NBER WORKING PAPER SERIES

EXPECTATIONS DURING THE U.S. HOUSING BOOM: INFERRING BELIEFS FROM ACTIONS

Itzhak Ben-David
Pascal Towbin
Sebastian Weber

Working Paper 25702
http://www.nber.org/papers/w25702

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2019, Revised November 2019

We are grateful to Olivier Blanchard, Alex Chinco, Hamid Faruqee, Andreas Fuster, Maja Ganarin, Ayhan Kose, Kenneth Kuttner, Xiaoji Li, Prakash Loungani, John Muellbauer, Chris Otrok, Antoinette Schoar, Emil Stavrev, Martin Straub, Susan Watcher, and Paul Willen as well as participants at the American Economic Association Conference, the Annual Paul Woolley Conference, the Bundesbank-IMF-DFG Conference, the SSES Congress, and the International Monetary Fund and SNB Research seminars for very helpful comments. This draft replaces an earlier version entitled "Price Expectations and the U.S. Housing Boom." Ben-David is with The Ohio State University and the National Bureau of Economic Research; Towbin is with the Swiss National Bank; and Weber is with the International Monetary Fund. The views expressed herein are those of the authors and do not necessarily reflect the views of International Monetary Fund, the Swiss National Bank, or the National Bureau of Economic Research.

Expectations During the U.S. Housing Boom: Inferring Beliefs from Actions
Itzhak Ben-David, Pascal Towbin, and Sebastian Weber
NBER Working Paper No. 25702
March 2019, Revised November 2019
JEL No. E32,E71,G12,R31

ABSTRACT

We infer the role of price expectations in forming the U.S. housing boom in the early- 2000s from examining housing inventories. We use a reduced form model to show that agents invest in vacant homes when they anticipate prices will increase. Empirically, vacancy can discriminate between price movements related to shocks to demand for housing services (low vacancy) and shocks to expectations (high vacancy). Using a structural vector autoregression with sign restrictions, we show that expectation shocks were a prime factor explaining the boom particularly in the Sand States, which experienced unprecedented booms.

Itzhak Ben-David
The Ohio State University
Fisher College of Business
606A Fisher Hall
Columbus, OH 43210-1144
and NBER
ben-david.1@osu.edu

Pascal Towbin
Swiss National Bank
Bundesplatz 1
CH-3003 Berne
Switzerland
Pascal.Towbin@snb.ch

Sebastian Weber
European Department International Monetary Fund
700 19th Street, NW
Washington DC 20431
sweber@imf.org

1 Introduction

Many economists agree that expectations played some role in forming the residential housing boom in the early 2000s. They disagree, however, on the degree of the contribution of expectations to the boom. Several studies propose that expectations were a prime factor in fueling the run-up in prices (Shiller, 2007; Foote, Gerardi, and Willen, 2012; Case, Shiller, and Thompson, 2015; Adelino, Schoar, and Severino, 2016; Kaplan, Mitman, and Violante, 2019). The macro-level empirical evidence supporting this assertion, however, often relies on a residual, that is, researchers attribute the unexplained variation in house prices to expectations (Dokko, Doyle, Kiley, Kim, Sherlund, Sim, and Heuvel, 2011; Glaeser, Gottlieb, and Gyourko, 2012). At the micro-level, researchers can identify the relation between speculative investment activity (e.g., out-of-town investors) and prices using granular data (Bayer, Geissler, and Roberts, 2011; Chinco and Mayer, 2015; DeFusco, Nathanson, and Zwick, 2017; Gao, Sockin, and Xiong, 2017; Griffin, Kruger, and Maturana, 2018; Bailey, Cao, Kuchler, and Stroebel, 2018), but it is unclear how this generalizes to the entirety of the boom in the U.S. Understanding the drivers of the boom in the residential market and its consequent bust is important from both an academic and a policy-making perspective. To date, the role of expectations in forming the housing boom is still unknown.

In this study, we infer the role of expectations from disparities between the demand for and the supply of housing in conjunction with price movements. The intuition is simple: When supply has a convex cost, investors gather vacant homes if they anticipate that prices will increase, often resulting in an increase in new construction (Topel and Rosen, 1988). We use a stylized model to formalize this argument. We then incorporate the model's predicted sign restrictions into a structural vector-autoregression estimation to identify expectation shocks. These tools allow us to estimate the contribution of expectation shocks to the housing price boom separately from shocks to factors such as demand for housing services, housing supply, and mortgage rate. To our best knowledge, no other study has attempted to harness vacancy to infer beliefs in a systematic manner to assess the role of price expectations in forming prices around the boom in the U.S. housing market in the early 2000s.

The idea that high inventories indicate that individuals anticipate that future prices will increase is not new to economists. Originally, this insight was implemented in the energy market, where investors amass large quantities of oil in anticipation of price increases (Kilian and Murphy, 2014; Juvenal and Petrella, 2015; Knittel and Pindyck, 2016). In the real estate market, vacancies stem from the search and matching frictions and the limited amount of housing that can be constructed in a given period at reasonable costs. Researchers have noticed that during boom times individuals often hold vacant properties (see, e.g., Malpezzi and Wachter, 2005; Mayer, 2011; Glaeser, Huang, Ma, and Shleifer, 2017) and that construction activity accelerates (Glaeser, Gyourko, and Saiz, 2008; Haughwout, Peach, Sporn, and Tracy, 2012; Glaeser and Nathanson, 2017; Glaeser et al., 2017). Topel and Rosen (1988) explain the key mechanism that causes vacancies to increase during boom times. Due to convex construction costs in a given period (Rosenthal, 1999), construction activity is front-loaded and occurs prior to the realization of the actual demand for housing services: "an anticipated transitory increase in future demand causes bubble-like price and investment responses." (Topel and Rosen, 1988, page 727) In practice, investors purchase existing homes from owners and homebuilders (Chinco and Mayer, 2015; DeFusco et al., 2017). The result is that construction activity increases to fill the gap in quantity due to the increased participation of investors in the housing market.[1] To illustrate the effect, in Figure 1 we present the annual change in the number of households formed, together with the flow of new homes sold in the U.S.. The figure shows that while the rate of household formation moves within a very narrow range over time, sales of new homes move in a cyclical manner, rising noticeably in the early 2000s. Of course, a large fraction of new homes replace existing dilapidated homes, but some homes remain vacant in the hands of investors.[2]

[1]Also move-up buyers may temporarily increase their participation in the market. Given expectations of rising prices, they may prefer to first buy a home before selling their current home, thereby owning two homes for a period. This drives up prices and triggers investment activity without a change in the demand for housing services.

[2]See Haughwout et al. (2012) for a discussion regarding the supply side during the boom.

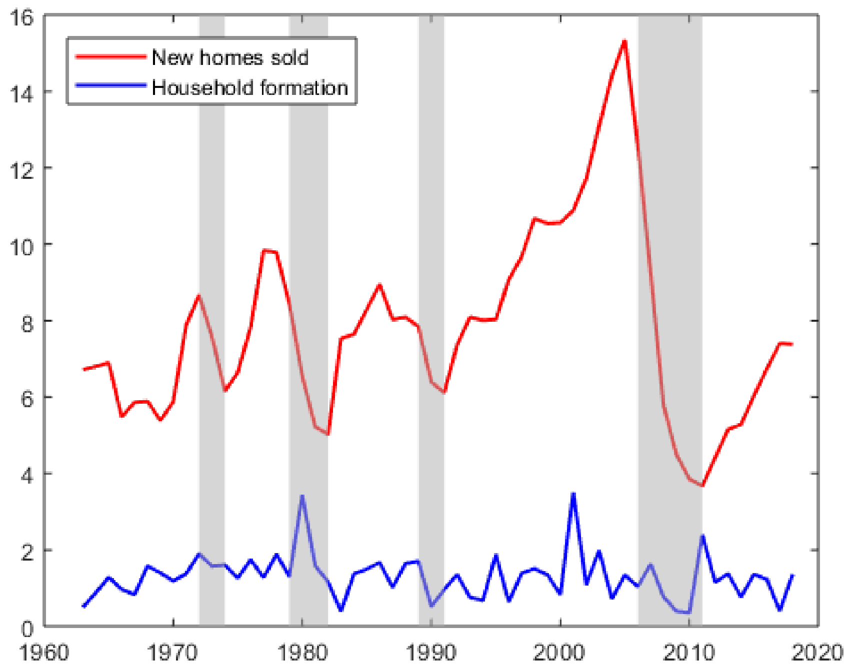

Figure 1. New Homes Sold and the Year-on-Year Change in Total Households

The figure shows the number of new single-family houses sold in the U.S. (series HSN1F) as well as the change in the number of households in the U.S. (series TTLHH) in millions. The units on the y-axis represent cumulative log differences. Both series are available on St. Louis Federal Reserve Bank's website (`https://fred.stlouisfed.org/`) and from the U.S. Census. The gray shaded areas reflect episodes of real housing price declines.

We flesh out the intuition that a high vacancy rate may be an indication of high price expectations in a stylized model that explores the determinants of aggregate housing prices by considering shocks to four distinct variables: demand for housing services, housing supply, mortgage rate, and expectations about future prices. We assume that current prices are determined by rental rates, current mortgage rates, and expectations about future prices. We make standard assumptions, such as that the demand for housing services slopes downward, the housing supply slopes upward with convex adjustment costs, and the credit supply slopes upward. In addition, we assume that suppliers of housing have less bargaining power over rental rates when the vacancy rate is high.

The main purpose of the model is to generate sign restrictions that describe the responses of the variables to the shocks. For example, the model predicts that a positive shock to demand for housing services will increase the housing supply, house prices, and the mortgage rate, and will decrease housing vacancy. Our interest lies in the effect of a shock to price expectations on the fundamental variables. The model predicts that a positive shock to price expectations will increase the housing supply, housing vacancies, house prices, and the mortgage rate. Thus, by observing the dynamics of the fundamental variables over time, we can back out the extent to which various shocks—and of particular interest, the expectations shock—could have shaped them.

Once derived, we use the sign restrictions to estimate the effects of expectations on housing prices. First, we estimate a Bayesian vector autoregression (BVAR) using observable variables for the entire country as well as for each state individually. Next, we simulate shocks to the different variables for each point in time (demand for housing services, housing supply, mortgage rate, expectations) and explore which set of shocks could account for the forecasting errors, i.e., deviations from the model's predictions. The idea behind the approach is that forecast errors occur because unexpected shocks pushed an observed variable away from its natural course. We collect only combinations of shocks that are consistent with the sign restrictions for the respective variables, as derived from the model. Overall, this analysis estimates the contribution of each shock to housing prices at any given point in time.

Our results show that expectation shocks were important both at the national level and across states. At the national level, we find that price expectation shocks were the second most important determinant of real estate prices during the boom, explaining about 23% of the magnitude of the boom, slightly less important than mortgage rates, which accounted for about 28% of the boom's magnitude.[3] In the cross-section of states, we find a high degree of heterogeneity. Expectation shocks particularly drove up prices in the Sand States (California, Florida, Nevada, and Arizona), which experienced unprecedented price booms.

To further validate our estimations, we contrast the size of the bust with the drivers of the boom. Among all shocks that hit the various states during the boom period, the one that explains the bust the best is the expectation shock. In other words, states that experienced the steepest expectation shocks during the boom also experienced the deepest subsequent bust. For the other identified shocks that we explore—demand for housing services, supply, and mortgage rates—the association is weaker. This finding is in line with the models of Burnside, Eichenbaum, and Rebelo (2016) and Chinco (2018), who propose that booms are fueled by both fundamental demand for housing services and expectations, and that busts occur when agents' expectations about future prices dissipate.

It is important to distinguish our goal of establishing the contribution of price expectation *shocks* from previous work that examined how expectations develop endogenously. For example, Glaeser and Nathanson (2017), Bordalo, Gennaioli, Ma, and Shleifer (2017), and Bordalo, Gennaioli, and Shleifer (2018) provide frameworks in which irrational individuals develop biased expectations based on observing past performance. Piazzesi and Schneider (2009), Adelino, Schoar, and Severino (2018b), De Stefani (2017), and Armona, Fuster, and Zafar (2019) use survey data to analyze real estate expectations and find that individuals form expectations based on extrapolations of past price changes. Similarly, Adam, Kuang, and Marcet (2012), Burnside et al. (2016), DeFusco et al. (2017), and Chinco (2018) explore

[3]This result is consistent with Levitin, Lin, and Wachter (2018), who show that risk premiums embedded in originated mortgages declined as housing prices increased.

mechanisms through which expectations develop endogenously, e.g., through social feedback among agents. In contrast to these studies, we are interested in the innovations to expectations, i.e., the part of expectations that *does not* arise endogenously, but rather is orthogonal to lagged observable variables. In our framework, endogenous expectations that develop in response to observable variables (e.g., past prices) would be absorbed in the autoregressive process.

