What Drives the Volatility and Persistence of House Price Growth?

Wieland Hoffmann*

February 15, 2016

(click here to access latest version)

Abstract

US house prices are characterized by strong persistence in annual growth rates but mean revert over longer horizons. I set up a calibrated life cycle model with heterogeneous agents and endogenous house price formation where both features emerge qualitatively in response to moderately persistent shocks to the aggregate income level. However, momentum is only forty percent of what is observed in the data, and while the presence of adjustment frictions in the model doubles the volatility of price growth rates, it remains at just twenty percent of its data counterpart. These quantitative conclusions change significantly when I incorporate subjective house price beliefs into the model. When agents use past observations to update their expectations regarding future price growth, the same shocks to aggregate income induce dynamics in house price growth rates that are very close to the data. The model fully matches the degree of persistence and mean reversion, and it generates eighty percent of the observed volatility. Notably, the subjective belief specification does not fully account for the strong price growth during the last housing boom.

Keywords: Housing, Heterogeneous Agents, Learning

JEL Classification: D83, D91, E21, G11, G12

---

*University of Mannheim, Graduate School of Economic and Social Sciences, L9 7, D-68161 Mannheim, wieland.hoffmann@gess.uni-mannheim.de; I thank my advisor Klaus Adam for his invaluable guidance and support; Michèle Tertilt, Matthias Doepke, and participants of the Mannheim Macro Seminar for helpful comments and suggestions.

1 Introduction

Residential real estate is the largest asset class on US households' balance sheets and subject to substantial price fluctuations. Figure 1 shows a recurring pattern where real house prices grow faster than trend for several years and then contract.[1] This results in considerable volatility and strong momentum in price changes. The standard deviation of annual price growth rates is five percent and the serial correlation is 0.7, meaning that a five percent increase in house prices is associated, on average, with another 3.5 percent increase in the following year. This high degree of predictability in price growth sets housing apart from other important asset classes. For the S&P 500 stock price index for example, the serial correlation between annual growth rates is only 0.27, while the standard deviation is thirteen percent. Glaeser et al. (2014) illustrate the difficulty of reconciling the strong momentum in annual house price growth with the observed mean reversion at lower frequencies in a rational, frictionless market.

Figure 1: Freddie Mac House Price Index

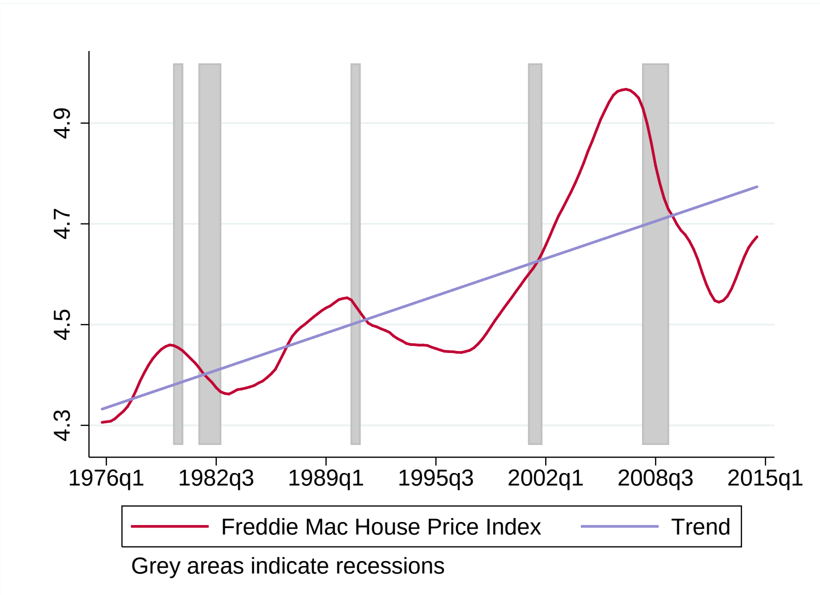

Notes: Log of the national price index deflated by the CPI urban excluding shelter.

In this paper I study the house price dynamics generated by a calibrated life cycle model where heterogeneous households trade owner-occupied real estate. To capture important frictions in the housing market, I introduce transaction costs and a downpayment requirement. Households are subject to idiosyncratic as well as aggregate income shocks that affect their purchase decisions. The consideration of aggregate income shocks is motivated by the observation that the *level* of detrended per capita income strongly correlates with the level of detrended house prices and also

[1] The data appendix gives a detailed description of the data that is used to construct all figures and statistics.

with house price *growth rates*.[2] Therefore, I use the model to ask to what extent the fluctuations in real per capita income in the US over the past forty years can account for the observed pattern in price growth rates. This is a nontrivial issue because the income process is moderately persistent and mean reverting. In order to replicate the observed momentum of annual price growth, a rise in the income level not only has to generate house price growth on impact but also in subsequent periods.

The modeling framework explicitly accounts for the fact that the market for single family homes is dominated by individuals infrequently trading in the homes they live in,[3] and my main contribution is the incorporation of subjective price beliefs in such a setting. Valuation of properties is subject to large uncertainty (Quan and Quigley (1991)), and in a series of surveys conducted since 1988 in different US markets, Case, Shiller, and Thompson (2012) find that past price movements play an important role in shaping both short and long term price expectations of homebuyers. Similarly, Piazzesi and Schneider (2009) document that during the last housing boom optimism regarding future house price growth peaked when prices had already reached their all time high. Such biases do not seem to be unique to real estate, but it is of interest to study their effect within the particular structure of the housing market. As chairman of the Federal Reserve, Greenspan (2004) argued that large transaction costs and the necessity to move when trading properties were important restraints on the development of house price bubbles. On the other hand, Adam, Marcet, Merkel, and Beutel (2015b) find that with such belief biases the introduction of a linear transaction tax would increase the volatility of stock prices and the likelihood of booms and busts.

In a first step of the analysis, I use a version of the model with fully elastic housing supply and constant prices to establish that the transaction costs and downpayment requirement themselves play an important role in shaping a persistent response of housing expenditures. In contrast to a model without these adjustment frictions, housing purchase volume rises steeply for several years when the aggregate income level is above trend. Adjustment frictions hence lead to positive serial correlation in the growth rate of housing expenditures. The result is reminiscent of Berger and Vavra (2015), who show that lumpy adjustment patterns make aggregate durable expenditures highly procyclical. During booms, the fraction of households close to adjusting increases over time, so that the response of expenditures to further positive shocks can be twice as large as during other periods. In my model, this applies in particular to rent-to-own transitions, as more and more

[2] The latter correlation is 0.47.

[3] Already Linneman (1986) and Case and Shiller (1989) point out that there is good reason to think that this makes the housing market less efficient than other financial markets. Regarding the difficulties of arbitrage by outside investors see also the detailed discussion in Glaeser and Gyourko (2007).

households pass the threshold to buying a home each year.

I then address the question how this pattern translates into equilibrium prices when there is a constant stock of owner occupied real estate. The assumption of a constant housing supply maximizes the potential impact of income shocks and is motivated by the fact that metropolitan areas with low supply elasticities are indeed an important driver of the the national index depicted in figure 1.[4] I find that if expectations are fully rational, both serial correlation and volatility are small despite the adjustment frictions. In high income periods, potential buyers become wealthier each period and prices rise to ensure market clearing. Yet, the calibrated income shock sequence implies that the serial correlation of annual price growth rates is 0.28 and the standard deviation is at most one percent - just twenty percent of the value observed in the data. While relatively small income shocks can elicit a strong nonlinear response in purchase volume under constant prices, small changes in equilibrium prices are also sufficient to offset the adjustment impulse. The majority of transactions is still carried out under those prices, but in conjunction with expected mean reversion they move a sufficient number of households to the non-adjustment region.

Nevertheless, in a model without adjustment frictions the standard deviation of price growth rates would be even lower at 0.4 percent. This is due to rent-to-own transitions in the benchmark model, where buying a house generally offers a superior return on the equity invested. Therefore, expected mean reversion in prices plays a smaller role for new homeowners than for existing homeowners, who are more apt to pay down relatively expensive mortgage debt and delay upgrades. Even if prices are likely to increase for several periods, homeowners can expect to upgrade at a lower price eventually. The trading dynamics in the model reveal that their purchase volume declines during high income periods, while that of new homeowners increases.

To explore the possibility that subjective house price beliefs contribute to house price dynamics, I build on the framework developed in Adam and Marcet (2011) and Adam, Marcet, and Beutel (2015a). The setup maintains internal rationality in the sense that decisions are optimal given a well-defined system of subjective probability beliefs about future house prices. It emphasizes agents' uncertainty about the exact relationship between house prices and economic fundamentals, giving rise to a learning problem, where optimal behavior implies that past observations are used to update beliefs regarding future prices. In line with the evidence from household surveys, this introduces an extrapolative component into price growth expectations. Given that beliefs are expressed as

[4]The correlation between the average price index for the 25 metropolitan areas with the lowest supply elasticity according to the measure developed by Saiz (2010) and the national index is 0.998, and the degree of persistence and mean reversion in growth rates are very similar. An example is the the San Francisco metropolitan area, where the housing stock on average grew by just 0.5 percent per year between 1984 and 2008 (Paciorek (2013)) which just offsets the rate of population growth.

an entire probability distribution over future prices, agents nevertheless assign some weight to the possibility of sharp price declines even during booms.

I find that with subjective beliefs that lead to extrapolation, income shocks give rise to price dynamics that are much closer to the data. In line with its data counterpart, the serial correlation of price growth rates is 0.75, and the standard deviation is 4.3 percent - more than eighty percent of the observed value. Overall, the correlation between actual price growth rates and the ones implied by feeding the time series of observed income shocks into the model is 0.7. A notable deviation occurs during the last housing boom, where the model only matches the high growth rates during the early stage. More generally, the extrapolative component in expectations does not generate explosive house price dynamics, but amplifies price responses to fundamental shocks. The presence of adjustment frictions means that only a fraction of households are in the market each period, and all of them plan large adjustments that do not change by much in response to an income shock. This leads to a dampened, drawn out price response compared to a frictionless scenario where households adjust by small amounts almost every period.

The remainder of this paper is organized as follows. Section 2 discusses some of the related literature. Section 3 describes the model setup and how expectations are formed. Section 4 presents the calibration approach and the fit between the model and cross sectional facts on savings, homeownership, and the composition of home buyers. Section 5 discusses the model results under both rational and extrapolative expectations. Section 6 concludes.

2 Related Literature

There is a number of studies emphasizing the role of systematic biases in beliefs for house price dynamics. Agents' uncertainty regarding fundamentals leads them to extrapolate from past prices in a way that is consistent with the survey evidence in Case et al. (2012) and Piazzesi and Schneider (2009). In Glaeser and Nathanson (2015) prices would be informative regarding the unobserved stochastic growth rate of fundamentals if all agents at all times used the rational filter, but a "convenient approximation" by agents induces significant momentum following innovations in the growth rate. Gelain, Lansing, and Mendicino (2013) introduce an adaptive component into forecasts for all variables in agents' decision problems which generates excess volatility in prices. They find that the nature of agents' expectations is important for policy assessment. Adam, Kuang, and Marcet (2011) also use the subjective beliefs framework from Adam and Marcet (2011) and study the impact of the decline in real interest rates starting in 2001.

My model adds to that literature by studying market clearing in a heterogeneous agent economy where agents trade for fundamental reasons, face significant transaction costs, and switch between renting and owning. It documents that in such a setting realistic level changes in income can explain much of the observed patterns in price growth rates, provided that one incorporates extrapolative elements into belief formation. The fact that the model generates the high price growth observed during the early years of the last housing boom is an interesting connection to Adam et al. (2011), where high initial price growth due to exogenous preference shocks amplifies the price response to the interest rate decline after the year 2000. My present setting abstracts from time variation in financial conditions.

