Speculative Dynamics of Prices and Volume*

Anthony A. DeFusco† **Charles G. Nathanson**‡ **Eric Zwick**§

February 4, 2017

Abstract

We present a dynamic theory of prices and volume in asset bubbles. In our framework, predictable price increases endogenously attract short-term investors more strongly than long-term investors. Short-term investors amplify volume by selling more frequently, and they destabilize prices through positive feedback. Our model predicts a lead-lag relationship between volume and prices that we confirm in the 2000-2008 US housing bubble. Using data on 50 million home sales from this episode, we document that much of the variation in volume arose from the rise and fall in short-term investment.

*We thank Stefano Giglio, Edward Glaeser, Sam Hanson, Stijn Van Nieuwerburgh, Alp Simsek, Johannes Stroebel, Richard Thaler, and Rob Vishny for helpful comments. Jessica Henderson, Saul Ioffie, Harshil Sahai, and Iris Song provided excellent research assistance. DeFusco and Nathanson thank the Guthrie Center for Real Estate Research for financial support, and Zwick gratefully acknowledges financial support from the Fama Miller Center and Booth School of Business at the University of Chicago.

†Kellogg School of Management, Northwestern University.

‡Kellogg School of Management, Northwestern University.

§Booth School of Business, University of Chicago and NBER.

The role of speculation in driving asset prices has long been a topic of debate among economists (Keynes, 1936; Fama, 1970; Shiller, 1981; Black, 1986).[1] Modern empirical work on "speculative dynamics" begins with Cutler et al. (1991), who document short-run momentum and long-run reversals in the prices of many diverse assets. These patterns are especially strong during asset bubbles, which have drawn attention due to their frenzied activity and subsequent social costs (Kindleberger, 1978; Shiller, 2005; Glaeser, 2013). Several distinct theories have been offered to explain these asset pricing facts.[2]

This paper explores a less studied feature of asset bubbles—the speculative dynamics of volume. Large movements in transaction volume consistently accompany price cycles (Stein, 1995; Genesove and Mayer, 2001; Hong and Stein, 2007), yet many theories of bubbles ignore implications for volume. As Cochrane (2011) writes:

> Every asset price "bubble"... has coincided with a similar trading frenzy, from Dutch tulips in 1620 to Miami condos in 2006... Is this a coincidence? Do prices rise and fall for other reasons, and large trading volume follows, with no effect on price? Or is the high price... explained at least in part by the huge volume... To make this a deep theory, we must answer why people trade so much.

Figure 1 plots time series patterns in prices and volume for four distinct bubble episodes: the 2000-2012 US housing market, the 1995-2005 market in technology stocks, the experimental bubbles studied by Smith et al. (1988), and the 1985-1995 Japanese stock market. During these episodes, prices and volume comove strongly. The figures also reveal a more nuanced feature of the data: in each case, volume peaks well before prices. Improving our understanding of these phenomena requires focus on this complex, joint dynamics of prices and volume.

To take up this challenge, we first present a simple model of the joint speculative dynamics of prices and volume during bubbles. Following past work, the model features extrapolative expectations: investors expect prices to increase after past increases.[3] The model departs from past work in two ways. First, instead of the standard dichotomy between feedback traders and rational arbitrageurs, investors differ only in their expected investment horizon.

[1] Harrison and Kreps (1978, p. 323) define speculation in the following way: "Investors exhibit speculative behavior if the right to resell a stock makes them willing to pay more for it than they would pay if obliged to hold it forever."

[2] These theories include Cutler et al. (1990), De Long et al. (1990), Barberis et al. (1998), Daniel et al. (1998), Hong and Stein (1999), Abreu and Brunnermeier (2003), Pástor and Veronesi (2006), Pástor and Veronesi (2009), Piazzesi and Schneider (2009), and Burnside et al. (2016).

[3] Key models with extrapolative expectations include Cutler et al. (1990), De Long et al. (1990), Hong and Stein (1999), Barberis and Shleifer (2003), Barberis et al. (2015), and Glaeser and Nathanson (2016).

Some buyers plan to sell after one year, while others plan to hold for many years. The second departure is to specify a term structure for extrapolation. In particular, extrapolation declines with the forecast horizon, so that short-run expectations display more sensitivity than long-run expectations to past prices. When this term structure holds, past price growth disproportionately attracts short-horizon investors, who in turn generate excess volume when they sell.

These modeling innovations benefit from much empirical support. Prior work estimating extrapolative future expectations finds that short-run expectations display more sensitivity to past price changes than long-run expectations (Graham and Harvey, 2003; Vissing-Jorgensen, 2004; Armona et al., 2016). In survey evidence from the National Association of Realtors, expected holding times vary considerably across buyers in the housing market. Furthermore, the share of respondents reporting an expected holding time less than 3 years comoves strongly with recent house price appreciation.

We model a housing market populated by extrapolative investors with heterogeneous horizons. Due to the generality of the stylized facts presented above, we abstract from several special features of the housing market, such as debt and new construction, in order to ease comparison to other asset markets. We focus on the housing market because data availability allows us to test directly the model's predictions about the composition of buyers and sellers over the course of a bubble episode. In the model, potential buyers arrive each instant and decide whether to buy a house. If they buy, they must hold the house for some period, after which they are free to sell. The expected duration of this period and the flow utility received during it vary across potential buyers. Potential buyers expect to sell immediately upon the period's expiration at the prevailing price. As in prior work (e.g., Hong and Stein (1999)), the asset price reacts sluggishly to changes in the number of potential buyers who wish to buy. This "price stickiness" combines with extrapolative expectations to generate positive feedback between price growth and demand that causes bubbles.

The model is analytically tractable, allowing us to characterize how prices, volume, and the composition of market participants respond to a one-time demand shock. Heterogeneity in expected holding periods interacts with extrapolative expectations to produce a rich volume dynamic. We partition the time following the shock into three epochs: a *boom* in which prices rise, volume rises, and all listings sell; a *quiet* in which prices continue to rise, but volume falls and unsold listings accumulate; and a *bust* in which prices fall on low volume.

This partition implies a lead-lag relationship between volume and prices that is consistent with the data in Figure 1 and absent from other theories in which trading volume comoves with prices.[4] In our model, the composition of buyers and sellers varies over the cycle because past buyers become future sellers. This time-varying investor composition links volume and prices across periods. In particular, volume begins to fall during the quiet when prices become so high that demand from long-horizon buyers falls and the large number of homes listed by previous short-term buyers do not sell. Prices, however, continue to rise as a result of near-term expected capital gains that continue to draw in short-term investors.

The prevalence of short-term investors is also important for the magnitude of the price and volume booms. In particular, we prove that the booms in prices and volume are larger when the frequency of potential buyers with short horizons is greater. This finding ties together prices and volume during bubbles by demonstrating that a single factor is responsible for movements in both. Using the empirical literature on extrapolative expectations and the survey evidence on expected holding times, we calibrate our model and find that the marginal effect of short-term potential buyers on the price and volume booms is quantitatively large and first-order relevant for explaining aggregate price and volume dynamics.

The second part of the paper documents new facts on the composition of buyers and sellers during the recent US housing bubble using transaction-level data between 1995 and 2014 for 115 cities that represent approximately 50% of the US population. While we organize the analysis around the model's key predictions, these facts provide a model-free depiction of the joint dynamics of prices and volume that highlights the role of speculation in driving the bubble.

We present four key facts, two about the composition of aggregate volume and two about the joint price-volume relationship. First, much of the variation in volume over time and across markets during this period arises from short-term investment. In the national time series, 42% of the rise in home sales between 2000 and 2005 can be attributed to the increase in sales of homes held for less than 3 years. Second, an equally large share of the increase in sales during this period came from buyers who did not occupy the property. Self-identified "investors" also report shorter expected holding times in surveys. These findings are consistent with the model's prediction that buyers with low flow utility from housing

[4]Alternative theories that generate comovement without a clear lead-lag relationship have focused on credit constraints (Stein, 1995; Ortalo-Magné and Rady, 2006), loss aversion (Genesove and Mayer, 2001), or disagreement (Scheinkman and Xiong, 2003; Piazzesi and Schneider, 2009).

respond most strongly to expected capital gains. The aggregate facts also describe the cross-section, as short-term and non-owner occupant buyers disproportionately influence volume across cities and across ZIP codes within cities. Third, the lead-lag relationship between prices and volume is robust not only at the aggregate level but also across cities, with the correlation maximized at a 15-month lag. Last, the 2000–2012 US house price cycle was larger in metropolitan areas in which the level of existing sales as a share of the housing stock was greater. As shown in our model, a higher frequency of short-term buyers increases both steady-state volume and the amplitude of the price response to the demand shock.

Related Literature. Our theoretical approach is most similar to a strand of papers beginning with Cutler et al. (1990) and De Long et al. (1990), which use the interaction of fundamental and feedback traders to explain asset price fluctuations. Hong and Stein (1999) offer a model of price underreaction and overshooting that provides similar intuition to the mechanism in our model. They explicitly model the slow diffusion of news about future dividends in order to microfound price stickiness. However, this literature does not consider the joint dynamics of volume and prices, nor the role of heterogeneous holding times.

Whether speculators anticipate wanting to sell in the near future distinguishes our approach from many in the literature. In particular, our approach is distinct from but complementary to models in which disagreement begets volume in financial markets (Hong and Stein, 2007; Simsek, 2013; Daniel and Hirshleifer, 2015). Market participants in our model always agree on the future price path of the asset, with differences in their willingness-to-pay coming only from their investment horizons. The seminal disagreement papers (Harrison and Kreps, 1978; Scheinkman and Xiong, 2003) link prices and volume through comparative statics and hence cannot explain the dynamics of volume over the course of a bubble episode in a single asset without also assuming that the extent of disagreement exogenously rises and falls during the bubble.

Burnside et al. (2016) present an epidemiological model in which disagreement may endogenously vary during a bubble and raise transaction volume. This mechanism does not produce the variation in expected holding periods present in the data. Our focus on *ex ante* short-term investors also contrasts with the disagreement model of Barberis et al. (2016), in which volume during bubbles involves investors whose beliefs unexpectedly change through a process called "wavering."[5] While disagreement is likely an important part of why there is so much trade in financial markets, our focus on holding periods offers a clear and measurable alternative way of explaining several of the most quantitatively important drivers of volume.[6]

The paper is related to a large literature beginning with Wheaton (1990) that uses search and matching models to understand housing market dynamics. In these models, the housing market is typically a closed system and the joint buyer-seller problem drives equilibrium dynamics. We show that the extensive margin of speculators is a crucial driver of volume. Moreover, we confirm the findings in Anenberg and Bayer (2013) that fewer than half of sellers return to the same market to buy homes around the time they sell. In our model and in the data, volume rises during the boom because owners are more likely to list their houses, whereas volume falls during the quiet because listings take longer to sell. Documented also by Ngai and Sheedy (2016), the apparent importance of the propensity to list and irrelevance of "time on the market" for explaining volume cycles matches our model but does not fit most search models, which rely heavily on fluctuations in time on the market.[7] Relative to our paper, search and matching models are often less concerned with price determination, usually set through Nash bargaining with simple aggregation rules. In these models, prices typically react quickly to shocks and do not exhibit the autocorrelation present in the data.[8]

Our empirical findings contribute to a growing empirical literature documenting the significant role of speculators in the 2000s US housing bubble. Haughwout et al. (2011) show that a third of purchase mortgages in 2006 were originated to buyers with first liens on multiple properties, Bhutta (2015) shows that inflows from real estate investors tripled during the boom, and Chinco and Mayer (2016) calculate that purchases of homes in Las Vegas, NV,

[5]Piazzesi and Schneider (2009) also present a model in which investor beliefs unexpectedly change. In their setting, optimistic homebuyers unexpectedly become pessimists upon buying a house, which produces short-term trading volume. See also Penasse and Renneboog (2016) for a theory of bubble volume based on exogenous movements in disagreement.

[6]Bailey et al. (2016) present empirical evidence that housing volume is higher in U.S. counties with larger disagreement about future house price growth. They measure an individual's house price growth expectations using the past house price growth experienced by the people in that individual's social network.

[7]See, for example, Díaz and Jerez (2013), Head et al. (2014), and Burnside et al. (2016). Ngai and Tenreyro (2014) and Ngai and Sheedy (2016) develop models that account for some of the cyclical joint dynamics of prices, volume, and listings, but do not generate overshooting in prices or the lead-lag relationship. Ngai and Sheedy (2016) argue the failure to generate overshooting and the lead-lag relationship likely owes to the omission of speculative forces.

[8]See, e.g., Caplin and Leahy (2011), Díaz and Jerez (2013), Head et al. (2014), and Guren (2016) for discussions of this point and proposed solutions. Caplin and Leahy (2011) propose a solution based on the slow diffusion of information that is similar in spirit to Hong and Stein (1999) and our microfoundation.

by people living outside of the metropolitan area exceeded 5% of Las Vegas's gross domestic product in 2004. Similarly, Bayer et al. (2015) estimate a three-fold increase during the boom in the share of Los Angeles, CA, home purchases resold within two years. Nathanson and Zwick (2015) document speculative behavior by builders and developers in land markets. Our work is unique in its focus on the aggregate implications of speculative behavior across a large set of cities and adds a new measure based on realized holding periods.[9]

1 Motivating Evidence

1.1 The Term Structure of Extrapolative Expectations

Much of the early theoretical work on extrapolative expectations is silent on the forward term structure of extrapolation. The two papers we are aware of that explicitly model how past price changes are extrapolated into expectations of future prices at varying horizons are Barberis et al. (2015) and Glaeser and Nathanson (2016).[10] In both papers, extrapolation is modeled in a way that leads short term expectations to exhibit more sensitivity to recent price changes than long term expectations. This approach, which we adopt in our model, is supported by a growing body of empirical evidence suggesting that past asset returns do indeed influence annualized expected capital gains more strongly over short versus long future horizons.

In the housing market, Armona et al. (2016) survey expected capital gains over 1- and 5-year horizons and relate those expectations to perceptions of recent local price changes. They find that 1-year ahead expectations are nearly five times more sensitive to perceived past price changes than annualized 2-5 year ahead expectations.[11] Furthermore, when provided with

[9] A complementary literature documents a large increase in the availability of mortgage credit during the early 2000s and argues that this expansion of credit caused the boom and subsequent bust in house prices (e.g., Mian and Sufi, 2009; Favilukis et al., 2015). Gao et al. (2016) find that speculation was higher in states without capital gains taxes.

[10] Hong and Stein (1999) assume that past price changes are extrapolated linearly into a fixed future horizon, but their model does not deliver predictions about the effects of recent price changes on expectations at differing future horizons.

[11] The coefficient when regressing 1-year expectations on past 1-year price growth perceptions equals 0.262 (0.029), whereas the coefficient when regressing the implied 2-5 year annualized expectations on 1-year perceptions equals 0.058 (0.012). Case et al. (2012) conduct a similar survey and report the 1- and 10-year annualized capital gains expectations for the housing market at the city-year level. Using CoreLogic county house price indices and their survey data, we find a coefficient of 0.25 (0.02) when regressing 1-year expectations on the past year's house price appreciation and a coefficient of 0.16 (0.04) when regressing the annualized 10-year expectations on the past year's house price appreciation.

new information about local changes in house prices over the last year, respondents in the survey update their forecasts of 1-year price gains more strongly than their 2-5 year forecasts. Both of these facts suggests that short-run house price expectations display significantly more sensitivity to past returns than do long-run expectations.

Similar evidence exists for the US stock market. Vissing-Jorgensen (2004) reports the average expectation of annualized stock market returns over 1- and 10-year horizons among respondents to the UBS/Gallup Index of Investor Optimism survey between 1998 and 2002. Over this period, 1-year expectations moved closely with recent price changes—first rising from 10% to 16% as stock prices increased, and then falling to 6% as prices fell. In contrast, 10-year expectations remained relatively constant over this period and were uncorrelated with the large contemporaneous movements in the stock market.

These patterns persist even in a sample of more sophisticated survey respondents. The Duke CFO Global Business Outlook, which surveys chief financial officers of US firms, provides data on annualized 1- and 10-year stock return expectations. Graham and Harvey (2003) use data from the 2000-2003 waves of this survey and find that the 1-year expected risk premium (expected return less treasury yield) is positively and significantly related to excess returns over the previous week, month, two months, and quarter, whereas the 10-year annualized expected risk premium is slightly negatively related to these past returns. In Appendix Table A1, we use the survey data from 2000-2011 and confirm that the 1-year expectations remain more sensitive to past returns than the 10-year expectations in the longer sample.[12] Thus, the available evidence all points toward a term structure for extrapolation in which short-run forecasts are more sensitive to recent prices changes than long-run forecasts.

1.2 Variation in Expected Holding Times

In the presence of a downward sloping term structure for extrapolative expectations, our model implies that recent price changes will differentially draw in short-term investors who amplify volume by selling more frequently and destabilize prices through positive feedback. The magnitude of these effects will depend on the degree of heterogeneity in the distribution

[12]Although the data are available from 2000-2016, we use only 2000-2011 to match the window used by Greenwood and Shleifer (2014), who also find a coefficient of about 0.03 when regressing the 1-year return expectation on the prior year's return. Interestingly, the coefficients decline considerably when the 2012-2016 sample is included, possibly because declines in interest rates over this time both increased lagged returns and decreased future return expectations.

of expected holding times among prospective investors. While not much data are available concerning the expected holding times of investors, the best data we are aware of, which come from the housing market, suggest that investment horizons vary considerably across individuals and commove strongly with recent price changes.

Each March, as part of the Investment and Vacation Home Buyers Survey, the National Association of Realtors (NAR) surveys a nationally representative sample of around 2,000 individuals who purchased a home in the previous year. The survey asks respondents to report the type of home purchased (investment property, primary residence, or vacation property) as well as the "length of time [the] buyer plans to own [the] property." Data on expected holding times and the share of purchases of each type are available for 2006-2015 (2008-2015 for primary residences).

Figure 2(a) documents the substantial cross-sectional heterogeneity in expected holding times among respondents to the survey.[13] Each bar reports an equal-weighted average of the share of recent buyers who report a given expected holding time across survey years. Averages are reported separately by property type. Two things stand out. First, the vast majority of recent homebuyers (roughly 80%) report knowing what their expected holding time will be. Second, there is wide variation in expected holding times among those who report. About half of the expected holding times are between 0 and 11 years and are distributed somewhat uniformly over that range. The survey question groups the remaining half of the responses into a single expected holding time of greater than or equal to 11 years; however, there may be substantial variation within that group as well. Expected holding times also vary in an intuitive way across property types. Recent buyers of investment properties report substantially shorter expected holding periods than recent buyers of primary residences or vacation homes.

There is also significant variation in the time series. To demonstrate this, we construct a "short-term buyer share," which is measured as the fraction of respondents who report an expected holding time of less than 3 years or had already sold their property by the time of the survey.[14] Across survey years, the short-term buyer share varies from 26% to 41% for investment properties, from 10% to 22% for primary residences, and from 12% to 34%

[13] The bins in the figure are those used by the NAR in its data release (we do not have access to less aggregated data). We reclassify respondents who have already sold their properties as having an expected holding time in [0,1).

[14] In constructing this measure, we leave out those who report that they do not know their expected holding time.

for vacation properties. The weighted average of the short-term buyer share across property types varies from 13% to 26%.