Overall, our analysis shows that shocks to expectations about future prices were major contributors to the price evolution around the boom, both nationally and for some of the states. This result is in contrast to survey-based expectations (Piazzesi and Schneider, 2009) and estimated structural models (Landvoigt, 2017) that show that expectations had little impact during the boom. We speculate that the source of the difference between the two expectation measures is the population in question. Recall that we identify the expectations from vacancies arising in conjunction with changes in housing prices. Therefore, the marginal buyers are those who hold or dispose of the vacant homes, i.e., the investors. In contrast, surveys estimate *average* expectations. Most researchers in the field use the Michigan Survey of Consumers (Piazzesi and Schneider, 2009; Lambertini, Mendicino, and Punzi, 2013; Cox and Ludvigson, 2018), and some of the respondents may not participate actively in the market. Conversely, our method measures the expectations of the *marginal* buyers who actually transacted. Those are the individuals whose demand and expectations set prices.

The rest of the paper is organized as follows. Section 2 introduces a simple model of a housing market that generates the sign restrictions. Sections 3 and 4 describe the empirical method and the data sets that are used in the study. Sections 5 and 6 describe the results of the analysis at the national and state level, respectively. Section 7 compares the model-based expectation measures to survey-based measures. Section 8 provides concluding remarks.

2 A Stylized Model

Our identification approach relies on five basic assumptions about the housing market. We embed these assumptions in a small reduced-form model. The model is kept deliberately simple to make the underlying mechanisms transparent. Each assumption is backed up by a large theoretical and empirical literature.

All endogenous variables are presented in capital letters and expressed in logs. Coefficients in small caps are restricted to be positive, with subscripts reflecting the respective variables. ε_X denotes exogenous drivers of the endogenous variable X that are shocked.[4]

1. **Prices are forward-looking and depend on the mortgage rate and expected future prices.**

$$P = p_R R - p_I I + \varepsilon_{P^e} \tag{1}$$

This assumption can be motivated by the present discount value relationship between current house prices (P) and future house prices (ε_{P^e}) as well as rental rates (R) and mortgage rates (I) (see for instance Poterba, 1984). A lower mortgage rate increases the present discount value of a given future price.[5] Importantly, to keep the problem tractable, we ignore the effects of depreciation and taxes.

At the same time, the above equation represents an arbitrage condition between renting and selling. Housing investors are indifferent between letting a house or selling it. Accordingly, we only consider the overall housing market and do not model the choice between selling and letting.[6] In that sense, the rental rate should be understood as the price to consume housing services in a given period.

[4]The expected value of ε_X is not necessarily zero. Alternatively, the equations can be interpreted as a deviation from a steady state with $E(\varepsilon_X) = 0$.

[5]This result is obtained if we consider the current house price P_t to be the expected present value of future rents R_{t+1}, discounted at rate I_{t+1}:

$$P_t = E_t \frac{R_{t+1} + P_{t+1}}{1 + I_{t+1}}.$$

[6]Vice versa, occupants of houses are indifferent between renting or buying.

Price expectations (ε_{P^e}) are treated as exogenous. The definition of an expectation shock does not specify whether expectations are "realistic" or "unrealistic." It is consistent both with a realistic response to new information about future housing fundamentals[7] or with unrealistic expectations about future house prices as emphasized by Case and Shiller (2003). A combination of the two is an overreaction to a signal ("kernel of truth") (Bordalo et al., 2017).

2. **Demand for housing services is downward sloping with respect to rents and depends negatively on the mortgage rate.**

$$D = -d_R R - d_I I + \varepsilon_D \tag{2}$$

Demand for housing services is understood as demand for the consumption of housing services. It is distinct from the demand to buy a house, since the buyer needs to live in the house to consume housing services. Most importantly, it excludes speculative demand for owning a house without actually living in it. Demand for housing services decreases with the rental rate (R).

Furthermore, the demand for housing services depends negatively on the mortgage rate. This assumption is motivated by borrowing constraints in the form of debt service coverage ratios (DSCR) that are relieved through lower mortgage rates (Adelino, Schoar, and Severino, 2018a; Greenwald, 2018; Levitin et al., 2018).

ε_D captures exogenous drivers of demand for housing services. Possible reasons for a positive shock to the demand for housing services are exogenous increases in the population, higher personal incomes, or shifts in tastes (Green and Hendershott, 1996; Zabel, 2004; Green and Lee, 2016).

[7] An example would be if housing investment activity today is triggered by the expectation that a large company (e.g., Amazon Inc.) moves its headquarters to a new location and brings additional future demand for housing services.

In addition, ε_D captures shocks to the loan-to-value (LTV) ratio or credit standards as emphasized, for example, in Mian and Sufi (2009), Duca, Muellbauer, and Murphy (2011), Favilukis, Kohn, Ludvigson, and Van Nieuwerburgh (2012), Favilukis, Ludvigson, and Van Nieuwerburgh (2017), and Greenwald (2018). Looser standards would shift the demand curve up. Holding interest rates constant, the demand for housing services increases with the availability of credit.[8] In an extension, we will also distinguish between shocks to the demand for housing services that are associated with a loosening of LTV standards and shocks that are associated with a tightening of LTV.

3. **Housing supply is upward sloping with respect to house prices and subject to convex adjustment costs.**

$$S = s_P P + \varepsilon_S \tag{3}$$

Supply (S) is increasing in the price of housing. The upward-sloping supply can be a result of various factors identified in the literature, including zoning regulations, land limitations, and rising construction costs (Glaeser et al., 2008; Huang and Tang, 2012). We consider house flippers to be part of the supply in our model; in general, they purchase derelict homes, improve their quality through investments, and sell them. Thus, flipping activity increases the supply of housing. A similar argument can be made for housing improvements (Choi, Hong, and Scheinkman, 2014).

ε_S represents exogenous determinants of supply. Negative supply shocks may arise from cost increases in the construction sector and changes in the regulatory environment

[8] The argument here is that looser credit standards would increase overall demand for housing services, thereby pushing up both rents and prices. Greenwald and Guren (2019) present a model in which rental and owner-occupied markets are allowed to be segmented. In segmented markets, looser credit standards can then encourage a shift from renting to owning, which will push up the prices of owner-occupied homes, but not rents. We acknowledge that our approach does not directly consider this possibility . However, Kaplan et al. (2019) provide evidence against strong segmentation, citing a large number owner-occupied properties being converted to rental properties. Furthermore, Begley, Loewenstein, and Willen (2019) show that during the boom the price of rental properties increased relative to owner-occupied properties, which goes against this hypothesis that loosened credit standards boosted the price of owner-occupied property only.

that reduce the provision of land (e.g., zoning restrictions) or make it more costly to construct on existing land (Gyourko and Saiz, 2004).

Note that the demand for housing services is decreasing with respect to the rental rate (the price for the current flow of rental services), whereas housing supply is increasing with respect to the price of housing (the price for current and discounted expected future flows of rental services, see Assumption 1). This is a simple way to represent the idea that the housing supply is more responsive than the demand for housing services to future developments. This higher sensitivity is a direct consequence of convex adjustment costs. Specifically, adjusting the housing supply in a given period leads to adjustment costs that increase with the size of the adjustment, creating an incentive to spread it over several periods and adjust in a forward-looking manner. Adjusting demand for housing services, instead, is costless, and the consumption of housing services can be re-optimized each period. Hence, convex adjustment costs motivate that the housing supply reacts both to current rents and to discounted future prices, whereas demand only responds to current rents.[9]

4. **Suppliers of housing have less bargaining power when vacancies are high.**

$$V = S - D = -v_r R \tag{4}$$

Markets do not clear fully, and there is a positive gap between the housing supply and the demand for housing services (Wheaton, 1990; Leung and Tse, 2012; Head, Lloyd-Ellis, and Sun, 2014). The vacancy rate is given by the (log) difference between the supply of and the demand for housing services. The rental rate declines as the stock

[9] We also considered an extended model in which switching houses is costly and occupants take the future into account. For example, one could imagine that young people move out earlier and household formation increases when prices are expected to increase. In such a model, demand for housing services depends negatively on current rents and positively on house prices. Our results still hold, as long as the demand for housing services' response to house prices is sufficiently small relative to the supply response, or put differently, the cost of switching houses should be sufficiently small compared to the cost of adjusting supply in a given period.

of vacant houses increases because high vacancy rates give housing suppliers relatively little bargaining power compared to a tight housing market with low vacancy rates. Rosen and Smith (1983) develop a model for a negative relationship between rents and vacancies. Recent empirical evidence is provided for example in Gabriel and Nothaft (2001).

5. **Credit supply is upward sloping.**

$$I = i_{SP}\,(S+P) + \varepsilon_I \tag{5}$$

Housing credit is not perfectly elastic (Adelino, Schoar, and Severino, 2012; Glaeser et al., 2012; Drechsler, Savov, and Schnabl, 2019). If demand for housing credit is strong, mortgage rates rise. An important driver of housing credit demand is the value of the outstanding stock of housing (i.e., the sum of log supply S and log house prices P). If the amount of housing increases or housing becomes more expensive, demand for credit increases and mortgage rates rise.

ε_I represents exogenous determinants of mortgage rates. Mortgage rates could experience a surprise drop for several reasons. The drop could be the result of an expansionary monetary policy shock, as emphasized in Taylor (2007), a lower term premium on risk-free long-term bonds (e.g., due to higher demand for long-term safe assets, as in Bernanke, 2005; Caballero, Farhi, and Gourinchas, 2008), or a lower lending spread above the risk-free rate (Justiniano, Primiceri, and Tambalotti, 2019) (e.g., because expected losses decline or bank intermediation becomes more efficient).[10] Our approach does not attempt to disentangle the various causes.

[10] Sá and Wieladek (2015) and Sá, Towbin, and Wieladek (2014) compare the importance of monetary policy and capital inflows shocks for the U.S. and a sample of countries in the Organisation for Economic Co-operation and Development (OECD), respectively.

The Model's Solution

The solution to the five-equation system is given by:[11]

$$
\begin{aligned}
P &= -\left[\left(1+d_I i_{SP}\right) p_R + \left(v_R + d_R\right) p_I i_{SP}\right] \cdot \varepsilon_S^* + p_R \cdot \varepsilon_D^* \\
&\quad - \left[\left(d_R + v_R\right) p_I + d_I \cdot p_R\right] \cdot \varepsilon_I^* + \left(v_R + d_R\right) \cdot \varepsilon_{Pe}^* \qquad (6)
\end{aligned}
$$

$$
\begin{aligned}
S &= \left[v_R + d_R + i_{SP}\left(d_I p_R + d_R p_I + v_R p_I\right)\right] \cdot \varepsilon_S^* + p_R s_P \cdot \varepsilon_D^* \\
&\quad - \left(d_I p_R + d_R p_I + v_R p_I\right) s_P \cdot \varepsilon_I^* + \left(v_R + d_R\right) s_P \cdot \varepsilon_{Pe}^* \qquad (7)
\end{aligned}
$$

$$
\begin{aligned}
V &= v_R\left[1 + i_{SP}\left(d_I + p_I\right)\right] \cdot \varepsilon_S^* - v_R\left[1 + (1+s_P) p_I i_{SP}\right] \cdot \varepsilon_D^* \\
&\quad \left[v_R\left(d_I - p_I s_P\right)\right] \cdot \varepsilon_I^* + v_R\left[s_P + (1+s_P) d_I i_{SP}\right] \cdot \varepsilon_{Pe}^* \qquad (8)
\end{aligned}
$$

$$
\begin{aligned}
I &= \left(d_R + v_R - p_R\right) i_{SP} \cdot \varepsilon_S^* + i_{SP} p_R\left(1+s_P\right) \cdot \varepsilon_D^* \\
&\quad + \left(d_R + v_R + p_R s_P\right) \cdot \varepsilon_I^* + \left[\left(d_R + v_R\right)\left(1+s_P\right)\right] i_{SP} \cdot \varepsilon_{Pe}^* \qquad (9)
\end{aligned}
$$

We use the model to identify structural shocks with sign restrictions (Canova and De Nicolo, 2002; Uhlig, 2005). We also solve for the rent-price ratio but impose no restrictions on its response; rather, we check empirically whether the predictions of the model are borne out by the data.

$$
\begin{aligned}
R - P &= \left[(p_R - 1)(1 + d_I i_{SP}) + (d_R + v_R - 1) p_I i_{SP}\right] \cdot \varepsilon_S^* + \left[1 + p_i i_{SP}(1+s_P) - p_R\right] \cdot \varepsilon_D^* \\
&\quad + \left[\left(d_R + v_R\right) p_I + p_I s_P + (p_R - 1) d_I\right] \cdot \varepsilon_I^* - \left[s_P + v_R + d_R + d_i i_{SP}(1+s_P)\right] \cdot \varepsilon_{Pe}^* .
\end{aligned}
$$

The main focus is on the identification of the price expectation shock. We compare this shock with shocks to the demand for housing services, to the supply of houses, and to the

[11] All exogenous factors are marked with a star because they are standardized by a positive constant:

$$\Omega = v_r + d_r + p_r s_P + \left(v_r + d_r\right) p_i i_{SP}\left(1+s_P\right) + d_I i_{SP} p_R\left(1+s_P\right).$$

mortgage rate and discuss how we can distinguish them. Table 1 summarizes the identification restrictions of the baseline specification. All structural shocks have been normalized to imply an increase in the real price of housing. To be able to identify the model using the data, our methodology requires a unique combination of sign restrictions for each shock.