The emphasis on trade that occurs for fundamental reasons is also of interest with respect to the literature that studies the effect of extrapolative expectations on stock prices. While there is no trade in the representative agent model of Adam et al. (2015a), the stock market model of Adam et al. (2015b) considers trading that is due to belief disagreement. They find that a linear transaction tax would increase price volatility during normal times and the probability of boom-bust scenarios. The reason is that, due to investors' inaction regions, it may take relatively large price adjustments to absorb small supply shocks. With richer heterogeneity there are always buyers and sellers in the market who trade for fundamental reasons and could absorb small disturbances. The main effect of adjustment frictions in my model is that the initial response to an income shock is dampened when there is just a small group of adjusters each period.

Even with a comprehensive set of shocks, many rational expectations models have difficulty matching the volatility of house prices. Iacoviello and Neri (2010), for example, require a volatile and very persistent exogenous process driving housing preferences. However, Favilukis, Ludvigson, and Van Nieuwerburgh (2015) present a rich general equilibrium model with heterogeneous agents and adjustment frictions that can match the volatility of HP filtered house prices along with other business cycle moments. They do not focus on the serial correlation in price growth rates but emphasize the importance of time-varying risk premia in their model. High house prices are hence associated with lower expected returns which corresponds well with the data but less well with survey evidence on expectations. While my model is less rich than Favilukis et al. (2015) along several dimensions, I do add rent-to-own transitions and find that they play a role in inducing additional price volatility under rational expectations. The trading patterns that arise also suggest a role for relaxing the common assumption that the size distribution of the housing stock can be rearranged each period. This is consistent with findings in the more stylized models by Ortalo-Magné and Rady (2006) and Landvoigt, Piazzesi, and Schneider (2015).

Favilukis et al. (2015) is also one of several papers that calculate the transition path to a new steady state following a change in fundamentals. Such transitions are another explanation for high price growth over several consecutive years. One factor is that agents typically take the change in fundamentals to be permanent.[5] While there can be overshooting along the transition path such as in Chu (2014), mean reversion requires a revision in fundamentals. The fact that this is not anticipated by agents can be seen as another type of belief bias driving house price dynamics, but the approach focusses on single episodes. Finally, models that take into account search frictions in the housing market have been successful at generating serial correlation in price growth rates. Head, Lloyd-Ellis, and Sun (2014) empirically document serial correlation and excess volatility in house price growth rates in response to mean reverting income shocks at the city level. Relocation of agents can generate this pattern for single cities in their model, but they do not consider the implications for national house price trends. In Hedlund (2015b) the combination of search frictions and long-term defaultable mortgages generates considerable volatility and momentum in house prices which is a significant achievement relative to the literature. It is nevertheless of interest to study the role of systematic belief biases, since policy conclusions may differ. For example, Hedlund (2015a) finds in related work that stricter borrowing limits reduce welfare, but extrapolative price expectations may produce a different assessment.

3 The Model

3.1 Environment

Time is discrete and infinite and measured in years. Each period, a constant mass of households enters the economy and starts working for N years. Income $y_{i,n,t}$ of household i aged $n \leq N$ in period t is measured in units of the numeraire good. It depends on a deterministic age-specific component $\bar{y}_n$, an idiosyncratic component $z_{i,t}$, an aggregate component Z_t, and the deterministic growth rate λ:

$$\ln y_{i,n,t} = \ln \bar{y}_n + z_{i,t} + Z_t + \lambda t.$$

Both $z_{i,t}$ and Z_t follow Markov chains defined on the finite sets Ω_z and Ω_Z, and $\bar{y}_n$ is hump shaped in age. Households have time-separable preferences and discount the future at rate β. Their

[5] In current work by Kaplan, Violante, and Mitman (2015), fundamentals such as credit conditions and housing preferences are expected to follow a unit root process, and a particular sequence of innovations causes a large swing in house prices as well.

per-period utility function is given by

$$u(c,h) = \frac{1}{1-\gamma}\left(c^{1-\nu}h^{\nu}\right)^{1-\gamma}$$

where c denotes how much of the numeraire good is consumed and h measures the quantity of housing services available to the household. Those are derived from either renting or owning a housing unit of size h in that period.[6]

Owner occupied housing and rental housing are traded in separate markets. The stock of owner occupied housing is fixed and owned by working age households in the economy. The menu of different sizes available for purchase is given by Ω_h, and households own at most one of these units. In that case they pay a maintenance cost $\mu_t h$ each period that fully offsets depreciation. To adjust their housing space, owners must sell the unit they currently own and newly purchase or rent a unit of the desired size. Each seller incurs a transaction cost that is a fraction φ of the sales price. The outlined life cycle income profile will nevertheless lead homeowners to move to larger or smaller units occasionally to adapt to income changes. In order to account for other reasons to reenter the housing market, I also include exogenous moving shocks. With probability π_ε a household must move to a different location and is subject to the transaction cost for selling her current house. I assume that the size distribution going into the market can be converted at no further cost into the size distribution that agents want to purchase. This reduces market clearing to the single condition that the total quantity of owner occupied housing does not change. It also means that the value of a home is directly proportional to its size and given by $p_t h$.

Rental housing is owned and operated by firms that fund themselves in the international financial market. In each period, they are able to supply the desired quantity of apartment space at a rental rate q_t. The rental rate is not explicitly linked to a market clearing condition, but it can fluctuate with the price for owner occupied housing and grows with per capita income. Moving between rental units is costless for households, and renters can choose any unit size between $h_{min} = \min(\Omega_h)$ and $h_{max} = \max(\Omega_h)$.

Households can invest their surplus funds in the international financial market at riskless rate r. They can also borrow at rate $r^b > r$ when posting housing collateral. Upon origination of the loan, the amount outstanding (including the first interest payment) cannot exceed a fraction θ of the current home value, and it must always be fully repaid whenever the house is sold. As long as debt is below the borrowing limit $\frac{\theta p_t h}{1+r^b}$, a homeowner is free to increase the outstanding amount. If prices

[6] While I will always refer to it as size, h could also be interpreted more broadly as a composite of size and location for example.

fall this threshold may be exceeded. In such a case, the only requirement regarding repayment is that the amount outstanding grows at most at the exogenous growth rate λ, but it is not necessary to supply additional equity. This setup reflects the pattern that equity extraction is often observed in response to house price increases, while households' leverage ratios can remain high after a price drop (compare for example Justiniano, Primiceri, and Tambalotti (2015)). In summary, the lower bound for financial assets a_{t+1} at the beginning of period $t+1$ is given by

$$a_{t+1} \exp(-\lambda) \quad \geq \quad \min\left(-\frac{\theta p_t o_t h_t}{1+r^b}, a_t \mathbb{1}_{[o_t h_t = o_{t-1} h_{t-1} \wedge \varepsilon_t = 0]}\right)$$

where o_t is one if the household owns a house in period t and zero otherwise, ε_t is one if the household experiences a moving shock in period t and zero otherwise, and $\mathbb{1}_x$ is the indicator function for condition x. Here, the indicator function is one if a household does not move and borrowing corresponds to a negative value for a.

Similar to Corbae and Quintin (2015), I assume that upon retirement all agents exit the housing market that is studied here. They immediately sell their home and derive utility from consuming out of their financial wealth and a pension benefit b_t that also grows at rate λ. Retirees' probability of dying is constant at π_d and there is an asset that allows them to annuitize financial wealth at the actuarially fair rate $\frac{1+r}{1-\pi_d} - 1$. They procure housing at a rental rate q_t^r that may be correlated with the rental market for working age households. This very stylized setup maximizes the potential for autocorrelation in house price growth by making an important component of net supply independent of price expectations, but it neglects the role that price fluctuations may play in the slow decumulation of housing wealth during retirement.

Finally, taxation corresponds to the benchmark setup in Diaz and Luengo-Prado (2008). Labor as well as interest income and retirement benefits are subject to a flat tax rate τ. Since the implicit rental income from owning a house is not taxed, this makes home ownership more attractive given a price-to-rent ratio. In addition, mortgage interest expenditure is fully deductible from taxable income and capital gains on housing are not taxed.

3.2 Decision problems

To bound the state space for y, all variables expressed in terms of the numeraire good are normalized by trend growth when setting up the decision problems. Specifically, the detrended analogues of

$y_t, a_t, p_t, q_t, q_t^r, \mu_t, b_t$, and c_t are defined by the following rule:

$$\grave{x}_t = x_t \exp(-\lambda t).$$

The resulting value functions can be normalized by a factor $\exp[-\lambda t(1-\nu)(1-\gamma)]$ without affecting the decision problem. The baseline specification of the retirees' problem also assumes that benefits and the rental rate grow deterministically at rate λ, so that $\grave{b}_t = b$ and $\grave{q}_t^r = q^r$. In that case, the normalized value of retirement is a function of just the beginning-of-period financial assets $\grave{a}_t$ which for notational convenience are defined net of the tax liability on accrued interest:

$$\begin{aligned}
V^r(\grave{a}_t) &= \max_{\{\grave{c}_t, h_t, \grave{a}_{t+1}\}} \left[u(\grave{c}_t, h_t) + \tilde{\beta}(1-\pi_d)V^r(\grave{a}_{t+1})\right] \\
\text{s.t. } \grave{c}_t &= (1-\tau)b + \grave{a}_t - \frac{\exp(\lambda)(1-\pi_d)}{1+r-\tau(r+\pi_d)}\grave{a}_{t+1} - h_t q^r \\
\grave{a}_{t+1} &\geq 0 \\
h_t &\in [h_{min}, h_{max}].
\end{aligned}$$

The adjusted discount factor $\tilde{\beta} = \beta \exp[\lambda(1-\nu)(1-\gamma)]$ reflects the normalization of future consumption.

An agent aged $n \leq N$ in period t will retire in period $R = t + N - n + 1$ and solves the following problem:

$$\max_{\{\grave{c}_s, \grave{a}_{s+1}, h_s, o_s\}_{s=t}^{R-1}} E_t\left[\sum_{s=t}^{R-1} \tilde{\beta}^{s-t} u(\grave{c}_s, h_s) + \tilde{\beta}^{N+1-n} V^r\left(\grave{a}_R + (1-\varphi)o_{R-1}h_{R-1}\grave{p}_R\right)\right] \tag{1}$$

s.t. $\forall s = t...R$

$$\begin{aligned}
\grave{c}_s &= (1-\tau)\grave{y}_s + \grave{a}_s - \left(\frac{\exp(\lambda)}{1+(1-\tau)r}\mathbb{1}_{\grave{a}_{s+1}\geq 0} + \frac{\exp(\lambda)}{1+(1-\tau)r^b}\mathbb{1}_{\grave{a}_{s+1}<0}\right)\grave{a}_{s+1} \\
&\quad -(1-o_s)h_s\grave{q}_s - o_s h_s \grave{\mu}_s - (o_s h_s - o_{s-1}h_{s-1})\grave{p}_s - \varphi o_{s-1}h_{s-1}\grave{p}_s \mathbb{1}_{[o_s h_s \neq o_{s-1}h_{s-1} \wedge \varepsilon_s = 1]} \\
\grave{y}_s &= \bar{y}_{n+s-t}\exp(z_s + Z_s) \\
\grave{a}_{s+1} &\geq \min\left(-\theta\frac{1+(1-\tau)r^b}{1+r^b}\grave{p}_s o_s h_s, \grave{a}_s \mathbb{1}_{[o_s h_s = o_{s-1}h_{s-1} \wedge \varepsilon_s = 0]}\right) \\
o_s h_s \in \Omega_h &\quad \wedge \quad (1-o_s)h_s \in [h_{min}, h_{max}];
\end{aligned}$$

$\grave{a}_t, \grave{p}_t, \grave{q}_t, \grave{\mu}_t, o_{t-1}, h_{t-1}, \varepsilon_t, z_t, Z_t$ given.