This variation over time is not random. As shown in Figure 2(b), the short-term buyer share moves closely with recent price appreciation in the housing market.[15] A simple regression of the pooled short term buyer share on the equal-weighted average year-over-year change in the nominal quarterly FHFA US house price index during the survey year yields a statistically significant coefficient estimate of 0.81. This implies that a recent nominal gain of 10% in house prices is associated with an increase in the short-term buyer share of roughly 8.1 percentage points. The nominal house price appreciation in the US in 2005 was equal to 11% and was much larger in some metropolitan areas. Thus, changes in house prices over the last cycle may have induced significant shifts in the distribution of expected holding times among homebuyers at different points in the cycle. We will investigate this hypothesis below using detailed micro-data that allows us to construct disaggreagated measures of holding times across markets and time periods. Before doing so, however, we first present our model, which is motivated by the evidence just discussed.

2 A Model of Investors with Heterogeneous Horizons

2.1 Primitives and Information Environment

We present an infinite-horizon, continuous-time model of a city with a fixed amount of perfectly durable housing, normalized to have measure one.[16] Agents go through a life cycle with three possible phases: potential buyer, stayer, and mover.

In each instant, agents arrive. Each agent begins as a potential buyer and must decide between buying a home immediately and leaving the city forever. A potential buyer who buys a home becomes a stayer. Stayers receive flow utility $\delta > 0$ from living in the city until receiving an idiosyncratic taste shock and becoming movers. This taste shock arrives with an instantaneous Poisson hazard $\lambda > 0$, which is distributed across potential buyers independently from δ and according to a time-invariant probability density function $f(\lambda)$. For any $\delta_0 > 0$, the measure of agents arriving at t for whom $\delta \geq \delta_0$ equals $A_t \delta_0^{-\epsilon}$, where

[15] Appendix Figure A1 shows that interest in "house flipping," an explicitly short-term investment strategy, has moved closely with the housing cycle in the US.

[16] Although we use terminology fitting the housing market, our model applies to any asset market.

$\epsilon > 0$ and $\int_0^\infty \lambda^\epsilon f(\lambda) d\lambda$ exists.[17]

Potential buyers and movers maximize the present value of lifetime utility by choosing whether to buy or list, respectively. Stayers do not sell their homes until becoming movers, at which point they choose whether to list their homes for sale at the current price P_t. Until selling, movers decide at each instant whether to list at the current price. Total flow utility is linear in consumption and in the flow benefit from living in the city, which is nonzero only for stayers. All agents may borrow or lend at the common discount rate r.

This environment contains three features common in the housing search literature: a lockup period in which homeowners do not sell, a Poisson hazard of the expiration of this lockup period, and a loss of housing flow benefits upon lockup expiration (Wheaton, 1990; Caplin and Leahy, 2011; Burnside et al., 2016). In contrast to the existing literature, our model allows this Poisson hazard to differ across agents. This variation captures *ex ante* specialization among agents in their investment horizons. For instance, short-term housing investors specialize in renovating houses while long-term housing investors specialize in property maintenance and tenant management. In the stock market, short-term investors are "high-frequency traders" who specialize in order-flow information while long-term investors specialize in learning about the balance sheets and growth opportunities of individual firms.

In addition to their individual types and the current price, all agents observe summary information about the complete history of prices. In particular, there exists a function $\omega(\cdot) \to \mathbb{R}$ that maps the history of prices into a single factor observed by market participants at time t. We denote this factor by $\omega_t \equiv \omega(\{P_{t'} \mid t' \leq t\})$. Potential buyers and movers use ω_t to form expectations of future prices, which govern their decisions of whether to buy or list, respectively. Agents form expectations regarding price growth between time t and $t + \tau$ in a manner that is consistent with the following assumption:

Assumption 1. *There exist functions γ and g such that for all $\delta, \lambda, \tau \geq 0$ and $P_t, \omega_t \in \mathbb{R}$*

$$\mathrm{E}[P_{t+\tau}/P_t \mid \delta, \lambda, P_t, \omega_t] = 1 + \gamma(\omega_t) g(\tau) \tag{1}$$

and the following properties hold:

[17] This heterogeneity in δ is necessary only to produce the demand curve in Lemma 3 that is a smoothly decreasing function of the current price. We believe that a similar demand curve would hold in a model with risk-averse agents and homogeneous δ. This alternate model may describe the stock market better because dividends are the same for all owners of a stock.

(a) $(g(\tau)/\tau)' < 0$ *for all* $\tau > 0$*;*

(b) $g(0) = 0$*;*

(c) $\gamma(\omega)g'(0) \leq r$ *for all* $\omega \in \mathbb{R}$*; and*

(d) $\int_0^\infty e^{-r'\tau} g(\tau) d\tau > 0$ *for all* $r' > r$*.*

Assumption 1(a) endows agents with extrapolative expectations that satisfy the empirical evidence on the forward term structure presented in Section 1. The decrease of $g(\tau)/\tau$ is necessary and sufficient for nontrivial increases in ω_t to raise short-term expected capital gains more strongly than long-term expected capital gains:[18]

Lemma 1. *Given* (1), *Assumption 1(a) holds if and only if*

$$\frac{\partial^2}{\partial\tau\partial\omega_t} \mathrm{E}\left[\frac{P_{t+\tau} - P_t}{\tau P_t} \mid \omega_t\right] < 0$$

for all $\tau > 0$ *and* $\omega_t \in \mathbb{R}$ *such that* $\gamma'(\omega_t) > 0$*.*

Assumption 1(b) imposes the weak constraint that $\mathrm{E}[P_{t+\tau}/P_t] = 1$ for $\tau = 0$. Assumption 1(c) is necessary and sufficient for the expected growth rate of prices to always fall below r:

Lemma 2. *Given* (1) *and Assumptions 1(a) and 1(b),* $\mathrm{E}[P_{t+\tau}/P_t \mid \omega_t] < e^{r\tau}$ *for all* $\tau > 0$ *and* $\omega_t \in \mathbb{R}$ *if and only if Assumption 1(c) holds.*

Finally, Assumption 1(d) guarantees that an increase to $\gamma(\omega_t)$ raises the present value of expected capital gains for all potential buyers.

2.2 Equilibrium Quantities

Solving the model requires knowing the number of agents of each type at each point in time, as well as the stock of previous listings that did not sell. In particular, we track the number of potential buyers who decide to buy D_t and the number of stayers S_t. These jointly determine the flow of listings L_t, the inventory of unsold listings I_t, and sales volume V_t.

Due to Lemma 2, $P_t > e^{-r\tau}\mathrm{E}[P_{t+\tau} \mid P_t, \omega_t]$ for all t and $\tau > 0$, so movers always choose to list their homes for sale. As a result, we may describe the evolution of listings using the

[18] Appendix A contains the proof of Lemma 1 as well as all other omitted proofs.

stock of stayers of each type λ, or $S_t(\lambda)$. Because all new movers list, the flow of listings is simply the total number of stayers receiving idiosyncratic mover shocks:

$$L_t = \int_0^\infty \lambda S_t(\lambda) d\lambda. \tag{2}$$

Given A_t, P_t, and ω_t, a measure D_t of potential buyers decide to buy. If D_t is less than the number of homes listed for sale, then all interested potential buyers are able to buy. Otherwise, homes are rationed randomly among the set of interested potential buyers. Let $D_t(\lambda)$ denote the measure of potential buyers of type λ who decide to buy. Then sales volume to potential buyers of type λ equals

$$V_t(\lambda) = \begin{cases} D_t(\lambda) & \text{if } L_t > D_t \text{ or } \quad I_t > 0 \\ \frac{D_t(\lambda)}{D_t} L_t & \text{if } L_t \leq D_t \text{ and } \quad I_t = 0, \end{cases} \tag{3}$$

where I_t is the inventory of unsold listings. The stayer stocks for each λ evolve according to the law of motion

$$\dot{S}_t(\lambda) = V_t(\lambda) - \lambda S_t(\lambda), \tag{4}$$

which is the number of new buyers of type λ less the number of stayers of type λ who become movers. Unsold listings follow the law of motion

$$\dot{I}_t = L_t - V_t, \tag{5}$$

where $V_t = \int_0^\infty V_t(\lambda) d\lambda$ equals total sales.

These expressions have two key implications for the dynamics of volume. First, volume today depends on volume before, as past buyers become current sellers. Second, the number of listings today depends both on the level of past volume and on the expected holding periods among past buyers. The larger the number of past buyers with short horizons, the larger the flow of current listings.

2.3 The Composition of Buyers

The interesting dynamics in the model concern how the composition of buyers—specifically, the composition of expected holding periods—varies over the cycle. This composition depends on the distribution of λ among buyers. The probability density function of λ among buyers at time t is given by the function $V_t(\lambda)/V_t$. By (3), this function coincides with $D_t(\lambda)/D_t$, the probability density function of demand across potential buyers at t. Thus to understand how the composition of buyers varies over time, we must calculate the distribution of demand across λ.

To derive $D_t(\lambda)$, we assume that potential buyers believe that they will sell their house as soon as they list it:[19]

Assumption 2. *Potential buyers believe that listing movers sell instantaneously.*

One consequence of Assumption 2 is that the expected holding time of a buyer of type λ equals $1/\lambda$. The other consequence is that a potential buyer of type (λ, δ) tries to buy if the current price is below the present discounted value of the flow utility she would receive as a stayer plus the expected resale value of the house at the time she anticipates becoming a mover. That is, a potential buyer tries to buy if and only if

$$P_t \leq \int_0^\infty \lambda e^{-\lambda\tau} \left(\int_0^\tau e^{-r\tau'} \delta d\tau' + e^{-r\tau} \mathrm{E}_t P_{t+\tau} \right) d\tau. \tag{6}$$

The value of δ at which a buyer of type λ is indifferent implies the equation for demand given in Lemma 3.

Lemma 3. *Demand from each λ type equals*

$$D_t(\lambda) = \underbrace{f(\lambda)A_t}_{\substack{\text{potential}\\ \text{buyer}\\ \text{measure}}} \times \underbrace{(rP_t)^{-\epsilon}}_{\substack{\text{fundamental}\\ \text{demand}}} \times \underbrace{\Sigma_\lambda(\omega_t)}_{\substack{\text{speculative}\\ \text{demand}}},$$

where

$$\Sigma_\lambda(\omega_t) = \left(1 - \frac{\lambda\gamma(\omega_t)}{r} \int_0^\infty e^{-(r+\lambda)\tau} g'(\tau) d\tau \right)^{-\epsilon}. \tag{7}$$

Demand for each λ is composed of three terms. The first term $f(\lambda)A_t$ equals the relative measure of potential buyers of type λ. The second term, which we call "fundamental demand," is a decreasing function of current prices with constant elasticity ϵ. This term reflects the relationship between demand and prices were prices to remain permanently at their current level. The third term, which we call "speculative demand," links current demand to

[19] Appendix B microfounds Assumption 2 by reducing it to other behavioral rigidities.

expected capital gains. If buyers expect prices to remain constant, then $\Sigma_\lambda(\omega_t) \equiv 1$ and speculative demand does not magnify total demand. Otherwise, speculative demand magnifies total demand when buyers expect capital gains and attenuates total demand when buyers expect capital losses, with the force of this multiplier depending on the buyer's horizon.

Only speculative demand matters for variation in the composition of buyers over time, as the other two demand components are always proportional to each other across λ types. As a consequence, variation in the composition in buyers depends entirely on changes in expected capital gains. Proposition 1 formally states this relationship:

Proposition 1. *At any ω_t such that $\gamma'(\omega_t) > 0$, the following hold for all $\lambda > 0$.*

(a) Expected capital gains increase demand from all types:

$$\frac{\partial \log D_t(\lambda)}{\partial \omega_t} > 0.$$

(b) Short-term (higher λ) buyers are more sensitive to expected capital gains:

$$\frac{\partial^2 \log D_t(\lambda)}{\partial \lambda \partial \omega_t} > 0.$$

(c) Expected capital gains skew the composition of buyers shorter-term:

$$\frac{\partial \int_0^\lambda V_t(\lambda')d\lambda'/V_t}{\partial \omega_t} \leq 0,$$

with equality if and only if $|\operatorname{supp} f| = 1$.

Part (a) formally states that greater expected capital gains increase demand. This effect appears in any user cost model of housing (e.g. Poterba, 1984). The focus in most user cost models is primarily the intensive margin demand for housing capital. Our model highlights the extensive margin instead, as the stimulative effect of expectations on demand operates entirely through drawing new buyers into the market. Consistent with this mechanism, Agarwal et al. (2015) document increased participation in the owner-occupied housing market in response to rising prices.

The extensive margin effect also predicts that expected capital gains stimulate demand more for buyers with low flow utility. Buyers with high flow utility are off the margin, so variation in ω does not impact their decision to buy. Consistent with that observation, in

Section 4 we document a large increase in investor participation in the US housing market during the 2000-2006 boom. Investors receive less flow utility than owner-occupants because in a competitive market, the rent received by landlords reflects the flow utility of the marginal renter; in contrast, most owner-occupants are infra-marginal. Investors also receive less flow utility than owner-occupants due to frictions that arise from the separation of ownership of control (Nathanson and Zwick, 2015).

Part (b) shows that the stimulative effect of capital gains is stronger for buyers with shorter expected holding times. Buyers looking to make a "quick buck" are drawn to rising prices more than those buying for the long run. The proof of Proposition 1 shows that this effect follows from the higher sensitivity of short-term expectations to ω_t embodied by Assumption 1(a).

Part (c) links expected capital gains to the composition of buyers. Because short horizon buyers are more sensitive to capital gains, an increase in expected capital gains skews the composition of buyers towards those with shorter holding times. Proposition 1(c) explains the evidence presented in Section 1 that both expected capital gains and the short-term buyer share respond strongly to recent home price appreciation.

Proposition 1 illustrates the key mechanism generating time-variation in volume in the model, namely, time-variation in the composition of holding periods for buyers and sellers driven by time-variation in expected returns. To close the model, we now specify how prices are determined.

2.4 Equilibrium Prices

As pointed out by Barberis et al. (2015), extrapolative expectations do not generate protracted booms and busts in asset prices without some additional source of sluggish adjustment. As the goal of this paper is to study volume dynamics during an extended price boom and bust, we specify a sluggish price adjustment process rather than assuming a Walrasian market in which prices equate the number of interested buyers and sellers at each instant.[20]

The log price $p_t = \log P_t$ changes according to the rule

$$\dot{p}_t = c \log(D_t/\overline{D}), \tag{8}$$

[20] An alternative source of sluggish adjustment is for agents to base extrapolative expectations on only lagged prices and not the current price (Glaeser and Nathanson, 2016; Barberis et al., 2016).

where $c > 0$ is a constant determining the speed of price adjustment, and $\overline{D}$ is given by

$$\overline{D} = \left(\int_0^\infty \lambda^{-1} f(\lambda) d\lambda \right)^{-1}. \tag{9}$$

As shown in Appendix A, (9) gives the unique value of $\overline{D}$ such that a steady state exists in which all listings sell instantaneously and prices are expected to and do remain constant. The rule in (8) brings prices towards such a steady state in a slow and predictable manner.

Appendix B microfounds (8) as the equilibrium outcome of a simple stock-flow matching model. In this microfoundation, some sellers are inattentive as in Hong and Stein (1999), Mankiw and Reis (2002), and Guren (2016) and do not adjust their listing prices in response to abnormal levels of demand, while other sellers observe current demand and choose their listing prices in response to it. The constant c is larger when the share of attentive sellers is higher and the demand elasticity ϵ is lower. In equilibrium, the allocation of listed houses to potential buyers coincides with the rule assumed in (3).

3 The Joint Dynamics of Prices and Volume

We use the model in Section 2 to explore the joint dynamics of prices and volume over the course of a boom-bust cycle. We provide a series of propositions that allow for a qualitative characterization of the relationship between prices and volume over the cycle. We then turn to a calibration of the model which allows us to study the potential quantitative relevance of the factors driving the model.

We focus on how volume and prices respond to a one-time permanent demand shock. In particular, we study an impulse response around the unique steady state that results from a single positive shock to the number of potential buyers, A_t, at a time we normalize to $t = 0$. Specifically, A_t follows the path where $A_t = A^i$ for $t < 0$ and $A_t = A^f > A^i$ for $t \geq 0$.

To characterize how prices and volume respond to this shock, we impose additional structure on how agents form expectations. We specify both the information that agents have available to them when they form forecasts of future prices and how that information influences these forecasts, which are governed in the model by $\omega(\cdot)$ and $\gamma(\cdot)$, respectively.

Following Barberis et al. (2015), we assume that agents observe only a weighted average of past price changes:

$$\omega_t = \int_{-\infty}^{t} \mu e^{-\mu(t-\tau)} \dot{p}_\tau d\tau, \tag{10}$$

where the parameter $\mu > 0$ measures the relative weight put on more recent price changes. To study dynamics, it is useful to know how this average changes in response to an instantaneous change in prices, which is simply the differential form of (10):

$$\dot{\omega} = \mu(\dot{p} - \omega). \tag{11}$$

Our specification of ω is sufficiently general that the impulse response may not always result in a well-behaved boom and bust in prices. In some cases, prices may rise without falling, or they may oscillate indefinitely. We restrict focus to a parameter region in which a boom is followed by a bust that asymptotes without indefinitely oscillating. In particular, we require $\gamma(\cdot)$, the function mapping past price information to future expected price changes, to satisfy the following assumption:

Assumption 3. *$\gamma(\omega) \equiv 0$ for $\omega \leq 0$, $\lim_{\omega \to 0^+} \gamma(\omega) = 0$, $\gamma'(\omega) > 0$ for $\omega > 0$, and*

$$\lim_{\omega \to 0^+} \gamma'(\omega) > \frac{r}{\min(c\epsilon, \mu)} \left(\int_0^\infty \int_0^\infty \lambda e^{-(r+\lambda)\tau} g'(\tau) d\tau f(\lambda) d\lambda \right)^{-1}.$$

The requirement that $\gamma \equiv 0$ for $\omega < 0$ rules out oscillations after the bust, as agents stop expecting capital gains once the historical average return becomes negative. The rest of Assumption 3 guarantees that price increases initially beget further increases and that prices eventually overshoot. With this additional structure, we are now able to provide a complete characterization of the joint dynamics of prices and volume following the one-time permanent demand shock.

3.1 The Boom

We begin with a series of three propositions which divide the boom-bust cycle into three distinct epochs. We refer to these epochs as "the boom," "the quiet," and "the bust." Proposition 2 presents the conditions that characterize the boom.

Proposition 2. *There exists $t_1 > 0$ such that for $0 < t \leq t_1$: prices rise ($\dot{p}_t > 0$), prices are convex ($\ddot{p}_t > 0$), expected capital gains increase ($\dot{\omega}_t > 0$), demand exceeds listings ($D_t > L_t$),* *and if* $|\operatorname{supp} f| > 1$ *volume rises* ($\dot{V}_t > 0$)*.*

As the proof shows, the shock to A at $t = 0$ causes demand D to jump above available listings. As a result, the price of housing begins to increase, and this increase raises ω. The rise in ω stimulates demand, further increasing prices and leading to convexity in the price path. Demand rises more sharply for potential buyers with higher values of λ, causing an increase in listings and hence volume. We define the boom as the time interval $(0, t_1^*]$, where t_1^* is the largest t_1 such that the conditions of Proposition 2 hold.

Proposition 2 fits the 2000-2005 US housing market remarkably well, as shown in Figure 3. Prices and volume rose during this time, with prices rising at an increasing rate. Unsold inventories increased at the same rate as volume, implying that demand was sufficient to exhaust the flow of listings and that a stock of unsold listings did not accumulate.