Table 1. Baseline Shock Identification

	Shock to:			
	Housing Supply (ε^*_S)	Demand for Housing Services (ε^*_D)	Mortgage Rate (ε^*_I)	Price Expectation ($\varepsilon^*_{P^e}$)
Housing Price (P)	> 0	> 0	> 0	> 0
Housing Supply (S)	< 0	> 0	> 0	> 0
Vacancy Rate (V)	< 0	< 0	·	> 0
Mortgage Rate (I)	·	> 0	< 0	> 0

All structural shocks have been normalized to imply an increase in the real price of housing.

Price expectation shocks: A positive house price expectation shock leads to increases in house prices, the housing supply, vacancies, and the mortgage rate. The restriction on house prices follows from Assumption 1: The prospect of being able to sell the house at a higher price in the future leads to higher prices now. The housing supply increases as a result of Assumption 3: As large adjustments to the supply of houses are more costly than small changes, the prospect of higher future prices also creates the incentive to start building now, causing the housing supply to increase. The increase in housing construction leads to higher demand for mortgage credit. Because the mortgage credit supply is not perfectly elastic, the increased demand for loans associated with higher prices and construction activity will lead to higher mortgage rates (Assumption 5). The increase in vacancies relies on Assumption 4: Markets do not fully clear because of search and matching frictions. As supply increases and the current demand for housing services is not affected by expectations, the vacancy rate rises.

The idea of supply increasing following an expectation shock while demand for housing services remains static is not new. Specifically, Topel and Rosen (1988) use a similar modeling assumption. They write, “For instance, an anticipated transitory increase in future demand causes bubble-like price and investment responses: House prices increase immediately in a rational market, and this signals increased construction activities prior to the time the change occurs” (p. 727). The rent-price ratio is predicted to increase, as prices increase but current rents do not (Assumptions 1 and 2).

Mortgage rate shocks: A negative mortgage rate shock is characterized by a decrease in the real mortgage rate and increases in house prices and the housing supply. Lower interest rates increase current prices through their impact on the present discount value (Assumption 1). Higher prices encourage supply (Assumption 3). The response of the vacancy rate is ambiguous, as the increase in supply in response to lower mortgage rates may be more than compensated by higher demand for housing services because of relieved credit constraints. The response of the price-rent ratio is again theoretically ambiguous: While a lower interest rate will drive up prices (Assumption 1), it may also drive up rents because of higher demand (Assumption 3). In practice, we expect the sensitivity of housing demand to interest rate d_I to be relatively small, such that vacancies increase and the rent-price ratio decreases. While prices and supply move in the same direction, opposite movement of mortgage rates enables us to distinguish the two shocks.

Demand for housing services shocks: A positive shock in the demand for housing services (i.e., occupying a house or moving to a bigger house) leads to an increase in house prices, an increase in the housing supply, a decrease in housing vacancies, and higher mortgage rates. The restriction on house prices and residential investment are as in Jarociński and Smets (2008) and follow from Assumptions 1, 2, and 3: An upward shift in the demand curve leads to higher house prices and increases in the housing supply. An upward-sloping mortgage credit supply will lead to higher interest rates, as mortgage demand increases with higher prices and investment in housing (Assumption 5). Because it takes time for the supply of houses to adjust, growth in the demand for housing services temporarily exceeds the growth in supply, which reduces the vacancy rate (Assumption 4). The restriction on vacancies is crucial to distinguish the demand for housing services shock from an expectation shock. The identified shock to the demand for housing services thereby focuses on exogenous variations in the demand for housing services but does not capture an increase in housing demand for investment purposes. The response of the price-rent ratio is ambiguous, as more demand pushes up both prices and rents.

Housing supply shocks: A negative housing supply shock is associated with a rise in house prices and decreases in the supply of housing and the vacancy rate. The restrictions on house prices and residential investment are again as in Jarociński and Smets (2008) and follow from Assumptions 1 and 3. An upward shift in the supply curve leads to higher prices and lower quantities. As there are now fewer houses for a given demand for housing services, the vacancy rate falls (Assumption 4). The response of mortgage rates is ambiguous: Credit demand may either increase due to higher house prices or decrease due to lower construction activity (Assumption 5). The response of the price-rent ratio is also ambiguous, as less supply pushes down both prices and rents.

3 Data

We conduct our analysis at both the national-quarter level and the state-year level.

3.1 National-level Data

The change in the real housing price (ΔP) is measured by the log difference of the national real Case-Shiller house price index. Housing supply is approximated by real private residential investment from the Bureau of Economic Analysis (ΔS). The real mortgage rate (I) is approximated by the nominal contract rate on the purchases of existing single-family homes provided by the Federal Housing Financing Agency (FHFA), less the long-term inflation expectations, measured by the 10-year-ahead forecast of the inflation rate

(Macroeconomic Advisers, downloaded from Haver). The rent-to-price ratio ($R - P$) is computed as the log of the ratio between the housing component of the Consumer Price Index (Bureau of Labor Statistics) and the nominal Case-Shiller house price index. Finally, the log difference of U.S. real gross domestic product (GDP; $\Delta RGDP$), which we use as a control for economic activity, is taken from the Bureau of Economic Analysis. In extensions in which we identify LTV shocks, we use the LTV ratio at the time of house purchase from FHFA.

Of special importance in our analysis is the series of vacancy rates (V). This series is given by the overall ratio of vacant houses that are part of the market relative to the total housing stock, excluding seasonal factors (Census Bureau). The Census uses the following definitions: Vacant homes are vacant year-round (i.e., excluding seasonal vacancies) and include "vacant for rent," "vacant for sale," "held off the market," and "rented or sold."[12]

Our national-quarter level data set covers the period from 1973Q1 to 2018Q1. Figure 2 shows the evolution of the respective variables.

3.2 State-level Data

The log changes in the real house price indexes (ΔP) are provided by the FHFA. The change in the supply of housing (ΔS) is measured as the log of new private housing permits, using data from the U.S. Census Bureau.[13] Vacancy (V) data series are available for owner-occupied homes as well as for rental homes from the U.S. Census Bureau. Because we are interested in the aggregate vacancy at the state-year level, we combine these data into a single number using the weights of homeownership, based on region-year homeownership

[12] *Vacant for rent* consists of vacant units offered for rent and those offered for both rent and sale. *Vacant for sale* comprises units for sale only; it excludes units for both rent and sale. *Vacant units held off the market* includes units held for occasional use, temporarily occupied by persons with a usual residence elsewhere, and vacant for other reasons. *Vacant units rented or sold* consists of year-round vacant units that have been rented or sold but the new renters or owners have not moved in as of the day of the Census interview. See further details at `https://www.census.gov/housing/hvs/definitions.pdf`.

[13] Data on residential investment are not available at the state level. Housing permits promise to be a fair proxy as they tend to lead residential investment by only a few months and are positively correlated with actual investment.

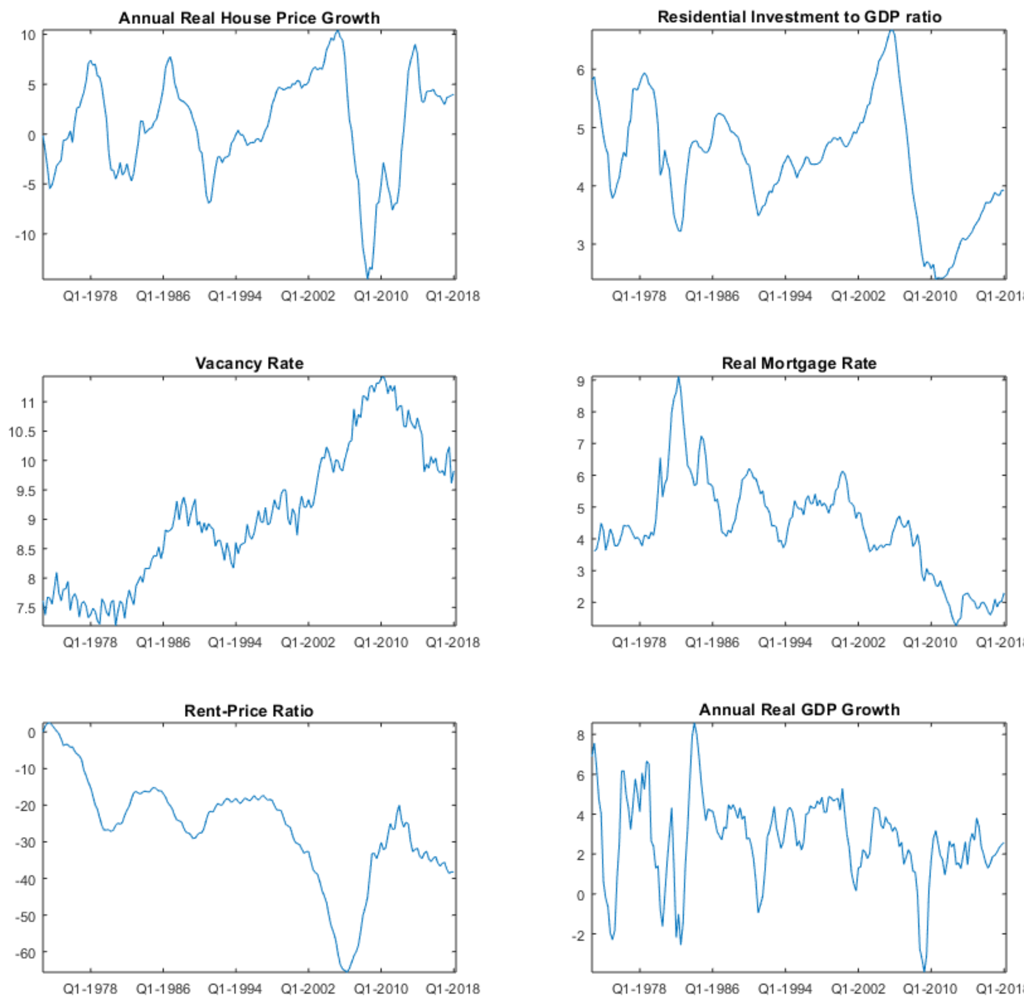

Figure 2. Evolution of Variables Over Time

The figure shows the time-series evolution of the main variables used in the analysis. The units on the y-axis represent cumulative log differences.

statistics from the U.S. Census Bureau.[14] State-year level mortgage rates (I) are available on the FHFA website, deflated using the national 10-year inflation expectation. To capture long-run dynamics in prices and to control for economic activity, we include the housing price-to-median income ratio, with median income data coming from the U.S. Census Bureau.[15] Data limitations (in particular regarding the vacancy rate) constrain the state-level analysis to the annual frequency. Our state-level data set covers the period 1988 to 2017 for 49 states[16] and Washington, D.C.

Ideally, we would perform the analysis at the metropolitan statistical area (MSA) or commuting zone (CZ) levels. Prior studies have found that some MSAs experienced unusual price increases in the early 2000s. For example, between 2000Q4 and 2006Q4, nominal housing prices increased in the Miami MSA by 163%, in the Phoenix MSA by 110%, and in the Las Vegas MSA by 115%. Unfortunately, housing vacancy data, a necessary component for identifying the expectation shock, are available only since 2005 for rental properties and since 2015 for owner-occupied properties at the MSA level (see Census website). A second-best approach is to examine state-level data, for which vacancy data are available since 1987. The disadvantage at this level of aggregation is that it masks within-state variation. The price variation during both the boom and bust at the state level was less dramatic than it was in some of the MSAs, yet it was notably higher than that of the aggregate U.S. For example, from 2000Q4 to 2006Q4, the national nominal housing price index increased by 70%. In contrast, Florida experienced a nominal price increase of 121%, Arizona 107%, and Nevada 103%.

[14] Due to limited data availability, we focus on the vacancy categories "vacant for rent" and "vacant for sale."

[15] There are no comparable data on rents at the state level. Therefore, we do not include the rent-price ratio in the state-level regressions. We do not include state-year real GDP merely in order to preserve degrees of freedom.

[16] Oklahoma is excluded due to missing data.

4 Estimation Method

This section introduces the empirical framework of the study. We first present the econometric model and then describe the identification approach and discuss the inference from the computational implementation.