Denoting the vector of normalized beginning-of-period state variables $(n, z, a, h, o, \varepsilon)$ by ω, the decision problem gives rise to time dependent decision rules $c_t^*(\omega), a_t^*(\omega), h_t^*(\omega)$, and $o_t^*(\omega)$ for all

$\omega \in \Omega = \Omega_n \times \Omega_a \times \Omega_h \times \Omega_o \times \Omega_\varepsilon \times \Omega_z$ where $\Omega_n = \{1, ..., N\}, \Omega_a = [a_{min}, a_{max}]$, and $\Omega_o = \Omega_\varepsilon = \{0, 1\}$. I index all objects by t to account for the potential time dependence of the expectations operator E_t. In the next section I outline how I model the evolution of agents' expectations regarding payoff relevant variables over time.

3.3 Expectations

Throughout, expectations for z and Z are formed according to the corresponding Markov processes. Due to aggregate shocks, also housing demand and hence $\grave{p}$ change over time. In principle, agents would therefore need to predict how the distribution of agents over Ω evolves conditional on the aggregate shocks and how it translates into a market clearing price for housing. This approach does not only put a large cognitive burden on agents, but it also makes the model computationally intractable. Instead, I assume that agents have a system of beliefs that allows them to consistently assign joint probabilities to all future sequences of payoff relevant variables based on a much smaller set of states. This system of beliefs governs how expectations are updated each period, and while it may not reflect the true probabilities, agents act optimally given their beliefs. I consider two established approaches from the literature to specify such a system of beliefs.

In order to approximate a rational expectations equilibrium, I follow a strategy based on Krusell and Smith (1998).[7] Specifically, I use a forecasting rule of the form

$$\ln \grave{p}_{t+1} - \ln \bar{p} = \alpha_1(Z_t, Z_{t+1}) + \alpha_2(Z_t, Z_{t+1})(\ln \grave{p}_t - \ln \bar{p}). \tag{2}$$

The forecasting rule is written in terms of deviations from a steady state value $\bar{p}$, and the coefficients are indexed by the contemporaneous and future realization of the aggregate shock. They are chosen in a way that maximizes the forecasting performance with respect to $\grave{p}$, given that all agents use this very rule to form expectations. The appendix reports more details on the iterative procedure employed to find such a set of coefficients. With an R^2 of at least 0.997, the forecasting performance of the rules employed here is high. Equation (2) implies that agents' decision rules depend on time only through the state variables Z_t and $\grave{p}_t$ which are sufficient to assign probabilities to all future sequences of payoff relevant variables.

The alternative approach I employ follows Adam et al. (2011). Agents are uncertain about the exact relationship between house prices and economic fundamentals, meaning that their beliefs are characterized by a non-degenerate distribution over future prices given any sequence of aggregate shocks. The underlying system of beliefs emphasizes uncertainty about the future growth rate of

[7] See Favilukis et al. (2015) for an application involving house prices.

house prices. More specifically, log house prices are believed to evolve as a random walk with an unobservable time-varying drift δ_t:

$$\begin{aligned}
\ln \grave{p}_t &= \ln \grave{p}_{t-1} + \ln \delta_t + \ln \epsilon_t && (3)\\
\ln \delta_t &= \ln \delta_{t-1} + \ln \phi_t && (4)\\
\begin{pmatrix} \ln \epsilon_t \\ \ln \phi_t \end{pmatrix} &\sim iiN\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_\epsilon^2 & 0 \\ 0 & \sigma_\phi^2 \end{pmatrix}\right).
\end{aligned}$$

Since neither the drift nor the innovations ϵ_t and ϕ_t can directly be observed, this setting gives rise to a learning problem. Agents use Bayesian updating to adjust their estimate of δ_t and hence their beliefs regarding future prices based on observed price growth. Adam et al. (2015a) show that under suitable assumptions regarding priors and the information structure the evolution of beliefs under the optimal Bayesian filter has the following recursive representation:

$$\begin{aligned}
\ln \delta_t &\sim N(\ln m_t, \sigma^2) && (5)\\
\ln m_t &= \ln m_{t-1} + g\left(\ln \grave{p}_{t-1} - \ln \grave{p}_{t-2} - \ln m_{t-1}\right). && (6)
\end{aligned}$$

The uncertainty regarding the estimate for $\ln \delta_t$ corresponds to its Kalman filter steady state value

$$\sigma^2 = \frac{1}{2}\left(-\sigma_\phi^2 + \sqrt{\left(\sigma_\phi^2\right)^2 + 4\sigma_\phi^2\sigma_\epsilon^2}\right).$$

and results from the unobservability of the initial drift and subsequent innovations. The learning gain is given by $g = \frac{\sigma^2}{\sigma_\epsilon^2}$.

Equation (6) is a key determinant of the model's price dynamics. Observed price growth feeds back into higher price growth expectations. This raises agents' willingness to pay for housing, so that higher price growth expectations can be justified ex post and may even be revised further upwards. The information structure implies that beliefs about m are only updated after agents have made their choices and observed equilibrium prices. This circumvents the possibility of multiple equilibria which may result from contemporaneous changes in the expected growth rate. Combining equations (3), (4), and (5), agents' believes in period t for $\grave{p}_t$ are given by

$$\ln \grave{p}_{t+1} = \ln \grave{p}_t + \ln m_t + \ln \tilde{\epsilon}_{t+1}. \qquad (7)$$

The uncertainty $\tilde{\epsilon}$ regarding future price growth combines the uncertainty regarding m_t and the uncertainty regarding future persistent and transitory innovations. It is normally distributed with variance $\tilde{\sigma}^2 = \sigma^2 + \sigma_\epsilon^2 + \sigma_\phi^2$. Even if agents are optimistic regarding regarding m_t, they always take into account the possibility of falling prices. Their decision rules depend on time only through $\grave{p}_t$ and the additional state variable m_t.

3.4 Equilibrium

Based on the decision problem, an equilibrium consists of an initial distribution $\pi_0(\omega)$ of working age households over individual states, the corresponding housing stock $H = \int_\Omega \omega_h \omega_o \pi_0(\omega) d\omega$, a system of beliefs over all possible sequences of payoff relevant variables with the associated probability measure P, a sequence of aggregate shocks $\{Z_t\}$, sequences of house prices $\{\grave{p}_t\}$, maintenance costs $\{\grave{\mu}_t\}$, and rental rates $\{\grave{q}_t\}$, sequences of decision rules, and a sequence of distributions $\{\pi_t(\omega)\}$ such that

1. Each period, the decision rules are optimal given $\grave{p}_t, \grave{q}_t, \grave{\mu}_t, Z_t$, and P.
2. Each period, the housing market clears
$$\int_\Omega h_t^*(\omega) o_t^*(\omega) \pi_t(\omega) d\omega = H.$$
3. The sequence of distributions $\{\pi_t(\omega)\}$ is implied by the sequence of decision rules, the initial distribution, and the distribution of new workers in each period.

4 Parameter Selection and Calibration

My objective is to produce a good fit between the model and cross sectional data on savings, homeownership, and the composition of home buyers. I choose a number of parameters in a standard way and in line with evidence presented in other studies. The remaining parameters are then jointly calibrated to match selected moments from the Panel Study of Income Dynamics (PSID). The specification for the aggregate shock process is based on real per capita income.

4.1 Parameters Selected Independently

Agents' working life goes from age 21 to age 65, so that $N = 45$. Like Iacoviello and Pavan (2013) I approximate their earnings profile $\bar{y}_n$ with a quadratic polynomial that is normalized at 1 at age 21, peaks at 1.82 at the age of 50, and reaches 1.61 at the age of 65. Based on the estimates in Storesletten et al. (2004), I construct a nine state Markov chain for the idiosyncratic component of earnings by discretizing an AR(1) process with persistence $\rho_z = 0.95$ and conditional volatility $\sigma_z = 0.17$ using the Rouwenhorst method. I set the risk aversion parameter γ in the utility function to the standard value of 2. The unit elasticity of substitution between housing and consumption is based on evidence in Davis and Ortalo-Magné (2011) that housing expenditure shares are approximately constant over time and across U.S. metropolitan statistical areas. From this study I also take the housing expenditure share parameter $\nu = 0.18$. The minimum house size is set such that the poorest agents in the economy spend fifty percent of their income on rents, and the maximum size is set sufficiently large to not be chosen by any agent. To parametrize the process for μ_t, I impose that maintenance makes makes up 2.4 percent of the home value during a phase of normal income shocks. This is average depreciation of the residential housing stock from BEA fixed asset tables between 1975 and 2014.

To obtain a measure of the riskless real rate, I adjust the 10-Year Treasury Constant Maturity Rate by CPI inflation and average over the sample period 1975 to 2014. This yields $r = 0.026$. I determine the borrowing premium $r^b - r$ as the average difference between the 30-Year Conventional Mortgage Rate and the 10-Year Treasury Constant Maturity Rate over the same period, yielding $r^b = 0.043$. The borrowing constraint is set at $\theta = 0.8$. The data in Bokhari et al. (2013) shows that there has been considerable bunching of mortgage origination at exactly this threshold for several decades. While in reality households have always used higher leverage ratios as well, this observation seems to indicate that for many it becomes more difficult or costly beyond this point.[8] The typical commission charged by real estate agents is six percent of the home value which I take as the parameter value for φ. A value of this magnitude is used in many studies and is similar to the transaction cost parameter estimated in Berger and Vavra (2015) based on the observed frequency of durable adjustments in the PSID. For the tax rate, I take $\tau = 0.2$ from the benchmark setup in Diaz and Luengo-Prado (2008). Finally, log trend growth in the deflated Freddie Mac House Price Index has been 0.01. This is the value I use for the growth rate λ. Because I have a constant housing stock, I cannot match the higher growth rate of per capita income at the same time.

4.2 Joint Parameter Selection

The remaining parameters are the discount rate β, the normalized retirement benefit b, and the probability of an exogenous moving shock π_ε. In addition, the process for the rental rate q_t has to be specified. While the house price can fluctuate in response to aggregate shocks, a long sequence of identical aggregate shocks will result in a constant normalized house price level. This also leads

[8] For example, they must buy private mortgage insurance in order to conform with loan requirements by government sponsored enterprises.

to a constant price-to-rent ratio, independent of what I assume on the comovement between house prices and rents. I denote the constant price-to-rent ratio following a long series of normal shocks by $\bar{p}$ and use it as a fourth calibration parameter. Normalizing the corresponding rental payment to one, this parameter and the degree to which rents comove with market clearing house prices pin down q_t.

I determine the four parameters by matching four selected cross-sectional moments. For the financial variables, I use median values since my model will not be able to generate a strong right tail in the wealth distribution, and I normalize by median income to make numbers comparable between the model and the data. Because the process for retirement is very stylized, I target median net worth of households aged 61 to 65 to have some comparability of wealth accumulation over the life cycle. For the other moments, I restrict calculations to the group between ages 26 and 60 since there are relatively few households below 26 in the data and because of the stylized nature of the retirement process. I target median net worth, the homeownership rate, and the value of houses acquired by new homeowners relative to the value of all houses acquired by movers within twelve months. The latter is included to ensure that both groups have plausible weights in the housing market.[9]

The cross-sectional moments are derived from the 1999 PSID interviews. The survey's panel structure has the benefit that I can determine the previous homeownership status of the household from its 1997 interview.[10] Together with the reported most recent moving date, this allows me to derive estimates of the fraction of homeowners who have bought a new residence within twelve months and whether they are new homeowners or not. The interviews were conducted during the first months of 1999, and I interpret the data as representative of the distribution arising after a long sequence of normal income shocks. According to my calibration of the aggregate income process described in the next section, the nineties have indeed the longest sequence of normal shock realizations in the sample. Even though this sequence ends in 1998 while 1999 already falls into the dot-com boom phase, I still consider the above interpretation reasonable.