3.2 The Quiet

To continue characterizing the cycle, we provide a lemma that guarantees that prices do not rise indefinitely:

Lemma 4. *There exists (a finite)* $t_2 > 0$ *such that* $\dot{p}_{t_2} = 0$*.*

Lemma 4 implies that the boom must end ($t_1^* < \infty$). It also precludes a price path that asymptotes to some level without overshooting. We define t_2^* to be the smallest $t_2 > 0$ such that $\dot{p}_{t_2} = 0$. Proposition 3 characterizes the time period between t_1^* and t_2^*.

Proposition 3. *Prices rise* ($\dot{p}_t > 0$) *for all* $t \in (t_1^*, t_2^*)$*, listings exceed demand* ($L_t > D_t$) *for some* $t \in (t_1^*, t_2^*)$*, and volume falls* ($\dot{V}_t < 0$) *for some* $t \in (t_1^*, t_2^*)$*.*

As can be seen from the price-adjustment equation (8), $D_{t_2^*}$ equals the long-run level of demand and listings $\overline{D}$. This level represents a decline from the magnitude of demand during the boom, so demand must fall sometime during (t_1^*, t_2^*). In contrast, listings remain above this steady-state value because past capital gains continue to lure short-term potential buyers disproportionately. The combined effect is an excess of listings over demand. Inventories of unsold listings accumulate by (5), and volume equals demand and hence falls with demand. Because trading activity is falling, we refer to the time interval $(t_1^*, t_2^*]$ as "the quiet."

Proposition 3 fits the time between 2005 and 2007 in Figure 3. In 2005, prices reached at inflection point and began to grow at a slower rate, and volume began to fall. Unsold inventories then sharply increased as volume continued to decline.

3.3 The Bust

The final proposition characterizes prices and volume after prices stop rising.

Proposition 4. $\lim_{t\to\infty} p_t = p_0 + \epsilon \log(A^f/A^i)$*, and for* $t > t_2^*$ *prices decline* ($\dot{p}_t < 0$) *and volume is below steady-state* ($V_t < \overline{D}$) *until this limit is reached.*

After the quiet, prices decline on low volume, eventually reaching the steady state level. We refer to the time interval (t_2^*, ∞) as "the bust."[21]

Genesove and Mayer (2001) document low volume during a bust in the Boston apartment market and ascribe it to high prices posted by loss-averse sellers trying to avoid nominal losses. In our setting, low volume during the bust derives from the sticky price adjustment that states that prices decline whenever demand is low. Low demand holds during the bust for two reasons: speculative demand is low because expected capital gains are small or absent, and fundamental demand is low because prices exceed the steady-state value.

The bust corresponds to the period after 2007 shown in Figure 3 during which prices fell and volume was lower than average.

3.4 Short-Term Buyers and the Size of the Cycle

Propositions 2–4 characterize the relative timing of the price and volume responses. We now describe the relative *magnitude* of these responses. In particular, we show that a common factor—the distribution of expected holding times $f(\cdot)$—determines the magnitude of the price and volume responses as well as the level of steady-state volume.

To discuss the size of the price and volume responses, we define $P^{max} = \max_{t\in(0,t_2^*)} P_t$ and $V^{max} = \max_{t\in(0,t_2^*)} V_t$ to be the largest values of prices and volume attained during the boom and quiet, and we denote the final steady-state price by P_∞. The presence of additional short-term buyers magnifies the price response and steady-state volume:

Proposition 5. *An increase to* $f(\cdot)$ *raises* V_0 *and* P^{max} *while keeping* P_0 *and* P_∞ *constant.*

[21] Prices do not overshoot on the decline because Assumption 3 rules out negative expected capital gains. This simplifying assumption allows us to highlight the dynamic interactions between the composition of buyers, volume, and prices during the periods leading to the bust, and how buyer composition can lead to overshooting as prices rise. Without this assumption, prices would overshoot on the way down, which is a standard feature of models with extrapolative expectations (e.g., Glaeser and Nathanson, 2016). Empirically, negative overshooting may occur for a variety of reasons, such as a continuing glut of listed houses (Shleifer and Vishny, 1992) or foreclosures (Guren and McQuade, 2015).

The increase to $f(\cdot)$ is under first-order stochastic dominance, which implies a greater frequency of short-term buyers (those with high λ).

As shown in the proof, a larger $f(\cdot)$ raises steady-state volume V_0 and the speculative demand function $\Sigma(\cdot)$. The demand of short-term buyers is more sensitive to expected capital gains, so a greater frequency of such buyers raises aggregate speculative demand. A larger $\Sigma(\cdot)$ then increases the maximal price reached during the cycle without changing the steady states. Proposition 5 makes the empirical prediction that price cycles are larger in markets with higher steady-state volume; we confirm this prediction in Section 4.

To demonstrate how $f(\cdot)$ jointly determines the magnitudes of the price and volume responses, we now apply Proposition 5 to a special case in which we compare two markets. In market A, all potential buyers have the same expected holding time. In market B, potential buyers differ in their expected holding times, none of which are longer than the expected holding time in A. The markets are otherwise identical.

Proposition 6. *If* $1 = |\operatorname{supp}(f_A)| < |\operatorname{supp}(f_B)|$ *and* $f_A < f_B$*, then* P^{max}/P_0*,* P^{max}/P_∞*,* V^{max}/V_0*, and* V_0 *are larger under* f_B *than* f_A*.*

Proposition 5 clarifies the importance of volume for understanding asset bubbles in a special case. Volume is not a sideshow to prices, but rather a manifestation of the speculative forces responsible for price dynamics. These speculative forces are captured by the heterogeneity in $f(\cdot)$. We conjecture that a mean-preserving spread in $f(\cdot)$ always increases P^{max} and V^{max} while leaving V_0 unchanged. Although we are not able to prove this conjecture, we verify it in our calibration below.

3.5 Calibration

To study predictions of our model quantitatively, we simulate the impulse response characterized in Propositions 2 through 4. We choose parameter values using surveys and prior literature to discipline the calibration and ask whether, subject to this parameterization, the changes in volume and buyer composition are large.

3.5.1 Parameter Choices

Table 1 lists the model parameters and their sources. We calibrate $f(\cdot)$ using the expected holding times reported in the NAR survey shown in Figure 2(a).[22] To select μ, we rerun the regression shown in Figure 2(b) by replacing the lagged four-quarter house price change with ω as given by (10). We choose μ to maximize the R^2 of this regression, leading us to $\mu = 1.19$ (standard error 0.65) and $R^2 = 67\%$. This estimate is noisy due to the small amount of data used in the estimation, but it is close to the value of $\mu = 0.5$ estimated by Barberis et al. (2015) in the context of the stock market.

To calibrate $g(\cdot)$, we adopt the functional form $g(\tau) = \rho(1 - e^{-\tau/\rho})$. For all ρ, $g'(0) = 1$, so ρ controls the extent to which the initial gain is extrapolated into the future. The half-life of the cumulative gains implied by $g(\cdot)$ equals $\rho \log 2$, so a larger ρ produces greater relative sensitivity of long-term expected gains to short-term expected gains.[23] In each year from 2014-2016, the New York Fed's Survey of Consumer Expectations reports the median expectation of house price growth in the United States over the next 1 and 5 years (see Fuster and Zafar, 2015 and Kuchler and Zafar, 2016 for information on this survey). The ratio of these expectations equals $(1 - e^{-5/\rho})/(1 - e^{-1/\rho})$, so we use the sample ratios to obtain a value of $\rho = 20.9$ (half-life of 14.5 years). An alternative method is to use the coefficient estimates from Armona et al. (2016) of the relative sensitivities of 1-year and 5-year forward capital gains expectations to the prior year's house price appreciation. By equating these estimates to $g(\tau) \lim_{\omega \to 0^+} \gamma'(\omega)$ for $\tau = 1$ and $\tau = 5$, we obtain a much smaller number of $\rho = 1.4$ (half-life of 1.0 year).[24] We use the average of these two numbers as our baseline, and return to each extreme in the sensitivity analysis.

To calibrate $\gamma(\cdot)$, we adopt the functional form $\gamma(\omega) = r(1 - e^{-\phi\omega/r})$ for $\omega > 0$. Here $\phi > 0$ is a free parameter governing the sensitivity of expectations to small increases in ω. As required by Assumption 1, $\gamma(\omega)g'(0) \leq r$ for all ω, and as required by Assumption 3, $\gamma'(\omega) >$

[22] We map the expected holding time for each bin to its median, except for $[11, \infty)$ which we map to 20. We then associate λ equal to the inverse of the expected holding time to each bin (e.g. $\lambda = 2$ for an expected holding time of 0.5 years). Using the share of sales to each property type in each year, we calculate $f(\cdot)$ for each year and then take an equal-weighted average across years to obtain the distribution used in the calibration.

[23] The half-life is the value τ_{hl} such that $g(\tau_{hl}) = (1/2) \lim_{\tau \to \infty} g(\tau)$.

[24] Armona et al. (2016) calculate the expected 1-year gain, which equals $\gamma(\omega)g(1)$, and the annualized 2-5 year expected gain, which equals $(1+\gamma(\omega)(g(5)-g(1))/(1+\gamma(\omega)g(1)))^{1/4}-1$. The right derivatives at $\omega = 0$ are $g(1) \lim_{\omega \to 0^+} \gamma'(\omega)$ and $(1/4)(g(5) - g(1)) \lim_{\omega \to 0^+} \gamma'(\omega)$. The ratio of these equals $(1/4)(g(5)/g(1) - 1)$, the value of which uniquely identifies ρ.

0 for all $\omega > 0$. We use the results of the survey of homeowner expectations conducted by Case et al. (2012) to estimate ϕ. Case et al. (2012) report the average expectation of the next year's price growth of homeowners in Alameda County (CA), Middlesex County (MA), Milwaukee County (WI), and Orange County (CA) in the spring of each year from 2003 to 2012. Using the CoreLogic monthly house price indices going back to 1976, we calculate ω_t for each county and year with (10) and our estimate of μ mentioned above. We then choose ϕ and a constant to minimize the mean-squared error of $g(1)\gamma(\omega)$ plus this constant versus the expectation reported by Case et al. (2012). The resulting value is $\phi = 0.98$ (standard error 0.27).[25] Our specification explains 70% of the variance in 1-year expectations across counties and years in this sample.

We set $r = 0.07$, which corresponds to a steady-state price-rent ratio of about $1/0.07 = 14$. We choose $c = 1$, which implies a half-life of price adjustment of about 8 months. We set the elasticity of demand, ϵ, equal to 0.6, a value in the range of estimates suggested by Hanushek and Quigley (1980). Finally, for round number convenience, we choose a demand shock size to match a long-run price impact of 10%. This price impact equals $(A^f/A^i)^{1/\epsilon}$, so we set $A^f/A^i = 1.06$: a demand shock of 6%.

3.5.2 Results

The differential equations given by our model allow us to solve for the impulse response in continuous time. In order to quantify the marginal effects of heterogeneous holding times, we supplement the baseline model with one in which λ is the same for all potential homebuyers. We set this value to $(\int_0^\infty \lambda^{-1} f(\lambda) d\lambda)^{-1}$, the unique value at which steady-state volume remains unchanged.

Figure 4 displays the resulting impulse responses. Panel (a) plots our two main objects of study: prices and volume. In the core model, prices significantly overshoot the long-run cumulative growth of 10%, more than doubling before decreasing to the new level. Initially, prices are convex, meaning that price changes beget further changes. In contrast, prices display a much less pronounced boom and bust when expected holding times equal the average.

In the baseline model, volume rises and then falls, beginning to decline 11 months before prices. This delay is very close to the empirical delay of 15 months we document below.

[25] The value of the constant is 0.022 (0.004).

The total rise of volume in our simulation equals 22%, a substantial fraction of the 34% rise in existing sales volume in the US between 2000 and 2005. As shown on the right, volume remains constant as prices rise when λ is homogeneous.

Panel (b) documents the changing composition of buyers over the cycle. At each time, we calculate the share of purchases going to buyers whose expected holding time is less than 3 years. In steady state, this share equals 20% (as identified by the NAR survey), and it rises to a high of 39%. This rise occurs as prices increase, and it drives the concomitant and subsequent surge in volume. In contrast, no homebuyers have expected holding times less than 3 years in the homogeneous simulation, as the average holding time (and hence universal holding time in that case) equals 10.5 years.

Finally, panel (c) documents the evolution of unsold listings over the cycle. Until the quiet begins, all listings sell, so unsold listings equal 0. The stock grows as volume begins to decline. Quantitatively, it reaches 6% of the housing stock in the baseline model, but only 2% when horizons are homogeneous.

In sum, Figure 4 shows that our calibrated model can quantitatively generate large swings in prices and volume during the boom and quiet periods, a dramatic shift in the composition of buyers, and a sharp rise in inventories in the period leading to the bust. These features depend critically on heterogeneity in the expected holding times of potential homebuyers.

3.5.3 Sensitivity Analysis

To provide intuition on how the parameters drive the results, we report key statistics of the simulation under parameters other than our baseline in Table 2. The three statistics we report are the excess price boom $P^{max}/P_\infty - 1$, the volume boom $V^{max}/V_0 - 1$, and the maximal inventory of unsold listings $\max_t I_t$. We vary each parameter to a low and high value while keeping the remaining parameters at the baseline values.

First we vary the degree of heterogeneity in $f(\cdot)$. In the "high" treatment, we keep steady-state volume $(\int_0^\infty \lambda^{-1} f(\lambda) d\lambda)^{-1}$ constant but put all the mass in $f(\cdot)$ on the most extreme values of λ in its support. The "low" treatment simply replicates the right panel of Figure 4 in which no heterogeneity exists. The booms in price and volume are much larger in the high treatment than in the baseline—prices more than quadruple, and volume almost doubles. Unsold inventories also rise to 18% of the housing stock. These results provide strong evidence tying together price booms, volume booms, and the distribution of expected holding times.

Varying the housing demand elasticity produces similar effects. Our low treatment sets $\epsilon = 0.3$ (half the baseline), whereas our high treatment sets $\epsilon = 1.8$ (an average of values calculated by Diamond, 2016). Prices more than quadruple and volume nearly doubles under the high elasticity, whereas prices and volume are much more stable under the low treatment. Short-term buyers enter the market more aggressively when housing demand is more elastic, so increasing the elasticity achieves similar results to increasing the frequency of short-term buyers. The simulation results are less sensitive to variations in the other parameters.

4 Speculative Dynamics in the U.S. Housing Bubble

While the calibration in the previous section allows for an assessment of the potential quantitative relevance of the factors that drive our model, their actual empirical relevance has yet to be established. In this section, we provide empirical evidence linking shifts in the distribution of realized holding periods over the course of the 2000–2008 US housing cycle to dynamic patterns in volume and prices that directly mirror the patterns implied by our model. We focus on the housing market both because of its macroeconomic relevance and because the availability of comprehensive, asset-level microdata permits a uniquely rich analysis of holding periods and the details of buyers and sellers.

4.1 Data

To conduct our analysis, we use data on individual housing transactions provided by CoreLogic, a private vendor which collects and standardizes publicly available tax assessments and deeds records from municipalities across the US. Our main analysis sample spans the years 1995–2014 and includes data from 115 Metropolitan Statistical Areas (MSAs), which together represent roughly 50 percent of the US population.

We include all transactions of single-family homes, condos, or duplexes that satisfy the following filters: (a) the transaction is categorized by CoreLogic as occurring at arm's length, (b) there is a non-zero transaction price, and (c) the transaction is not coded by CoreLogic as being a nominal transfer of title between lenders following a foreclosure. We also drop a small number of duplicate transactions where the same property is observed to sell multiple times at the same price on the same day or where multiple transactions occur between the same buyer and seller at the same price on the same day. These restrictions leave us with a final sample of 49,709,319 transactions. Given the geographic coverage of these data and their source in administrative records, our analysis sample serves as a rough proxy for the population of transactions in the US during our sample period.

We supplement these data with national and MSA-level housing stock counts from the US Census and national counts of sales and listings of existing homes from the NAR.

4.2 The Composition of Buyers

The key mechanism that generates time-variation in transaction volume in our model is that changes in expected capital gains over the course of the housing cycle differentially attract buyers with shorter-vs-longer expected holding periods. This phenomenon was stated formally in Proposition 1 and implies that large swings in volume should be accompanied by equally large changes in the distribution of realized holding periods among those who choose to sell their homes at various points in the cycle.

As evidence for this prediction, Figure 5 presents a simple yet compelling illustration of the time variation in realized holding periods during the 2000–2008 US housing cycle. We define the holding period of each transaction as the number of days since the last transaction of the same property. We then group all transactions with holding periods less than or equal to 5 years into bins of 1, 2, 3, 4, or 5 years and count the number of transactions falling into each bin. Figure 5 plots these bin counts by year for each year between 2000 and 2008.

During the boom years of 2000–2005, there is a clear compression in the distribution of realized holding periods toward shorter holding periods. This pattern then reverses as national house prices peak in 2006 and begin to fall in the subsequent years. The increase in transaction volume at short holding periods during the boom years represents a non-trivial portion of the overall increase in volume during this period. For example, total volume across all holding periods (including those greater than 5 years) increased from 2,735,490 transactions in 2000 to 3,817,122 transactions in 2005. During the same period, total volume in the 1-, 2-, and 3-year bins increased from 471,057 transactions to 921,766, which implies that these three groups alone can account for 42 percent of the total increase in volume between 2000 and 2005.[26]

[26]This finding is in line with the evidence provided by Bayer et al. (2015), who document a similar increase in volume among short holding period buyers in the Los Angeles MSA during this period.

This shift in the composition of buyers and sellers toward shorter holding periods during the boom years correlates highly with changes in total volume across local markets. This correlation can be seen clearly in Figure 6, which presents scatter plots of the percent change in total volume at the MSA-level from 2000–2005 versus the percent change in volume for short holding periods ($\leq$ 2 years) in Panel (a) and long holding periods ($>$ 2 years) in Panel (b).[27] Not only does the growth in volume of short holding period transactions correlate strongly with the increase in total volume across MSAs during this period, but this relationship is much stronger for short holding periods relative to long holding periods.

Panel (c) further shows that these cross-sectional differences in the growth rate of short holding period volume explain a significant portion of the differences in the growth in total volume across cities during this period. For each city, we plot the change in short holding period volume divided by initial total volume on the y-axis against the percent change in total volume on the x-axis. The slope of this line provides an estimate of how much of a given increase in total volume during this period came in the form of short holding period volume. The answer is approximately 33 percent. Thus, as predicted by Proposition 1, shifts in the distribution of holding periods of buyers and sellers over the course of the cycle appear to be a major determinant of changes in total transaction volume.

Because expected capital gains increase demand through the extensive margin, Proposition 1 additionally predicts that volume increases more strongly in groups of buyers with low flow utility. While we do not observe flow utility in our data, we do observe whether each purchased property is owner-occupied. Under the assumption that non-occupants receive less flow utility than occupants, we further test Proposition 1 by examining whether non-occupant purchases rose more than occupant purchases from 2000 to 2005.

To track the participation of non-owner-occupants in the market over time, we follow Chinco and Mayer (2016) by marking buyers as non-owner-occupants when the transaction lists the buyer's mailing address as being distinct from the property address.[28] While this proxy may misclassify some non-owner-occupants as living in the home if they choose to list the property's address for property tax collection purposes, we believe it to be a useful gauge

[27] For visual clarity, we group MSAs into 25 equal-sized bins based on their percent change in total volume during this period and calculate the average percent change in short and long holding period volume in each of these bins.