4.1 Econometric Model

For the national-level analysis, we estimate a Bayesian vector autoregressive (BVAR) model of the following form:

$$\mathbf{y}_t = \sum_{i=1}^{L} \mathbf{A}_i \mathbf{y}_{t-i} + \mathbf{e}_t, \quad \text{with} \quad \mathbf{e}_t \sim \mathrm{N}(\mathbf{0}, \mathbf{\Sigma}) \quad \forall\, t = 1, ..., T \tag{10}$$

$\mathbf{y}_t$ is a vector of seven variables

$$\mathbf{y}_t = \begin{pmatrix} \triangle P_t & \triangle S_t - RGDP_t & V_t & I_t & R_t - P_t & \triangle RGDP_t \end{pmatrix}^T$$

$\mathbf{e}_t$ is a reduced-form error term with variance-covariance matrix $\mathbf{\Sigma}$; L is the lag length; and $\mathbf{A}_i$ are coefficient matrices.[17]

The first four variables in the BVAR are required for identification.[18] The inclusion of the rent-to-price ratio is motivated by the co-integrating relationship between rents and prices. It allows us to capture long-run dynamics in prices, while only including stationary variables (see, for example, King, Plosser, Stock, and Watson, 1991). Real GDP growth is included to capture general economic conditions. The combined responses of the residential investment to GDP ratio and GDP growth allow us to compute the level response of residential investment and GDP. Furthermore, the responses of real GDP and the rent-to-price ratio allow us to assess the consistency of the responses to theoretical arguments.

[17] For the state-level regression, $S_t - RGDP_t$ is proxied by new private building permits; $RGDP_t$ is dropped; and $R_t - P_t$ is replaced by dividing the housing price at the state level with the state median income. Data come from the U.S. Census. L equals 2 for the national-level regression and 1 for the state-level regressions.

[18] We introduce the housing price in the first difference and use the ratio of residential investment to GDP to account for co-integrating relationships and deterministic trends. However, we apply the sign restrictions on the levels of house prices and residential investment.

4.2 Computational Implementation

We sample the regression coefficients A_i and covariance matrix Σ from the posterior distribution, with an uninformative prior distribution.[19] Given the parameter draws, we implement the identification based on sign restrictions. We can think of the one-step-ahead prediction error e_t as a linear combination of orthonormal structural shocks $e_t = B \cdot v_t$, with $E(v_t' v_t) = I$, where the matrix B describes the contemporaneous response of the endogenous variables to structural shocks, $\Sigma = E(e_t e_t') = E(B v_t v_t' B') = BB'$. To sample candidate matrices B, we compute the Cholesky factorization V of the draws of the covariance matrix Σ. We then multiply V using a random orthonormal matrix Q $(B = VQ)$. Q is sampled as in Rubio-Ramírez, Waggoner, and Zha (2010).[20] The Q matrices are orthonormal random matrices. Given matrix Q and the impact matrix B, we compute candidate impulse responses. If the impulse response functions implied by B are consistent with the sign restrictions in Table 1 for all shocks, we keep the draw. We constrain the sign restriction to hold for the first two periods for all of the variables' responses. We repeat the procedure until we accept 5,000 models.

In contrast to exact identification schemes (e.g., zero restrictions), error bands for standard vector-autoregression (SVAR) models based on sign restrictions reflect two types of uncertainty: parameter and identification uncertainty. Parameter uncertainty occurs both in models with exact restrictions and in models with sign restrictions: With a limited amount of data, there is uncertainty about the true parameters of the model. Identification uncertainty is specific to models with sign restrictions. When applying sign restrictions, there is a set of impulse response functions that satisfy the restriction for a given parameter draw.

We report the pointwise mean of accepted impulse response functions for each variable. We proceed similarly for the historical decomposition and the variance forecast error decom-

[19] Σ is drawn from an inverted-Wishart distribution $IW(\Sigma_{OLS}, T)$, and the coefficient matrices A_i from a normal distribution $N(A^k_{OLS}, \Sigma_{OLS})$, where T is the number of observations and subscript OLS stands for the ordinary least squares estimates.

[20] We compute Q by drawing an independent standard normal matrix X and applying the QR decomposition $X = QR$.

position and use the pointwise mean as our baseline measure. As error bands, we report the pointwise 16th and 84th percentiles. As is standard in the literature, historical decompositions are constructed using point estimates, i.e., discarding parameter uncertainty. Doing so facilitates the interpretation of results, as it ensures that the sum of the individual contributions add up to the total.

5 Results: National-level Analysis

We begin our discussion of national-level results with a discussion of the impulse response functions of the identified shocks. This discussion prepares the ground for the main interest of the study: the contribution of price expectation shocks to the housing boom of the 2000s, which is examined in Section 5.2. Section 5.3 presents results from the forecast error variance decomposition, i.e., the importance of the different shocks over the entire sample period.

5.1 Impulse Response Functions

Figure 3 depicts the response of the six variables in the VAR to the four identified shocks. In each case, the size of the shock is normalized to one standard deviation and the sign of the shock is normalized such that the response of the house price is positive (i.e., a positive expectation shock and a negative mortgage rate shock). The responses of real house prices, real residential investment, and real GDP are displayed in levels.

In response to a positive one standard price expectation shock, house prices rise in the first three years by about 1.2% and then start to decline slowly. The rent-to-price ratio declines initially, consistent with the (unrestricted) sign prediction from the model, implied by Assumption 1 (i.e., house prices reflect the present discount value of future rents and the sale price). Residential investment increases on impact by close to 1% and follows a hump-shaped pattern, peaking at about 2%. The hump-shaped response of investment is consistent with convex adjustment costs (Assumption 3). Investment peaks six quarters

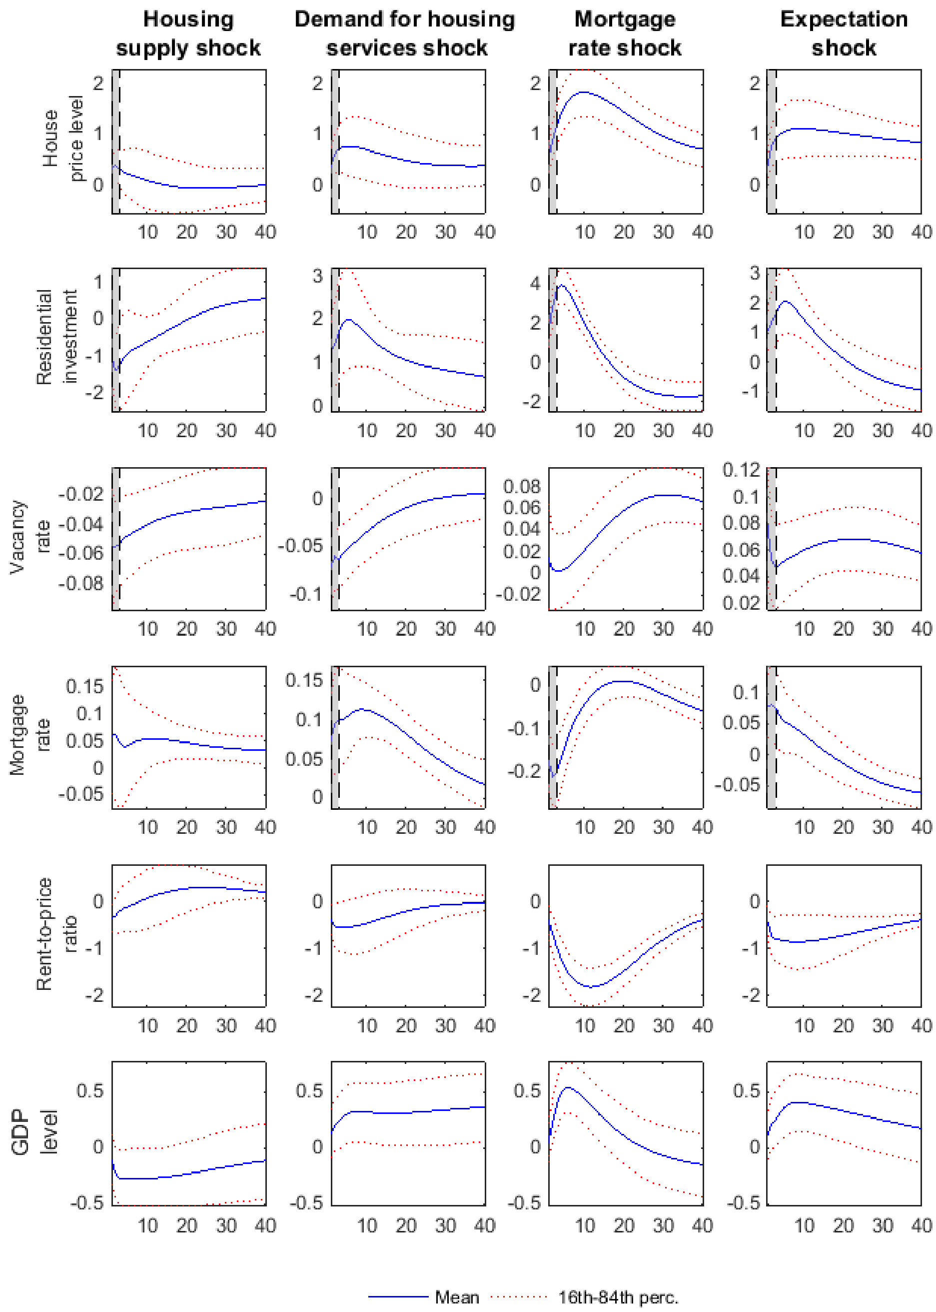

Figure 3. Baseline Model: Impulse Response Functions

This figure shows the pointwise mean of accepted impulse response functions and is the main summary measure, along with pointwise 16th and 84th percentile error bands. These measures are described in Section 4.2. The identification assumptions are summarized in Table 1. A gray shaded area marks periods for which sign restrictions have been imposed. The units on the y-axis represent cumulative log differences.

after the expectation shock. In later quarters, it persistently falls below its pre-shock path after about six years. Real GDP increases only temporarily, a pattern that is consistent with the hypothesis that overly optimistic expectations about future housing conditions are compensated in the medium run with persistently lower residential investment. The real mortgage rate increases by roughly 8 basis points in the first year due to higher mortgage demand. The real mortgage rate then starts to decline and eventually undershoots initial levels, consistent with persistently low residential investment and low demand for mortgages. At the same time, the vacancy rate is increasing for an extended period, reverting slowly only after about six years. This suggests that a persistent excess supply follows the expansion in construction, underpinning the need for residential investment to decline below its pre-shock path for a prolonged period.

A negative mortgage rate shock leads to qualitatively similar responses of house prices and residential investment as a positive price expectation shock. Quantitatively, however, a mortgage rate shock is associated with substantially stronger responses in residential investment and output, and the house price increase is also somewhat larger. The response of variables are, however, less persistent. As in the case of the price expectation shock, residential investment falls below zero over the medium term, and the temporary output (GDP) increase dissipates. The rent-to-price ratio initially decreases, as we would expect from the present value relationship discussed under Assumption 1, if interest rates fall. The vacancy rate increases persistently. Both these responses (unconstrained) imply that the sensitivity of demand for housing services (d_I) to interest rates is limited.

Positive shocks to the demand for housing services are associated with a persistent increase in residential investment, real housing prices, and output. Residential investment and output rise initially by about the same amount as in response to the price expectation shock, but the response is more persistent. House prices rise by less than in response to price expectation shocks. Different from the price expectation shock, the response of the rent-to-price ratio is weak: Although it falls initially, the response turns quickly insignificant. Hence, house prices and rents grow by about the same amount. The pattern is also consistent with the simple model described above, which predicts that an expectation shock unambiguously reduces the rent-price ratio, while the response of the rent-price ratio to a shock to the demand for housing services is necessarily smaller and could even be positive. The mortgage rate rises by about 10 basis points in response to increased housing activity and remains relatively elevated for an extended period, in line with the persistent increase in residential investment. The vacancy rate drops initially but within five years returns to its pre-shock path, suggesting that increased residential investment closes the gap between the demand for housing services and the supply of housing.

Negative supply shocks are associated with an increase in house prices and a contraction of residential investment and output. As in the case of the demand shock for housing services, the response of the rent-to-price ratio is insignificant at most horizons, as we would expect from Assumption 5 if the shock mainly affects current fundamentals and is not driven by expectations.

5.2 Contribution of Price Expectation Shocks to the Housing Boom

This section explores how the four identified shocks contributed historically to housing dynamics at specific points in time. Figure 4 displays the decomposition of real house prices over time. The solid black line is the log real house price (normalized to zero at the starting point of the boom period in 1996Q4). The log house price is presented in deviation from its deterministic path (i.e., the path house prices would have taken if no shock had occurred since the starting point of the sample). The colored bars indicate the contribution of the four shocks to the observed path. Finally, there is an unexplained residual that occurs because the model is only partially identified. Quantitative results are shown in Table 2.