I generate the parameter estimates from a version of the model with constant normalized house prices and then use the same parameter values for all specifications in the next chapter. As documented in the appendix, especially for the rational expectations models the changes in moments are quite small. Calibrating each of the flexible price models separately would be much more challenging computationally however. Table 1 shows the fit between targeted moments and the

[9] As documented in the next chapter, their reaction to price fluctuations is quite different.

[10] The data appendix discusses how I match interview data from both survey waves and how I handle cases where previous data is missing.

constant price model. While all parameters are determined jointly, each row lists the one that mainly drives the corresponding statistic.

Table 1: Calibration Results

Targeted Moment	*Data*	*Model*	*Parameter*	*Value*
Homeownership rate	63.7%	63.7%	$\bar{p}$	18.3
Home value new HO / Home value movers	63.9%	63.9%	π_ε	0.014
Median net worth	101.3%	101.3%	β	0.977
Median net worth before retirement	343.0%	343.4%	b	58.8%

Note: Median net worth before retirement is calculated based on households between age 61 and age 65, the other rows are calculated based on households between age 26 and age 60. All financial variables including the retirement benefit b are given as a percentage of median income.

Overall, the parameters look very reasonable. The discount factor β falls in the range of values routinely employed in the literature. The retirement benefit is about fifty percent of average income.[11] The estimate for π_ε implies that fifty percent of all moves are not induced by the exogenous moving shock. In the PSID, about 44 percent of moving homeowners report as their primary reason for moving the desire to expand or contract housing or to move to a better neighborhood. The price-to-rent ratio is within the range documented by Davis et al. (2008) for the period from 1960-1995. To also assess life cycle patterns implied by the calibration, I compare homeownership rates and median net worth by five year age bins. Figure 2 shows that the model matches their increase over the life cycle quite well.

4.3 Aggregate Income Shocks

My calibration of the aggregate shock process is based on personal income from the national income and product accounts (NIPA). I adjust the per capita series by taking out income receipts on assets and rental income of persons. Then I deflate the series with the consumer price index excluding shelter, convert it to per capita terms and take logs. The deviations of this series from a linear trend can be interpreted as the Z_t process in my model. I group them into three categories. I assign all values between -0.02 and 0.02 to the normal state with $Z_{normal} = 0$ which is almost exactly their average value. The average value of all values above 0.02 is 0.04 which I take as the realization for Z_{high}. The low realization is set to $Z_{low} = -0.05$, again based on the average of the corresponding deviations. Figure 3 shows the original series and my discretization. Notably, income remains high during the 2001 recession. The process would be similar when working with all components of per

[11] While this is above the typical replacement rate for Social Security alone, one has to consider the additional healthcare and spousal benefits and preferential tax treatment not modeled here.

Figure 2: Evolution of assets by age group

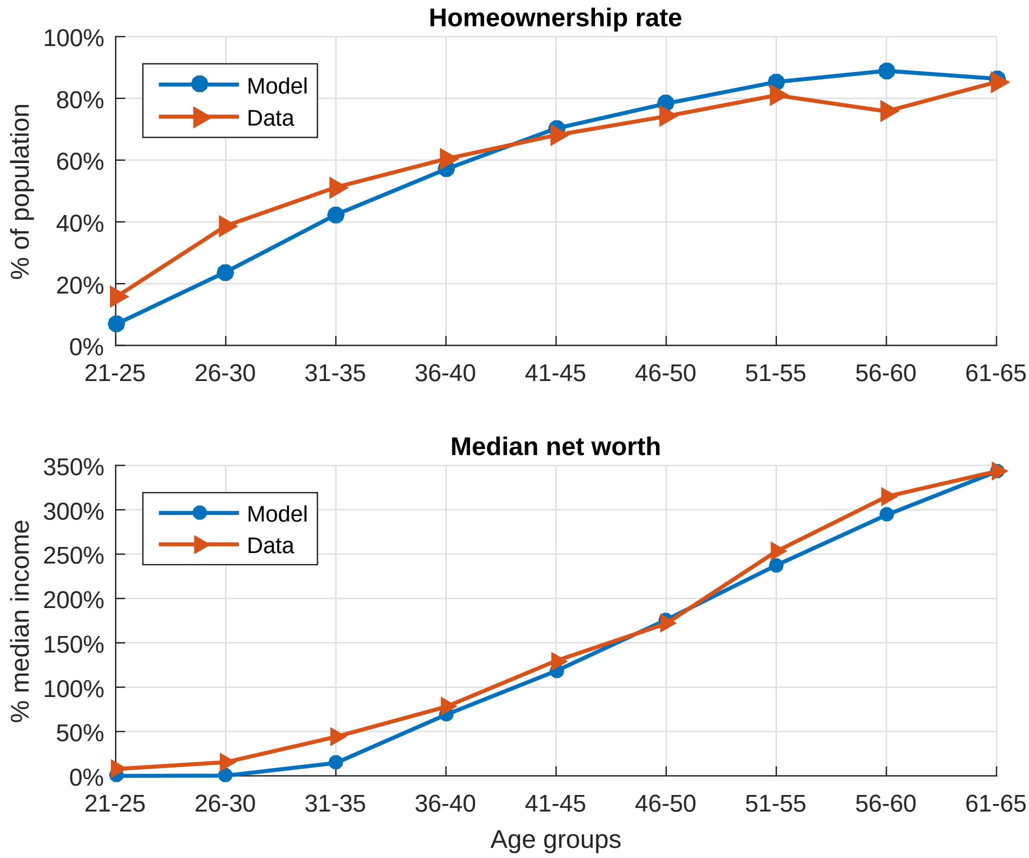

Notes: Homeownership rate is the fraction of all households of a given age group owning a house. Net worth is the sum of the value of all assets net of liabilities and expressed as a percentage of median income in either the data or the model. The data is taken from the 1999 PSID survey.

capita income instead.

From the discretized series, I estimate transition probabilities between states under the assumption that the last sequence of income realizations below trend ends in 2015. Based on the estimation results, I use the following matrix as Markov transition probabilities Π_Z in my model:

Table 2: Aggregate income states

		Transition probabilities		
Income state	*Value*	to Z_{low}	to Z_{norm}	to Z_{high}
Z_{low}	-0.05	0.8	0.2	0
Z_{norm}	0	0.15	0.65	0.2
Z_{high}	0.04	0	0.2	0.8

Figure 3: Aggregate income process

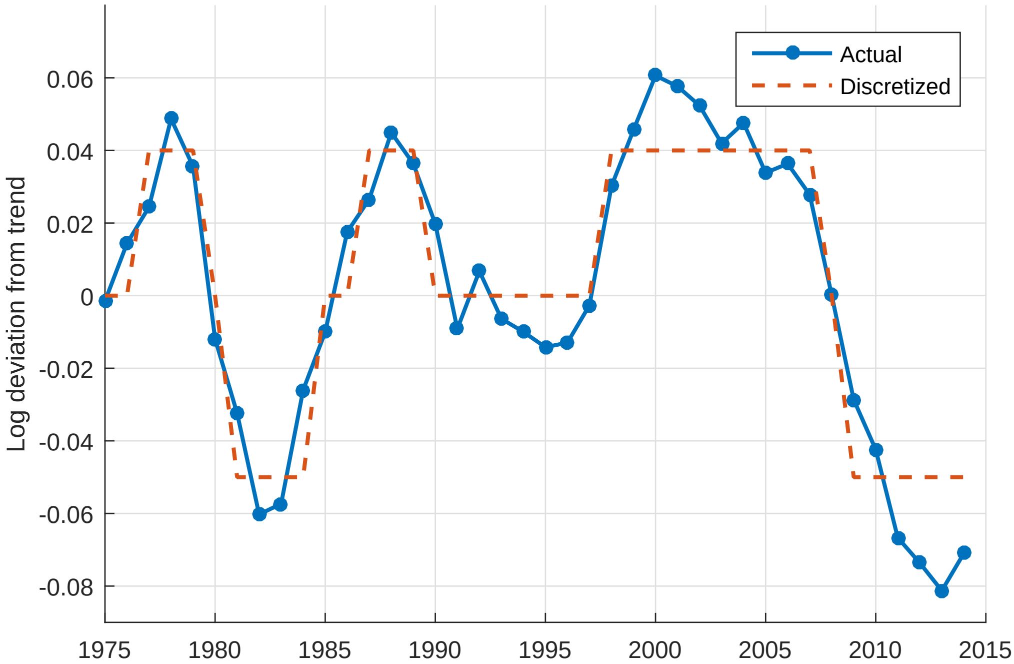

Notes: Log deviations of the real per capita income series from trend and a three state discretization.

5 Results

5.1 Rational Expectations Benchmark

The benchmark setting is chosen in a way that ensures the maximum possible price impact of innovations in the aggregate income process. To that end, I impose that rental rates fully fluctuate with house prices and that maintenance costs remain constant. Figure 9 shows the resulting price path, expressed as log deviations from trend, in response to the aggregate income shock series described in section 4.3. The initial distribution is the one resulting from a long series of normal income realizations. The largest innovations in the price series happen in direct response to changes in the income level. Nevertheless, also in subsequent periods where the income level remains constant the price level continues to drift further in the direction of the initial change. Overall, the correlation in annual price growth rates is 0.28, about forty percent of what is observed in the data.

It is apparent that fluctuations in the price level are of smaller magnitude than fluctuations in the aggregate income level. The standard deviations are 0.018 and 0.036 respectively. The standard deviations in the data counterparts of the price and income series are 0.113 and 0.040, so that aggregate income shocks of plausible magnitude can only generate a small fraction of the observed house price volatility. The low volatility in the level of house prices also means that the

Figure 4: Simulated house price series

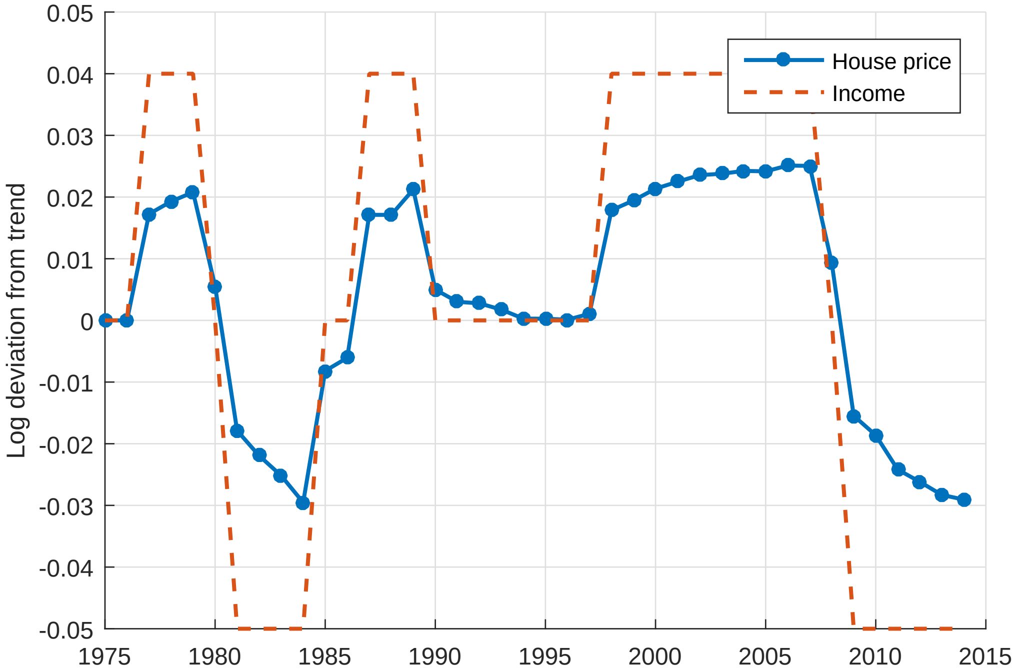

Notes: Deviations of log house prices and log aggregate income from trend.

volatility in price growth rates is small. At 0.01 it is about twenty percent of its data counterpart.

