their estimated values from table 5. First, I will consider the hypothetical case that the only variation in model inputs over time is the estimation sample. In other words, all model parameters, including the realized price and income growth, are set to their long-run values for each year. Figure 1 shows data and model-generated moments for this case. The model almost perfectly matches ownership rates over time, but significantly misses data value-to-income ratios and loan-to-value ratios.

The next step is to feed in the non-estimated time-varying parameters from table 1. These are interest rates, rent-to-price ratios, mortgage spreads, and realized price and income growth for each 3-year period. The resulting model-generated moments are shown in figure 2 by the solid red line. Simulating the model using the realized price gains and low interest rates and spreads in 2004 to 2010 drives value-ot-income ratios up significantly. For the 2001-2004 period, interest rates are already lower than the long-term average, but rent-to-price ratios are still close to the long-term average, which results in a model-predicted ownership rate that is too high. From 2004 to 2010, the drop in interest rates is counteracted by a simultaneous drop in rent-to-price ratios, keeping the ownership rate stable. Leverage is too low for the 2001-2007 period, but somewhat too high for 2007-10. Note that in 2001 and 2004, even though households substantially increase their mortgage debt, leverage stays roughly constant due to the large realized rise in house values. In the last model period, the large realized drop in house prices pushes leverage above the data value.

We can think of the solid red line in figure 2 as the starting point for the estimation of the time-varying parameters, $\hat{m}_t$, δ_t, and $\hat{\sigma}_{H,t}$. The estimation procedure can choose these parameters for each period to make the model-implied moments match the data.

Figure 3 shows the model fit if only the mean expectation parameters are set to their estimates from specification (1), but volatilities and credit constraints are held fixed at their long run values (the solid green line). The lower expected price growth from 2001-04 reduces the model-predicted ownership rate for that period to a value slightly below the data. This has the additional effect of also lowering model leverage and VTI ratios for that period. The slightly more optimistic expectation for 2001-04 and 2007-10 cause the ownership rate to move slightly above the data values.

Next, I will also set the down payment parameters to their estimated values from specification (1), so that the only parameters still held fixed at their long run values are the volatilities $\hat{\sigma}_{H,t}$. The solid blue lines in figure 4 show the result. By relaxing the constraints in 2001 and 2004, and tightening the constraints in 2007 and 2010, the estimation matches home ownership rates and house value-to-income ratios closely. These moments have a lower variance in the data, and the GMM objective function penalizes deviation for loan-to-value ratios less. The end result of specification (1) with fixed volatility parameters therefore predicts too much leverage in 2001 and too little leverage in 2007-10.

Freeing up the volatility parameters allows the estimation to also match loan-to-value ratios, as can be seen in figure 5. This requires lower volatility from 1998-2001 to decrease leverage, and higher volatility from 2004-2010 to increase leverage, which is what we see in specification (2) in table 5.

To summarize, the estimates of mean expected house price growth are mainly identified from the extensive margin – the home ownership rate – in the data. As one can see in figure 2, however, model-predicted ownership rates are close to their data counterparts without any deviation from long run parameter values for all but the 2001-04 period. This is because time variation in interest rates and rent-to-price ratios has opposing effects, keeping the ownership rate roughly stable. Hence the estimated values of $\hat{m}_t$ are close to the long-run value of 2.5% annually for all periods but 2001.

The estimated values of the credit constraints, on the other hand, are mainly identified from the intensive margin – house value-to-income ratios. Here the estimation finds lax constraints for the 1998-2004 period and a tightening of constraints for the 2004-10. The ultimate reason for this finding is the inertia of household choices due to transaction costs. If the majority of households stay in their current house, then housing quantities are mostly fixed, and a change in the price per unit of housing (as measured by a repeat sales index) is directly reflected in values. Low interest rates and mortgage spreads during the boom years induce an additional expansion in housing quantities of new home buyers. This force would inflate average model-predicted house values beyond their data counterparts at persistently low down payment constraints. Thus the model matches the data by slightly restricting the maximum feasible house size of credit constrained buyers during the boom period.

4.3 Model Fit

As an "out-of-sample" test for the model, this section evaluates the fit of the model for more narrowly defined age and wealth groups. Since the estimation only targets averages of home ownership rate, value-to-income ratio and loan-to-value ratio, the model's cross-sectional fit is useful to understand how well the model's mechanisms capture the actual heterogeneity in choices in the data.

Table 7 shows the three main model outcomes, broken down by age and net worth, ans comparing the data to the model-generated sample. For all three choice margins, the model is able to qualitatively match the pattern in the data.

With respect to the home ownership rate, the pattern in the data is best described by the statement that young/poor households rent, whereas old/wealthy households own their homes. While the model generally matches this pattern, the model is not able to explain the steep increase of the ownership rate in net worth for the oldest group of households. The credit constraint in the model mainly affects young households in connection with their upward-sloping labor income profile and the transaction cost. As the income process of the oldest households is much less variable, the credit constraint is not a good reason for these households to avoid ownership. Since most of the poor old households also have a small income, the optimal choice for these agents according to the model is simply to also own a small house[19]. The flip side of this issue is that the model's ownership rate among wealthy old households is too low. The need of some wealthy old households to save additional funds for their bequest causes them to downsize to a smaller house and save in bonds. Since downsizing requires them to sell their house and incur the transaction cost, they optimally decide to rent for their remaining life span to avoid incurring the transaction cost again when their final wealth is liquidated for the bequest.

[19] Some other studies impose a minimum house size for owner occupancy to deal with this issue. While such a restriction would improve the cross-sectional fit of this model as well, it would not significantly alter the estimation results based on targeting data averages.

The model matches the cross-sectional distribution of house value-to-income ratios reasonably well, again with the exception of the data values for poor and old households, who report very low ratios in the data.

Similar to the ownership rate, the model generates the overall pattern for leverage across age and wealth groups that we observe in the data. However, it somewhat overstates the leverage of young households, and it undershoots for older households. The reason for the too low LTV ratios of old households is the lack of other assets that can function as savings devices in the model. Specifically, not a small fraction of older households in the data have both a mortgage and a portfolio of liquid assets such as stocks and mutual funds. These households then have both a non-zero LTV and a non-zero portfolio of liquid funds. Given the two-asset structure of the portfolio choice, this portfolio composition cannot represented within the model. Put differently, home owners with positive savings automatically have an LTV of zero in the model. To match average LTV ratios, the model has to compensate for this low leverage of old households by overutilizing the life-cycle borrowing motive of young households.

Overall, the model matches the cross-sectional distribution of ownership, house values, and leverage reasonably well. Table 7 demonstrates that the inference about the estimated parameters does not come at the expense of highly counterfactual cross-sectional implications.

5 Conclusion

This paper estimates short-run expectations of house price appreciation and minimum down payment requirements during the recent housing boom. The inference is based on structural estimation of a life-cycle dynamic program that encompasses housing demand choices at the extensive and intensive margin, applied to repeated cross-sections from the Survey of Consumer Finances. I implement the estimation by constructing a pseudo-panel and applying a Simulated Method of Moments estimator.