The four identified shocks explain a substantial share of the house price increase in the run up to the crisis. About 70% of the increase between 1996Q4 and 2006Q1 is explained by the four identified shocks in the baseline model. The largest contribution comes from mortgage

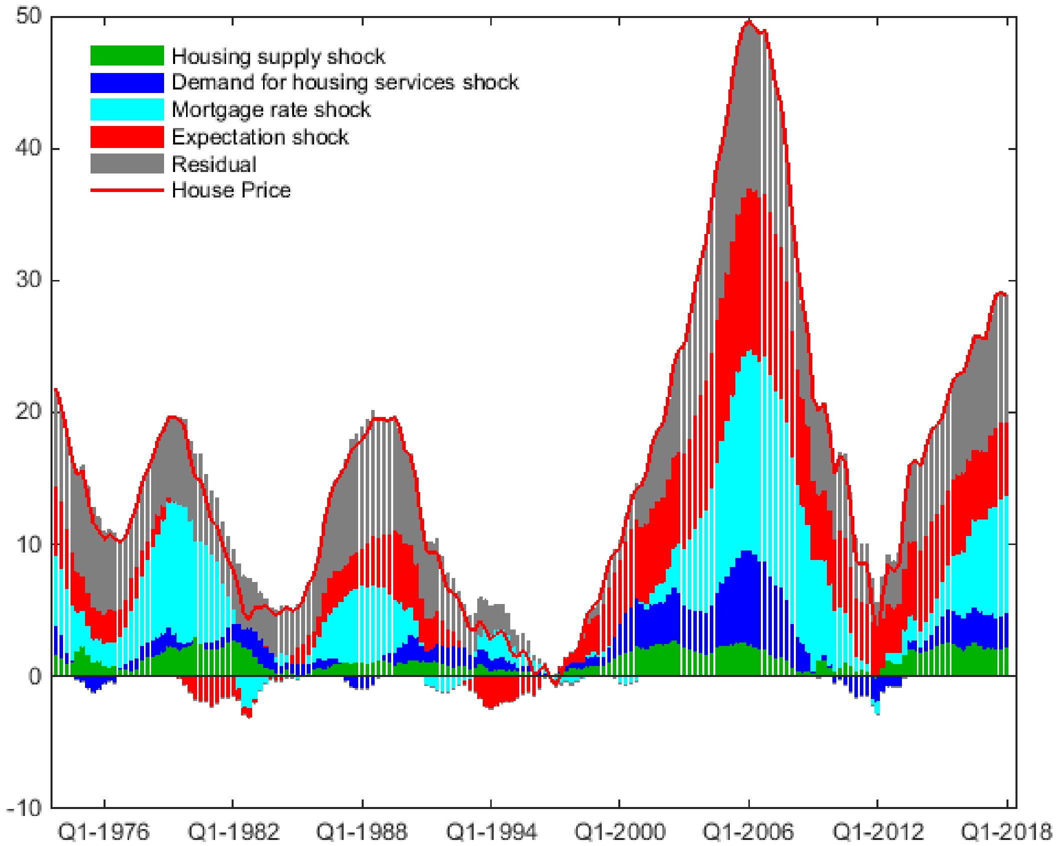

Figure 4. Decomposition of Real House Price Developments

This figure shows the contribution of the different shocks to the national home prices between 1973Q1 and 2018Q1. The solid red line is the log real house price (normalized to zero at the starting point of the boom period in 1996Q4). It is presented as the deviation from its deterministic path (i.e., the path house prices would have taken if no shock occurred since the starting point). The bars indicate the contribution of the four shocks and the unexplained residual to the observed path. The units on the y-axis represent cumulative log differences.

rate shocks, explaining 28% of the increase. The second most important contribution comes from price expectation shocks, accounting for about 23% of the rise. The price path generated by these two shocks increases monotonically over the boom period. The contribution from the mortgage rate shock gains in importance after the 2001 recession, when monetary policy is widely perceived as accommodating. Shocks to the demand for housing services and the supply of housing account only for a small fraction of the boom (5% and 13%). Finally, as our model is only partially identified, there is a sizable unexplained residual of about 22%. The deterministic component explains a small part, less than 10%.

A model that accounts only for the three traditional shocks (housing supply, demand for housing services, and mortgage rate),[21] explains about 50% of the house price increase in the 2000s boom. The contribution ascribed in our model to the expectation shock would be partly assigned to the shock to the demand for housing services (as the higher demand for housing services can coexist with a vacancy rate increase under the traditional identification assumptions) and remains partly unexplained (increases the residual contribution). In our view, these results suggest that accounting for price expectation shocks is not only relevant by itself, but also relevant to better estimate the contribution of other shocks to the boom. At the same time, attributing the residual that cannot be explained by conventional shocks to expectations would overestimate the contribution of expectation shocks to the 2000s house price boom.

In a model in which we distinguish between housing demand shocks associated with LTV increases (labeled LTV shock) and those that are not (labeled shock to demand for housing services), the contribution of the price expectation shock is somewhat smaller but remains the most important contributor to the boom after the mortgage rate shock. The contribution of the LTV shock is limited. This result is in line with empirical studies showing that the LTV distribution and mortgage approval rates did *not* change materially during the boom

[21]Hence, the identification would be identical to the one described in Table 1 without any constraints on the vacancy rate and no identification of the price expectation shock.

Table 2. Contribution of Shocks to Price Boom and Bust

Model:	Shock to: Housing Supply	Demand for Housing Serv.	Mortgage Rate	Expect-ations	LTV	Deter-ministic	Residual
			Contribution to Boom (1996Q4–2006Q1)				
Baseline	4.8	12.8	28.0	22.9		9.5	22.0
Baseline							
− excl. Expectation	6.5	21.9	22.6			9.5	39.5
− with LTV	3.6	13.6	26.7	18.8	7.8	10.8	18.8
			Contribution to Bust (2006Q2–2012Q1)				
Baseline	6.4	20.3	37.1	20.9		-7.7	23.0
Baseline							
− excl. Expectation	7.6	27.4	30.7			-7.7	42.0
− with LTV	5.2	17.9	36.4	16.3	12.1	-10.0	16.3

The table shows the share of the change in house prices explained by the respective shock. The baseline identification assumptions are summarized in Table 1.

and bust years and across the United States (Glaeser et al., 2012; Adelino et al., 2018a).[22]

Turning now to the decline in real house prices that started in 2006Q2 and ended in 2012Q1, the historical decomposition reveals that the decline was again mainly driven by mortgage rate and price expectation shocks, as they explain more than 50% of the path (see Table 2). However, the contribution of the mortgage rate shock (about 37%) and demand for housing services shock (about 20%) mattered more in the bust than in the boom.

As shocks to the demand for housing services have persistent effects on house prices, the decline can mainly be attributed to new negative shocks, consistent with a decline in income during the recession. The high contribution from mortgage rate shocks to the house price decline partly reflects the reversion of house prices implied by the impulse response function.

[22]For example, Adelino et al. (2018a) report that lenders indeed provided larger mortgages relative to income during the boom period; however, LTV remained constant over the boom and bust periods because mortgage size increased with the value of assets, keeping LTV constant. In their Figure 10, the authors show that between 1996 and 2012 LTV distribution stayed almost constant in both boom and non-boom states.

Table 3. Variance Decomposition

	Shock to:				
Quarter Ahead	Housing Supply	Demand for Housing Services	Mortgage Rate	Expectations	Residual
			Real Housing Price		
1	12	16	28	19	25
4	9	15	32	22	22
16	9	15	33	21	22
40	9	14	34	21	22

The table displays the share of the forecast error variance explained by the respective shocks at various forecast horizons. Shocks are identified as summarized in Table 1.

It also reflects new contractionary mortgage rate shocks.[23] The contribution of shocks to the housing supply remains minor.

5.3 Forecast Error Variance Decomposition

While the previous section looked at the importance of the various shocks during boom and bust periods, this section analyzes their importance over the entire sample, using a forecast error variance decomposition (see Table 3).

The mortgage rate and price expectation shocks are again the most important drivers of the variation in housing prices. Price expectation shocks account for about 20% and mortgage rate shocks for more than 30% of house price variation at the 10-year horizon. Hence, the contribution of the mortgage rate shock is slightly higher than during the boom period. The contribution of the price expectation shock is still substantial, but lower than the contribution during the 2000s boom period.

Shocks to the supply of housing account for around 10% and the shocks to the demand for housing services about 15% of the variation in house prices. The contribution in the boom and subsequent bust period in the 2000s is tilted somewhat more to demand for

[23] The mortgage rate fell less than predicted by the VAR given the other variables' developments. Thus, a positive mortgage rate shock materialized in the period directly after the housing price peak.

housing services than supply variation. About 22% of the variation in house prices remains unexplained by the four shocks, which is in line with the residual's contribution during the boom period.

6 Results: State-level Analysis

Next, we expand the analysis to the state level. This extension allows us to verify that our inference about the roles of the various shocks agrees with observations of prior research about the sources of the boom in different states. Our analysis uses state-year data from 1988 to 2017, and employs a similar BVAR analysis and simulations at the state level. We conduct the analysis independently for each state.

We begin the discussion of the results by presenting charts of the components contributing the evolution of prices averaged across states. Mayer (2011) documents that there was significant heterogeneity in the magnitude of the U.S. housing boom. The locations that experienced unprecedentedly large booms were cities in the Sand States: Las Vegas (Nevada), Miami (Florida), Phoenix (Arizona), Los Angeles (California), and San Diego (California). At the state level, the magnitudes of these booms in our data were on the order of 50% to 80% (in real terms, log differences, from 1996 to 2006). The busts in these locations were correspondingly deep: 60% to 90% from peak to trough (2006–2012). We therefore distinguish between Sand States (Arizona, California, Florida, and Nevada) and non-Sand States (all other states and Washington, D.C.). In each case, we show the average contributions of the various shocks to the boom and bust in log real house prices, equally weighted across states.[24]

Panel (a) of Figure 5 shows the shock contributions for the non-Sand States. The figure shows that at the peak of the boom (2006), the main shock contributing to the average magnitude of the boom was the mortgage shock (about 33% of the average magnitude of the boom explained by shocks), followed by the expectation shock (about 22% of the boom's magnitude). Shocks to the demand for housing services and to the supply of housing are smaller in magnitude (jointly about 29% of the boom's magnitude), and the remainder is accounted for by an unidentified residual. We find a similar ordering of the contribution of the shocks in the bust period (2006–2012): The mortgage shock is the prime contributor, followed by the expectation shock.

Panel (b) of Figure 5 shows the contribution of shocks to the boom in the Sand States (Arizona, California, Florida, and Nevada). Clearly, the magnitude of the boom is almost three times larger than the average magnitude of the boom in other states (Panel (a)). In addition, the importance of expectation shocks increases (to about 26% of the boom's magnitude) and is now comparable to the mortgage rate shock contribution (about 29% of the boom magnitude).[25] Shocks to the demand for housing services and to the supply of housing play a minor role in the boom (together, accounting for about 28% of the boom's magnitude). Finally, since 2014, the average contribution of expectation shocks to price changes has again been positive in Sand States, while it remained subdued in other states.

Next, we present the absolute contributions of the different shocks to the boom in Figure 6 for the various states. States are sorted by the size of the contribution of the shock. We

[24]At the national level, the house price path is shown in deviation from the deterministic path, i.e., the path house prices would have taken if no shock occurred since the starting point. Different from the national-level analysis, a sizable share of the entire boom (a bit less than half for non-Sand States and about one-fourth for Sand States) is accounted for by the deterministic component. As a result, the boom and bust appear less dramatic in the charts (where the deterministic component is omitted) than in the raw data. The contribution of the higher deterministic component may be a result of a short sample period and the larger influence of initial conditions, i.e., unobserved shocks that occurred before the sample starts (Hamilton, 1994). The role of the these initial conditions disappears as the sample length increases, as is the case in our national data that start about 15 years earlier. In some states, there is also a structural upward trend in real house prices.

[25]With the exception of Florida, expectation shocks were more important than mortgage rate shocks during the boom in the Sand States.