[28] In related work, Haughwout et al. (2011) use an alternative proxy for non-owner-occupancy based on the number of first-lien mortgages present on an individual's credit report. They also find a large increase in non-owner-occupant purchases during this time period.

of the level of non owner-occupied home sales in the market.

Using this proxy, Figure 7 displays plots that are analogous to those in Figure 6 but use non-owner occupancy as the sorting variable rather than holding periods. Similar to the patterns we documented for short holding periods, we find that non owner-occupant volume is an important driver of total volume during the cycle. The top panels compare volume growth for non owner-occupants and owner-occupants, and show that the relationship between total volume growth and non owner-occupant volume growth is much stronger. The bottom panel shows that this growth is also quantitatively important in accounting for total volume growth. Non owner-occupant volume accounts for more than half of the growth in total volume across cities.

The city-level results from Figures 6 and 7 also hold across local neighborhoods *within* cities. To show this, in Table 3 we repeat the cross-city analysis using a zip-code-level panel data set constructed from the same underlying sample of transactions. The first three columns report results that are directly analogous to panels (a)-(c) of Figure 6. In column 1, we regress the percentage change in the number of short ($\leq$ 2 years) holding period transactions in each zip code on the corresponding percentage change in total volume for that zip code. Column 2 runs the same regression, using the percentage change in the number of long ($>$ 2 years) holding period transactions as the outcome. Both regressions include a full set of MSA-fixed effects, so that the coefficient estimates reflect only within city variation in transaction volume. Echoing the cross-city results from Figures 6, the change in short holding period volume is strongly correlated with changes in total volume and is much stronger than the corresponding relationship for long holding periods. Column 3, which is analogous to panel (c) of Figure 6, shows that these changes in short holding period volume are also quantitatively important for explaining cross-zip-code differences in the change in total volume with in MSA. In this specification, we regress the change in short holding period volume measured as a fraction of the level of total volume in 2000 on the change in total volume for the zip code. The coefficient estimate implies that roughly 28% of the cross-zip code variation within an MSA in the growth of volume over this period can be attributed to differences in the number of short holding period transactions. The results for owner-occupants versus non-owner occupants are also similar. Columns 4–6 of Table 3 report estimates from similar regressions where we instead group borrowers by their occupancy status. As in Figure 7, we find an important role for non-owner-occupants. The coefficient estimate in column 6 implies that roughly a third of the within city variation in the growth rate of total volume between 2000 and 2005 can be attributed to changes in the number of purchases by owner occupants.

4.3 The Joint Dynamics of Prices and Volume

Propositions 2–4 predict that prices and volume both go through a boom and bust cycle, with the volume cycle leading the price cycle. In Figure 8, we present evidence that this relationship holds on average across MSAs in our sample. To do so, we search for the horizon over which a given change in volume has the most predictive power for the contemporaneous change in prices at the MSA level. Changes in volume generally lead changes in prices if the correlation between prices and volume is maximized at a positive lag.

To implement this search, we construct a monthly panel of house prices and total transaction volume at the MSA level running from January 1995 to December 2014. Prices equal the log of CoreLogic MSA-level house price indices, demeaned at the MSA level. Volume equals the total number of transactions in our data in a given month and MSA divided by that MSA's housing stock in the 2000 Census. Because the CoreLogic price indices are seasonally adjusted, we adjust our volume series by subtracting the MSA-specific average volume for the relevant calendar month from each observation. Using this panel, we then run a series of simple regressions of the form

$$p_{i,t} = \beta_\tau v_{i,t-\tau} + \eta_{i,t}, \tag{12}$$

where p is price, v is volume, i indexes MSAs, and time is measured in months.

The coefficient β_τ provides an estimate of how movements in volume around MSA-calendar month averages at a τ-month lag are correlated with contemporaneous movements in prices around MSA averages. We run these regressions separately for up to 4 years of lags ($\tau = 48$) and one year of leads ($\tau = -12$). Figure 8 plots the implied correlation from each regression along with its 95% confidence interval.[29] The correlation is positive at all leads and lags, but reaches its maximum at a positive lag of approximately 15 months. Thus, changes in volume generally lead changes in prices by a little over a year.

[29]The implied correlation equals $\rho_\tau = \beta_\tau \operatorname{std}(v_{i,t-\tau}) / \operatorname{std}(p_{i,t} - \overline{p_i})$. Confidence intervals for each estimate are calculated from standard errors that are clustered at the month level.

4.4 Listing versus Selling Probabilities

Proposition 2 predicts that during the boom, volume increases entirely because owners are more likely to list their houses for sale. The speed at which listing sell remains constant, as all listing sell instantaneously as they do in the steady state. Conversely, Proposition 3 shows that volume falls during the quiet because the speed at which listings sell begins to decline, as listings outstrip demand.

To determine the relative importance of listing and selling rates for explaining empirical movements in volume, we use the methodology presented by Ngai and Sheedy (2016) to measure each of these rates over the course of the housing cycle. In discrete time, the equations determining these rates are $V_t = s_t I_t$, $I_t = I_{t-1} - V_{t-1} + L_t$, and $L_t = n_t(K_t - I_t)$; K denotes the housing stock, s denotes the rate at which listings sell, and n denotes the rate at which owners list unlisted houses.[30] Using monthly data on V, I, and K, we calculate $s_t = V_t/I_t$ and $n_t = (I_t - I_{t-1} + V_{t-1})/(K_t - I_t)$ for each month. We average these probabilities within each year and plot the resulting averages in Figure 9.[31]

Figure 9 shows that the listing rate n_t rose 40% during the 2000-2005 boom, whereas the rate s_t at which listings sold remained flat. In contrast, the decline in volume during the 2005-2006 quiet is driven primarily by s_t, which contracts by about 25% while n_t remains high. As is suggested by the steady-state equation $V = Ksn/(s + n)$, these rates multiplicatively determine volume, so their proportional changes determine their relative importance. This evidence confirms the model's predictions about the drivers of volume during the boom and the quiet.

4.5 Initial Volume and the Size of the Cycle

Proposition 5 shows that an increase to the distribution of expected holding times $f(\cdot)$ raises both the level of steady-state volume as well as the magnitude of the increase in prices during the boom and subsequent decreases in prices during the bust. Rather than measure $f(\cdot)$ directly, we test whether steady-state volume and the amplitude of the price cycle are

[30] In discrete time, V and L are the integrals of their continuous time counterparts, whereas I equals the integral of continuous-time listings.

[31] We cannot impute L_t in January of 1999, as it is the first month in our sample, so the average of n_t shown in Figure 9 is taken over February through December of 1999. The remaining data points in the figure represent averages taken over all 12 months of each corresponding year. Burnside et al. (2016) plot the same series for s_t.

positively correlated across cities. Our measure of steady-state volume equals the number of existing home sales in 2000 as a share of the housing stock. The boom in each city equals the percentage change in prices between January 2000 and the month in which prices peaked, and the bust equals the percentage change between the month of the peak and the month in which prices reached their lowest level subsequent to the peak month.[32]

Figure 10 plots the relationship between steady-state volume and the magnitude of the boom and bust in prices across cities. Panel (a) shows that there is a clear positive relationship between initial volume and the magnitude of the price boom. Cities with higher initial volume experienced significantly larger house price booms. As shown in Panel (b), these cities also experienced more drastic drops in prices following the boom.

Columns 1 and 3 of Table 4 quantify this relationship by reporting coefficient estimates from simple linear regressions of the price boom (column 1) and bust (column 3) on steady-state volume. A one percentage point increase in the share of the existing housing stock that turned over in 2000 is associated with a 15 percentage point higher increase in prices from January 2000 to peak and a four percentage point larger fall in prices from peak to trough. In columns 2 and 4, we report analogous and nearly identical estimates from regressions which instead assume that the boom ended in January of 2006 for all cities. These results are strongly consistent with the prediction of our model that steady-state volume should be correlated with the magnitude of swings in house prices during boom-bust episodes.

5 Conclusion

This paper shows that short-term investors have the capacity to destabilize financial markets. This observation raises two lines of inquiry.

First: do the expansions in credit that accompany asset price booms appeal disproportionately to short-term investors? Barlevy and Fisher (2011) document a strong correlation across US metropolitan areas between the size of the 2000s house price boom and the take-up of interest-only mortgages. These mortgages back-load payments by deferring principal repayment for some amount of time, and thus might appeal especially to buyers who expect to resell quickly. The targeting of credit expansions to short-term buyers might explain the

[32] We restrict the price peak to occur prior to January 2012 since prices in some markets had already recovered to levels higher than those experienced during the boom by the end of our sample.

amplification effects of credit availability on asset price booms documented by Di Maggio and Kermani (2015) Favara and Imbs (2015), and Rajan and Ramcharan (2015).

Second: do policies that aim to achieve financial market stability work better if they discourage the participation of short-term investors? For instance, consider the financial transactions tax proposed by Tobin (1978), supported by Stiglitz (1989) and Summers and Summers (1989), and analyzed theoretically by Dávila (2015). If the incidence of this tax falls entirely on buyers, then the tax burden is independent of the investment horizon; if the incidence falls entirely on sellers, then the tax burden is larger in present-value terms for short-term investors who plan to resell quickly. Our model suggests that transaction taxes discourage bubbles more powerfully when their incidence falls more strongly on sellers. One policy that discourages short-term investors more directly is the short-term capital gains tax, and our model provides a rationale for this policy. Any tax that discourages short-term investors will also discourage the liquidity provision and, in the case of the housing market, the residential investment they provide. We hope that future work will weigh all of these effects to guide policy carefully.

A Omitted Proofs of Mathematical Statements

Lemma 1

By (1), $\mathrm{E}[(P_{t+\tau} - P_t)/(\tau P_t) \mid \omega_t] = \gamma(\omega_t)g(\tau)/\tau$. The cross-partial in Lemma 1 equals $\gamma'(\omega_t)(g(\tau)/\tau)'$. Because $\gamma'(\omega_t) > 0$, this cross-partial is negative for all $\tau > 0$ if and only if Assumption 1(a) holds.

Lemma 2

If Assumption 1(c) fails, then we can find ω_t such that $\gamma(\omega_t)g'(0) > r$. Then $\mathrm{E}[P_{t+\tau}/P_t \mid \omega_t] - e^{r\tau}$ equals 0 at $\tau = 0$ and has a positive derivative with respect to τ at $\tau = 0$, which means that $\mathrm{E}[P_{t+\tau}/P_t \mid \omega_t] > e^{r\tau}$ for some $\tau > 0$. Now suppose Assumption 1(c) holds. For $\tau > 0$, $g(\tau) = \int_0^\tau g'(\tau_0)d\tau_0 < \int_0^\tau g(\tau_0)/\tau_0 d\tau_0 < g'(0)\tau$. As a result, for all $\omega_t \in \mathbb{R}$ $\mathrm{E}[P_{t+\tau}/P_t \mid \omega_t] = 1 + \gamma(\omega_t)g(\tau) < 1 + \gamma(\omega_t)g'(0)\tau \leq 1 + r\tau < e^{r\tau}$. The last inequality follows because $1 + r\tau$ and $e^{r\tau}$ coincide for $\tau = 0$ and the derivative of the latter exceeds that of the former for all $\tau > 0$.

Lemma 3

By (6), a potential buyer buys if and only if

$$\delta \geq rP_t \left(1 - \frac{\lambda}{r}\int_0^\infty (r+\lambda)e^{-(r+\lambda)\tau}\left(\mathrm{E}\left[\frac{P_{t+\tau}}{P_t} \mid \omega_t\right] - 1\right)d\tau\right).$$

Substituting (1) reduces the integral to $\int_0^\infty (r+\lambda)e^{-(r+\lambda)\tau}\gamma(\omega_t)g(\tau)d\tau$ and then integrating by parts further reduces it to $\int_0^\infty e^{-(r+\lambda)\tau}\gamma(\omega_t)g'(\tau)d\tau$. The measure of potential buyers at t of type λ whose flow utility exceeds some $\delta_0 > 0$ equals $f(\lambda)A_t\delta_0^{-\epsilon}$, so we are done.

Proposition 1

We prove the stronger statement of Proposition 1 in which each derivative with respect to ω_t is replaced by the right-sided derivative $\partial_+/\partial\omega_t$ or the left-sided derivative $\partial_-/\partial\omega_t$ throughout. We write the proof in terms of $\partial_+/\partial\omega_t$; the identical proof holds replacing those partials with $\partial_/\partial\omega_t$. We use this more general form of the proposition in the proof of Proposition 2.

From Lemma 3, $\log D_t(\lambda) = \log(f(\lambda)A_t(rP_t)^{-\epsilon}) + \log\Sigma_\lambda(\omega_t)$. Only $\Sigma_\lambda(\omega_t)$ depends on ω. We write $\log\Sigma_\lambda(\omega_t) = -\epsilon\log(1 - \gamma(\omega_t)i(\lambda))$. From Lemma 2, $1 + \gamma(\omega_t)g(\tau) < e^{r\tau}$ for $\tau > 0$, so $\gamma(\omega_t)i(\lambda) < \lambda/r\int_0^\infty (r+\lambda)e^{-(r+\lambda)\tau}(e^{r\tau} - 1)d\tau = 1$. For ω_t such that $\partial_+\gamma(\omega_t)/\partial\omega_t > 0$, $\partial_+\log\Sigma_\lambda(\omega_t)/\partial\omega_t = \epsilon(\partial_+\gamma(\omega_t)/\partial\omega_t)i(\lambda)/(1 - i(\lambda)) > 0$ because $i(\lambda) > 0$ by Assumption 1(d). Differentiating again yields $\partial\partial_+\log\Sigma_\lambda(\omega_t)/\partial\omega_t\partial\lambda = \epsilon(\partial_+\gamma(\omega_t)/\partial\omega_t)i'(\lambda)/(1 - i(\lambda))^2$. Integrating by parts and applying Assumption 1(a) yields

$$ri'(\lambda) = \int_0^\infty e^{-(r+\lambda)\tau}g'(\tau)d\tau - \int_0^\infty \lambda e^{-(r+\lambda)\tau}\tau g'(\tau)d\tau > \frac{r}{r+\lambda}\int_0^\infty e^{-(r+\lambda)\tau}g'(\tau)d\tau > 0,$$

where the final inequality follows because the last integral equals $r^2/(\lambda(r+\lambda))i(\lambda) > 0$.

To prove (c), we first note that $V_t(\lambda)/V_t = D_t(\lambda)/D_t$ by (3). By Lemma 3, $D_t(\lambda)/D_t = f(\lambda)\Sigma_\lambda(\omega_t)/\int_0^\infty f(\lambda')\Sigma_{\lambda'}(\omega_t)d\lambda'$. Thus, we must show that

$$\frac{\partial_+}{\partial \omega_t}\frac{\int_0^\lambda \Sigma_{\lambda'}(\omega_t)f(\lambda')d\lambda'}{\int_0^\infty \Sigma_{\lambda'}(\omega_t)f(\lambda')d\lambda'} \leq 0$$

for all λ. Differentiating, subtracting parts of the integral that appear on each side, and multiplying and dividing by $\sigma_{\lambda'}$ reduces this inequality to

$$\int_\lambda^\infty \Sigma_{\lambda'}(\omega_t)f(\lambda')d\lambda' \int_0^\lambda \frac{\partial_+ \log \Sigma_{\lambda'}(\omega_t)}{\partial \omega_t}\Sigma_{\lambda'}(\omega_t)f(\lambda')d\lambda' \leq$$
$$\int_0^\lambda \Sigma_{\lambda'}(\omega_t)f(\lambda')d\lambda' \int_\lambda^\infty \frac{\partial_+ \log \Sigma_{\lambda'}(\omega_t)}{\partial \omega_t}\Sigma_{\lambda'}(\omega_t)f(\lambda')d\lambda'.$$

The first part of the proof showed that $\partial_+ \log \Sigma_\lambda(\omega_t)/\partial\omega_t > 0$ and $\partial\partial_+ \log \Sigma_\lambda(\omega_t)/\partial\omega_t\partial\lambda > 0$ for all λ. Thus $\partial_+ \log \Sigma_{\lambda'}(\omega_t)/\partial\omega_t$ increases in λ' and is positive, allowing us to reduce the inequality to

$$\int_\lambda^\infty \Sigma_{\lambda'}(\omega_t)f(\lambda')d\lambda' \int_0^\lambda \Sigma_{\lambda'}(\omega_t)f(\lambda')d\lambda' \leq$$
$$\int_0^\lambda \Sigma_{\lambda'}(\omega_t)f(\lambda')d\lambda' \int_\lambda^\infty \Sigma_{\lambda'}(\omega_t)f(\lambda')d\lambda',$$

which holds with equality. Strict inequality results if and only if $\partial_+ \log \Sigma_\lambda(\omega_t)/\partial\omega_t$ is distinct across two points in the support of f. Because it strictly increases, strict inequality results if and only if the support of f consists of more than a single point.

Demand Target

Because prices remain constant, $D_t = \overline{D}$ by (8). Because potential buyers expect prices to remain constant, $\sigma_\lambda(\omega_t) = 1$ for all λ and ω_t by (7). Thus, by Lemma 3, $D_t(\lambda)/D_t = f(\lambda)$ for all t. Because $L_t = D_t$, $V_t = L_t$ by (3), so $V_t(\lambda) = f(\lambda)\overline{D}$. In a steady state, each $S_t(\lambda)$ remains constant, so $S_t(\lambda) = V_t(\lambda)/\lambda = \overline{D}f(\lambda)/\lambda$. Because $V_t = L_t$, $I_t = 0$. As a result, the housing stock is comprised entirely of stayers, so $1 = \int_0^\infty S_t(\lambda)d\lambda = \overline{D}\int_0^\infty f(\lambda)/\lambda d\lambda$. Solving for $\overline{D}$ provides (9).

Proposition 2

By Lemma 3, total demand across all potential buyers equals

$$D_t = A_t(rP_t)^{-\epsilon}\Sigma(\omega_t), \tag{A1}$$

where $\Sigma(\omega_t) \equiv \int_0^\infty \Sigma_\lambda(\omega_t)f(\lambda)d\lambda$ equals aggregate speculative demand. From (10), ω_t depends on price changes before t, so $\omega_0 = 0$. By (A1), $D_0 = A^f/A^i\overline{D}$, so by (8) $\dot{p}_0 =$

$c\log(A^f/A^i) > 0$. It follows that $\dot{p}_t > 0$ for $t \in (0, t_1)$ for some $t_1 > 0$.

By (11), $\dot{\omega}_0 > 0$, so $\dot{\omega}_t > 0$ for $t \in (0, t_1)$ for some $t_1 > 0$.

Substituting (A1) into (8), differentiating, and using (11) yields $\ddot{p}/c = -\epsilon\dot{p} + \sigma'(\omega)\mu(\dot{p} - \omega)$ for $\omega > 0$, where $\sigma(\omega) \equiv \log\Sigma(\omega)$. Differentiating yields

$$\sigma'(\omega) = \frac{\epsilon\gamma'(\omega)\int_0^\infty(1-\gamma(\omega)i(\lambda))^{-\epsilon-1}i(\lambda)f(\lambda)d\lambda}{\int_0^\infty(1-\gamma(\omega)i(\lambda))^{-\epsilon}f(\lambda)d\lambda},$$

where $i(\lambda)$ is as defined in the proof of Proposition 1. Taking limits yields $\lim_{\omega\to0^+}\sigma'(\omega) = (\lim_{\omega\to0^+}\gamma'(\omega))\epsilon\int_0^\infty i(\lambda)f(\lambda)d\lambda > \epsilon/\min(c\epsilon, \mu)$ by the bound in Assumption 3. It follows that $\ddot{p}_0 > 0$ because $\dot{p}_0 > 0$ and $\omega_0 = 0$, so there exists $t_1 > 0$ such that $\ddot{p}_t > 0$ for $t \in (0, t_1)$.