Table 3: Moments of price growth rates

		Model results		
		$\varphi = 0.06$	$\varphi = 0$	$\varphi = 0$
Moment	*Data*	$\theta = 0.8$	$\theta = 0.8$	$\theta = 1$
One year correlation	0.70	0.28	0.34	0.38
Three year correlation	0.08	-0.18	-0.14	-0.11
Five year correlation	-0.48	-0.30	-0.26	-0.31
Standard deviation	0.05	0.01	0.008	0.004

Notes: Price growth rates are changes in the detrended log price index. Data: Freddie Mac House Price Index deflated by CPI urban excluding shelter. Model: rational expectations benchmark calibration with different parametrizations for the leverage ratio θ and the transaction cost φ.

One important feature of the model studied here are micro-level adjustment frictions in combination with a rent-or-own decision. How do those affect market clearing and house prices? Table 3 compares the volatility and momentum of price growth rates in the benchmark model to a frictionless version where the borrowing constraint is relaxed and the sales fee φ is zero. The absence of those frictions implies that all households become homeowners immediately. The volatility of price growth rates is then less than half of what I find in the benchmark version. The results for a hybrid model, where only the sales fee is zero, indicates that both frictions contribute to the increased volatility. The main difference induced by adjustment frictions is a larger price increase on impact of the shock, while the subsequent drift is very similar. This means that serial correlation actually slightly decreases with adjustment frictions.

The role that demand by new homeowners plays is illustrated in figure 5. Panel a) shows the impulse response of net housing purchases to a sequence of high income realizations in a model with fully elastic housing supply and constant house prices. Absent the price response, net housing purchases would be up to thirty percent above their steady state level and the growth rate of purchases remains high for several years. As the split by buyer groups shows, the latter is due to the rent-or-own decision. More and more households pass the adjustment threshold each period, which leads to the depicted increase in purchase volume and in the homeownership rate. This pattern would be very different in a frictionless model, where all households are homeowners. There, the bulk of the adjustment happens directly on impact and absent an extensive margin the overall increase in owner-occupied housing is less pronounced.

Figure 5: Changes in net purchase volume

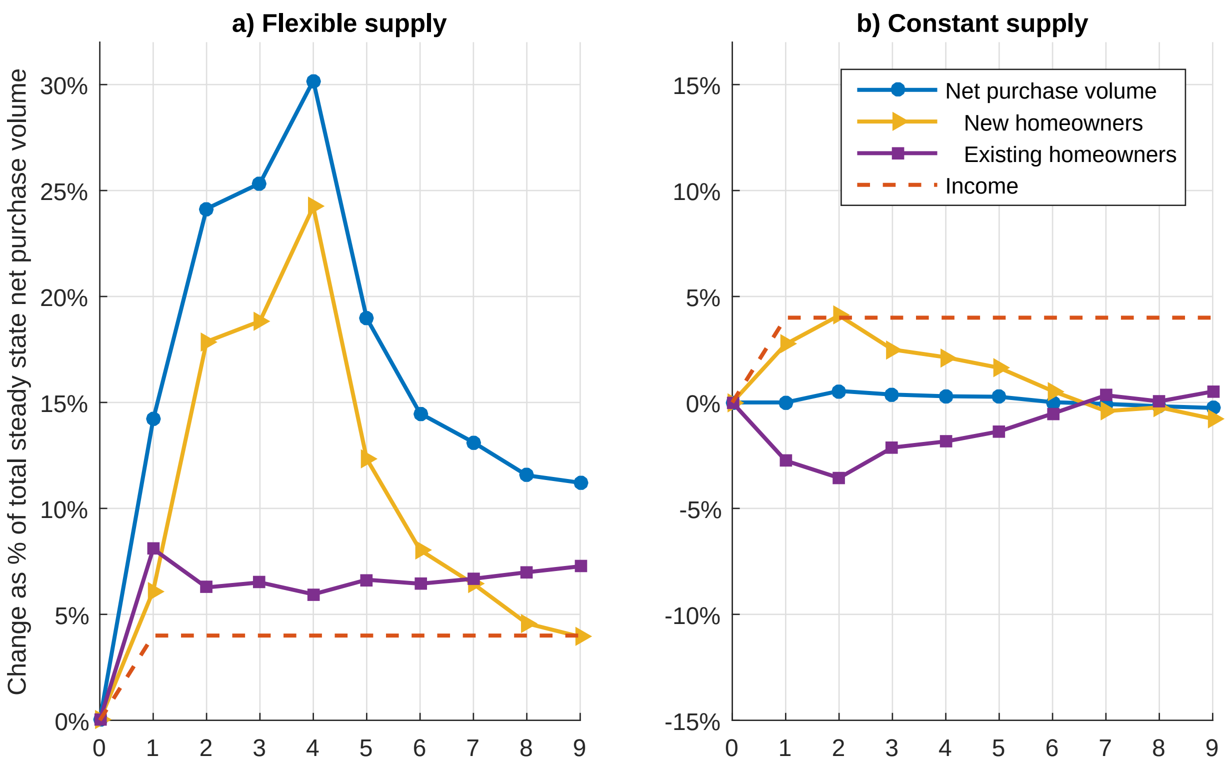

Notes: Net purchase volume is the size difference between the new housing unit and the one owned previously, where rental units are assigned a zero size. The benchmark is net purchase volume following a long period of normal income shocks (steady state). The graph shows the deviations of net purchase volume from this quantity split by whether a households transitions from renting or initially owns a home. For both groups these deviations are normalized by total steady state net purchase volume so that the depicted series add up to the change in total net purchase volume.

Panel b) depicts the impulse responses with a constant housing stock. Total net purchases are then also essentially constant since the only source of net supply in the model is the largely predetermined housing stock of retirees. Nevertheless, the composition varies considerably in response to income shocks as shown in figure 5. Now the only group that increases net purchases during the boom state are new homeowners. The endogenous price increase is sufficient to induce existing homeowners to delay their upgrades, so that the group's decrease in net purchase volume accomodates increased purchases by new homeowners. Since existing homeowners are the wealthier group and less likely to hit the downpayment constraint, this different pattern is not explained by how price changes affect households' ability to afford home purchases. Instead, it can be seen as a manifestation of different return considerations.

In order to generate a large homeownership rate in this model, owning a house must yield a greater total return than investing in financial assets.[12] Buying a home enables a household to put all their equity into this superior investment technology. Existing homeowners on the other hand have already invested all their equity and can use surplus funds to pay down mortgage debt which carries an interest rate premium. They are hence less likely to invest in more housing in times where higher prices and expected mean reversion lower expected returns. In a frictionless model, upgrades by existing homeowners make up more than ninety percent of the net purchase volume, so that their return considerations dominate. The benchmark version, on the other hand, has been calibrated to have a plausible ratio of purchases by new and existing homeowners, so that new homeowners play a much larger role in determining the market clearing price. Nevertheless, the more elastic response by existing homeowners continues to limit the price response.

To assess the relatively small fluctuations in house prices, one can also look at innovations in life-time income caused by the aggregate income shock. As an approximation I compare the discounted earnings difference between the different states for Z for an infinitely lived agent with no idiosyncratic earnings risk. Similar to Diaz and Luengo-Prado (2010) I calculate the "permanent" earnings shock $\hat{Z}$ as

$$\hat{Z} = Z + \frac{\exp(\lambda)}{1+r^b} \sum_{Z'} \Pi_Z(Z, Z')\hat{Z}'.$$

While Z_{high} is about nine percent greater than Z_{low}, expected mean reversion in the income process means that the difference in discounted total earnings is just 1.2 percent between the two states. This is a similar order of magnitude as the price fluctuations observed in the frictionless model.

In summary, under rational expectations only a fraction of the observed house price fluctuations

[12]This is also true if housing yields a utility premium which makes owning a house equivalent to renting a unit of a larger size. It is not necessary if rental units come only in sizes too small for the majority of households.

could be attributed to aggregate income shocks. While adjustment frictions and a rent-or-own decision lead to volatile net housing purchases with strongly correlated growth rates, the small price responses are sufficient to make net purchases constant over time. Nevertheless, adjustment frictions and a rent-or-own decision double the volatility of price growth rates. Despite some serial correlation in price growth rates, expected mean reversion dominates, so that expected returns of investing in housing are lower when prices are high.

5.2 Extrapolative Expectations

Under extrapolative expectations, the expected price dynamics are described by equations (3) and (6) as fluctuations around a stochastic trend. They are parametrized by the standard deviation of the expected fluctuations $\tilde{\sigma}$ and by the gain parameter g. I set $\tilde{\sigma} = 0.04$ which is slightly below the volatility of price growth rates in the data, but close to the volatility ultimately implied by the model. The baseline value for g is 0.022 which is in line with survey evidence on the speed of learning in other asset markets. The effect of choosing other values for g is shown in table 4. I deviate from the setup in the previous section by imposing that normalized rental rates are constant, i.e rents grow only as fast as productivity. This specification reflects the empirical fact that rental rates fluctuate much less than prices. While the choice of rental rates that fully comove with house prices was motivated by maximizing price movements in the rational expectations solution, it would exaggerate volatility in this setup.

Figure 6 shows the growth rate of prices generated by the extrapolative specification together with its data counterpart. Both series are strongly correlated (0.69) and also similar in terms of magnitude. As shown in table 4, the extrapolative expectations model therefore does not only match the serial correlation of price growth rates very well, but it also matches the degree of mean reversion and comes close to the volatility observed in the data. The table also shows results for different values for g, indicating that similar dynamics can be observed for a range of values. One notable difference between the two series in figure 6 are growth rates during the most recent housing boom. Personal income is above trend over the entire period, but price growth slowly declines in the model while it remains exceptionally high in the data. There is a lot evidence that several factors beyond the model's scope contributed to these developments. The model illustrates however that a moderate increase in aggregate income could have contributed to the early stages of the boom and to elevated price growth expectations.