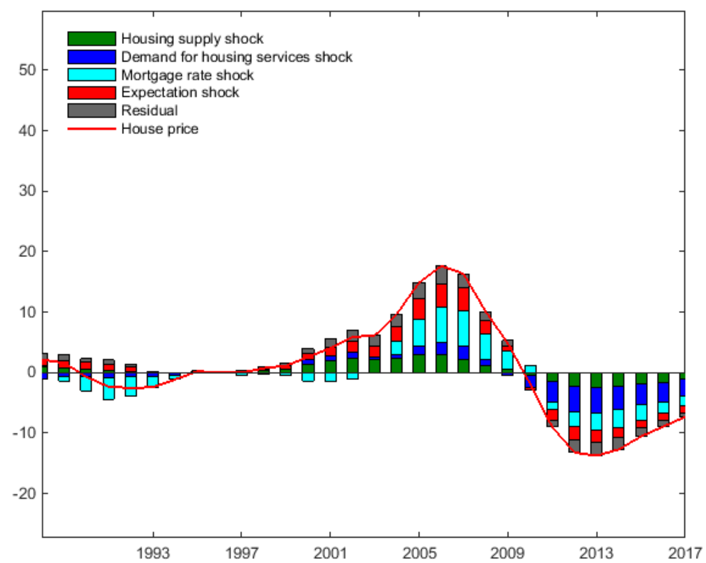

(a) All States, Excluding Sand States

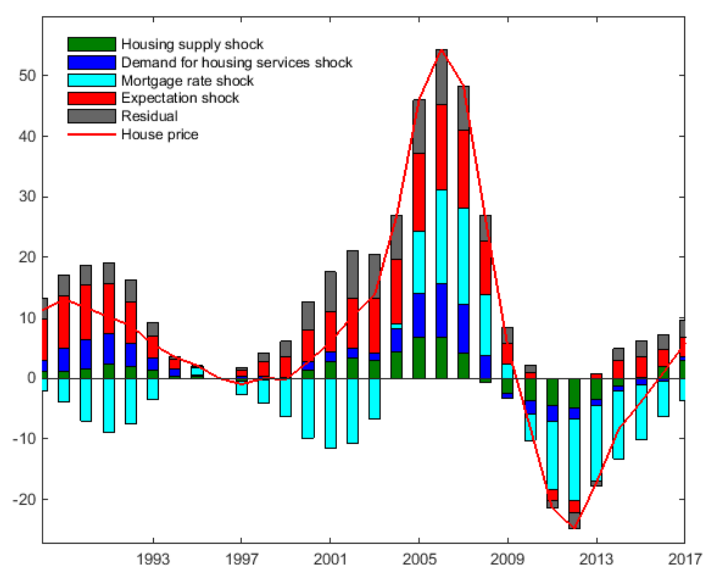

(b) Sand States

Figure 5. State-level Contribution of Shocks to the Price Evolution

This figure presents the contribution of the different shocks to house prices at the state level. Panel (a) shows the average across contributions for all states excluding the Sand States. Panel (b) shows the average across contributions for the four Sand States (Arizona, California, Florida, and Nevada). In each chart, the analysis was performed at the state level and equally averaged across all participating states. The solid red line is the log real house price (normalized to zero at the starting point of the boom period in 1996). The house price path is shown in deviation from its deterministic path, i.e., the path house prices would have taken if no shock occurred since the starting point. The bars indicate the contribution of the four shocks and the unexplained residual to the observed path. The units on the y-axis represent cumulative log differences.

show the top 20 states as well as the minimum, median, and maximum for the remaining 30 states. Panel (a) depicts the contribution of the expectation shock to the price increase during the boom. While the contribution rises gradually for most states from left to right, a marked increase in the importance of the expectation shock can be observed for three of the Sand States—Arizona, California, and Nevada—and to a lesser extent for Idaho. The magnitude of the contribution of this shock is 11 percentage points in the case of Idaho, up to nearly 19 percentage points for Nevada. The rest of the states experienced significantly lower expectation shocks of between –1 and 10 percentage points.

Panel (b) of Figure 6 shows the contribution of the mortgage shock for the top 20 states. Mortgage shocks contributed the most to price increases in Maine, Florida, Nevada, Vermont, and Michigan. The magnitude of the contributions to these states' booms is on the order of 15 to 22 percentage points. The contribution of the mortgage shock to the boom in the rest of the states ranges from –4 to 16 percentage points.

In Panel (c) of Figure 6, we present the contribution of shocks to the demand for housing services to the boom. The panel shows that California has the highest contribution of 11 percentage points, while for all other states, the contribution of the shock to the demand for housing services ranges between –6 and 9 percentage points. The result that the shock to the demand for housing services had the most significant contribution to California is consistent with the interdependent analysis of Ferreira and Gyourko (2017), which shows that California's boom was disproportionately driven by fundamental demand for housing services.

Panel (d) of Figure 6 shows the contribution of the supply shock to the boom of the various states. It shows that several of the most densely populated states have a high supply shock contribution. Maryland, Nevada, Washington, D.C., Massachusetts, Rhode Island, Washington, Hawaii, Arizona, and New Jersey had the highest contribution, of close to 10 percentage points each. This finding is consistent with studies pointing to the importance of geography in housing supply elasticities and house price growth (Green, Malpezzi, and Mayo, 2005; Saiz, 2010). The contribution to other states' booms ranges from –3 to 8 percentage points.

(a) Expectation Shock

(b) Mortgage Shock

(c) Demand for Housing Services Shock

(d) Supply Shock

Figure 6. State-level Contribution of Shocks to the Boom (1996–2006)

The figure shows the contribution of the four types of shocks to the evolution of prices during the boom in various states. The bars show the contribution of the top 20 states (log real house prices). The dashed lines show the minimum (black), maximum (black), and median (red) of the remaining 30 states. The units on the y-axis represent cumulative log differences.

Another useful way to look at the results is through heat maps of the boom (Figure 7) and bust (Figure 8).[26] Figure 7 shows that the expectation shock was most important in the Sand States, as discussed earlier. The mortgage shock shows no particular geographic clustering:

[26]Hawaii and Alaska are excluded from these charts only for visual ease.

Its effect is the strongest in the Sand States as well as the eastern coast, Michigan, and Idaho. The shock to the demand for housing services was important primarily for California. The supply shock was important mostly in the Northeast as well as the western states.

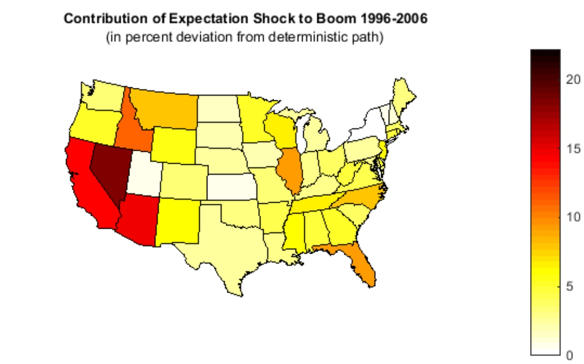

(a) Expectation Shock

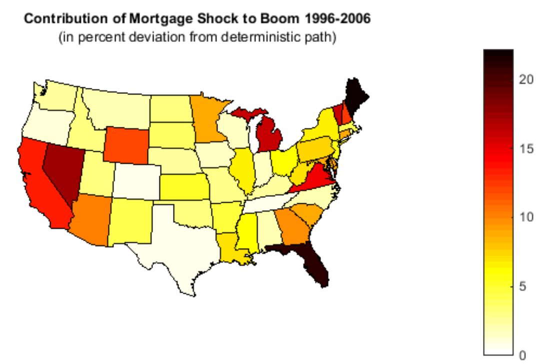

(b) Mortgage Shock

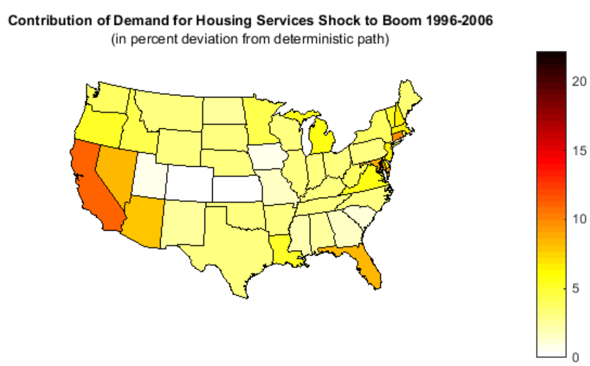

(c) Demand for Housing Services Shock

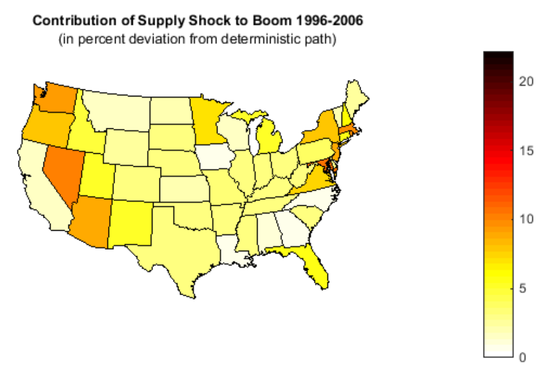

(d) Supply Shock

Figure 7. State-level Contribution of Shocks to the Boom (1995–2006)

The figure shows the contribution of the various shocks to the evolution of prices. The maps show the contribution of each type of shock to the states' booms for the period 1996–2006. The color scheme shows the magnitude of the contribution of the shock to the size of the boom, in cumulative log differences.

When examining the contribution to the bust (2006–2012) (Figure 8), the picture mirrors that of the boom in most, albeit not all, cases. Expectation shocks had the greatest contribution in California, Arizona, and Nevada, but was also high in Montana. The mortgage shock had the greatest contribution to the bust in Nevada and Florida, followed by

California, Michigan and Arizona. Mortgage shocks also made important contributions in Washington, Maryland, Minnesota, Georgia, Maine, New Hampshire, and Virginia. Shocks to the demand for housing services had a mild contribution to the bust nationwide, with the strongest impact in Nevada and Florida. Supply shocks had little bearing on the bust, with the largest effects in Nevada and Arizona.

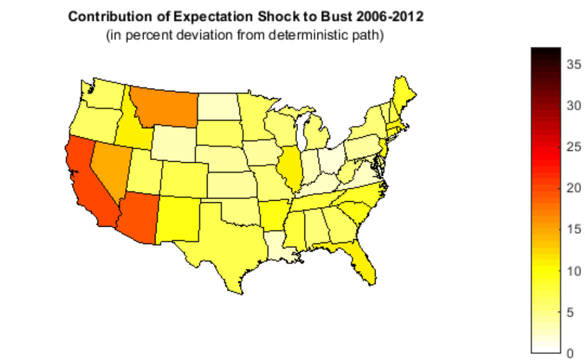

(a) Expectation Shock

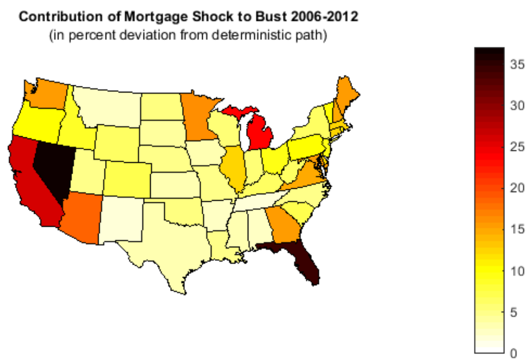

(b) Mortgage Shock

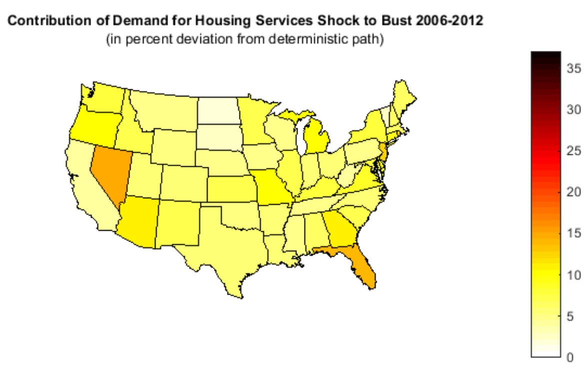

(c) Demand for Housing Services Shock

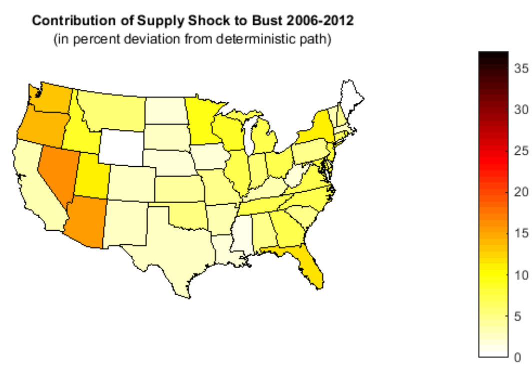

(d) Supply Shock

Figure 8. State-level Contribution of Shocks to the Bust (2006–2012)

The figure shows the contribution of the various shocks to the evolution of prices. The maps show the state-level contribution for the different states for the period 2006–2012. The color scheme shows the magnitude of the contribution of the shock to the size of the bust, in cumulative log differences.

Overall, the state-level analysis supplements our observations from examining the data at the national level. As in the national level, we observe that expectation and mortgage shocks were the most important to the evolution of housing prices during both the boom and bust periods. Yet, we observe considerable geographic variation in the effects. The boom and bust were most severe in the Sand States, where expectation shocks contributed the most to the large magnitude of the price fluctuations. In other states, mortgage shocks generally had the largest effect, followed by expectation shocks. Supply shocks appear to have mattered most during the boom in densely populated states, while contributing little to lower prices during the bust, when the decrease in demand for housing services dragged prices down nationwide.

The bust in the residential market cannot be seen as an event independent from the preceding boom. We use the cross-section of states to explore which of the four shocks that caused the boom can also explain the bust. In particular, we measure the correlation between the contribution of the shocks to the *boom*, with the depth of the following bust. Our examination is to some extent a test of Burnside et al. (2016), who propose that booms are built based on some good news and some expectations about future growth, and that busts reflect the reversal of these expectations.