The main result is that model-implied aggregate expectations of future price growth were very close to the long-term average throughout the period from 1998 to 2010, with slightly higher expectations at the beginning and the end of the boom period. The estimation also finds that down payment constraints were less strict at the beginning of the boom, and then tightened toward its end and during the bust. These findings are driven by the need for the model to jointly rationalize the dynamics of the home ownership rate, loan-to-value, and house value-to-income ratios, in an environment of high house prices, low interest rates, and high price-to-rent ratios. The stability in expected house price gains is needed to explain the stability in the ownership rate – with very high expectations, the model would predict a counterfactually large rise in the ownership rate. Since expectations are pinned down in this way by the ownership rate, the initially loose and then subsequently tighter credit constraints are required to explain the evolution of value-to-income ratios over time.

The realistically large transactions costs associated with selling houses and moving in the model are crucial for identifying these estimates. The transactions costs cause most existing home owners to simply stay in the same house despite large swings in prices (as is confirmed by the data), rent-to-price ratios, and interest rates. Therefore changes in house prices translate directly into changes in house values (prices times quantities), and a further relaxation of credit constraints would lead to an even higher rise of value-to-income ratios than in the data.

The estimation further finds that high uncertainty about future house prices during the boom years may have contributed to the increase in mortgage debt. Everything else equal, higher uncertainty leads to an increase in the value of the call option on housing that is implied by the combination of owning a house with a defaultable mortgage. If debt is more easily adjusted than housing, households will optimally consume part of this higher option value through higher debt.

Overall, the quantitative results are consistent with households beliefs in mean-reverting house prices, in the sense that they expect long-run future price growth to follow the historical average of roughly 2.5% real growth per year.

Apart from the quantitative results, the paper's main methodological contribution is the way in which it infers subjective short-run expectations by closely tracking household choices over time using a life-cycle portfolio choice model. An important aspect of the method is that it takes into account short-term movements of other variables that characterize the economic environment for housing choices, while keeping household expectations of the long-term values of these variables set to their long-run averages.

The results of this paper do not contradict the notion that the recent housing boom was partly fueled by overly optimistic expectations of house price appreciation. Only a relatively small share of the housing stock gets traded per year, and these transactions form the basis for price measurement. It is hence possible that few very optimistic households caused the price spike by increasing short-term demand in local housing markets - Piazzesi and Schneider (2009) make this point using a simple search model. Aggregate beliefs, however, are identified from observing the majority of households who did not substantially adjust their housing choices during the boom.

It would be an interesting extension to allow heterogeneous beliefs across age or income groups to see whether the approach taken here could identify those subgroups with optimistic beliefs that possibly were the driving force behind the strong price movement.

References

Bajari, P., P. Chan, D. Krueger, and D. Miller (2013): "A Dynamic Model Of Housing Demand: Estimation And Policy Implications," *International Economic Review*, 54(2), 409–442.

Browning, M., A. Deaton, and M. Irish (1985): "A profitable approach to labor supply and commodity demands over the life-cycle," *Econometrica*, 53(3), 503–543.

Burnside, C., M. Eichenbaum, and S. Rebelo (2011): "Understanding Booms and Busts in Housing Markets," NBER Working Papers 16734, National Bureau of Economic Research, Inc.

Campbell, J. R., and Z. Hercowitz (2005): "The Role of Collateralized Household Debt in Macroeconomic Stabilization," NBER Working Papers 11330, National Bureau of Economic Research, Inc.

Campbell, J. Y., and J. F. Cocco (2003): "Household risk management and optimal mortgage choice," *Quarterly Journal of Economics*, Nov, 1449–1494.

Case, K. E., J. M. Quigley, and R. J. Shiller (2003): "Home-buyers, Housing and the Macroeconomy," in *Asset Prices and Monetary Policy*, ed. by A. Richards, and T. Robinson, RBA Annual Conference Volume. Reserve Bank of Australia.

Case, K. E., and R. J. Shiller (2005): "Is there a bubble in the housing market?," *Brookings Papers on Economic Activity*, 2, 299–361.

Cocco, J. F. (2004): "Portfolio Choice in the Presence of Housing," *Review of Financial Studies*, 18(2), 535–567.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005): "Consumption and Portfolio Choice over the Life-Cycle," *Review of Financial Studies*, 18(2), 492–533.

Corbae, D., and E. Quintin (2014): "Leverage and the Foreclosure Crisis," *Journal of Political Economy*, forthcoming.

Davis, M. A., A. Lehnert, and R. F. Martin (2008): "The Rent-Price Ratio For The Aggregate Stock Of Owner-Occupied Housing," *Review of Income and Wealth*, 54(2), 279–284.

Demyanyk, Y., and O. V. Hemert (2011): "Understanding the Subprime Mortgage Crisis," *Review of Financial Studies*, 24(6), 1848–1880.

Economist, T. (2005): "The global housing boom: In come the waves," *The Economist*.

Favilukis, J., S. C. Ludvigson, and S. V. Nieuwerburgh (2013): "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium," NBER Working Papers 15988, National Bureau of Economic Research, Inc.

Flavin, M., and T. Yamashita (2002): "Owner-Occupied Housing and the Composition of the Household Portfolio," *American Economic Review*, 92(1), 345–362.

Foote, C. L., K. Gerardi, and P. S. Willen (2008): "Negative equity and foreclosure: Theory and evidence," *Journal of Urban Economics*, 64(2), 234–245.

Geanakoplos, J. (2010): *The Leverage Cycle*pp. 1–65. University of Chicago Press.

Glaeser, E. L., J. D. Gottlieb, and J. Gyourko (2010): "Can Cheap Credit Explain the Housing Boom?," Working Paper 16230, National Bureau of Economic Research.

Hall, G., and J. Rust (2002): "Econometric Methods for Endogenously sampled time series: the Case of Commodity Price Speculation in the Steel Market," *Working paper*.

Himmelberg, C., C. Mayer, and T. Sinai (2005): "Assessing high house prices: bubbles, fundamentals and misperceptions," *Journal of Economic Perspectives*, 19(4), 67–92.

Iacoviello, M., and M. Pavan (2013): "Housing and debt over the life cycle and over the business cycle," *Journal of Monetary Economics*, 60(2), 221–238.

Kiyotaki, N., A. Michaelides, and K. Nikolov (2011): "Winners and Losers in Housing Markets," *Journal of Money, Credit and Banking*, 43(2-3), 255–296.

Landvoigt, T., M. Piazzesi, and M. Schneider (2014): "The Housing Market(s) of San Diego," *American Economic Review*, forthcoming.

Li, W., H. Liu, and R. Yao (2009): "Housing over Time and over the Life Cycle," *Federal Reserve Bank of Philadelphia Working Paper*, NO. 09-7.

Lustig, H. N., and S. G. V. Nieuwerburgh (2005): "Housing Collateral, Consumption Insurance, and Risk Premia: An Empirical Perspective," *Journal of Finance*, 60(3), 1167–1219.

Mian, A., and A. Sufi (2009): "The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis," *The Quarterly Journal of Economics*, 124(4), 1449–1496.

——— (2011): "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis," *American Economic Review*, 101(5), 2132–56.