It should be noted that the presence of an adjustment cost plays a different role for model dynamics under extrapolative expectations. While equation (3) implies that households always

Figure 6: Price growth rates: extrapolative expectations and data

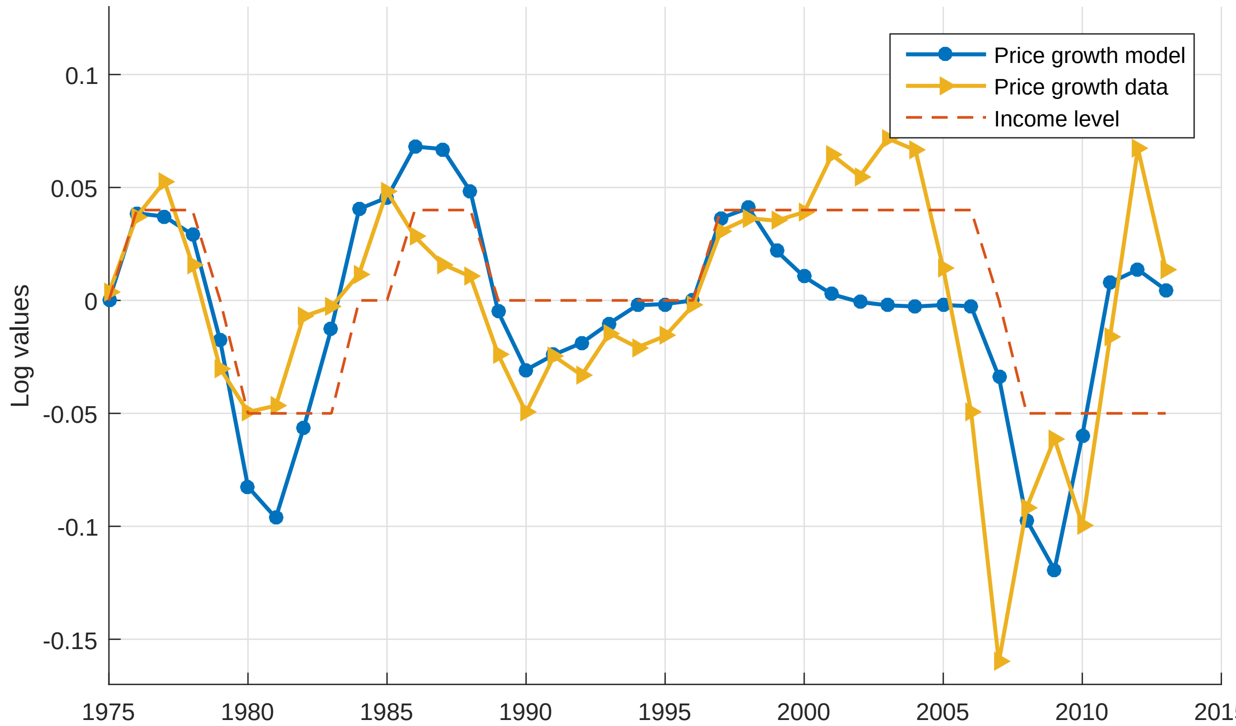

Price growth rates are changes in the detrended log price index. The income series is a discretization of the level of real per capita income relative to trend. Data: Freddie Mac House Price Index deflated by CPI excluding shelter. Model: extrapolative expectations benchmark calibration

take the possibility of falling prices into account, an innovation in the price level shifts their entire probability distribution over future prices upward. This lack of expected mean reversion particularly affects how existing homeowners react to price increases. Many of them are leveraged, so that their net worth increases by more than prices. The disproportionate increase in net worth means that it becomes easier to buy a larger house despite the price increase - especially for relatively recent homeowners who were at the downpayment constraint when buying. Absent a transaction cost many homeowners would be in the housing market each period, and in that case their total purchase volume can even be locally increasing in the price. This causes a large but transient price increase on impact of the income shock. With transaction costs homeowners will upgrade only very infrequently and hence make up a smaller fraction of the market, especially those with recent purchases that were constrained by net worth. Non-homeowners on the other hand do not experience an increase in net worth following a house price increase. Their total demand for housing is always falling in the price, and the price increase on impact of the income shock is less pronounced when they make up a larger share of the market.

Figure 7 shows the impulse response of prices and net housing purchases to a sequence of high income realizations in the baseline extrapolative expectations model. The initial price increase is an

Table 4: Moments of price growth rates

		Model results				
Moment	*Data*	$g = 0$	$g = 0.01$	$g = 0.02$	$g = 0.022$	$g = 0.03$
One year correlation	0.70	0.63	0.71	0.74	0.75	0.78
Three year correlation	0.08	-0.21	-0.21	-0.16	-0.15	-0.06
Five year correlation	-0.48	-0.44	-0.47	-0.46	-0.47	-0.49
Standard deviation	0.05	0.026	0.031	0.041	0.043	0.066

Notes: Price growth rates are changes in the detrended log price index. Data: Freddie Mac House Price Index deflated by CPI excluding shelter. Model: extrapolative expectations benchmark calibration with different parametrizations for the gain parameter g.

equity injection for existing homeowners whose net purchase volume increases.[13] In contrast to the rational expectations model net purchases of new homeowners give way. For them, houses are less affordable and also less profitable due to the high price-to-rent ratio. Because price growth is above trend growth the perceived profitability of housing increases over time however as m grows according to equation (6). Importantly, these dynamics are not explosive. Housing requires maintenance and an important part of the return is still the implicit rental income, so that homeowners' demand can be saturated. In addition, potential new homeowners face rising downpayment requirements. Eventually, realized price growth falls below expected price growth and m shrinks. Over a very long period of high income realizations prices hence first overshoot and then slowly decline.

To assess the importance of changes in the expected price growth rate I also plot the price path for a model with $g = 0$. In this case m is constant and equation (3) implies that house prices are a random walk, an assumption that is also routinely used in the literature.[14] Figure 8 shows that without innovations in m the overall price increase is only half of the benchmark specification and price growth seizes after three periods. Panel b) shows that the observed growth in prices after the income shock is accompanied by an increase in purchases by new homeowners. As was shown in in table 4 this pattern already generates considerable momentum in response to the calibrated income shock series.

In summary, the extrapolative expectations setup generates price dynamics that differ significantly from the rational expectations solution. It implies that a calibrated series of aggregate income shocks can generate dynamics in house prices that are much closer to the data. The in-

[13] One part of this is that old homeowners with low net worth are less likely to sell their home before reaching retirement. For some of them, the risk associated with falling prices is very significant because having to supply additional equity upon retirement would leave them with hardly any savings. The price shift reduces this risk and results in an increase in the housing stock of retirees, so that net purchases by other households can be higher than in steady state.

[14] See for example Campbell and Cocco (2015) or Berger, Guerrieri, Lorenzoni, and Vavra (2015).

Figure 7: Impulse responses with $g = 0.022$ (benchmark)

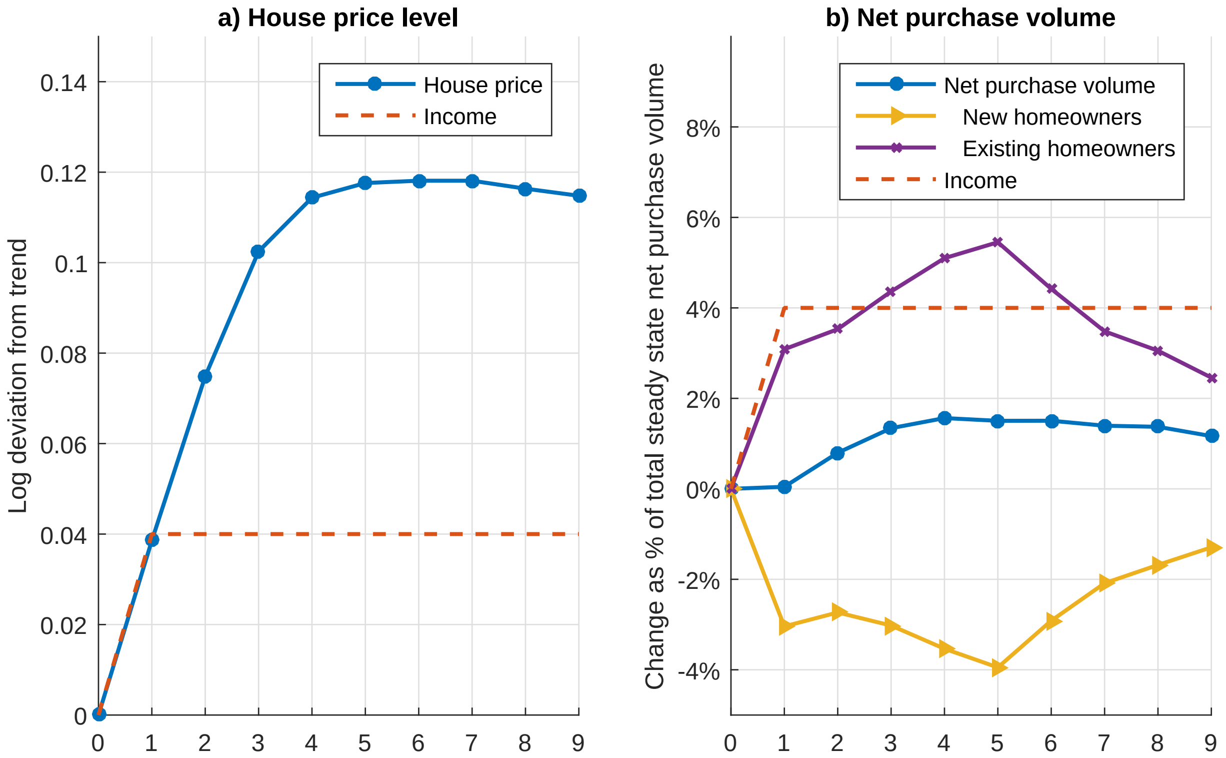

Notes: Panel a) shows deviations of log house prices and log aggregate income from trend. Panel b) shows changes in net purchase volume as a fraction of total steady state net purchase volume as in figure 5

crease in volatility is driven both by the lack of expected mean reversion that is already implied by a unit root process and by the learning dynamics that are triggered by observed price increases. These learning dynamics do not lead to an explosive path for prices because agents' demand for housing can be saturated, and eventually realized price growth fails to keep up with expectations. Nevertheless, at that turning point growth expectations are still very high and further stimulus to housing demand, such as easing credit conditions, would likely fuel the boom in prices.

6 Conclusion

I present a calibrated life cycle model and study the house price dynamics generated by purchase decisions of optimizing households. In line with survey evidence pertinent to the housing market as well as other asset markets, I consider a specification where households extrapolate observed house price growth when forming their expectations. This significantly alters the model's predictions compared to the rational expectations solution. In particular, the observed fluctuations in US per capita income over the past forty years generate price dynamics that pick up much of the fluctuations observed in the data. House price growth rates are volatile and match both the observed degree of serial correlation and mean reversion. The combination of income shocks and learning dynamics

Figure 8: Impulse responses with $g = 0$ (random walk)

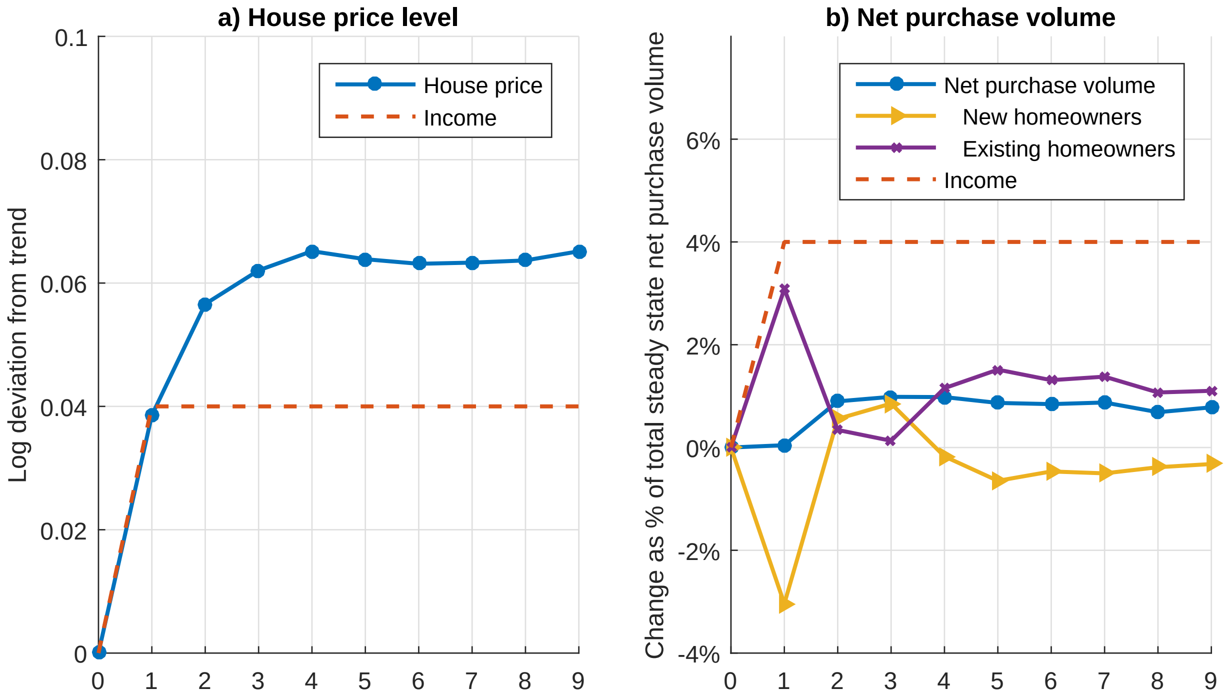

Notes: Panel a) shows deviations of log house prices and log aggregate income from trend. Panel b) shows changes in net purchase volume as a fraction of total steady state net purchase volume as in figure 5

also matches the early stage of the last housing boom and leads to high growth expectations that could have amplified the impact of further stimulus to housing demand.