In Figure 9, we present scatter plots of the depth of the state-level bust, as a function of the contribution of the various shocks to the boom. The figure shows that the price expectation shock explains the bust the best ($R^2 = 0.46$). The contributions of the other shocks are also positively correlated with the depth of the bust; however, the explanatory power of each is weaker (R^2s range between 0.26 and 0.33). Overall, we view these results as corroborating the model of Burnside et al. (2016) about how real estate booms develop and why they collapse.

7 Model-based vs. Survey-based Expectations

The expectation shocks that we measure were identified through variation in the fraction of vacant properties. As such, they reflect the expectations of actual home buyers and

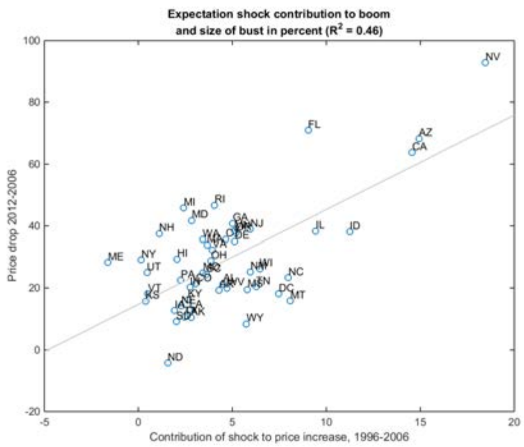

(a) Expectation Shock

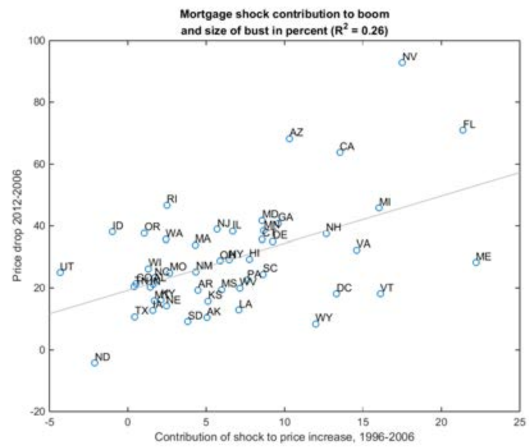

(b) Mortgage Shock

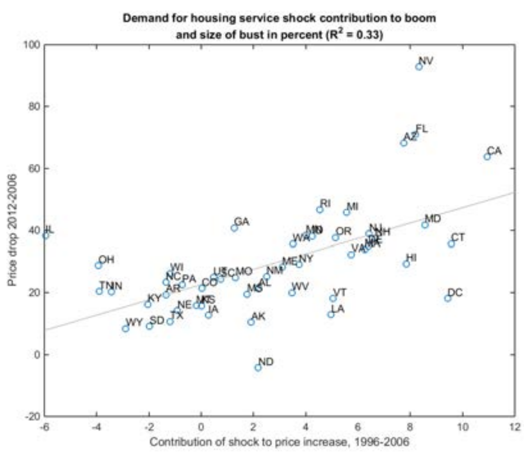

(c) Demand for Housing Services Shock

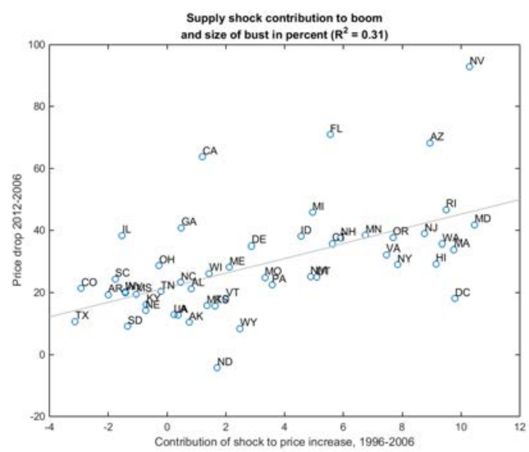

(d) Supply Shock

Figure 9. Contribution of Shocks to the Boom and the Subsequent Bust

The figure shows the state-level depth of the bust (2006–2012) as a function of the contribution of the four shocks to the evolution of the boom (1996–2006). The units on the x- and y-axes represent cumulative log differences.

investors, who are the marginal buyers of these vacant properties.

An important question is how these expectations measures fit the evidence on expectations derived from other sources. In particular, the literature that studies house price cycles draws expectations data from surveys (e.g., Piazzesi and Schneider, 2009; Lambertini et al., 2013; Cox and Ludvigson, 2018). The main source of survey expectations about housing is the Michigan Survey of Consumers, which, to our best knowledge, is the only survey of house price expectations that covers the entire boom period. The survey asks respondents to comment on whether this is generally a good time or bad time to buy a house.[27] Then, respondents are asked as an open question the reason for their answer; responses are then categorized. The following survey choices are related to expectations about future prices: "because current prices are low" and "because prices will increase."

In Figure 10, we overlay the time series of the survey responses with our measure of expectations. Figure 10, Panel (a) shows the response that it is a good time to buy because prices are expected to increase, and Panel (b) shows the response that it is a good time to buy because prices are currently low. As noted previously by Piazzesi and Schneider (2009), survey responses do not line up well with housing prices during the boom (peaking in 2006). In Figure 10, Panel (a), if we abstract from the very early and volatile part of the survey, the share of survey respondents that expect prices to increase tends to rise during housing boom phases. However, this share increased only moderately during the most recent and largest boom and only at the very end of the boom period. Furthermore, the fraction of survey respondents who perceive prices to be low (Figure 10b) seems to lag the house price cycle. In particular, an unprecedented fraction perceived prices to be low following the housing bust in 2008. Our measure of expectations, in contrast, captures very high expectation shocks during the boom years and negative expectation shocks during the bust.

One explanation for why survey responses do not align with the boom is that respondents were asked the wrong question. In particular, survey respondents are only asked indirectly

[27]See survey instrument at `https://data.sca.isr.umich.edu/fetchdoc.php?docid=24776`.

about their price expectations (i.e., the reason why it is a good time to buy); respondents may have answered differently if asked directly. However, since 2007, the survey also has asked more directly whether respondents expect prices to increase, decrease, or stay about the same (at the one-year and five-year horizon). The share of respondents that expect prices to increase is strongly correlated with the share of respondents that answer it is a good time to buy a house "because prices will increase" (0.96 for the one-year and 0.68 for the five-year measure). This correlation indicates that the "because prices will increase" measure is a reasonable proxy for survey-based price expectations and that the indirect question about price expectations is not the reason for the lack of correlation between respondents' recorded price expectations and actual price boom.[28]

Another way to reconcile the difference between model-based and survey-based expectations is to look at the difference between the surveyed population and the population our model captures. The goal of the survey is to provide a snapshot of expectations of a representative sample of consumers. Because most individuals are not actively looking to buy a home at any given time, their views have little weight in setting prices. [29] In contrast, investors in vacant homes, whose expectations we capture, participate in the market. These are people who transact on the margin; therefore, their beliefs are instrumental in setting prices.

8 Conclusion

A number of observers have suggested that shifts in house price expectations were an important driver of the U.S. house price boom in the early 2000s. Researchers have gathered evidence in favor of this hypothesis indirectly by either using deviations from benchmark models (residuals), by relying on measures of house price expectations (survey, media), or

[28]By contrast the "because price are low" is negatively correlated with the share of respondents who expect prices to increase, which indicates that it may not be an appropriate measure.

[29]Benítez-Silva, Eren, Heiland, and Jiménez-Martín (2015) provide evidence that self-assessed housing wealth depends on the time when the house was initially bought.

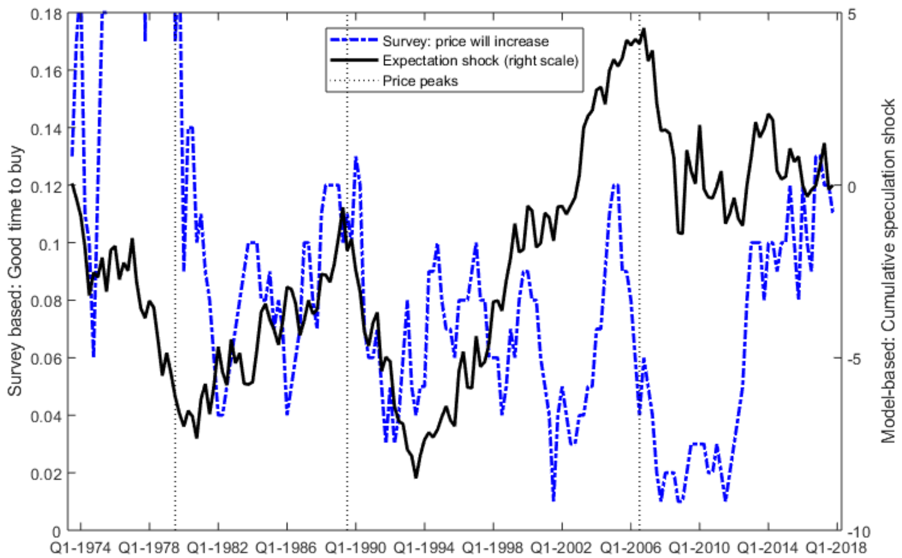

(a) Survey: “Prices Will Increase”

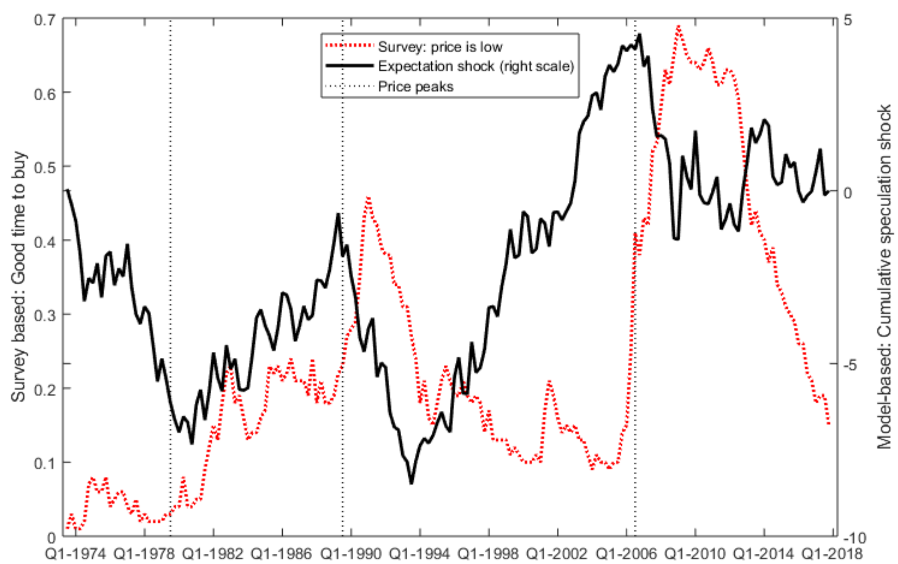

(b) Survey: “Prices Are Low”

Figure 10. Survey-based vs. Model-based Expectations

The figure compares survey-based expectations from the Michigan Survey of Consumers and expectation shocks resulting from our model. The survey expectations show the fraction of people who indicate the reasons why they think it is a good time to buy homes: future prices will increase (blue dashed line) (Panel (a)), or current prices are low (red dashed line) (Panel (b)). The cumulative expectation shock based on our model is indicated with a solid black line. The units on the y-axis represent fraction of respondents (left-hand scale) and cumulative log differences (right-hand scale).

by using localized data about speculators.

We present a novel approach to quantify the contribution of price expectations to the housing boom, based on agents' actual actions. Our main identification tool is the insight that high price expectations result in agents holding unused capacity so that they can benefit from higher future prices. This approach allows us to separate price expectation shocks from shocks to mortgage rates, shocks to the demand for housing services, and shocks to the supply of housing.

We find that the contribution of price expectation shocks to the U.S. housing boom in the 2000s was substantial. In our baseline specification, price expectation shocks explain roughly 23% of the increase, close to the the contribution explained by mortgage shocks (28%). After accounting for shocks to the demand for housing services and the supply of housing (5% and 13%, respectively), 22% of the price increase still remains unaccounted for by the four identified shocks. This finding indicates that attributing to price expectation shocks the entire residual that cannot be explained by traditional shocks will overestimate the former's contribution.

Our analysis also provides new insights into the contributions of these various shocks in the cross-section. We find that price expectation shocks were main contributors to the booms in the Sand States (Arizona, California, Florida, and Nevada). Furthermore, we observe that mortgage rate shocks were important contributors to the booms across all states. Specifically, mortgage rate shocks drove up prices in states that otherwise did not experience particularly large booms (e.g., Virginia). In regard to the shock to the demand for housing services, California's boom stands out as the one state influenced the most by it, whereas most states experienced a drag from shocks to the demand for housing services during the bust period. Densely populated states were among those most affected by the supply shock during the boom. In total, these cross-sectional observations add to the validity of our results, as they corroborate the observations of other independent research (Mayer, 2011; Chinco and Mayer, 2015; Ferreira and Gyourko, 2017).