Pakes, A., and D. Pollard (1989): "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 57(5), 1027–1057.

Piazzesi, M., and M. Schneider (2009): "Momentum Traders in the Housing Market: Survey Evidence and a Search Model," *American Economic Review*, 99(2), 406–11.

Tauchen, G. (1986): "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics Letters*, 20(2), 177–181.

Wooldridge, J. M. (2002): *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge.

Yao, R., and H. H. Zhang (2005): "Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints," *Review of Financial Studies*, 18(1), 197–239.

A Dynamic Programming Solution

First, denote the value function of the household who did not own a house or has sold its house as $v_t^M(w_t)$, where $w_t = W_t/Y_t$ and W_t is defined as in equation 10 in the main text. $v_t^M(\cdot)$ is further defined as the value conditional on survival until age a_{t+1}, and after all shocks are realized. Thus one gets

$$v_t^M(w_t) = \max_{c_t, l_t, \tau_t, h_t} u(c_t, h_t) + \beta \, \mathrm{E}_t \left[v_{t+1}(x_{t+1}) \, \left(G_{t+1}^Y (R_{t+1}^H)^{-\rho} \right)^{1-\gamma} \right] \tag{18}$$

subject to

$$w_t = c_t + l_t + (1-\tau_t)\alpha_t h_t + \tau_t(1+\psi)h_t, \tag{19}$$

$$l_t \geq -\tau_t \, (1-\delta_t) h_t. \tag{20}$$

where $v_t(x_t)$ is defined as in equation 15, the constraints 19 and 20 are obtained by normalizing the budget constraint 5 and the downpayment constraint 4 by income Y_t, and the utility function $u(c,h)$ is defined in equation 9. Secondly, denote the value function of a home owner who is *forced* to stay in the same house as $v_t^S(w_t, \bar{h}_{t-1})$. Again, I define $v_t^S(\cdot)$ as the value conditional on survival until a_{t+1}, and after realization of the mobility shock. This yields

$$v_t^S(w_t, \bar{h}_{t-1}) = \max_{c_t, l_t} u(c_t, \bar{h}_{t-1}) + \beta \, \mathrm{E}_t \left[v_{t+1}(x_{t+1}) \, \left(G_{t+1}^Y (R_{t+1}^H)^{-\rho} \right)^{1-\gamma} \right] \tag{21}$$

subject to

$$w_t + \nu \bar{h}_{t-1} = c_t + l_t + (1+\psi)\bar{h}_{t-1}. \tag{22}$$

$$l_t \geq -(1-\delta_t)\bar{h}_{t-1}. \tag{23}$$

Due to the definition of w_t as including the house value net of the transaction cost, the household that stays in the same house receives this cost back on the LHS of budget constraint 22. Thus the house value $\bar{h}_{t-1}$ cancels on both sides of the constraint, i.e. the constraint becomes

$$1 + (R_t + 1_{[l_{t-1}<0]}\zeta_t) l_{t-1} = c_t + l_t + \psi \bar{h}_{t-1}.$$

One can now express the normalized value function $v_t(\cdot)$ in terms of $v_t^M(\cdot)$ and $v_t^S(\cdot)$

$$\begin{aligned} v_t(\tau_{t-1}, M_t, w_t, \bar{h}_{t-1}) =& \lambda_{t+1} \left[\tau_{t-1}(1-M_t) \max\left\{ v_t^M(1-\kappa), v_t^M(w_t), v_t^S(w_t, \bar{h}_{t-1}) \right\} \right. \\ & \left. + (1 - \tau_{t-1} + \tau_{t-1} M_t) \max\left\{ v_t^M(1-\kappa), v_t^M(w_t) \right\} \right] \\ & + (1 - \lambda_{t+1})\, b(w_t). \end{aligned} \tag{24}$$

The first term in square brackets represents the choices of the household who enters the period owning a house ($\tau_{t-1} = 1$) and does not have to move ($M_t = 0$). This household can either stay in the same house ($v^S(\cdot)$), sell the current house and face the same optimization problem as a renter ($v^M(w_t)$), or default on its mortgage and face the problem of a renting household with cash equal to $1 - \kappa$ percent of its income.

The second term represents the choices of the households entering the period without owning a home ($\tau_{t-1} = 0$), or the home owner who has to move ($\tau_{t-1} M_t = 1$). This household also has the choice of defaulting, which can of course only be optimal for owners who are forced to move.

The two endogenous state variables of the model are w_t, and $\bar{h}_{t-1}$. To express their transition laws, it is useful to define the discrete choice variable $d \in \{0, 1\}$, with $d = 1$ indicating that the household defaults on its mortgage. Then the transitions for w_t and $\bar{h}_{t-1}$ are

$$w_{t+1} = (1-d)[(R_{t+1} + 1_{[l_t<0]}\zeta_{t+1}) + \tau_t(1-\nu) h_t\, R_{t+1}^H] \frac{1}{G_{t+1}^Y} + 1 - d\kappa, \tag{25}$$

$$\bar{h}_t = (1-d)\tau_t h_t \frac{R_{t+1}^H}{G_{t+1}^Y}, \tag{26}$$

where l_t, h_t, and τ_t denote the optimal savings and housing policies for period t.

The dynamic program specified by equations 18 to 26 can be solved recursively starting in period T, where

$$v_T(x_T) = b(w_T)$$

since $\lambda_{T+1} = 0$. To compute the value functions $v_t^M(\cdot)$ and $v_t^S(\cdot)$ in practice, I discretize the continuous state variables w_t and $\bar{h}_{t-1}$ on grids with 80 points each. The spacing of the grid points for w_t and $\bar{h}_{t-1}$ is chosen with the goal of estimation in mind such that the points are denser on intervals where more households from the SCF are located. Further, the boundaries are chosen such that almost all observations fall within the state space. The innovations to the income an house price processes ϵ_{t+1}^Y and ϵ_{t+1}^H are assumed to be jointly normally distributed, and are discretized using the method of Tauchen (1986). I use 3 nodes for the income innovation and 7 nodes for the house price innovation. Since for house price growth, the variance is also estimated, it is important to have enough nodes in the discretization. Increasing the number of nodes above 7 did not affect the estimation results. The shock forcing a home owner to move M_{t+1} is independent of both the house price and income shocks, and distributed as a $\{0, 1\}$-Bernoulli random variable with $\Pr(M_{it+1} = 1 \,|\, a_{it})$ set according to table 3. I use linear interpolation to compute the continuation value in case the the next period state variables do not lie on the grid.

B Estimation Procedure

Define the year-t sample of variables from the SCF that correspond to the model's normalized state variables

$$s_t = \left\{a_{it}, \tau_{it-1}, w_{it}, \bar{h}_{it-1}\right\}_{i=1}^{N_t},$$

which are as defined in section 3. The goal is to create the model-implied sample $\hat{s}_{t+1}(s_t, \theta)$ of simulated year-$t+1$ state variables. From the transition equations for the state variables 25 and 26 it is clear that the required ingredients for this step are

- the model policy functions for housing, bonds, and home ownership related choices,
- simulated random variables for the move shock M_t and the shocks to income and house price growth $(\epsilon_{t+1}^Y, \epsilon_{t+1}^H)$ for each observation,
- and the realized aggregate return to housing ΔP_{t+1} and realized aggregate income growth ΔY_{t+1}.