The model structure would accomodate the introduction of housing production and further shocks that affect housing demand. It can also be used to study the welfare impact of belief driven price fluctuations and to assess macroprudential policies such as countercyclical caps on mortgage debt. More broadly, it would be of interest to investigate hybrid specifications with both rational and extrapolative components in expectations.

References

Adam, K., P. Kuang, and A. Marcet (2011): "House Price Booms and the Current Account," *NBER Macroeconomics Annual*, 26, 77 – 122.

Adam, K. and A. Marcet (2011): "Internal rationality, imperfect market knowledge and asset prices," *Journal of Economic Theory*, 146, 1224–1252.

Adam, K., A. Marcet, and J. Beutel (2015a): "Stock Price Booms and Expected Capital Gains," Working paper, University of Mannheim.

Adam, K., A. Marcet, S. Merkel, and J. Beutel (2015b): "Can a financial transaction tax prevent stock price booms?" *Journal of Monetary Economics, forthcoming.*

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2015): "House Prices and Consumer Spending," Working Paper 21667, National Bureau of Economic Research.

Berger, D. and J. Vavra (2015): "Consumption Dynamics During Recessions," *Econometrica*, 83, 101–154.

Bokhari, S., W. Torous, and W. Wheaton (2013): "Why did Household Mortgage Leverage Rise from the mid-1980s until the Great Recession?" Working paper, Massachusetts Institute of Technology.

Campbell, J. Y. and J. F. Cocco (2015): "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium," *Journal of Finance, forthcoming.*

Case, K. E. and R. J. Shiller (1989): "The Efficiency of the Market for Single-Family Homes," *American Economic Review*, 79, 125–37.

Case, K. E., R. J. Shiller, and A. K. Thompson (2012): "What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets," *Brookings Papers on Economic Activity*, 45, 265–315.

Chu, Y. (2014): "Credit constraints, inelastic supply, and the housing boom," *Review of Economic Dynamics*, 17, 52–69.

Corbae, D. and E. Quintin (2015): "Leverage and the Foreclosure Crisis," *Journal of Political Economy*, 123, 1–65.

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, 279–284.

Davis, M. A. and F. Ortalo-Magné (2011): "Household Expenditures, Wages, Rents," *Review of Economic Dynamics*, 14, 248–261.

Diaz, A. and M. J. Luengo-Prado (2008): "On the User Cost and Homeownership," *Review of Economic Dynamics*, 11, 584–613.

——— (2010): "The Wealth Distribution With Durable Goods," *International Economic Review*, 51, 143–170.

Favilukis, J., S. C. Ludvigson, and S. Van Nieuwerburgh (2015): "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium," *Journal of Political Economy, forthcoming.*

Gelain, P., K. J. Lansing, and C. Mendicino (2013): "House Prices, Credit Growth, and Excess Volatility: Implications for Monetary and Macroprudential Policy," *International Journal of Central Banking*, 9, 219–276.

Glaeser, E. L. and J. Gyourko (2007): "Arbitrage in Housing Markets," Working Paper 13704, National Bureau of Economic Research.

Glaeser, E. L., J. Gyourko, E. Morales, and C. G. Nathanson (2014): "Housing dynamics: An urban approach," *Journal of Urban Economics*, 81, 45–56.

Glaeser, E. L. and C. G. Nathanson (2015): "An Extrapolative Model of House Price Dynamics," NBER Working Papers 21037, National Bureau of Economic Research, Inc.

Greenspan, A. (2004): "The mortgage market and consumer debt," Remarks by Chairman Alan Greenspan at the Americas Community Bankers Annual Convention, Washington, D.C.

Head, A., H. Lloyd-Ellis, and H. Sun (2014): "Search, Liquidity, and the Dynamics of House Prices and Construction," *American Economic Review*, 104, 1172–1210.

Hedlund, A. (2015a): "Illiquidity and its Discontents: Trading Delays and Foreclosures in the Housing Market," Working Papers 1417, Department of Economics, University of Missouri.

——— (2015b): "The Cyclical Dynamics of Illiquid Housing, Debt, and Foreclosures," Working Papers 1416, Department of Economics, University of Missouri.

Iacoviello, M. and S. Neri (2010): "Housing Market Spillovers: Evidence from an Estimated DSGE Model," *American Economic Journal: Macroeconomics*, 2, 125–64.

Iacoviello, M. and M. Pavan (2013): "Housing and Debt over the Life Cycle and over the Business Cycle," *Journal of Monetary Economics*, 60, 221 – 238.

Justiniano, A., G. Primiceri, and A. Tambalotti (2015): "Household leveraging and deleveraging," *Review of Economic Dynamics*, 18, 3–20.

Kaplan, G., G. L. Violante, and K. Mitman (2015): "Consumption and House Prices in the Great Recession: Model Meets Evidence," Manuscript.

Krusell, P. and A. A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896.

Landvoigt, T., M. Piazzesi, and M. Schneider (2015): "The Housing Market(s) of San Diego," *American Economic Review*, 105, 1371–1407.

Linneman, P. (1986): "An empirical test of the efficiency of the housing market," *Journal of Urban Economics*, 20, 140–154.

Ortalo-Magné, F. and S. Rady (2006): "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints ," *Review of Economic Studies*, 73, 459–485.

Paciorek, A. (2013): "Supply constraints and housing market dynamics," *Journal of Urban Economics*, 77, 11 – 26.

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

Quan, D. and J. Quigley (1991): "Price formation and the appraisal function in real estate markets," *The Journal of Real Estate Finance and Economics*, 4, 127–146.

Saiz, A. (2010): "The Geographic Determinants of Housing Supply," *The Quarterly Journal of Economics*, 125, 1253–1296.

Storesletten, K., C. I. Telmer, and A. Yaron (2004): "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112, 695–717.

A Appendix 1: Data Definitions

A.1 Aggregate Data

Consumption good prices: I use the monthly values for the "Consumer Price Index for All Urban Consumers: All items less shelter" published by the US Department of Labor (FRED Series ID CUUR0000SA0L2). Since the consumption good is the numeraire in my model, I deflate house prices as well as income data with this series. Since I focus on annual values, none of the series I use are seasonally adjusted.

House prices: I use the monthly values from the national Freddie Mac House Price Index available at http://www.freddiemac.com/finance/fmhpi/archive.html which start in January 1975. This is a commonly used repeat sales index covering single family homes with a mortgage that has been purchased by the government sponsored enterprises Freddie Mac or Fannie Mae. I deflate the series by the consumption price index described above, normalize so that December 2000 equals 100, and take logs. Annual values used to calculate the growth rate statistics are the average value for each year. Figure 1 plots a quarterly series of twelve month rolling averages.

Stock prices: I use the monthly closing values of the series "S&P 500 Composite Price Index (w/GFD extension)" from https://www.globalfinancialdata.com/. I deflate the series by the consumption price index described above, normalize so that December 2000 equals 100, and take logs. Annual values used to calculate the growth rate statistics are the average value for each year.

Income data: I use annual values from NIPA table 2.1 published by the U.S. Bureau of Economic Analysis and deflate by the average consumption price index for the year. For my preferred specification, I add "Compensation of employees" (FRED Series ID W209RC1A027NBEA), "Proprietors' income with inventory valuation and capital consumption adjustments" (FRED Series ID A041RC1A027NBEA), and "Personal current transfer receipts" (FRED Series ID A577RC1A027NBEA), thus excluding capital and (imputed) rental income. To calculate per capita values I use the ratio between the reported values for personal income (FRED Series ID A065RC1A027NBEA) and personal income per capita (FRED Series ID A792RC0A052NBEA). I take logs and calculate a linear trend for the log series based on the time window 1965-2014. The correlation between the detrended income level and price growth rates is not affected by choosing a shorter time window for detrending. It is at least 0.36 for different specifications of the income series I employ.

Interest rates and depreciation: The savings rate I use is the "10-Year Treasury Constant Maturity Rate" reported by the Federal Reserve (FRED Series ID DGS10) which I adjust by comsumer price index inflation. For the borrowing rate I use the "30-Year Conventional Mortgage Rate" reported by the Federal Reserve (FRED Series ID MORTG). I calculate depreciation as the ratio of "Current-Cost Depreciation of Fixed Assets: Residential" (FRED Series ID M1R53101ES000) and the beginning of year value of "Current-Cost Net Stock of Fixed Assets: Residential" (K1R53101ES000), both of which are provided by the US Bureau of Economic Analysis.

A.2 PSID Data

The PSID is a long established, nationally representative panel study of US households. After 1997, the interviews have been conducted at biannual frequency and all contain details on wealth, housing, and debt at the family level. The dataset provides family longitudinal weights, which I use throughout. In general, families are linked between periods by having the same person classified as "Head" in all interviews. Given the formal criteria by which the role is assigned, in most instances this will allow to make the connection between interviews for the same household. The interview data also indicates composition changes within the family. For the analysis of movers, I require for existing households that head and spouse remained the same between interviews, so that the link via the family head is comprehensive.

To calculate the homeownership rate, I use the answer to the corresponding question A19 from all 1999 interviews (where the head falls in the specified age range). Net worth (WEALTH2) is imputed by the PSID staff in case of missing data, so that I use all interviews as well. I identify movers as households who report having moved in response to question A42. I require that the calendar month and year are provided and count moves that have occurred in the twelve months prior to the interview month. I also require that the value of the home is reported in order to calculate overall purchase volumes for each group. To classify whether a household newly became homeowners, I first look at the same household's ownership status in the 1997 interview. In such cases I require that head and spouse did not change, because otherwise it is not clear what kind of move is reported. For some households there is no previous interview. I cannot discard this group, since so called splitoffs from existing families (e.g. children moving out) are a key feature of maintaining the panel and potentially also a relevant source of new homeowners. Therefore, I include in the analysis all households who either have a previous interview and where the head and spouse have not changed, or who are splitoffs according to the composition change indicator.[15] All splitoffs who report moving within the twelve months prior to the interview and now own a home are counted as new homeowners. Together, they make up about twelve percent of all new homeowners. A large majority of them indeed reports a recent mortgage origination date, so that further refining the approach would not yield materially different results. There may be a downward bias in the fraction of movers among all homeowners, because I do not count moves where the date has not been reported, but for the ratio between rent-to-own and own-to-own transitions which I calibrate to this should not be relevant. I check that the ratio is also not materially affected by the other selection criteria.