By (2), L_t depends directly $S_t(\lambda)$, and by (4), $S_t(\lambda) = \int_{-\infty}^t e^{-\lambda(t-\tau)}V_\tau(\lambda)d\tau$. Therefore L_t is continuous at 0, so $L_0 = \overline{D} < D_0$ and there exists $t_1 > 0$ such that $D_t > L_t$ for $t \in (0, t_1)$.

By (3), $V_t = L_t$ when $D_t > L_t$ and $I_t = 0$, so to prove that $\dot{V}_t > 0$ we must prove that $\dot{L}_t > 0$. We have $\dot{L}_0 = \int_0^\infty \lambda\dot{S}_0(\lambda)d\lambda$. But $\dot{S}_0(\lambda) = V_0(\lambda) - \lambda S_0(\lambda) = 0$ because $V(\lambda)$ and $S(\lambda)$ are continuous at $t = 0$. So $\dot{L}_0 = 0$. Taking another derivative yields $\ddot{L}_0 = \int_0^\infty \lambda\ddot{S}_0(\lambda)d\lambda$. We have $\ddot{S}_0(\lambda) = \dot{V}_0(\lambda) - \lambda\dot{S}_0(\lambda) = \dot{V}_0(\lambda)$. Thus $\ddot{L}_0 = \int_0^\infty \lambda\dot{V}_0(\lambda)d\lambda$. Because $\dot{\omega}_0 > 0$, by Proposition 1(c) $\int_0^\lambda \dot{V}_0(\lambda')d\lambda' \le 0$ for all $\lambda > 0$, so $V_0(\lambda) + \dot{V}_0(\lambda)$ first-order stochastically dominates $V_0(\lambda)$ as distributions (strictly iff $|\operatorname{supp} f| > 1$). The former must have a larger mean as a result, so $\int_0^\infty \lambda\dot{V}_0(\lambda) > 0$, giving $\ddot{L}_0 > 0$. Because $\dot{L}_0 = 0$, $\dot{L}_t > 0$ for $t \in (0, t_1)$ for some $t_1 > 0$.

Lemma 4

Define $\overline{p} = \log(A^f)/\epsilon - \log(r) - \log(\overline{D})/\epsilon$. Substituting (A1) into (8) yields

$$\dot{p} = c\epsilon(\overline{p} - p) + c\sigma(\omega). \tag{A2}$$

Substituting (A2) into (11) gives

$$\dot{\omega} = \mu c\epsilon(\overline{p} - p) + \mu c\sigma(\omega) - \mu\omega. \tag{A3}$$

With the initial conditions $p_0 = \overline{p} - \log(A^f/A^i)/\epsilon$ and $\omega_0 = 0$, (A2) and (A3) specify the joint dynamics of p and ω. Figure A2 illustrates the corresponding phase diagram.

The $\dot{p} = 0$ locus is given by $p = \overline{p} + \sigma(\omega)/\epsilon$. For a given ω, $\dot{p} < 0$ for p above this locus and $\dot{p} > 0$ for p below this locus. By Assumption 3, $\sigma(\omega) = 0$ for $\omega \le 0$ and $\sigma'(\omega) > 0$ for $\omega > 0$, so the $\dot{p} = 0$ locus equals $p = \overline{p}$ for $\omega < 0$ and increases for $\omega > 0$. The $\dot{\omega} = 0$ locus is given by $p = \overline{p} + \sigma(\omega)/\epsilon - \omega/(c\epsilon)$. For a given ω, $\dot{\omega} < 0$ for p above this locus and $\dot{\omega} > 0$ for p below this locus. For $\omega < 0$, this locus equals a decreasing line, and for $\omega > 0$, this locus lies beneath the $\dot{p} = 0$ locus. The right slope of this locus at 0 equals $\lim_{\omega\to0^+}\sigma'(\omega)/\epsilon - 1/(c\epsilon)$, which > 0 because $\lim_{\omega\to0^+}\sigma'(\omega) > \epsilon/\min(c\epsilon, \mu)$ as shown in the proof of Proposition 2.

We have drawn the phase diagram such that the $\dot{p} = 0$ locus is bounded for large ω and that the $\dot{\omega} = 0$ locus asymptotes to a decreasing linear function for large ω. These features hold as long as $\sigma(\omega)$ is bounded. As shown in the proof of Lemma 2, $g(\tau) < g'(0)\tau$ for $\tau > 0$, so $i(\lambda) < (\lambda g'(0)/r)\int_0^\infty(r+\lambda)e^{-(r+\lambda)\tau}\tau d\tau = (g'(0)/r)\lambda/(r+\lambda)$. Therefore Assumption 1(c) implies that $\sigma(\omega) < \log \int_0^\infty (r/(r+\lambda))^{-\epsilon} f(\lambda)d\lambda$, which exists because $\int_0^\infty \lambda^\epsilon f(\lambda)d\lambda$ exists.

Tracing the system from the initial point makes it clear that p must decrease in finite time. At first, p and ω increase. Eventually, the right $\dot{\omega} = 0$ locus is reached because this locus goes to $-\infty$ for large ω due to its linearity. Next, ω begins to increase while p continues increasing. The right $\dot{p} = 0$ locus is then reached because it is bounded from above. After this point, ω continues decreasing while p begins to decrease, as desired.

Proposition 3

By the definition of t_2^*, $\dot{p}_t > 0$ for $t \in (t_1^*, t_2^*)$. By (8), $D_{t_2^*} = \overline{D}$. By (3), $V_{t_2^*} \leq D_{t_2^*} = \overline{D}$. Because $V_0 = \overline{D}$ and $\dot{V}_t > 0$ for $t \in (0, t_1^*)$, there must exist $t \in (t_1^*, t_2^*)$ such that $\dot{V}_t < 0$.

Suppose for a contradiction that $L_t < D_t$ for all $t \in (0, t_2^*)$. Because $D_{t_2^*} = \overline{D}$, by continuity $L_{t_2^*} \leq \overline{D}$, so $t^* \equiv \inf\{t > 0 \mid L_t = \overline{D}\}$ exists and $t^* > 0$ because $L_t > \overline{D}$ for $t \in (0, t_1^*)$. Because $L_{t^*} < D_{t^*}$, $V_{t^*} = L_{t^*}$, so (2) and (4) imply that $\int_0^\infty \dot{S}_{t^*}(\lambda)d\lambda = 0$. If there exists $\lambda > 0$ such that $\int_0^\lambda \dot{S}_{t^*}(\lambda')d\lambda' < 0$, then the mean value theorem for integrals implies that $\dot{L}_{t^*} = \int_0^\lambda \lambda' \dot{S}_{t^*}(\lambda')d\lambda' + \int_\lambda^\infty \lambda' \dot{S}_{t^*}(\lambda')d\lambda' > \lambda \int_0^\lambda \dot{S}_{t^*}(\lambda')d\lambda' + \lambda \int_\lambda^\infty \dot{S}_{t^*}(\lambda')d\lambda' = 0$, an impossibility due to the infimum property of t^*. Thus $\int_0^\lambda \dot{S}_{t^*}(\lambda)d\lambda < 0$ for all $\lambda > 0$. Given that $\dot{p}_t > 0$ for $t \in (0, t^*)$, $\omega_{t^*} > 0$, so $V_t(\lambda)/V_t$ stochastically dominates $f(\lambda)$ by Proposition 1(c). Using (4) then yields $\int_0^\lambda \lambda' S_{t^*}(\lambda')d\lambda' < \int_0^\lambda V_{t^*}(\lambda')d\lambda' < (\int_0^\lambda f(\lambda')d\lambda')(\int_0^\infty \lambda' S_{t^*}(\lambda')d\lambda')$ for all $\lambda > 0$, so $\lambda S_{t^*}(\lambda)/\int_0^\infty \lambda' S_{t^*}(\lambda')d\lambda'$ stochastically dominates $f(\lambda)$. Because $1/\lambda$ decreases, it follows that $\int_0^\infty S_{t^*}(\lambda)d\lambda / \int_0^\infty \lambda S_{t^*}(\lambda)d\lambda < \int_0^\infty f(\lambda)/\lambda d\lambda$, so $L_{t^*} = \int_0^\infty \lambda S_{t^*}(\lambda)d\lambda > (\int_0^\infty \lambda^{-1} f(\lambda)d\lambda)^{-1} = \overline{D}$ because $\int_0^\infty S_{t^*}(\lambda)d\lambda = 1$ as $L_t < D_t$ for $t \in (0, t^*)$. We have reached the necessary contradiction.

Proposition 4

The proof of Lemma 4 showed that prices begin to decline when the right $\dot{p} = 0$ locus shown in Figure A2 is reached. After this point, the system moves left so that $\dot{p} < 0$. From (8), $D_t < \overline{D}$ during this movement. By (3) $V_t \leq D_t$, so $V_t < \overline{D}$ while $\dot{p} < 0$.

The price decline continues until either the $\omega = 0$ axis is reached or convergence to $(0, \overline{p})$ occurs. Before either of these events, the system can never fall below the right $\dot{p} = 0$ locus, and any intersection with it occurs for but an instant. As a result, p continues to decrease until convergence to the steady-state or until $\omega = 0$ but $p > \overline{p}$. In the first case, we are done. In the second case, we know that $\omega < 0$ right after $\omega = 0$ from examining the phase diagram. It is clear that ω remains below 0 for the rest of time until steady-state is reached, as the system remains weakly above the left $\dot{p} = 0$ locus. In this region, $\sigma(\omega) = 0$, so (A2) becomes $\dot{p} = c\epsilon(\overline{p} - p)$. Thus, $\dot{p} < 0$ for the rest of time, as $p > \overline{p}$ when $\omega = 0$.

Proposition 5

Because $V_0 = \overline{D} = (\int_0^\infty \lambda^{-1} f(\lambda)d\lambda)^{-1}$, an increase to $f(\cdot)$ increases V_0 because $1/\lambda$ decreases.

We next show that $\sigma(\cdot)$ increases pointwise if $f(\cdot)$ increases. We have $\partial\Sigma_\lambda/\partial\lambda = \epsilon\gamma(\omega)i'(\lambda)(1 - i(\lambda)\gamma(\omega))^{-\epsilon-1}$, which equals 0 for $\omega \leq 0$ and is positive otherwise. An increase to $f(\cdot)$ thus increases $\sigma = \log \int_0^\infty \Sigma_\lambda f(\lambda)d\lambda$ for $\omega > 0$ and keeps it constant at $\omega \leq 0$.

We now show that P^{max} increases in σ. $P_t = P^{max}$ when the system first intersects the $\dot{p} = 0$ locus shown in Figure A2. Because this locus is given by $p = \overline{p} + \sigma(\omega)/\epsilon$, it shifts up for $\omega > 0$ due to an increase in σ. If the path of (p, ω) also shifts to the right, then the intersection with this locus must occur at a higher value of p. Thus we prove that the path does shift to the right. This shift occurs if $dp/d\omega$ decreases as σ increases. This derivative equals $\dot{p}/\dot{\omega} = (\mu(1 - \omega/\dot{p}))^{-1}$. Because $\omega > 0$ and $\dot{p} > 0$ before the $\dot{p} = 0$ locus is reached, we must show that $\omega/\dot{p}$ decreases in σ. By (A2), $\dot{p}$ increases in σ, so we are done.

Finally, we show that P_0 and P_∞ are independent of f. As shown in the proof of Proposition 4, $\lim_{t\to\infty} p_t = \overline{p}$, which does not depend on f. The initial condition p_0 given in the proof of Lemma 4 also does not depend on f. Thus σ has no bearing on P_0 or P_∞.

Proposition 6

Because $f_A < f_B$, by Proposition 5 V_0 and P^{max} are larger under f_B while the steady-state prices P_0 and P_∞ stay unchanged. Under f_A, $L_t = \lambda^* S_t(\lambda^*) \leq \lambda^*$, where λ^* is the sole member of $\operatorname{supp} f_A$, and $V_0 = \lambda^*$. Therefore $V^{max}/V_0 = 1$ under f_A. By Proposition 2, $V^{max}/V_0 > 1$ under f_B because $|\operatorname{supp} f_B| > 1$.

B Microfoundation of Price Adjustment Rule

This appendix presents one possible microfoundation for the price adjustment rule (8). This microfoundation is consistent with (3), which specifies how housing is allocated between sellers and interested buyers at each time.

Market Mechanism

At each time, movers may post a listing price; we call listing movers "sellers." A centralized mechanism matches potential buyers to sellers. We assume there exists a cutoff price $P_t^c(\delta, \lambda)$ such that a potential buyer of type (δ, λ) purchases from a matched seller if and only if the seller's listing price is no greater than $P_t^c(\delta, \lambda)$. We verify the existence of $P_t^c(\cdot, \cdot)$ below.

The matching mechanism works as follows. When the stock of sellers equals 0 (but the flow may still be positive), the potential buyer with the highest cutoff price matches to the seller with the lowest listing price, then the potential buyer with the second highest cutoff price matches to the seller with the second lowest listing price, et cetera until either the pool of potential buyers or the pool of sellers has been exhausted.[1] When the stock of sellers is positive, this mechanism is applied first to the oldest cohort of movers and then is applied successively to younger cohorts of movers until either the pool of potential buyers or the entire pool of sellers has been exhausted.[2]

Each matched pair trades at the seller's listing price if and only if that price does not exceed the potential buyer's cutoff price. If trade does not occur, the potential buyer permanently exits the market and the seller remains a mover.

Information

Because trade may occur at different prices at the same time, some care is needed to define "the market price" at a given time. Letting i index sellers at t, we define the market price by $P_t = e^{\int \log P_{i,t} di}$, where $P_{i,t}$ is seller i's price. Under this definition, the log of the market price equals the average of the logs of the listing prices.

Before the posting of list prices, movers and potential buyers at t receive two pieces of information: the price history $\{P_{t'} \mid t' \leq t - dt\}$ and the number $\widetilde{D}_t$ of current potential buyers whose cutoff price is at least P_{t-dt}. We will take the limit as $dt \to 0$ to describe the continuous-time model.

Preferences

Movers act to maximize their utility after the matching process has occurred. We assume that they have lexicographic preferences over holding a house, preferring any outcome in which they do not own a house to all outcomes in which they do own a house. This extreme assumption strongly motivates movers to sell their houses at each time—selling a house is the only way to divest ownership of it (there is no free disposal)—and can be thought of as the limit as the "holding costs" assumed in housing search models (Piazzesi and Schneider,

[1] This mechanism presumes that δ and λ are observable.

[2] The assumption that older cohorts of sellers are matched before younger cohorts is made in the stock-flow matching literature (Ebrahimy and Shimer, 2010).

2009; Guren and McQuade, 2015; Han and Strange, 2015) go to infinity. Within each class of outcomes, utility is linear in consumption.

Movers understand the market mechanism described above, but there does not exist common knowledge of the rationality of movers, leading to uncertainty over the prices that other movers will post. Following Gilboa and Schmeidler (1989), we assume that each mover displays "ambiguity aversion" over this uncertainty, acting to maximize her utility contingent on the pricing decisions made by other movers that would minimize her utility. In this setting, the mover's decision is robust to the range of possible actions by other movers.

Potential buyers maximize the present value of utility discounted at a constant rate r, where utility is linear in consumption and the flow utility δ received while a stayer.

Rigidities

A fraction $e^{-\beta dt}$ of movers are inattentive to $\widetilde{D}_t$ and assume that it equals its steady-state value $\overline{D}$. The remaining $1-e^{-\beta dt}$ of movers use the true value of $\widetilde{D}_t$ when setting prices. The allocation of movers to each group is independent of their prior histories, so in continuous time attentiveness arrives with Poisson intensity β. This attention rigidity is similar to the slow diffusion of news in Hong and Stein (1999) and the sticky information sets in Mankiw and Reis (2002).

Potential residents and movers hold rigid beliefs concerning the the number of movers at each moment. Instead of inferring how this number may vary over time, agents assume that the flow of new movers and the stock of existing movers remain constant over time at their values in the unique steady state in which the stock of movers equals 0. As shown in Appendix A, the flow of new movers in this steady state equals the value $\overline{D}$ defined by (9).

As stated in Assumption 2, potential buyers assume that upon becoming a mover, they sell instantaneously at the prevailing market price. A potential buyer's expectation of future market prices is given by $\mathrm{E}_t[P_{t+\tau}|\delta, \lambda, P_{i,t}, \omega_{t-dt}, \widetilde{D}_t] = (1+\gamma(\omega_{t-dt})g(\tau))P_{i,t}$ for $\tau > 0$; $P_{i,t}$ is the listing price to which the potential buyer is matched, $\omega_{t-dt} = \omega(\{P_{t'} \mid t' \leq t-dt\})$, and γ and g satisfy the conditions in Assumption 1.

Cutoff Prices

The following lemma shows that, as assumed above, a potential buyer purchases the house to which she is matched if and only if the listing price does not exceed a time- and type-dependent cutoff.

Lemma B1. *A potential buyer of type (δ, λ) buys if and only if $P_{i,t} \leq \delta/r\Sigma_\lambda(\omega_{t-dt})^{-1}$.*

Proof. The potential buyer anticipates selling at the market price with certainty when becoming a mover, so her expected utility of becoming a mover at $t+\tau$ equals $(1+\gamma(\omega_{t-dt})g(\tau))P_{i,t}$, where $P_{i,t}$ is the listing price of the house to which she is matched at t. The expected utility of buying this house equals $\int_0^\infty \lambda e^{-\lambda\tau}(\int_0^{\tau'} e^{-r\tau'}\delta d\tau' + e^{-r\tau}(1+\gamma(\omega_{t-dt})g(\tau))P_{i,t})d\tau$. She buys if and only if this quantity is at least $P_{i,t}$, which occurs exactly when $rP_{i,t}(1-\lambda\gamma(\omega_{t-dt})/r\int_0^\infty(r+\lambda)e^{-(r+\lambda)\tau}g(\tau)d\tau) \leq \delta$. As shown in the proof of Lemma 2, $\gamma(\omega_{t-dt})g(\tau) < e^{r\tau}-1$ for all $\tau > 0$ and ω_{t-dt}. It follows that the term in parentheses exceeds $1-\lambda/r\int_0^\infty(r+\lambda)e^{-(r+\lambda)\tau}(e^{r\tau}-1)d\tau = 0$, so it can be divided by to obtain the decision rule $P_{i,t} \leq \delta/r(1-\lambda\gamma(\omega_{t-dt})/r\int_0^\infty(r+\lambda)e^{-(r+\lambda)\tau}g(\tau)d\tau)^{-1}$. Differentiating by parts as in the proof of Lemma 3 and then applying (7) yields the cutoff rule in Lemma B1. $\square$

Lemma B1 shows that the cutoff assumed above equals $P_t^c(\delta,\lambda) = \delta/r\Sigma_\lambda(\omega_{t-dt})^{-1}$. This cutoff leads to the following demand curve that is useful for proving the results below:

Lemma B2. *For any $\widetilde{P} > 0$, the number of potential residents for whom $P_t^c(\delta,\lambda) \geq \widetilde{P}$ equals $A_t(r\widetilde{P})^{-\epsilon}\int_0^\infty \Sigma_\lambda(\omega_{t-dt})^{-\epsilon}f(\lambda)d\lambda$.*

Proof. $P_t^c(\delta,\lambda) \geq \widetilde{P}$ if and only if $\delta \geq r\widetilde{P}\Sigma_\lambda(\omega_{t-dt})$. For each λ, the measure of such potential buyers equals $A_t f(\lambda)(r\widetilde{P}\Sigma_\lambda(\omega_{t-dt}))^{-\epsilon}$. Integrating over λ yields the result. $\square$

Listing Prices

The following lemma characterizes the price P_t^* chosen by attentive movers.