In the following, I will state the algorithm to be applied to each observation of sample s_t in order to construct the simulate sample $\hat{s}_{t+1}$.

1. Take observation i from sample s_t. Dropping the observation and time subscripts, denote age by a, liquid funds by w, homeownership status by τ, and the value of the house owned or rented by $\bar{h}$. Denote the realized aggregate growth in house prices and income from t to $t+1$ as g^H and g^Y.

2. If $\tau = 1$, that is the current observation is a home owner, draw a uniform random variable u_1 and simulate move shock $M = 1[u_1 < \Pr(M = 1\,|a)]$, where $1[\cdot]$ denotes the indicator function.

3. Using the optimal policies from the model's computational solution with parameter vector θ, calculate the model-implied household choices

$$\begin{aligned} l &= \hat{l}_a(M, w, \tau, \bar{h}; \theta) \\ \tau' &= \hat{\tau}_a(M, w, \tau, \bar{h}; \theta) \\ d &= \hat{d}_a(M, w, \tau, \bar{h}; \theta) \\ h &= \hat{h}_a(M, w, \tau, \bar{h}; \theta), \end{aligned}$$

   where $\hat{d}(\cdot)$ denotes the optimal mortgage default decision.

4. Draw a pair of normally distributed random variables (ϵ^Y, ϵ^H) for the innovations to house price and income growth.

5. Apply the transition equation for the state variables to get next period's implied states, i.e. compute

$$\begin{aligned} w' &= (1-d)[Rl + \tau'\,(1-\nu)\exp(g^H + \epsilon^H)\,h]\exp(-(f(a+1) + g^Y + \epsilon^Y)) + 1 - d\kappa \\ \bar{h}' &= (1-d)\tau'\,h\exp(g^H + \epsilon^H - f(a+1) - g^Y - \epsilon^Y). \end{aligned}$$

6. Set $a_{it+1} = a+1$, $\tau_{it} = \tau'$, $w_{it+1} = w'$, and $\bar{h}_{it+1} = \bar{h}'$ to obtain the simulated state variables for this observation for year $t+1$.

By applying this algorithm to each observation in the sample s_t once, one obtains the simulated sample $\hat{s}_{t+1}(s_t, \theta)$. In practice, one needs to repeat the whole procedure multiple times with different seeds for the random number generator and obtain multiple simulated samples. During the calculation of the moments for the distance function, one then uses an average over the moments calculated from each simulation. Denote the number of simulations conducted by Z. All estimation results reported in this paper are computed with $Z = 5$, which turns out to be sufficient to get stable results.

After repeating this procedure for each pair of consecutive years, one has the complete set of data and simulated samples, s_{t+1} and $\hat{s}_{t+1}(s_t, \theta)$.

Given these samples, the computation of moments and the construction of the GMM objective function proceeds as outlined in the main text. Denote by g_q the TK-vector of data means for birth cohort-education cell q, and by $\hat{g}_q$ the corresponding vector of simulated means. The aggregate moment conditions are computed as in equation 16, and the estimator for θ is defined based on equation 17 as

$$\hat{\theta} = \text{argmin}_{\theta}(G_Q - \hat{G}_Q(\theta))' D (G_Q - \hat{G}_Q(\theta)). \tag{27}$$

To minimize this distance function, I employ a global pattern search algorithm over a range of model parameters for which the dynamic programming solution is well-behaved. This algorithm is essentially an intelligent grid search that only uses direct function evaluation and does not compute any numerical derivatives. Once the search seems close to a minimum, I employ a simplex search algorithm until convergence. In the case of the exactly identified model (with free volatility parameters) the search succeeds to find a local minimum as it manages to reduce the objective function to a value very close to zero.

The objective is sufficiently smooth to be numerically differentiable using bi-directional differentiation at a delta of 0.01. This should be sufficient to calculate approximate gradients for standard errors once the minimum is found.

Based on the results of Pakes and Pollard (1989), the SMM estimator's asymptotic distribution is normal with

$$\sqrt{Q}(\hat{\theta} - \theta^*) \xrightarrow{d} N(0, (1 + 1/Z)\Lambda_1^{-1}\Lambda_2\Lambda_1^{-1}).$$

To write the expressions for Λ_1 and Λ_2, first define the Jacobian matrix of the population moments with respect to the parameters evaluated at $\theta*$

$$A = \mathrm{E}\left[\nabla_\theta (g_q - \hat{g}_q(\theta^*))\right].$$

Then one gets

$$\Lambda_1 = A' D A,$$

and

$$\Lambda_2 = A' D \Omega D A,$$

where Ω is the variance-covariance matrix of the population moments

$$\Omega = \mathrm{E}\left[(g_q - \hat{g}_q(\theta^*))(g_q - \hat{g}_q(\theta^*))'\right].$$

A consistent estimator for the variance-covariance matrix of $\hat{\theta}$ is thus

$$\frac{1}{Q}\,(\hat{A}' D \hat{A})^{-1}\,(\hat{A}' D \hat{\Omega} D \hat{A})\,(\hat{A}' D \hat{A})^{-1},$$

where $\hat{A}$ and $\hat{\Omega}$ are consistent estimators of A and Ω and are calculated as

$$\hat{A} = \nabla_\theta (G_Q - \hat{G}_Q(\hat{\theta})) \text{ and}$$
$$\hat{\Omega} = \frac{1}{Q} \sum_{q=1}^{Q} (g_q - \hat{g}_q(\hat{\theta}))(g_q - \hat{g}_q(\hat{\theta}))'.$$

Figure 1: Target and model-generated moments, all parameters fixed at initial values