B Appendix 2: Computation

B.1 Rational Expectations

Given the grids for housing and income realizations, I solve the households' problem on a fine grid for assets a and the house price p. The asset grid consists of 1000 points where the distribution of the 1000 points is chosen differently for each of the possible house sizes.[16] It is denser where the utility of a typical owner of the particular house size has more curvature and in particular around the borrowing limit. The price grid is made up of 71 points where the range of the 71 points is chosen differently for each level of the aggregate income shock. This means that there are 175 distinct points covering the relevant range of prices. This implies a step size of only about 0.05 percent between grid points. The advantage of the setup is that the large number of grid points helps to address the non-convexities induced by the linear adjustment cost. However, it also means that I restrict the savings policy to values on the asset grid.

For each grid point, I derive the housing policy based on comparing the value of saving optimally given the current house and the value of selling the current house. The value of selling the house is the same as the value of beginning the period as a renter with the level of cash that would result after selling the house.[17] This is in turn given by the maximum of either remaining a renter and saving optimally or buying a house of any size. The value of buying a house can be derived from the problem of a household who already enters the period with that house and the level of cash that

[15] There is no explicit splitoff indicator for 1999, and I use the condition that family composition change is either 5 or 6, codes which are primarily used for splitoffs.

[16] The grid for renters has 1300 points.

[17] This value may not be a point on the asset grid for renters. I randomly assign the household to the left or the right point based on the relative distance and calculate the expected value of selling accordingly.

remains after the purchase. I ensure that new buyers always at least meet the leverage constraint however, and, because the level of cash after the purchase will not exactly be on the grid, I compare the policies for the left and right grid point. Overall, the approach means that the savings problem only has to be solved for each (a, h) combination, but not for each (a, h, h') combination separately, which is a large computational advantage.

I simulate the economy with 12,000 households in each cohort, i.e. 540,000 households living at each point in time. All households start with no assets or housing and with a draw from the stationary distribution for z. The initial draw and all subsequent draws are the same for each cohort. Each simulation is initialized by calculating choices for one cohort from birth to retirement given the specified value for $\bar{p}$ and normal income realizations. The state of that cohort in each year of their life makes up the initial economy. The same distribution will reemerge after a long sequence of normal income realizations. The aggregate housing stock H is the sum over the cohort's housing choices in each period.

Starting from the initial distribution that would result from a long sequence of normal income realizations, I calculate market clearing prices in response to a sequence of aggregate income shocks. The tolerance for market clearing is h_{max}, since any household will either buy a house around a given size or not buy it, but only buying a fraction of the desired house would never be optimal. While the price grid is very fine, no point on the grid will imply exact market clearing. Instead, I place all households randomly on the interval between the two points with the lowest positive and highest negative excess demand. I then locate a threshold such that all households to the left are assigned to the left grid point and all households on the right are assigned to the right grid point and the market clears. The position of the threshold is interpreted as the exact market clearing price for the period. However, all households to the left act as if the price was slightly lower and all households to right act as if the price was slightly higher. Given the that the points are only 0.05 percent apart this seems acceptable.[18] It also matches the way expectations are formed: for any expected price households linearly interpolate between the values of the grid points to the left and to the right.

To determine the forecasting rule (2), I start from an initial guess, solve for all policy functions given these expectations and use the policy functions to simulate the economy for 2,600 periods. I discard the initial 600 periods, and for the simulated data from the remaining 2000 periods I run

[18] The alternative would be to interpolate all value functions on a much coarser price grid.

the following regressions

$$\ln \grave{p}_{t+1} - \ln \bar{p} = \hat{\alpha}_1(Z_t, Z_{t+1}) + \hat{\alpha}_2(Z_t, Z_{t+1})(\ln \grave{p}_t - \ln \bar{p}).$$

I use the estimated coefficients $\hat{\alpha}$ to update the guess for α and resolve and resimulate the model to obtain new coefficient estimates. After a number of iterations the estimated forecasting rules fit the simulation results well with an R^2 of at least 0.997. The R^2 and the aggregate statistics remain virtually constant across iterations, and I stop the iteration when the difference between estimated and used coefficients is below 0.001. The table below reports the coefficient estimates for the three different scenarios discussed in the text.

Table 5: forecasting rule coefficients

	Model specifications		
Coefficient	$\varphi = 0.06$ $\theta = 0.8$	$\varphi = 0$ $\theta = 0.8$	$\varphi = 0$ $\theta = 1$
$\alpha_1(Z_{low}, Z_{low})$	-0.0092	-0.0062	-0.0028
$\alpha_2(Z_{low}, Z_{low})$	0.7360	0.7903	0.8968
$\alpha_1(Z_{low}, Z_{norm})$	0.0108	0.0134	0.0075
$\alpha_2(Z_{low}, Z_{norm})$	0.6297	0.9046	0.7987
$\alpha_1(Z_{norm}, Z_{low})$	-0.0224	-0.0179	-0.0105
$\alpha_2(Z_{norm}, Z_{low})$	0.8057	0.7823	1.0419
$\alpha_1(Z_{norm}, Z_{norm})$	0.0001	0.0001	-0.0003
$\alpha_2(Z_{norm}, Z_{norm})$	0.7252	0.8441	0.9173
$\alpha_1(Z_{norm}, Z_{high})$	0.0174	0.0138	0.0071
$\alpha_2(Z_{norm}, Z_{high})$	0.6783	0.8155	0.5748
$\alpha_1(Z_{high}, Z_{norm})$	-0.0112	-0.0101	-0.0098
$\alpha_2(Z_{high}, Z_{norm})$	0.8048	0.8903	1.3304
$\alpha_1(Z_{high}, Z_{high})$	0.0074	0.0048	0.0047
$\alpha_2(Z_{high}, Z_{high})$	0.7160	0.8030	0.6616
$\frac{\alpha_1(Z_{low}, Z_{low})}{1-\alpha_2(Z_{low}, Z_{low})}$	-0.0349	-0.0295	-0.0271
$\frac{\alpha_1(Z_{high}, Z_{high})}{1-\alpha_2(Z_{high}, Z_{high})}$	0.0259	0.0245	0.0140

Notes: Coefficient values for forecasting rule (2) for specifications of the model with different degrees of adjustment frictions as given by the leverage ratio θ and the transaction cost φ. Coefficients are indexed by the current and future aggregate state.

One can see that the coefficient estimate for $\alpha_1(Z_{norm}, Z_{norm})$ is very close to zero as should be expected. To interpret the other coefficients, I report the log difference between the fixed point of (2) and $\bar{p}$ for the other states which is given by $\frac{\alpha_1(Z_s, Z_s)}{1-\alpha_2(Z_s, Z_s)}$. For $s = high$ prices are expected to rise till they are that much above $\bar{p}$ and for $s = low$ they fall until they reach that level. Consistent with the reported volatilities, the model with adjustment frictions generates the largest absolute values for the fixed points. The coefficients $\alpha_1(Z_{norm}, Z_{high})$ and $\alpha_1(Z_{norm}, Z_{low})$ indicate how much of the distance to the fixed points is covered on impact of the income shock. For the model with adjustment frictions this is a slightly larger fraction, which contributes to the lower serial correlation in growth rates.

To assess the quality of the solution, I also compare the result of simply iterating on the forecast rule given a sequence of income shocks to the model solution for the same sequence of shocks. Figure x shows that both series remain close the entire time. The largest deviation is 0.4 percent. The data is for the model with both frictions. Without lumpy adjustment, forecasting errors are somewhat smaller.

Figure 9: Simulated house price series

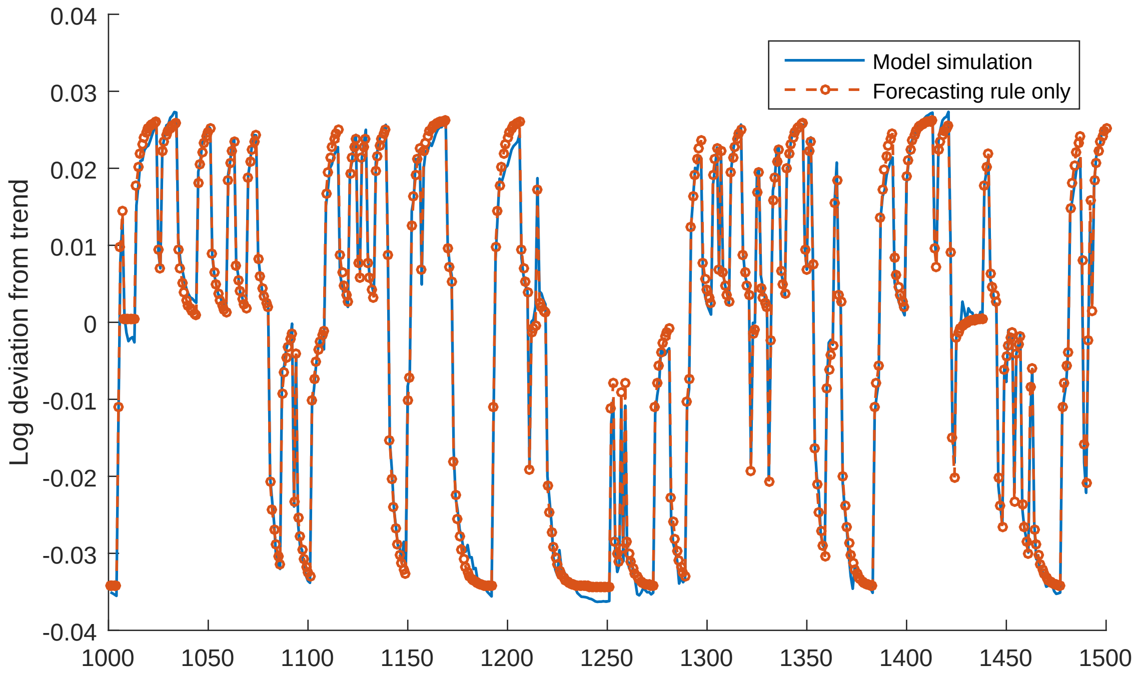

Comparison of market clearing prices in a long model simulation to a prediction that just uses the sequence of aggregate shocks and iterates on the forecasting rule.

B.2 Extrapolative Expectations

As shown in equation (7), extrapolative price expectations depend on an additional state m_t. Due to the additional state and because a larger price range has to be accomodated, I use smaller grids and interpolate on value and policy functions. The asset grid now has 60 points that again have values specific to the different house sizes. The price grid has 19 points and the grid for m has 9 points. I treat these three dimensions as continuous and use linear interpolation in all steps. I continue to use discrete values for h and y based on the the same grids as before. The logic to for calculating the grid point values for value and policy functions is then essentially the same.

For $\ln \tilde{\epsilon}_{t+1}$ which expresses agents' uncertainty regarding future prices I use five Gaussian quadrature nodes. Given a standard deviation of four percentage points, this means that agents always explicitly take a 12 percent price drop for the next period into account. This is higher than any observed price drops, and because there is a positive probability of having to move next period agents always choose sustainable debt positions.

Importantly, the choice whether to buy or change owner occupied housing remains a discrete choice which cannot be interpolated. In simulations, I therefore explicitly compare the values of remaining a renter and buying a house. I consider all house sizes implied by the policy function for adjacent grid points. For existing homeowners, I compare the value of keeping the house to the value of selling the house. Sellers can then be treated the same as initial renters to determine whether they buy a new house.

In a version of the model where prices are constant, I can directly compare the results for the different algorithms I use for the rational and extrapolative expectations framework. In terms of aggregate outcomes and individual histories the results are almost identical. The policies derived from the entirely grid based approach seem to be slightly more accurate, because realized utility is higher. However, the difference is only about 0.01 percent of per period consumption.