Lemma B3. *Each attentive mover chooses to list at the price $P_t^* = P_{t-dt}(\widetilde{D}_t/\overline{D})^{1/\epsilon}$.*

Proof. Suppose the measure of listing movers equals L_t and the cumulative distribution function of their list prices equals $F_t^p(\cdot)$. Suppose a given mover i chooses a listing price $P_{i,t}$. Given Lemma B2, the number of potential buyers whose cutoff is at least $P_{i,t}$ equals $\widetilde{D}_t(P_{i,t}/P_{t-dt})^{-\epsilon}$. Under the mover's assumption that the stock of movers equals 0, the mover sells with certainty at $P_{i,t}$ if and only if $L_tF_t^p(P_{i,t}) \leq \widetilde{D}_t(P_{i,t}/P_{t-dt})^{-\epsilon}$. We define $P_t^{sup}(L_t, F_t^p)$ to be the supremum of $P_{i,t}$ satisfying this inequality. The listing mover can guarantee selling her house and attain a payoff arbitrarily close to $P_t^{sup}(L_t, F_t^p)$ by choosing a price arbitrarily close to but below this value.

Due to ambiguity-averse preferences, the attentive mover sets $P_t^* = \min_{L_t,F_t^p} P_t^{sup}(L_t, F_t^p)$. Because this function decreases in L_t, the minimizing value of L_t is $\overline{D}$, the value that attains when all of the movers thought to exist decide to list. It follows that $P_t^* = \min_{F_t^p} P_t^{sup}(\overline{D}, F_t^p)$.

We claim that this minimum equals $P_{t-dt}(\widetilde{D}_t/\overline{D})^{1/\epsilon}$. If $F_t^p(P_{t-dt}(\widetilde{D}_t/\overline{D})^{1/\epsilon}) = 1$, then this assertion trivially holds. If $F_t^p(P_{t-dt}(\widetilde{D}_t/\overline{D})^{1/\epsilon}) < 1$, then $P_t^{sup}(\overline{D}, F_t^p) \geq P_{t-dt}(\widetilde{D}_t/\overline{D})^{1/\epsilon}$. It follows that $P_t^* = \min_{F_t^p} P_t^{sup}(\overline{D}, F_t^p) = P_{t-dt}(\widetilde{D}_t/\overline{D})^{1/\epsilon}$, as claimed. Because the mover is certain to sell at this price, she always chooses to list. $\square$

Lemma B3 allows us to characterize the evolution of the market price P_t over time. Because inattentive movers believe that $\widetilde{D}_t = \overline{D}$, Lemma B3 implies that they list at P_{t-dt} at time t. With lower-case p denoting $\log P$, the log market price at t equals $p_t = e^{-\beta dt}p_{t-dt} + (1-e^{-\beta dt})p_t^*$. Substituting the expression for p_t^* from Lemma B3 yields $p_t - p_{t-dt} = (1-e^{-\beta dt})\log(\widetilde{D}_t/\overline{D})/\epsilon$. Dividing by dt and taking the limit $dt \to 0$ yields

$$\dot{p}_t = (\beta/\epsilon)\log(D_t/\overline{D}), \tag{B1}$$

where we have used the definitional fact that $\lim_{dt\to 0}\widetilde{D}_t = D_t$. In continuous time, movers are inattentive almost surely and post at P_t, so transactions take place at P_t almost surely. As a result, (B1) reproduces (8) with $c = \beta/\epsilon$.

The fact that inattentive movers post at P_t rationalizes Assumption 2. Because potential buyers believe that the stock of movers remains equal to 0 at all times, they must believe that all movers sell instantaneously, and all but a measure 0 of such movers sell at P_t. In this way, Assumption 2 follows from the assumed rigidities in attention to $\widetilde{D}_t$ and to the variation in the number of movers.

The price chosen by inattentive movers coincides with a rigidity assumed in Guren (2016) in which some sellers copy the most recent listing price plus an increase proportional to recent price growth. In continuous time, under this rigidity inattentive movers copy the latest market price, which is the behavior they follow in our microfoundation.

Rationing

We now show that (3) holds given the market mechanism and the chosen prices of movers. In continuous time, all but a zero measure of matched potential buyers see a price of P_t.

If $L_t > D_t$ or $I_t > 0$, all D_t potential buyers whose cutoff price is at least P_t are matched, and all of them buy. Volume $V_t(\lambda)$ to potential buyers of type λ in this case equals $D_t(\lambda)$, the number of potential buyers of type λ whose cutoff price is at least exceeds P_t (according to the proof of Lemma B2).

If $L_t \leq D_t$ and $I_t = 0$, then only the L_t potential buyers with the highest cutoff prices buy. As shown in the proof of Lemma B2, the number of potential buyers of type λ whose cutoff price is at least $\widetilde{P}_t$ equals $D_t(\lambda)(\widetilde{P}_t/P_t)^{-\epsilon}$. The market-clearing cutoff for $\widetilde{P}_t$ solves $L_t = \int_0^\infty D_t(\lambda)(\widetilde{P}_t/P_t)^{-\epsilon} f(\lambda) d\lambda = D_t(\widetilde{P}_t/P_t)^{-\epsilon}$, so $\tilde{P}_t = P_t(D_t/L_t)^{1/\epsilon}$. It follows that $V_t(\lambda) = D_t(\lambda)(\widetilde{P}_t/P_t)^{-\epsilon} = L_t D_t(\lambda)/D_t$, as given by (3).

References

Abreu, Dilip and Brunnermeier, Markus K. (2003). 'Bubbles and Crashes', *Econometrica* 71(1), 173–204.

Agarwal, Sumit, Hu, Luojia and Huang, Xing. (2015). 'Rushing into the American Dream? House Prices Growth and the Timing of Homeownership'. Forthcoming, *Review of Finance*.

Anenberg, Elliot and Bayer, Patrick. (2013), Endogenous Sources of Volatility in Housing Markets: The Joint Buyer-Seller Problem, Technical report, National Bureau of Economic Research.

Armona, Luis, Fuster, Andreas and Zafar, Basit. (2016). 'Home Price Expectations and Behavior: Evidence from a Randomized Information Experiment'. Working Paper.

Bailey, Michael, Cao, Ruiqing, Kuchler, Theresa and Stroebel, Johannes. (2016), Social Networks and Housing Markets. Working Paper.

Barberis, Nicholas, Greenwood, Robin, Jin, Lawrence and Shleifer, Andrei. (2015). 'X-CAPM: An Extrapolative Capital Asset Pricing Model', *Journal of Financial Economics* 115(1), 1–24.

Barberis, Nicholas, Greenwood, Robin, Jin, Lawrence and Shleifer, Andrei. (2016). 'An Extrapolative Model of Bubbles'. Working Paper.

Barberis, Nicholas and Shleifer, Andrei. (2003). 'Style Investing', *Journal of Financial Economics* 68, 161–199.

Barberis, Nicholas, Shleifer, Andrei and Vishny, Robert. (1998). 'A Model of Investor Sentiment', *Journal of Financial Economics* 49(3), 307–343.

Barlevy, Gadi and Fisher, Jonas D.M. (2011). 'Mortgage Choices and Housing Speculation'. Working Paper, Federal Reserve Bank of Chicago.

Bayer, Patrick, Geissler, Christopher, Mangum, Kyle and Roberts, James W. (2015). 'Speculators and Middlemen: The Strategy and Performance of Investors in the Housing Market'. Working Paper.

Bhutta, Neil. (2015). 'The Ins and Outs of Mortgage Debt during the Housing Boom and Bust', *Journal of Monetary Economics* 76, 284–298.

Black, Fischer. (1986). 'Noise', *Journal of Finance* 41(3), 528–543.

Burnside, Craig, Eichenbaum, Martin and Rebelo, Sergio. (2016). 'Understanding Booms and Busts in Housing Markets', *Journal of Political Economy* 124(4), 1088–1147.

Caplin, Andrew and Leahy, John. (2011). 'Trading Frictions and House Price Dynamics', *Journal of Money, Credit, and Banking* 43(2), 283–303.

Case, Karl E., Shiller, Robert J. and Thompson, Anne K. (2012). 'What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets', *Brookings Papers on Economic Activity* Fall, 265–298.

Chinco, Alex and Mayer, Christopher. (2016). 'Misinformed Speculators and Mispricing in the Housing Market', *Review of Financial Studies* 29(2), 486–522.

Cochrane, John H. (2011). 'Discount Rates', *Journal of Finance* 66, 1047–1108.

Cutler, David M., Poterba, James M. and Summers, Lawrence H. (1990). 'Speculative Dynamics and the Role of Feedback Traders', *American Economic Review, Papers and Proceedings* 80(2), 63–68.

Cutler, David M., Poterba, James M. and Summers, Lawrence H. (1991). 'Speculative Dynamics', *Review of Economic Studies* 58(3), 529–546.

Daniel, Kent and Hirshleifer, David. (2015). 'Overconfident Investors, Predictable Returns, and Excessive Trading', *Journal of Economic Perspectives* 29(4), 61–88.

Daniel, Kent, Hirshleifer, David and Subrahmanyam, Avanidhar. (1998). 'Investor Psychology and Security Market Under- and Overreactions', *Journal of Finance* 53(6), 1839–1885.

Dávila, Eduardo. (2015). 'Optimal Financial Transactions Taxes'. Working Paper, NYU Stern.

De Long, J. Bradford, Shleifer, Andrei, Summers, Lawrence H. and Waldmann, Robert J. (1990). 'Positive Feedback Investment Strategies and Destabilizing Rational Speculation', *Journal of Finance* 45(2), 379–395.

Diamond, Rebecca. (2016). 'The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000', *American Economic Review* 106(3), 479–524.

Díaz, Antonia and Jerez, Belén. (2013). 'House Prices, Sales, and Time on the Market: A Search-Theoretic Framework', *International Economic Review* 54(3), 837–872.

Di Maggio, Marco and Kermani, Amir. (2015). 'Credit-Induced Boom and Bust'. Working Paper, Columbia Business School.

Ebrahimy, Ehsan and Shimer, Robert. (2010). 'Stock-Flow Matching', *Journal of Economic Theory* 145, 1325–1353.

Fama, Eugene F. (1970). 'Efficient Capital Markets: A Review of Theory and Empirical Work', *Journal of Finance* 25(2), 383–417.

Favara, Giovanni and Imbs, Jean. (2015). 'Credit Supply and the Price of Housing', *American Economic Review* 105(3), 958–992.

Favilukis, Jack, Ludvigson, Sydney C. and Van Nieuwerburgh, Stijn. (2015). 'The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium'. Forthcoming, *Journal of Political Economy*.

Fuster, Andreas and Zafar, Basit. (2015). 'The Sensitivity of Housing Demand to Financing Conditions: Evidence from a Survey'. Working Paper.

Gao, Zhenyu, Sockin, Michael and Xiong, Wei. (2016). 'Economic Consequences of Housing Speculation'. Working Paper.

Genesove, David and Mayer, Christopher. (2001). 'Loss Aversion and Seller Behavior: Evidence from the Housing Market', *Quarterly Journal of Economics* 116(4), 1233–1260.

Gilboa, Itzhak and Schmeidler, David. (1989). 'Maxmin Expected Utility with Non-Unique Prior', *Journal of Mathematical Economics* 18(2), 141–153.

Glaeser, Edward L. (2013). 'A Nation of Gamblers: Real Estate Speculation and American History', *American Economic Review* 103(3), 1–42.

Glaeser, Edward L. and Nathanson, Charles G. (2016). 'An Extrapolative Model of House Price Dynamics'. Working Paper.

Graham, John R. and Harvey, Campbell R. (2003). 'Expectations of Equity Risk Premia, Volatility and Asymmetry'. Working Paper.

Greenwood, Robin and Shleifer, Andrei. (2014). 'Expectations of Returns and Expected Returns', *Review of Financial Studies* 27(3), 714–746.

Guren, Adam M. (2016). 'House Price Momentum and Strategic Complementarity'. Working Paper.

Guren, Adam M. and McQuade, Timothy J. (2015). 'How do Foreclosures Exacerbate Housing Downturns?'. Working Paper.

Han, Lu and Strange, William C. (2015), The Microstructure of Housing Markets: Search, Bargaining, and Brokerage, *in* **Gilles Duranton, J. Vernon Henderson and William C. Strange**., eds, 'Handbook of Regional and Urban Economics', Vol. 5B, Elsevier, chapter 13, pp. 813–886.

Hanushek, Eric A. and Quigley, John M. (1980). 'What is the Price Elasticity of Housing Demand?', *Review of Economics and Statistics* 62(3), 449–454.

Harrison, J. Michael and Kreps, David M. (1978). 'Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations', *Quarterly Journal of Economics* 92(2), 323–336.

Haughwout, Andrew, Lee, Donghoon, Tracy, Joseph and van der Klaauw, Wilbert. (2011). 'Real Estate Investors, the Leverage Cycle, and the Housing Market Crisis'. Federal Reserve Bank of New York Staff Report no. 514.

Head, Allen, Lloyd-Ellis, Huw and Sun, Hongfei. (2014). 'Search, Liquidity, and the Dynamics of House Prices and Construction', *American Economic Review* 104(4), 1172–1210.

Hong, Harrison and Stein, Jeremy C. (1999). 'A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets', *Journal of Finance* 54(6), 2143–2184.

Hong, Harrison and Stein, Jeremy C. (2007). 'Disagreement and the Stock Market', *Journal of Economic Perspectives* 21(2), 109–128.

Keynes, John Maynard. (1936), *The General Theory of Employment, Interest and Money*, Macmillan, New York.

Kindleberger, Charles P. (1978), *Manias, Panics, and Crashes: A History of Financial Crises*, Wiley, 4th Edition.

Kuchler, Theresa and Zafar, Basit. (2016). 'Personal Experiences and Expectations About Aggregate Outcomes'. Working Paper.

Mankiw, N. Gregory and Reis, Ricardo. (2002). 'Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve', *Quarterly Journal of Economics* 117(4), 1295–1328.

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

Nathanson, Charles G. and Zwick, Eric. (2015). 'Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market'. Working Paper.

Ngai, L. Rachel and Sheedy, Kevin D. (2016). 'The Decision to Move House and Aggregate Housing-Market Dynamics'. Working Paper.

Ngai, Rachel and Tenreyro, Silvana. (2014). 'Hot and Cold Seasons in the Housing Market', *American Economic Review* 104(12), 3991–4026.

Ortalo-Magné, François and Rady, Sven. (2006). 'Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints', *Review of Economic Studies* 73(2), 459–485.

Pástor, L'uboš and Veronesi, Pietro. (2006). 'Was There a Nasdaq Bubble in the Late 1990s?', *Journal of Financial Economics* 81(1), 61–100.

Pástor, L'uboš and Veronesi, Pietro. (2009). 'Technological Revolutions and Stock Prices', *American Economic Review* 99(4), 1451–1483.

Penasse, Julien and Renneboog, Luc. (2016), Bubbles and Trading Frenzies: Evidence from the Art Market. Working Paper.

Piazzesi, Monika and Schneider, Martin. (2009). 'Momentum Traders in the Housing Market: Survey Evidence and a Search Model', *American Economic Review* 99(2), 406–411.

Poterba, James M. (1984). 'Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach', *Quarterly Journal of Economics* 99(4), 729–752.

Rajan, Raghuram and Ramcharan, Rodney. (2015). 'The Anatomy of a Credit Crisis: The Boom and Bust in Farm Land Prices in The United States in the 1920s', *American Economic Review* 105(4), 1439–1477.

Scheinkman, Jose A. and Xiong, Wei. (2003). 'Overconfidence and Speculative Bubbles', *Journal of Political Economy* 111(6), 1183–1220.

Shiller, Robert J. (1981). 'Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?', *American Economic Review* 71(3), 421–436.

Shiller, Robert J. (2005), *Irrational Exuberance*, Princeton University Press.

Shleifer, Andrei and Vishny, Robert W. (1992). 'Liquidation Values and Debt Capacity: A Market Equilibrium Approach', *Journal of Finance* 47(4), 1343–1366.

Simsek, Alp. (2013). 'Speculation and Risk Sharing with New Financial Assets', *Quarterly Journal of Economics* p. qjt007.

Smith, Vernon L., Suchanek, Gerry L. and Williams, Arlington W. (1988). 'Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets', *Econometrica* 56(5), 1119–1151.

Stein, Jeremy C. (1995). 'Prices and Trading Volume in the Housing Market: A Model with Down-Payment Effects', *Quarterly Journal of Economics* 110(2), 379–406.

Stiglitz, Joseph E. (1989). 'Using Tax Policy To Curb Speculative Short-Term Trading', *Journal of Financial Services Research* 3, 101–115.

Summers, Lawrence H. and Summers, Victoria P. (1989). 'When Financial Markets Work Too Well: A Cautious Case For a Securities Transaction Tax', *Journal of Financial Services Research* 3, 261–286.

Tobin, James. (1978). 'A Proposal for International Monetary Reform', *Eastern Economic Journal* 4(3), 153–159.

Vissing-Jorgensen, Annette. (2004), Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions, *in* 'NBER Macroeconomics Annual 2003, Volume 18', The MIT Press, pp. 139–208.

Wheaton, William C. (1990). 'Vacancy, Search, and Prices in a Housing Market Matching Model', *Journal of Political Economy* 98(6), 1270–1292.

FIGURE 1
The Joint Dynamics of Prices and Volume

(a) US Housing Market (2000–2012)

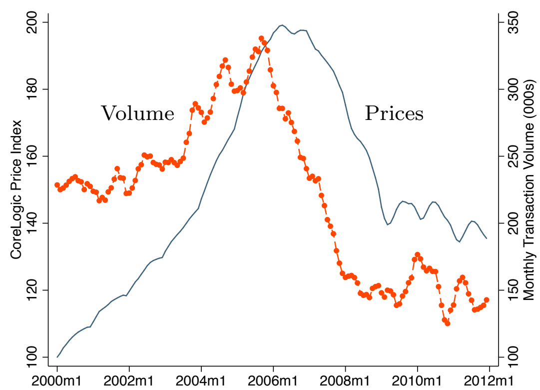

(b) US Equities, Tech (1995–2005)

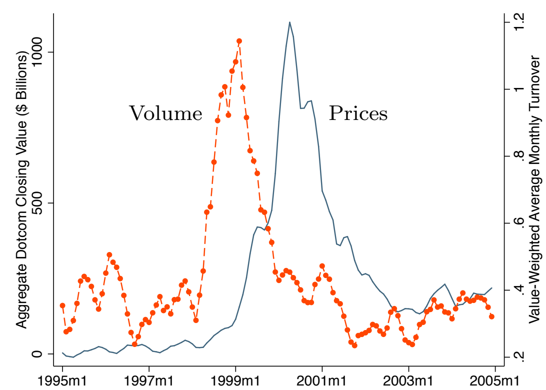

(c) Experimental Markets, SSW (1988)

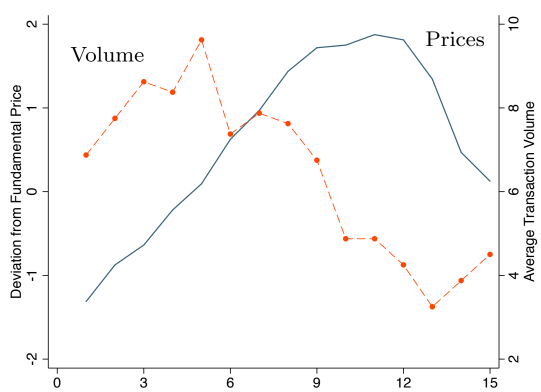

(d) Japan Equities (1985-1995)

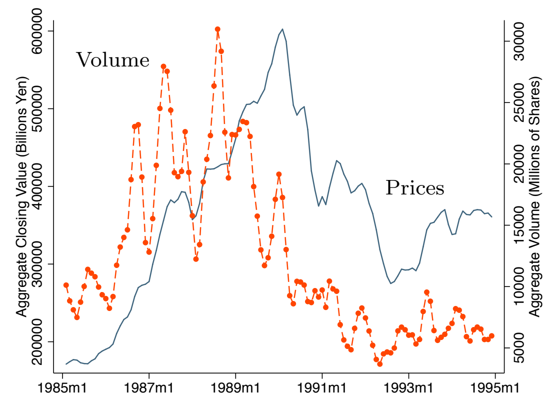

Notes: These figures display the dynamic relationship between prices and transaction volume for four distinct bubble episodes: (a) the 2000-2012 US housing market, (b) the 1995-2005 market in technology stocks, (c) the experimental bubbles studied by Smith et al. (1988), and (d) the 1985-1995 Japanese stock market.