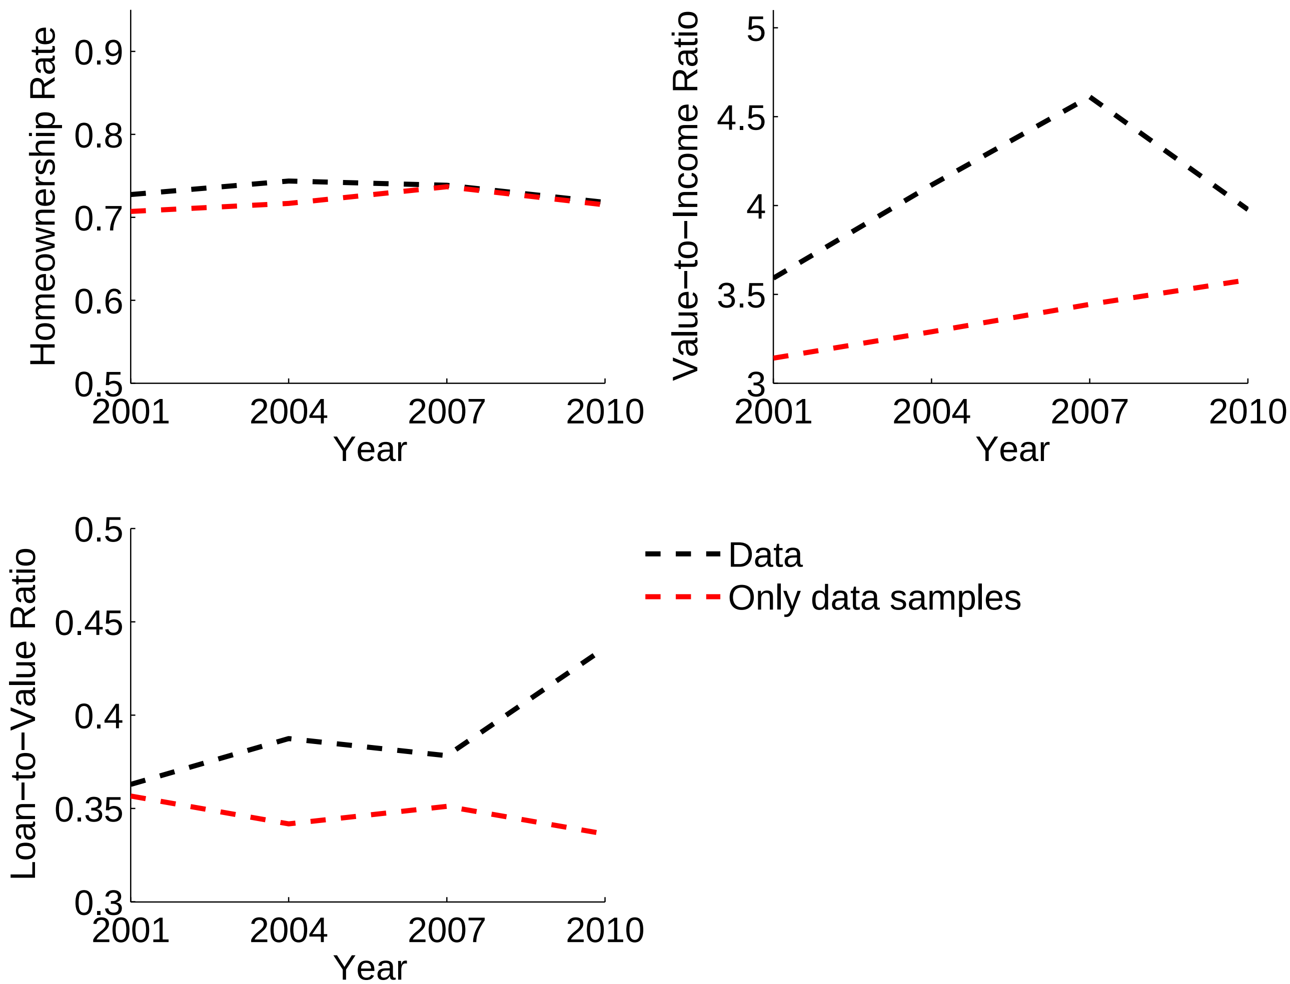

The figure compares model-generated to data moments when all model parameters are fixed at their base level (for the 1992/1995 samp). The only model input varying over time are the SCF samples that are fed into the model as distribution of households' state variables. Panel (A) shows home ownership rates, panel (B) value-to-income ratios, and panel (C) loan-to-value ratios. Optimal model policies cause the majority of existing home owners to stay in their current houses. Without taking into account realized price growth, this implies counterfactually low model-implied value-to-income ratios.

Figure 2: Target and model-generated moments, with time-varying prices and interest rates.

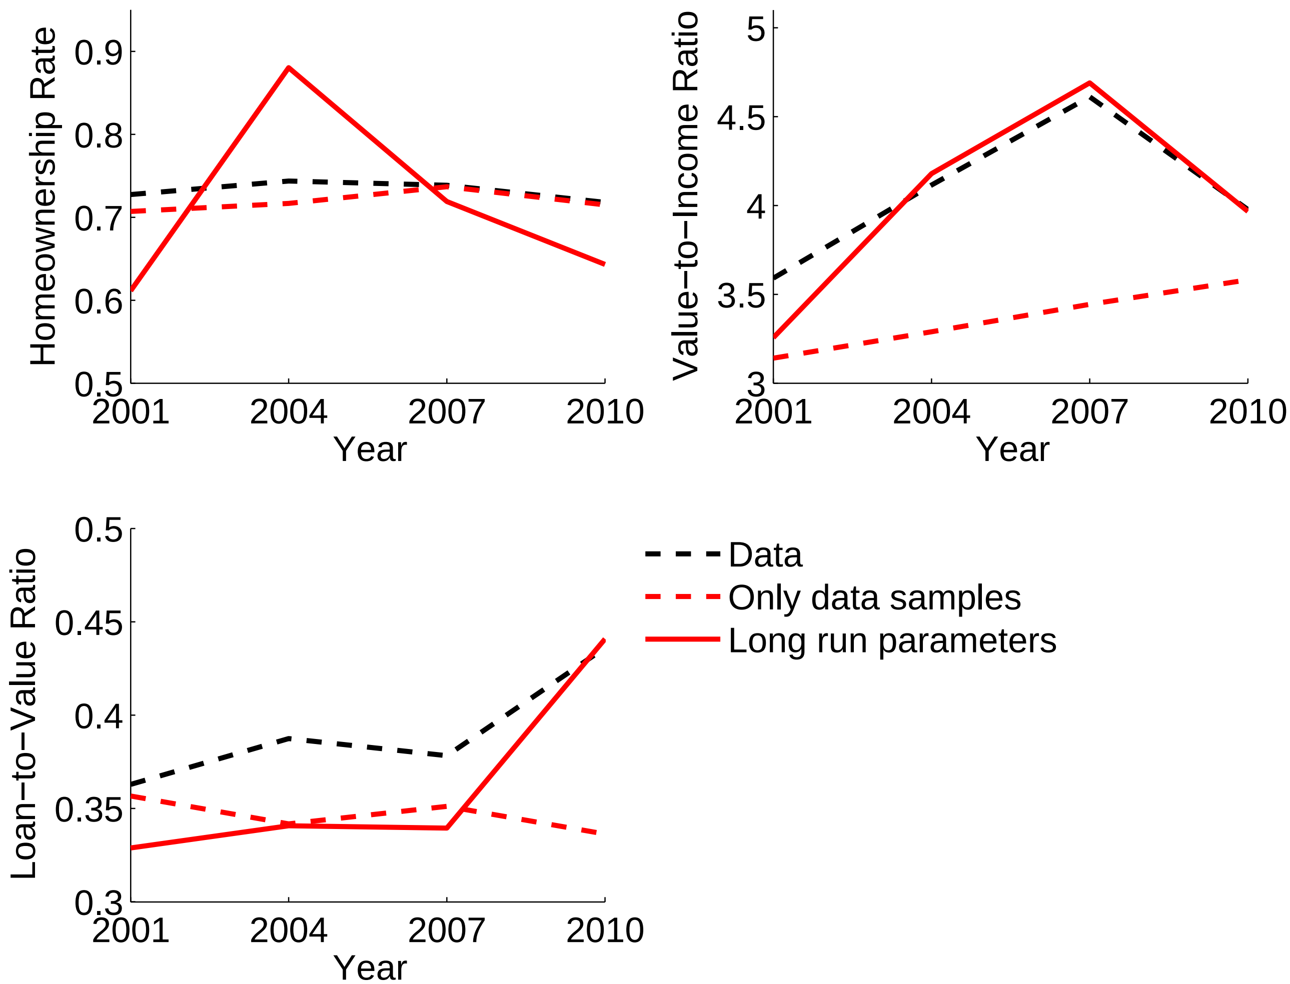

The figure compares model-generated to data moments when all estimated model parameters are fixed at their base level (for the 1992/1995 samp). The only model input varying over time are the SCF samples and realized rent-to-price ratios, interest rates, and price and income growth rates. Panel (A) shows home ownership rates, panel (B) value-to-income ratios, and panel (C) loan-to-value ratios. Low interest rates and rent-to-price ratios cause time variation in the model-predicted moments. Taking into account realized price growth moves model-predicted value-to-income ratios close to their data value.