Among the various shock contributions to the boom that we explored, positive price expectation shocks during the boom best predict the depth of the subsequent house price bust. This fact is consistent with a model of housing boom and bust by Burnside et al. (2016), suggesting that optimistic views fuel a boom, and when they are reversed, a bust occurs.

Our expectation measure produces a different picture from the one arising from commonly used survey data. As some researchers have noted (Piazzesi and Schneider, 2009; Landvoigt, 2017), the average sentiment during the boom period across the U.S. was not necessarily speculative. However, our results show that in pockets across the U.S., *marginal* investors behaved in a speculative manner, gathering excess capacity in anticipation of future price increases. Their actions were sufficient to drive up home prices during the boom period and depress them in the subsequent bust.

References

Adam, Klaus, Pei Kuang, and Albert Marcet, 2012, House price booms and the current account, *NBER Macroeconomics Annual* 26, 77–122.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, 2012, Credit supply and house prices: Evidence from mortgage market segmentation, Working paper, Duke University.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, 2016, Loan originations and defaults in the mortgage crisis: The role of the middle class, *Review of Financial Studies* 29, 1635–1670.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, 2018a, Dynamics of housing debt in the recent boom and great recession, *NBER Macroeconomics Annual* 32, 265–311.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, 2018b, Perception of house price risk and homeownership, Technical report, National Bureau of Economic Research.

Armona, Luis, Andreas Fuster, and Basit Zafar, 2019, Home price expectations and behavior: Evidence from a randomized information experiment, *Review of Economic Studies* forthcoming.

Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel, 2018, The economic effects of social networks: Evidence from the housing market, *Journal of Political Economy* 126, 2224–2276.

Bayer, Patrick, Christopher Geissler, and James W. Roberts, 2011, Speculators and middlemen: The role of intermediaries in the housing market, NBER Working Papers 16784, National Bureau of Economic Research.

Begley, Jaclene, Lara Loewenstein, and Paul S. Willen, 2019, The price-rent ratio in the boom and bust: Measurement and implications, mimeo.

Benítez-Silva, Hugo, Selcuk Eren, Frank Heiland, and Sergi Jiménez-Martín, 2015, How well do individuals predict the selling prices of their homes?, *Journal of Housing Economics* 29, 12–25.

Bernanke, Ben, 2005, The global saving glut and the US current account deficit, Working paper, Board of Governors of the Federal Reserve System (U.S.).

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer, 2017, Overreaction in macroeconomic expectations, Working paper, Harvard University.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2018, Diagnostic expectations and credit cycles, *Journal of Finance* 73, 199–227.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, 2016, Understanding booms and busts in housing markets, *Journal of Political Economy* 124, 1088–1147.

Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas, 2008, An equilibrium model of "global imbalances" and low interest rates, *American Economic Review* 98, 358–393.

Canova, Fabio, and Gianni De Nicolo, 2002, Monetary disturbances matter for business fluctuations in the G-7, *Journal of Monetary Economics* 49, 1131–1159.

Case, Karl E., and Robert J. Shiller, 2003, Is there a bubble in the housing market?, *Brookings Papers on Economic Activity* 2003, 299–342.

Case, Karl E., Robert J. Shiller, and Anne K. Thompson, 2015, What have they been thinking? Home buyer behavior in hot and cold markets–a 2014 update, Working paper, Yale University.

Chinco, Alex, 2018, The madness of the crowds and the likelihood of bubbles, Working paper, University of Illinois at Urbana-Champaign.

Chinco, Alex, and Christopher Mayer, 2015, Misinformed speculators and mispricing in the housing market, *Review of Financial Studies* 29, 486–522.

Choi, Hyun-Soo, Harrison Hong, and Jose Scheinkman, 2014, Speculating on home improvements, *Journal of Financial Economics* 111, 609–624.

Cox, Josue, and Sydney C. Ludvigson, 2018, Drivers of the great housing boom-bust: Credit conditions, beliefs, or both?, Working paper, New York University.

De Stefani, Alessia, 2017, Waves of optimism: House price history, biased expectations and credit cycles, Working paper, University of Edinburgh.

DeFusco, Anthony A., Charles G. Nathanson, and Eric Zwick, 2017, Speculative dynamics of prices and volume, Working paper, University of Chicago.

Dokko, Jane, Brian M. Doyle, Michael T. Kiley, Jinill Kim, Shane Sherlund, Jae Sim, and Skander Van Den Heuvel, 2011, Monetary policy and the global housing bubble, *Economic Policy* 26, 233–283.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl, 2019, How monetary policy shaped the housing boom, Working paper, New York University.

Duca, John V., John Muellbauer, and Anthony Murphy, 2011, House prices and credit constraints: Making sense of the us experience, *Economic Journal* 121, 533–551.

Favilukis, Jack, David Kohn, Sydney C Ludvigson, and Stijn Van Nieuwerburgh, 2012, International capital flows and house prices: Theory and evidence, in *Housing and the financial crisis*, 235–299 (University of Chicago Press).

Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh, 2017, The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium, *Journal of Political Economy* 125, 140–223.

Ferreira, Fernando V., and Joseph Gyourko, 2017, Anatomy of the beginning of the housing boom across U.S. metropolitan areas, Working paper, University of Pennsylvania.

Foote, Christopher L., Kristopher S. Gerardi, and Paul S. Willen, 2012, Why did so many people make so many ex post bad decisions? The causes of the foreclosure crisis, Working Paper 18082, National Bureau of Economic Research.

Gabriel, Stuart A., and Frank E. Nothaft, 2001, Rental housing markets, the incidence and duration of vacancy, and the natural vacancy rate, *Journal of Urban Economics* 49, 121–149.

Gao, Zhenyu, Michael Sockin, and Wei Xiong, 2017, Economic consequences of housing speculation, Working paper, Princeton University.

Glaeser, Edward L., Joshua D. Gottlieb, and Joseph Gyourko, 2012, Can cheap credit explain the housing boom?, in *Housing and the Financial Crisis*, 301–359 (University of Chicago Press).

Glaeser, Edward L., Joseph Gyourko, and Albert Saiz, 2008, Housing supply and housing bubbles, *Journal of Urban Economics* 64, 198–217.

Glaeser, Edward L., Wei Huang, Yueran Ma, and Andrei Shleifer, 2017, A real estate boom with Chinese characteristics, *Journal of Economic Perspectives* 31, 93–116.

Glaeser, Edward L., and Charles G. Nathanson, 2017, An extrapolative model of house price dynamics, *Journal of Financial Economics* 126, 147–170.

Green, Richard, and Patric H. Hendershott, 1996, Age, housing demand, and real house prices, *Regional Science and Urban Economics* 26, 465–480.

Green, Richard K, and Hyojung Lee, 2016, Age, demographics, and the demand for housing, revisited, *Regional Science and Urban Economics* 61, 86–98.

Green, Richard K., Stephen Malpezzi, and Stephen K. Mayo, 2005, Metropolitan-specific estimates of the price elasticity of supply of housing, and their sources, *American Economic Review* 95, 334–339.

Greenwald, Daniel L., 2018, The mortgage credit channel of macroeconomic transmission, Working paper, Massachusetts Institute of Technology.

Greenwald, Daniel L., and Adam Guren, 2019, Do credit conditions move house prices?, mimeo.

Griffin, John M., Samuel Kruger, and Gonzalo Maturana, 2018, What drove the 2003–2006 house price boom and subsequent collapse? Disentangling competing explanations, Working paper, University of Texas at Austin.

Gyourko, Joseph, and Albert Saiz, 2004, Reinvestment in the housing stock: The role of construction costs and the supply side, *Journal of Urban Economics* 55, 238–256.

Hamilton, James Douglas, 1994, *Time Series Analysis*, volume 2 (Princeton University Press, Princeton, N.J.).

Haughwout, Andrew, Richard W. Peach, John Sporn, and Joseph Tracy, 2012, The supply side of the housing boom and bust of the 2000s, in *Housing and the Financial Crisis*, 69–104 (University of Chicago Press).

Head, Allen, Huw Lloyd-Ellis, and Hongfei Sun, 2014, Search, liquidity, and the dynamics of house prices and construction, *American Economic Review* 104, 1172–1210.

Huang, Haifang, and Yao Tang, 2012, Residential land use regulation and the U.S. housing price cycle between 2000 and 2009, *Journal of Urban Economics* 71, 93–99.

Jarociński, Marek, and Frank Smets, 2008, House prices and the stance of monetary policy, *Federal Reserve Bank of St. Louis Review* 339–366.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti, 2019, Credit supply and the housing boom, *Journal of Political Economy* 127, 1317–1350.

Juvenal, Luciana, and Ivan Petrella, 2015, Speculation in the oil market, *Journal of Applied Econometrics* 30, 621–649.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante, 2019, The housing boom and bust: Model meets evidence, Working paper, National Bureau of Economic Research.

Kilian, Lutz, and Daniel P. Murphy, 2014, The role of inventories and speculative trading in the global market for crude oil, *Journal of Applied Econometrics* 29, 454–478.

King, Robert G., Charles I. Plosser, James H. Stock, and Mark W. Watson, 1991, Stochastic trends and economic fluctuations, *American Economic Review* 81, 819–840.

Knittel, Christopher R., and Robert S. Pindyck, 2016, The simple economics of commodity price speculation, *American Economic Journal: Macroeconomics* 8, 85–110.

Lambertini, Luisa, Caterina Mendicino, and Maria Teresa Punzi, 2013, Expectation-driven cycles in the housing market: Evidence from survey data, *Journal of Financial Stability* 9, 518–529.

Landvoigt, Tim, 2017, Housing demand during the boom: The role of expectations and credit constraints, *Review of Financial Studies* 30, 1865–1902.

Leung, Charles Ka Yui, and Chung-Yi Tse, 2012, Flippers in housing market search, in *Annual Meeting of the Society for Economic Dynamics, SED 2012*.

Levitin, Adam J., Desen Lin, and Susan M. Wachter, 2018, Mortgage risk premiums during the housing bubble, *Journal of Real Estate Finance and Economics* 1–48.

Malpezzi, Stephen, and Susan M. Wachter, 2005, The role of speculation in real estate cycles, *Journal of Real Estate Literature* 13, 143–164.

Mayer, Christopher, 2011, Housing bubbles: A survey, *Annual Review of Economics* 3, 559–577.

Mian, Atif, and Amir Sufi, 2009, The consequences of mortgage credit expansion: Evidence from the us mortgage default crisis, *Quarterly Journal of Economics* 124, 1449–1496.

Piazzesi, Monika, and Martin Schneider, 2009, Momentum traders in the housing market: Survey evidence and a search model, *American Economic Review: Papers and Proceedings* 99, 406–411.

Poterba, James M., 1984, Tax subsidies to owner-occupied housing: An asset-market approach, *Quarterly Journal of Economics* 99, 729–752.

Rosen, Kenneth T., and Lawrence B. Smith, 1983, The price-adjustment process for rental housing and the natural vacancy rate, *American Economic Review* 73, 779–786.

Rosenthal, Stuart S., 1999, Residential buildings and the cost of construction: New evidence on the efficiency of the housing market, *Review of Economics and Statistics* 81, 288–302.

Rubio-Ramírez, Juan F., Daniel F. Waggoner, and Tao Zha, 2010, Structural vector autoregressions: Theory of identification and algorithms for inference, *Review of Economic Studies* 77, 665–696.

Sá, Filipa, Pascal Towbin, and Tomasz Wieladek, 2014, Capital inflows, financial structure and housing booms, *Journal of the European Economic Association* 12, 522–546.

Sá, Filipa, and Tomasz Wieladek, 2015, Capital inflows and the U.S. housing boom, *Journal of Money, Credit and Banking* 47, 221–256.

Saiz, Albert, 2010, The geographic determinants of housing supply, *Quarterly Journal of Economics* 125, 1253–1296.

Shiller, Robert J., 2007, Understanding recent trends in house prices and home ownership, Working paper, National Bureau of Economic Research.

Taylor, John, 2007, Housing and monetary policy, Working paper, Stanford Institute for Economic Policy Research.

Topel, Robert H., and Sherwin Rosen, 1988, Housing investment in the United States, *Journal of Political Economy* 96, 718–740.

Uhlig, Harald, 2005, What are the effects of monetary policy on output? Results from an agnostic identification procedure, *Journal of Monetary Economics* 52, 381–419.

Wheaton, William C., 1990, Vacancy, search, and prices in a housing market matching model, *Journal of Political Economy* 98, 1270–1292.

Zabel, Jeffrey E., 2004, The demand for housing services, *Journal of Housing Economics* 13, 16–35.