FIGURE 2
Expected Holding Times of Homebuyers, 2008-2015

(a) Response Heterogeneity

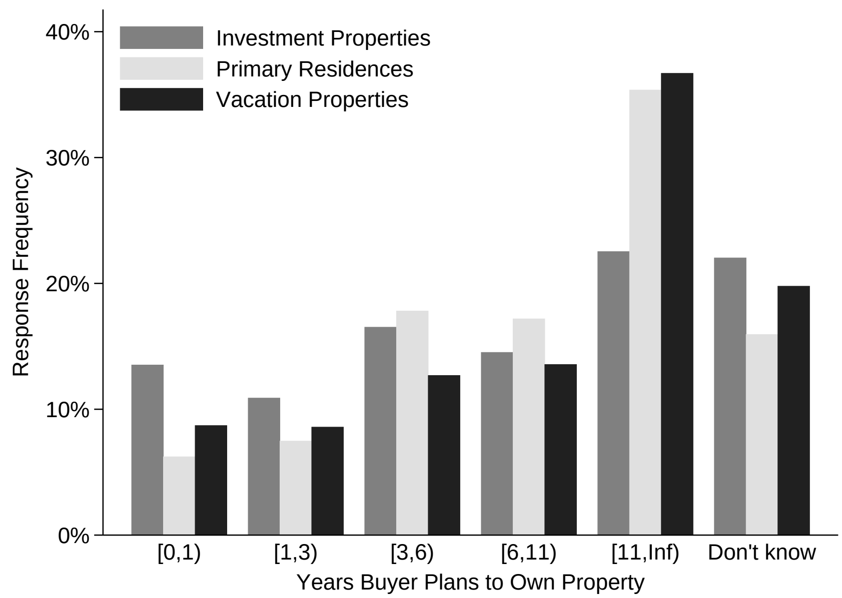

(b) Short-Term Buyers and Recent House Price Growth

Notes: Data come from the annual Investment and Vacation Home Buyers Survey conducted by the National Association of Realtors. We reclassify buyers who have already sold their properties by the time of the survey as having an expected holding time in [0,1). Panel (a) plots the response frequency averaged equally over each year from 2008 to 2015. In Panel (b), "annual house price growth" equals the average across that year's four quarters of the log change in the FHFA US house price index from four quarters ago, and "short-term buyer share" equals the share of respondents other than those reporting "don't know" who report a horizon less than three years.

FIGURE 3
The Dynamics of Prices, Volume, and Inventories

(a) Prices and Volume (2000–2012)

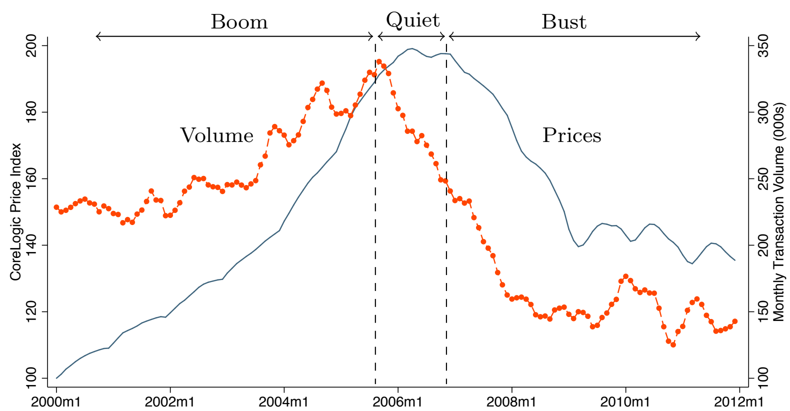

(b) Prices and Inventories (2000–2012)

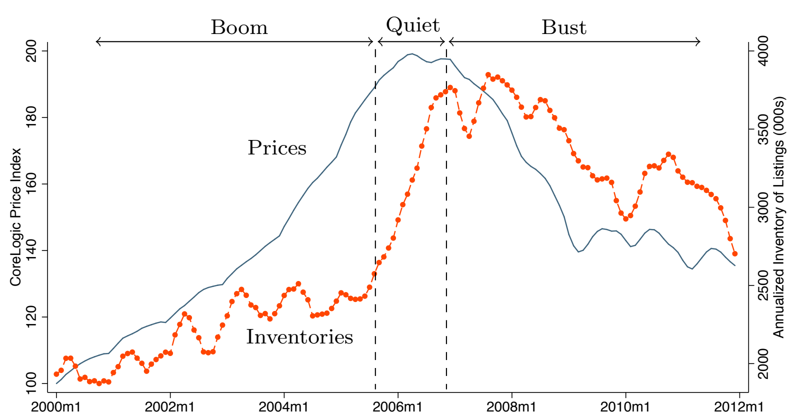

Notes: These figures display the dynamic relationship between prices, transaction volume, and the inventory of listings in the US between 2000 and 2015. Panel (a) plots monthly prices and sales volume and panel (b) plots monthly prices and inventory. For prices we use the national CoreLogic single family home price index, which is based on repeat sales. Inventory information comes from the National Association of Realtors. For volume, we plot the smoothed, seasonally-adjusted count of transactions in our sample of 115 Metropolitan Statistical Areas. We seasonally adjust volume by removing calendar month fixed effects and smooth the subsequent series using a three-month moving average. We apply the same smoothing and seasonal adjustment to inventory levels.

FIGURE 4
Simulation Results

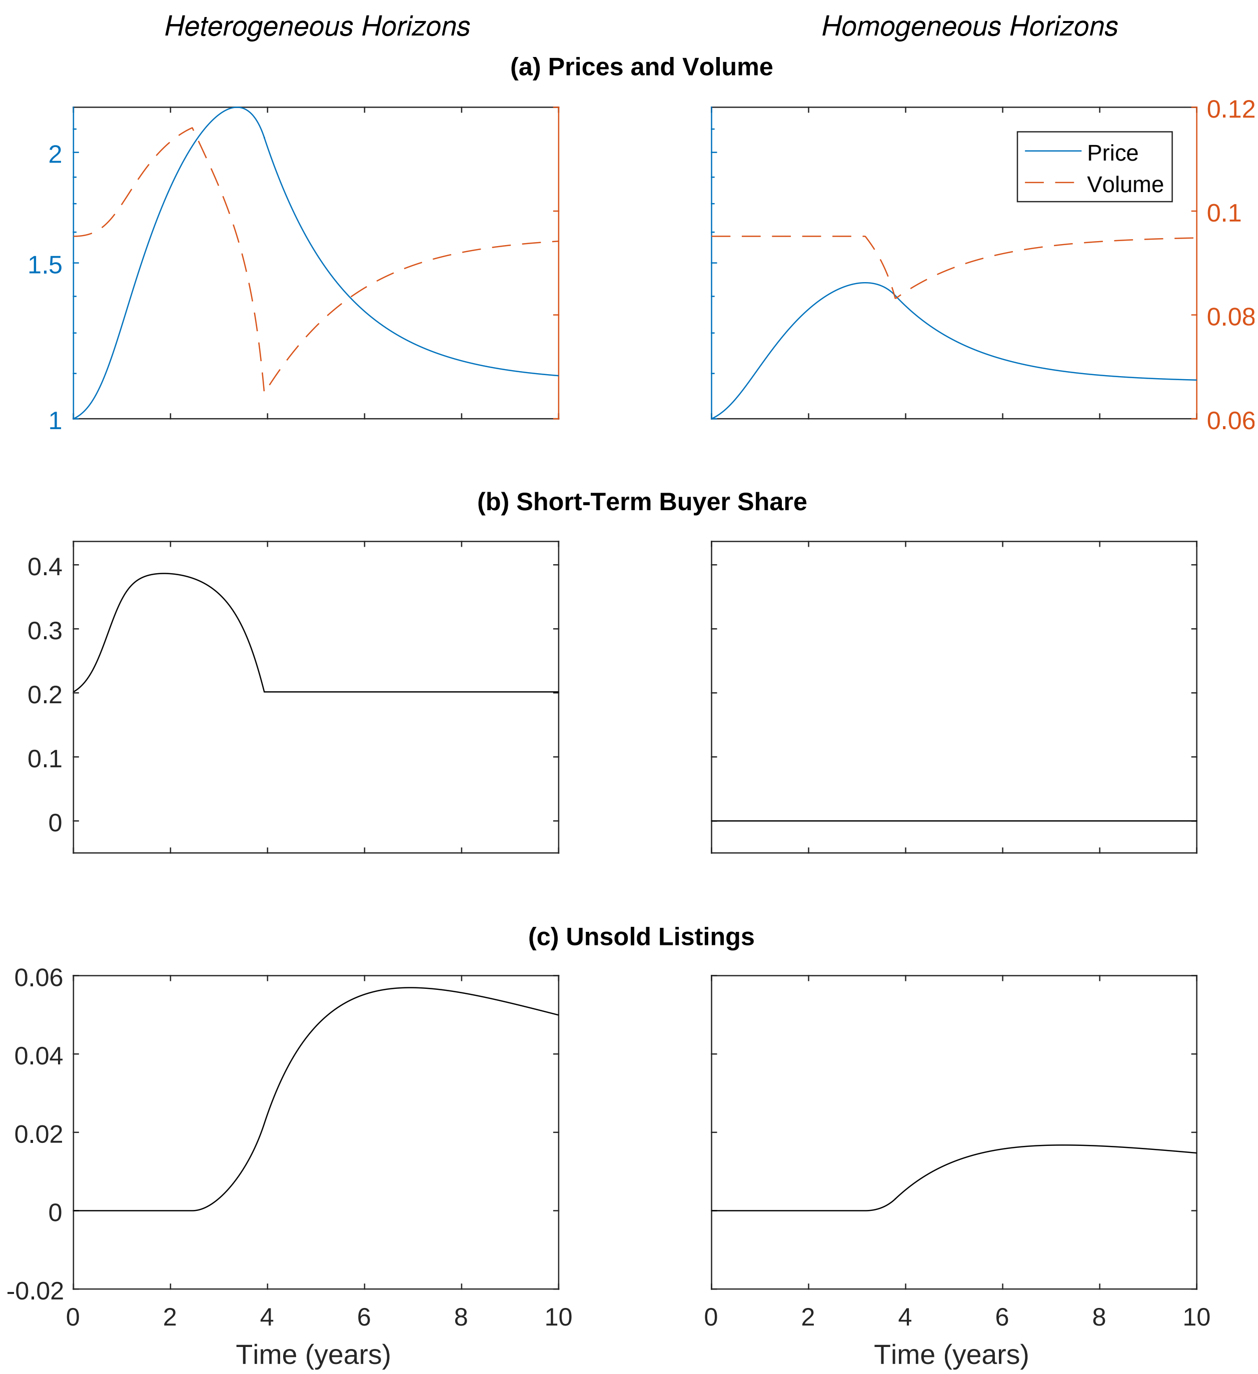

Notes: (a) The price when the demand shock occurs is normalized to 1. Volume is in units of share of the housing stock. (b) Here, "short-term buyer" is defined as one whose expected holding time is less than 2 years. The panel plots the share of such individuals among all buyers. (c) Unsold listings are in units of share of the housing stock.

FIGURE 5

The Dynamics of Holding Times in the Housing Market

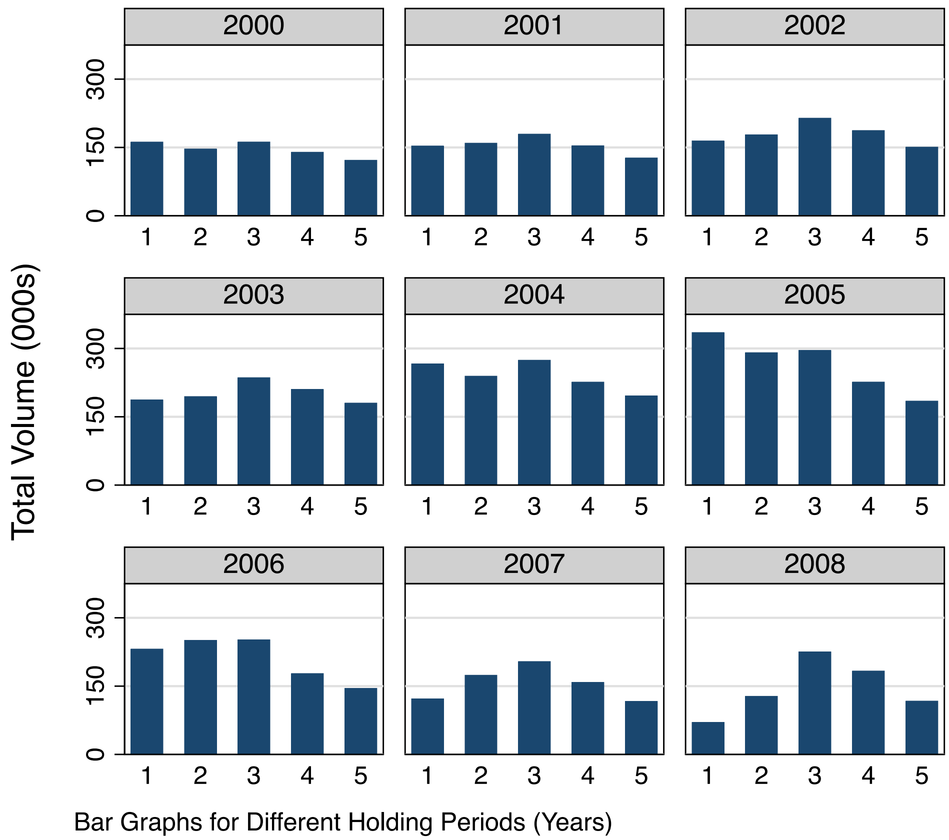

Notes: This figure illustrates the time variation in holding periods during the 2000-2008 housing cycle in the US. For each transaction, we define the holding period as the number of days since the last transaction of the same property. We then group all transactions with holding periods less than or equal to 5 years into bins of 1, 2, 3, 4, or 5 years, respectively. For each year between 2000–2008, we plot aggregate transaction counts in each of these five holding period groups.

FIGURE 6

The Role of Short Holding Period Volume Growth for Total Volume Growth

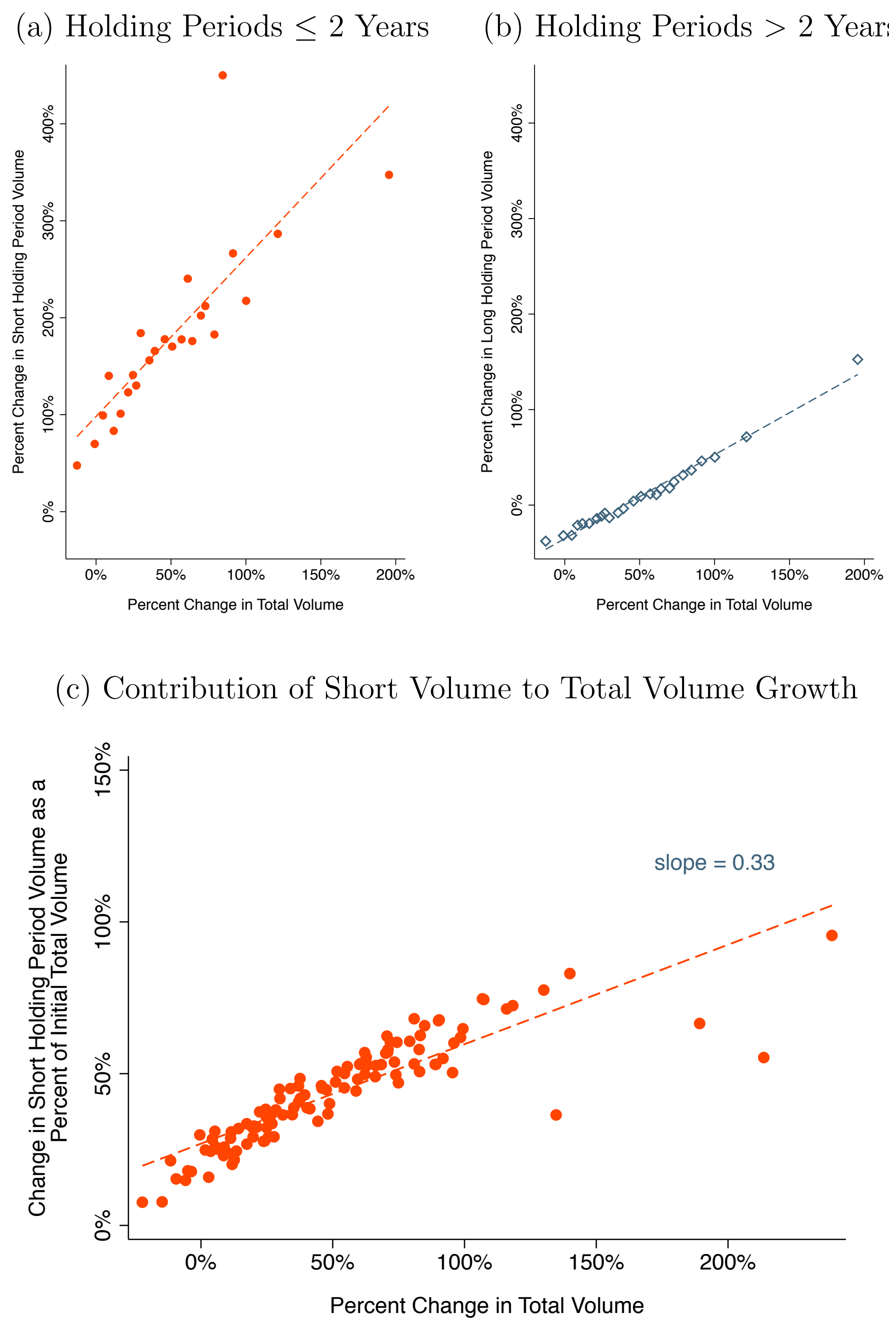

Notes: This figure illustrates of the quantitative importance of short holding period volume in accounting for the increase in total volume between 2000 and 2005. We present binned scatter plots ("binscatters") of the percent change in total volume from 2000–2005 versus the percent change in volume for short holding periods (≤ 2 years) in Panel (a) and long holding periods (> 2 years) in Panel (b). Panel (c) shows that the growth in short holding period volume is a quantitatively important component of the growth in total volume across cities. For each city, we plot the change in short holding time volume divided by initial total volume on the y-axis against the percent change in total volume on the x-axis.