Figure 3: Target and model-generated moments, estimated mean growth rate.

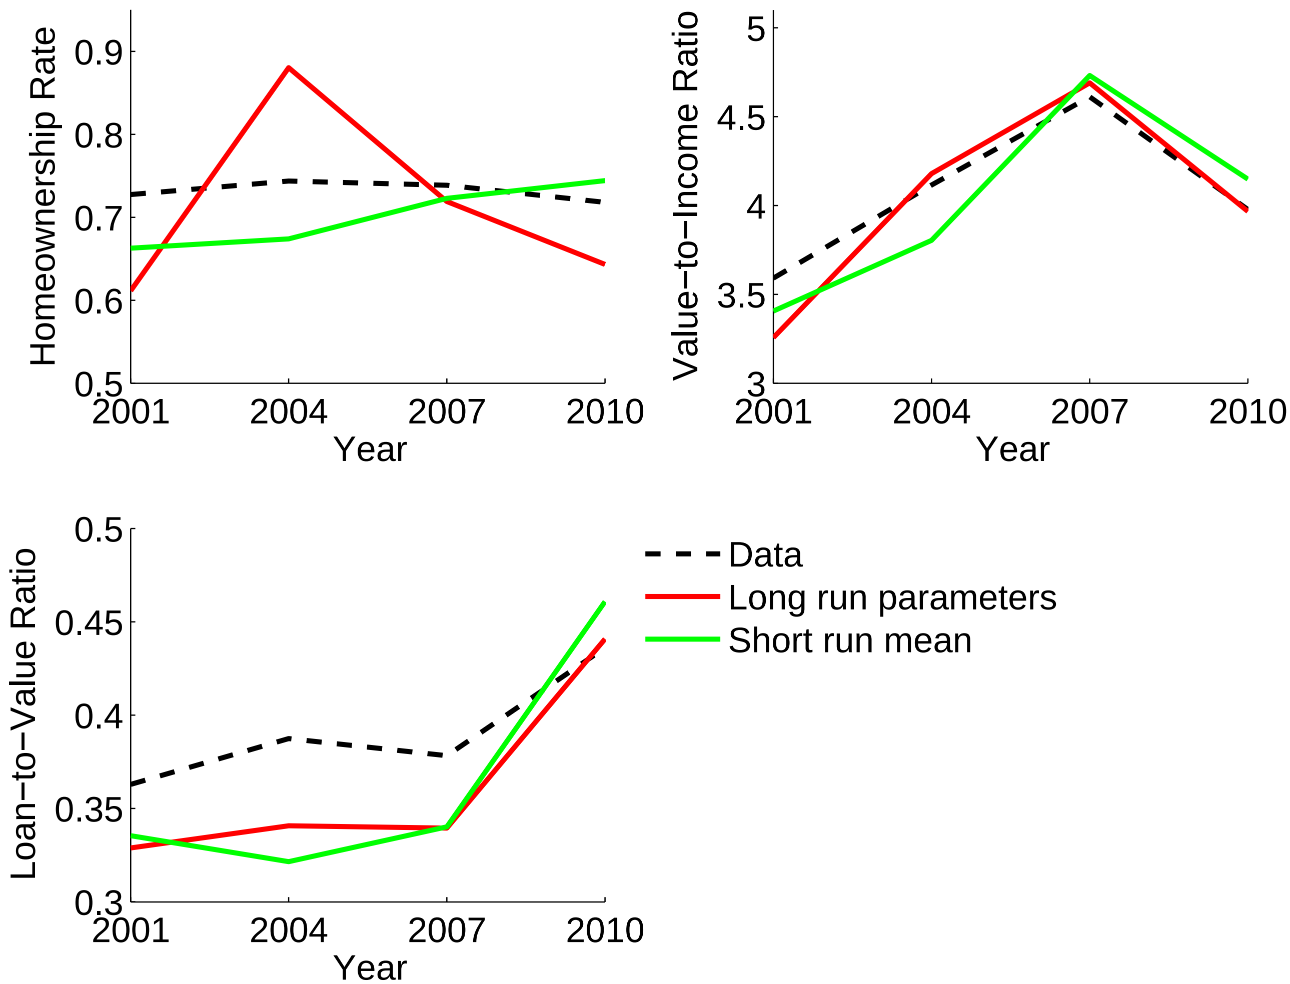

The figure compares model-generated to data moments when the expected growth rate of house prices is set to its estimated value for each period in addition to the time-varying inputs from 2. Panel (A) shows home ownership rates, panel (B) value-to-income ratios, and panel (C) loan-to-value ratios. The mean estimates are mainly identified from the home ownership rate in the data.

Figure 4: Target and model-generated moments, estimated mean and down payment requirements.

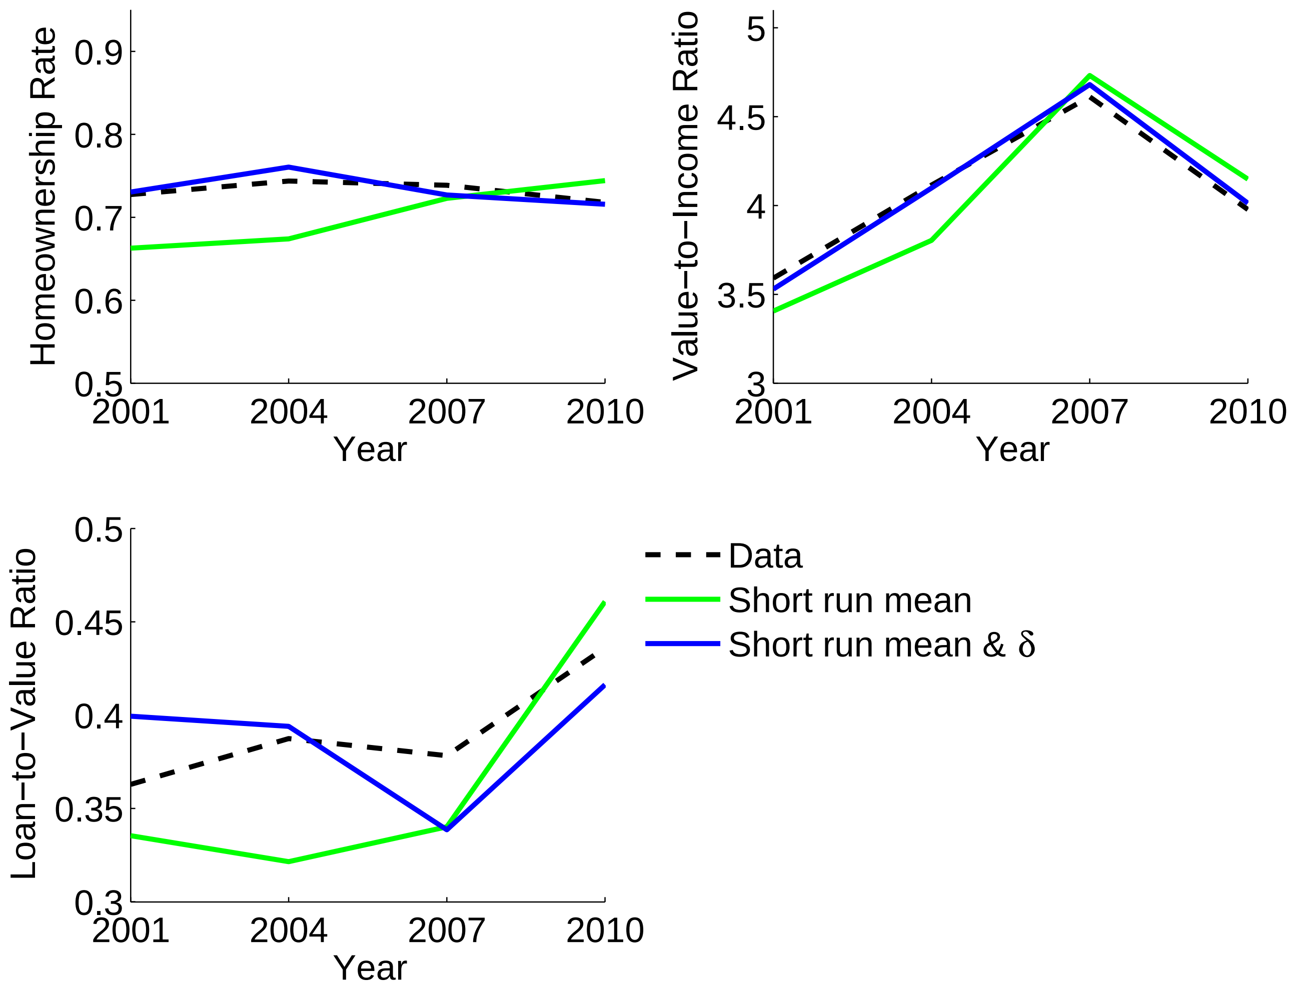

The figure compares model-generated to data moments when the expected growth rate of house prices and down payment requirements are set to their estimated values for each period in addition to the time-varying inputs from 2. Panel (A) shows home ownership rates, panel (B) value-to-income ratios, and panel (C) loan-to-value ratios. The estimates of δ_t are mainly identified from house value-to-income ratios.

Figure 5: Target and model-generated moments, all estimated parameters.

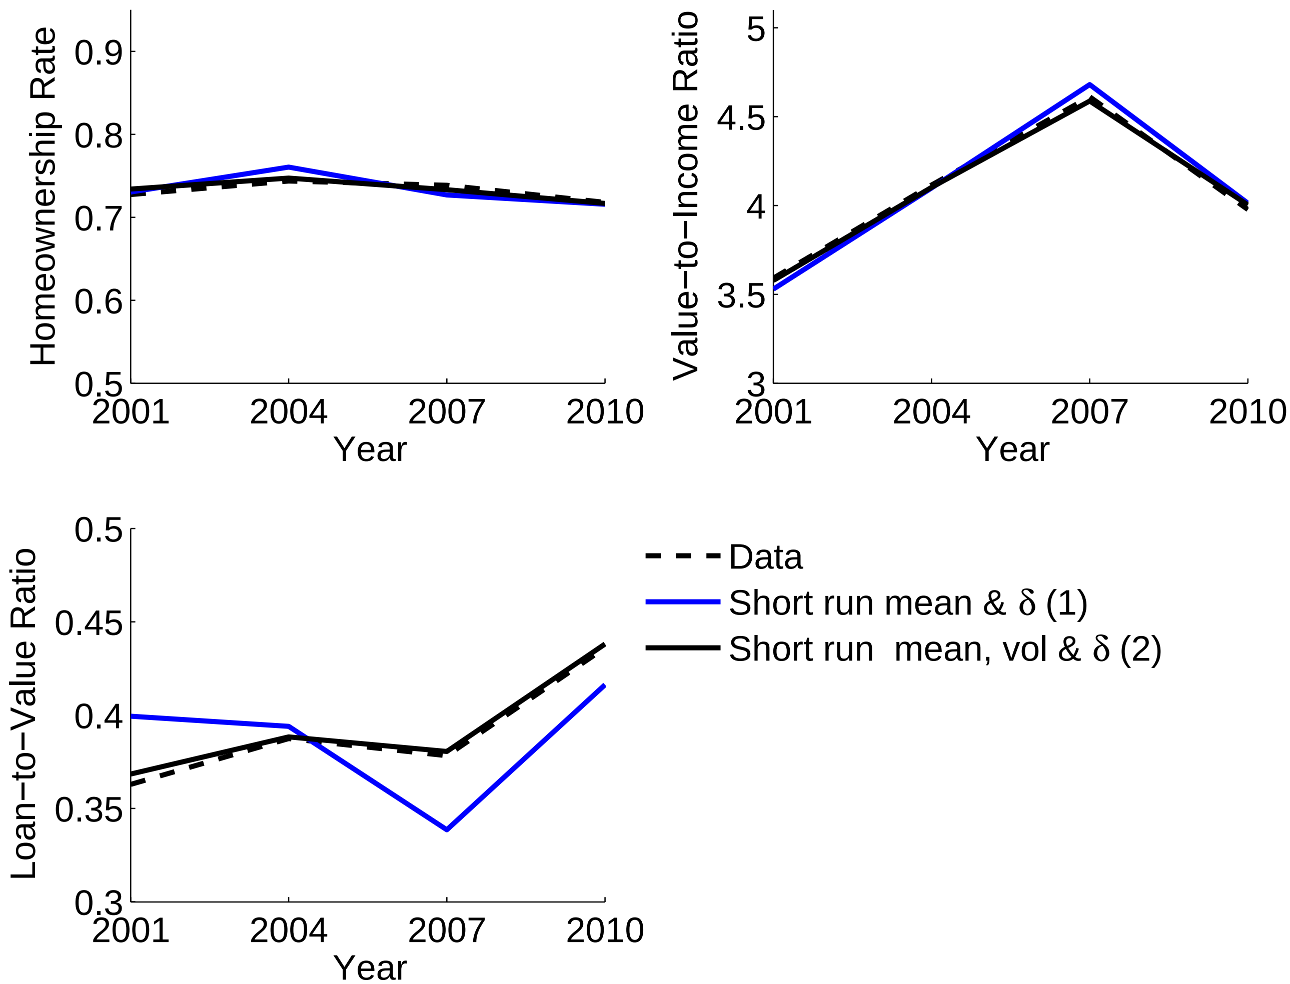

The figure compares model-generated to data moments when the expected growth rate of house prices, its volatility, and down payment requirements are set to their estimated values for each period in addition to the time-varying inputs from 2. Panel (A) shows home ownership rates, panel (B) value-to-income ratios, and panel (C) loan-to-value ratios. The volatility estimates are mainly identified from loan-to-value ratios.

Table 1: Short-run Model Inputs

Year t	r_t	ζ_t	α_t	δ_t	$\hat{m}_t$	$\hat{\sigma}_{H,t}$	ΔP_{t+1}	ΔY_{t+1}
1998	3.24	1.43	5.40	*	*	*	10.0	8.0
2001	0.86	1.10	5.19	*	*	*	14.0	6.0
2004	0.73	0.93	4.02	*	*	*	16.0	5.0
2007	0.73	1.18	3.30	*	*	*	−19.0	−2.0

All values annual in percent. r_t is real average annualized interest rate over the three-year period based on 3-year treasury yields. α_t is rent-to-price ratio calculated by rescaling 1992 base value of 0.06 over time. ζ_t is calculated as the difference of the 30-year fixed conventional mortgage rate and 20-year treasury yields.

* Minimum down payment (as percentage of house value) δ_t, expected house price growth $\hat{m}_t$ and volatility $\hat{\sigma}_{H,t}$ are to be estimated.

The last two columns contain realized aggregate house price and income growth from t to $t+1$.

Table 2: Long-run Beliefs ($s > 0$)

Parameter	Value
Interest rate r_{t+s}	3
Rent-to-price ratio α_{t+s}	5.5
Mortgage spread ζ_{t+s}	1.5
Minimum down payment δ_{t+s}	15
Expected price growth $\hat{m}_{t+s}$	2.5
Volatility of price growth $\hat{\sigma}_{H,t+s}$	17

All values annual in percent. The volatility includes both aggregate regional (7%) and idiosyncratic (10%) components of housing return risk.

Table 3: Time-invariant Parameters

Parameter	Value
Risk aversion γ	3
Sales transaction cost ν	10%
Maintenance share ψ	2%
Std.Dev.(ϵ_{it}^{Y})	13%
Corr($\epsilon_{it}^{Y}, \epsilon_{it}^{H}$)	0%
Income growth $\hat{g}$	2.5%
Cost of default κ	12%

Table 4: Target Moments

Year	HOR	VTI	LTV	LTV(>58)
1992/1995	0.708	3.174	0.369	0.148
1998	0.712	3.168	0.380	
2001	0.728	3.591	0.363	
2004	0.744	4.114	0.387	
2007	0.739	4.612	0.378	
2010	0.718	3.977	0.436	

VTI and LTV ratios computed for the subsample of homeowners.

Table 5: Estimation Results

Parameter	$\sigma = 0.18$		σ flexible	
ρ	0.124	(0.007)	0.124	(0.007)
β	0.824	(0.013)	0.824	(0.023)
$\bar{B}$	2.585	(0.186)	2.113	(0.454)
ϕ	0.722	(0.222)	0.671	(0.114)
δ_{1998}	0.127	(0.047)	0.098	(0.065)
δ_{2001}	0.113	(0.067)	0.106	(0.044)
δ_{2004}	0.152	(0.038)	0.152	(0.021)
δ_{2007}	0.183	(0.043)	0.189	(0.055)
$\hat{\mu}_{1998}$	0.033	(0.004)	0.034	(0.003)
$\hat{\mu}_{2001}$	0.008	(0.003)	0.007	(0.005)
$\hat{\mu}_{2004}$	0.026	(0.004)	0.027	(0.003)
$\hat{\mu}_{2007}$	0.033	(0.004)	0.034	(0.006)
$\hat{\sigma}_{H,1998}$			0.124	(0.106)
$\hat{\sigma}_{H,2001}$			0.171	(0.077)
$\hat{\sigma}_{H,2004}$			0.259	(0.116)
$\hat{\sigma}_{H,2007}$			0.199	(0.065)

Asymptotic standard errors in parenthesis. The estimates of $\hat{m}_t$ and $\hat{\sigma}_{H,t}$ are annual

Table 6: Fraction of home owners who did not purchase their house during last three years

	Data			Model		
	Net worth[a]			Net worth		
	≤ 30	> 30 ≤ 150	> 150	≤ 30	> 30 ≤ 150	> 150
	1992/1995					
Age ≤ 40	0.025	0.360	0.691	0.060	0.474	0.759
$> 40 \leq 61$	0.025	0.547	0.826	0.045	0.585	0.922
> 61	0.045	0.782	0.924	0.070	0.905	0.997
	1998					
Age ≤ 40	0.020	0.322	0.693	0.026	0.448	0.867
$> 40 \leq 61$	0.026	0.502	0.819	0.053	0.608	0.961
> 61	0.072	0.868	0.928	0.111	0.917	0.986
	2001					
Age ≤ 40	0.002	0.297	0.638	0.050	0.276	0.504
$> 40 \leq 61$	0.006	0.515	0.836	0.014	0.433	0.939
> 61	0.087	0.714	0.913	0.093	0.835	0.992
	2004					
Age ≤ 40	0.010	0.345	0.563	0.093	0.393	0.464
$> 40 \leq 61$	0.056	0.481	0.803	0.027	0.377	0.837
> 61	0.024	0.631	0.875	0.054	0.736	0.995
	2007					
Age ≤ 40	0.095	0.396	0.592	0.171	0.551	0.603
$> 40 \leq 61$	0.012	0.487	0.831	0.006	0.416	0.762
> 61	0.043	0.723	0.877	0.005	0.730	0.980

All averages computed for the subsample of home owners in each cell, using SCF sampling weights.

[a] Net worth is the SCF variable with that name measured in thousands of 2001 dollars.

Table 7: Cross-sectional model fit in base period (1992/1995)

	Data			Model		
	Net worth[a]			Net worth		
	≤ 30	> 30 ≤ 150	> 150	≤ 30	> 30 ≤ 150	> 150
	Home Ownership Rate					
Age ≤ 40	0.165	0.588	0.946	0.090	0.442	0.790
$> 40 \leq 61$	0.139	0.698	0.961	0.687	0.942	0.933
> 61	0.098	0.806	0.966	0.634	0.778	0.799
	Value-to-Income Ratio					
Age ≤ 40	2.449	2.056	2.458	2.174	2.491	2.665
$> 40 \leq 61$	2.469	2.394	2.629	2.590	2.901	2.703
> 61	0.998	3.737	5.585	2.766	3.821	5.072
	Loan-to-Value Ratio					
Age ≤ 40	0.678	0.632	0.493	0.847	0.718	0.284
$> 40 \leq 61$	0.631	0.459	0.353	0.817	0.480	0.098
> 61	0.197	0.174	0.086	0.781	0.144	0.009

All averages computed using SCF sampling weights.
[a] Net worth is the SCF variable with such name measured in thousands of 2001 dollars.