FIGURE 7

The Role of Investor Volume Growth for Total Volume Growth

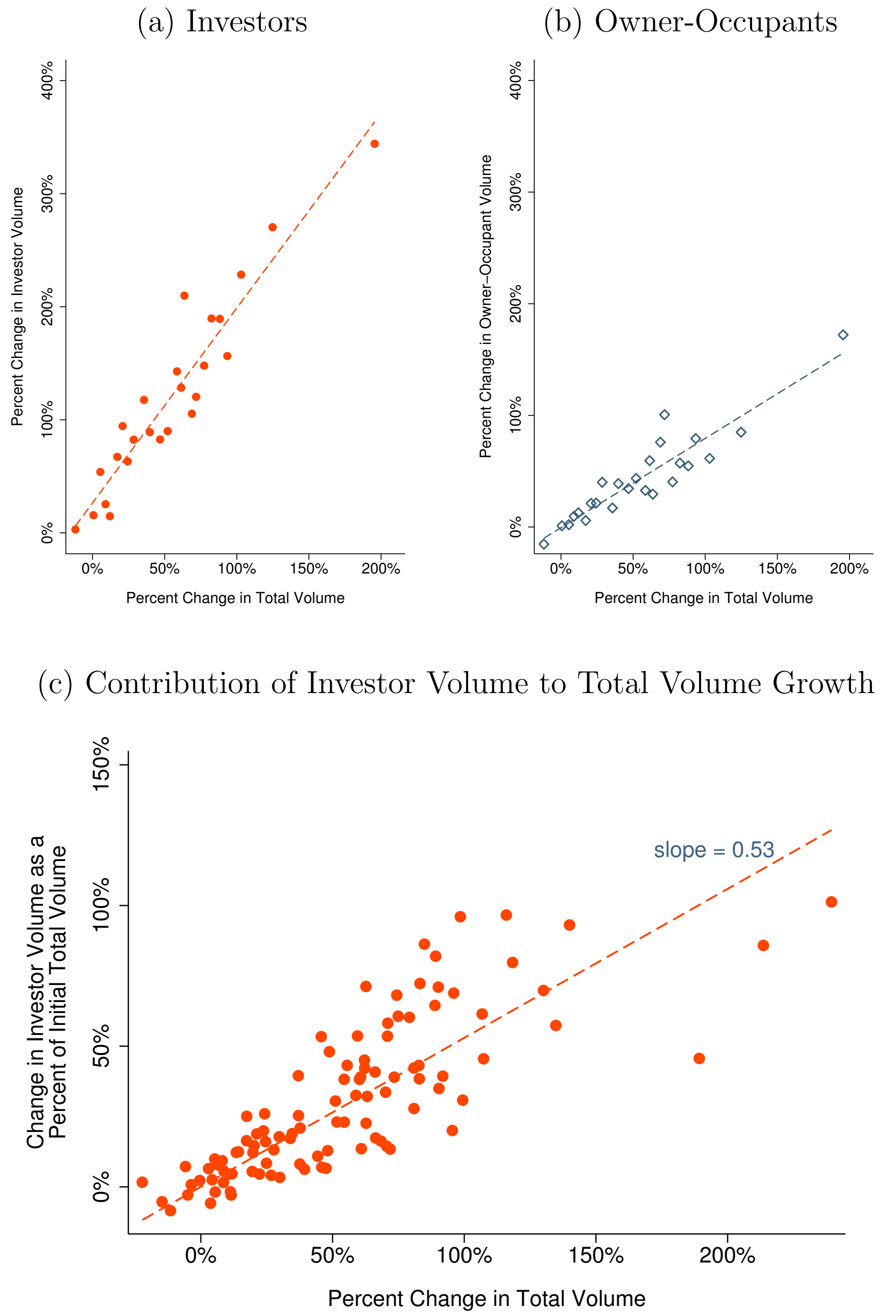

Notes: This figure illustrates of the quantitative importance of investor volume in accounting for the increase in total volume between 2000 and 2005. We present binned scatter plots ("binscatters") of the percent change in total volume from 2000–2005 versus the percent change in volume for investors (defined as transactions with distinct mailing and property addresses) in Panel (a) and owner-occcupants (defined as non-investors or having a missing mailing address) in Panel (b). Panel (c) shows that the growth in investor volume is a quantitatively important component of the growth in total volume across cities. For each city, we plot the change in investor time volume divided by initial total volume on the y-axis against the percent change in total volume on the x-axis.

FIGURE 8
The Correlation between Prices and Volume at Various Lags

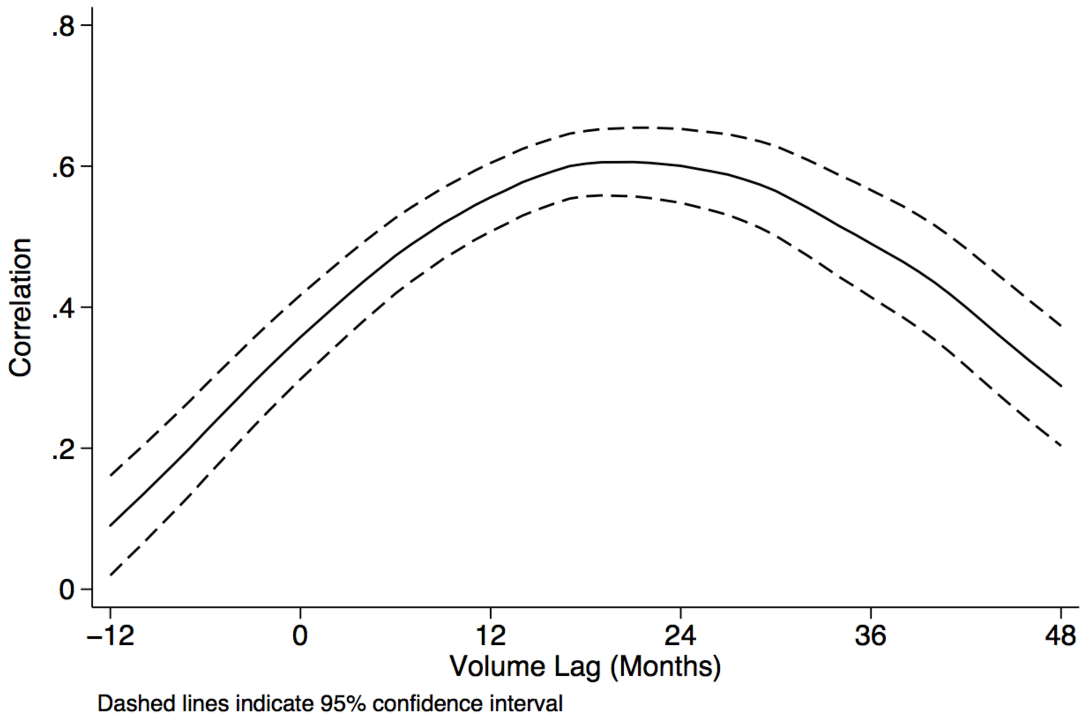

Notes: This figure shows that the correlation between prices and lagged volume is robust across cities and maximized at a positive lag of approximately 15 months. We regress the demeaned log of prices on seasonally adjusted lagged volume divided by the 2000 housing stock for each lag from -12 months to 48 months and plot the implied correlation and its 95% confidence interval calculated using standard errors that are clustered by month.

FIGURE 9
Sale and Listing Probabilities for U.S. Existing Homes, 1999-2015

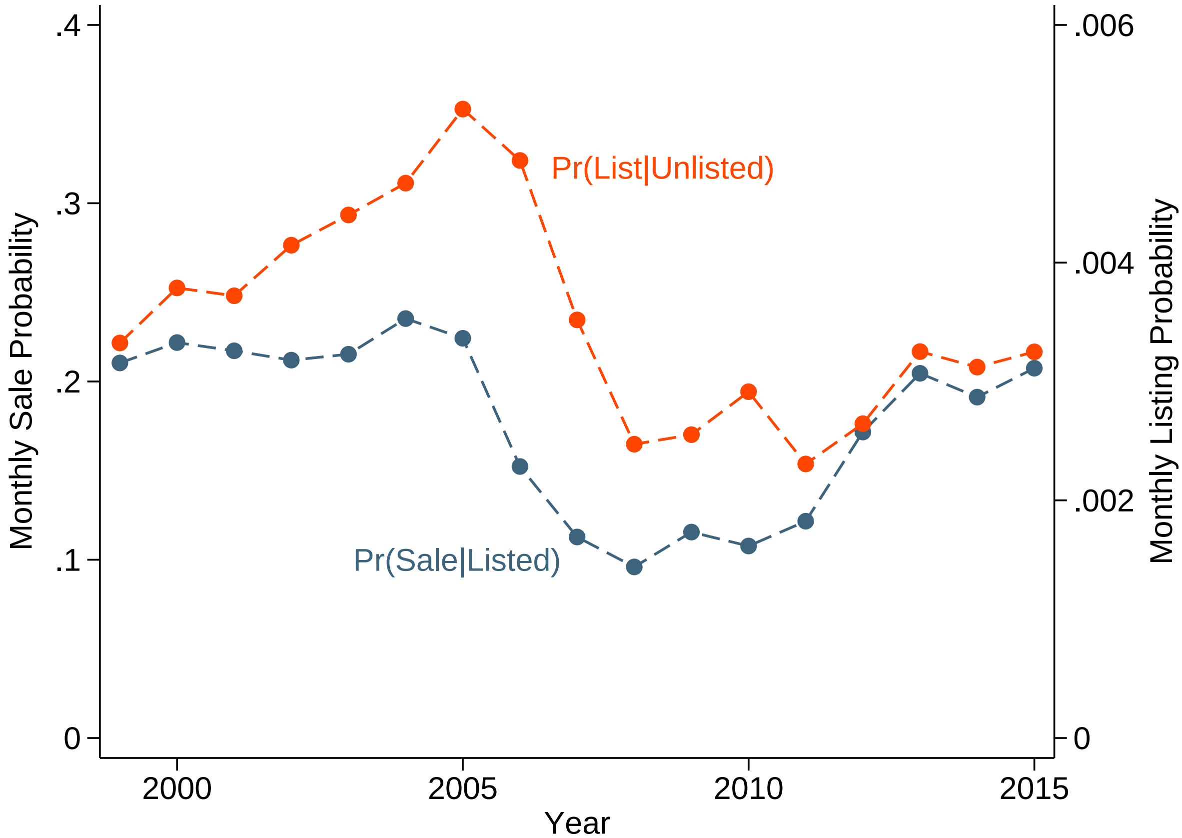

Notes: The figure displays the average monthly sale probability and listing probability for each year. For each month, the sale probability equals sales divided by listed inventory, and the listing probability equals new listings divided by the stock of unlisted houses. Monthly data on the U.S. housing stock are interpolated from quarterly estimates provided by the U.S. Census, and monthly sales and inventory numbers come from data provided by the National Association of Realtors; new listings are calculated using the inventory and sales data.

FIGURE 10

Initial Volume and the Magnitude of the Housing Boom and Bust

(a) Boom

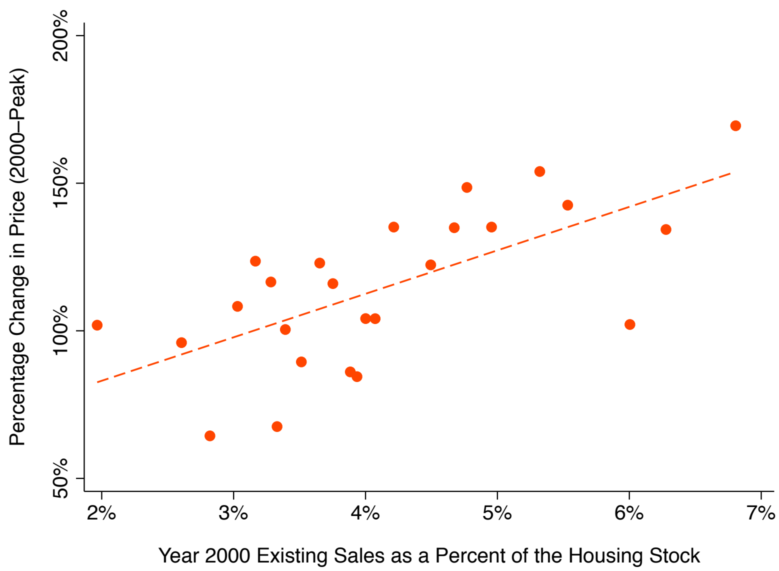

(b) Bust

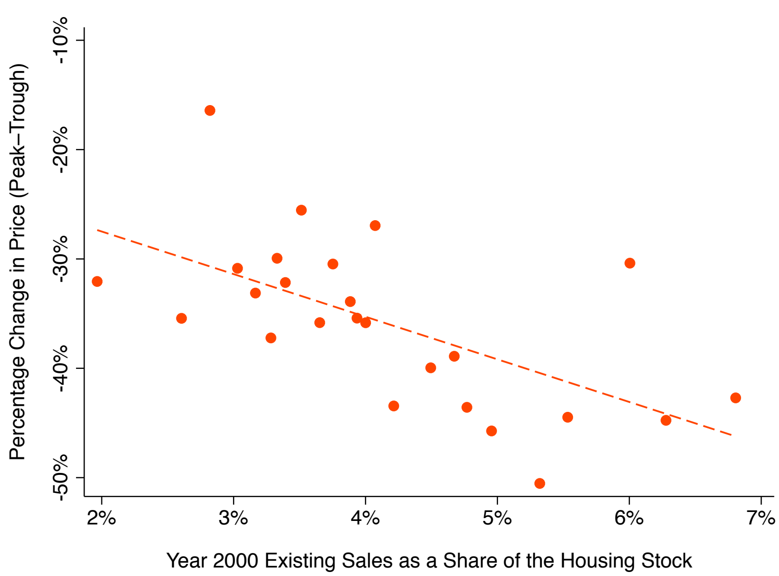

Notes: This figure provides empirical support for the cross-sectional prediction that the magnitude of price swings during boom-bust episodes should be correlated with the level of steady-state transaction volume across markets. We present binned scatter plots ("binscatters") of the percent change in prices from January 2000 to peak (Panel (a)) and from peak to trough (Panel (b)) versus total existing homes in 2000. To facilitate comparisons across cities of different sizes, we normalize existing sales by the size of the housing stock in 2000 for each city. House prices are measured using the monthly CoreLogic repeat-sales house price indicies. The price peak for each MSA is measured as the highest price recorded for that MSA prior to January, 2012. The trough is measured as the lowest price subsequent to the month in which the peak occurred.

TABLE 1
Calibration Sources

Calibrated Quantity	Role in Model	Source or Assumed Value
$f(\cdot)$	Distribution of expected holding times	NAR Investment and Vacation Home Buyers Survey
μ	Relative weight in expectations on recent price changes versus those in distant past	Estimation using NAR survey
$g(\cdot)$	Forward term structure of expectations	Survey of Consumer Expectations; Armona et al. (2016)
$\gamma(\cdot)$	Extrapolation function	Estimation using Case et al. (2012) survey
r	Discount rate	0.07
c	Price stickiness	1
ϵ	Housing demand elasticity	0.6
A^f/A^i	Size of demand shock	1.06

Notes: This table lists the model quantities we calibrate to produce Figure 4. Further details are provided in Section 3.5.

TABLE 2
Sensitivity of Simulation Results to Parameters

	f	ρ	μ	c	ϕ	ϵ	A^f/A^i	r
				A. Excess Price Boom				
Low	0.30	0.11	0.83	0.40	0.00	0.36	0.92	1.37
Baseline	1.05	1.05	1.05	1.05	1.05	1.05	1.05	1.05
High	3.31	1.46	0.99	1.48	0.01	3.46	1.17	0.68
				B. Volume Boom				
Low	0.00	0.04	0.15	0.11	0.00	0.05	0.21	0.26
Baseline	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22
High	0.94	0.26	0.23	0.19	0.00	0.86	0.23	0.17
				C. Maximal Unsold Listings				
Low	0.02	0.01	0.05	0.04	0.00	0.02	0.06	0.07
Baseline	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06
High	0.18	0.07	0.05	0.05	0.00	0.16	0.06	0.04

Notes: The excess price boom equals $P^{max}/P_\infty - 1$, the volume boom equals $V^{max}/V_0 - 1$, and maximal unsold listings equal $\max_t I_t$. The alternate values for f ρ, and ϵ are described in the text. The low and high values for the remaining parameters are half and double their baseline values (we half and double $\log(A^f/A^i)$).

TABLE 3
Within-MSA Role of Short Holding Period and
Investor Volume Growth for Total Volume Growth (2000–2005)

	Holding Period			Occupancy Status		
	%Δ Short Volume	%Δ Long Volume	Δ Short/ Total Volume	%Δ Investor Volume	%Δ Occupant Volume	Δ Investor/ Total Volume
%Δ Total Volume	1.195***	0.987***	0.275***	1.523***	0.953***	0.340***
	(0.105)	(0.026)	(0.013)	(0.136)	(0.091)	(0.032)
R^2	0.51	0.91	0.73	0.41	0.55	0.73
Observations	6,756	6,771	6,759	5,825	5,758	6,677

Notes: This table reports estimates of the quantitative importance of short holding period and investor volume in accounting for the increase in total volume between 2000 and 2005 at the zip-code level. Each column reports estimates from a separate regression of the change in a given component of volume on the percent change in total volume in the zip-code. All specifications include MSA fixed effects. A short holding period is defined as any holding period less than or equal to two years. Investors and owner-occupants are identified using information on the mailing address of the property as described in the text. In columns 3 and 6, we divide the level change in short holding period and investor holding period volume by total volume in the zip-code in 2000. All other changes are expressed as percent changes from the 2000 level for the indicated type of transaction. Zip codes are weighted according to their total volume in 2000. Standard errors are reported in parentheses and are clustered at the MSA level. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE 4
Initial Volume and the Magnitude of the Housing Boom and Bust

	Percentage Change in Prices			
	2000–Peak	2000–2006	Peak–Trough	2006–Trough
Existing Sales/Stock (2000)	15.046***	14.933***	–4.059***	–4.508***
	(3.655)	(3.573)	(1.120)	(1.218)
R^2	0.12	0.12	0.12	0.11
Observations	115	115	115	115

Notes: This table reports estimates of the cross-sectional relationship between the magnitude of the housing boom and bust and initial transaction volume at the MSA level. Each column reports estimates from a separate regression where the dependent variable is the percentage change in prices measured over the indicated horizon. Initial transaction volume is measured as total year 2000 existing home sales in each MSA scaled by the total number of housing units in the MSA as reported in the 2000 Census. House prices are measured using the monthly CoreLogic repeat-sales house price indicies. The price peak for each MSA is measured as the highest price recorded for that MSA prior to January, 2012. The trough is measured as the lowest price subsequent to either the month in which the peak occurred (column 3) or January, 2006 (column 4). In columns 1, 2, and 4, price changes are calculated using the January price level in 2000 and 2006. Heteroskedasticity robust standard errors are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

FIGURE A1
Interest in "House Flipping" over the Housing Cycle

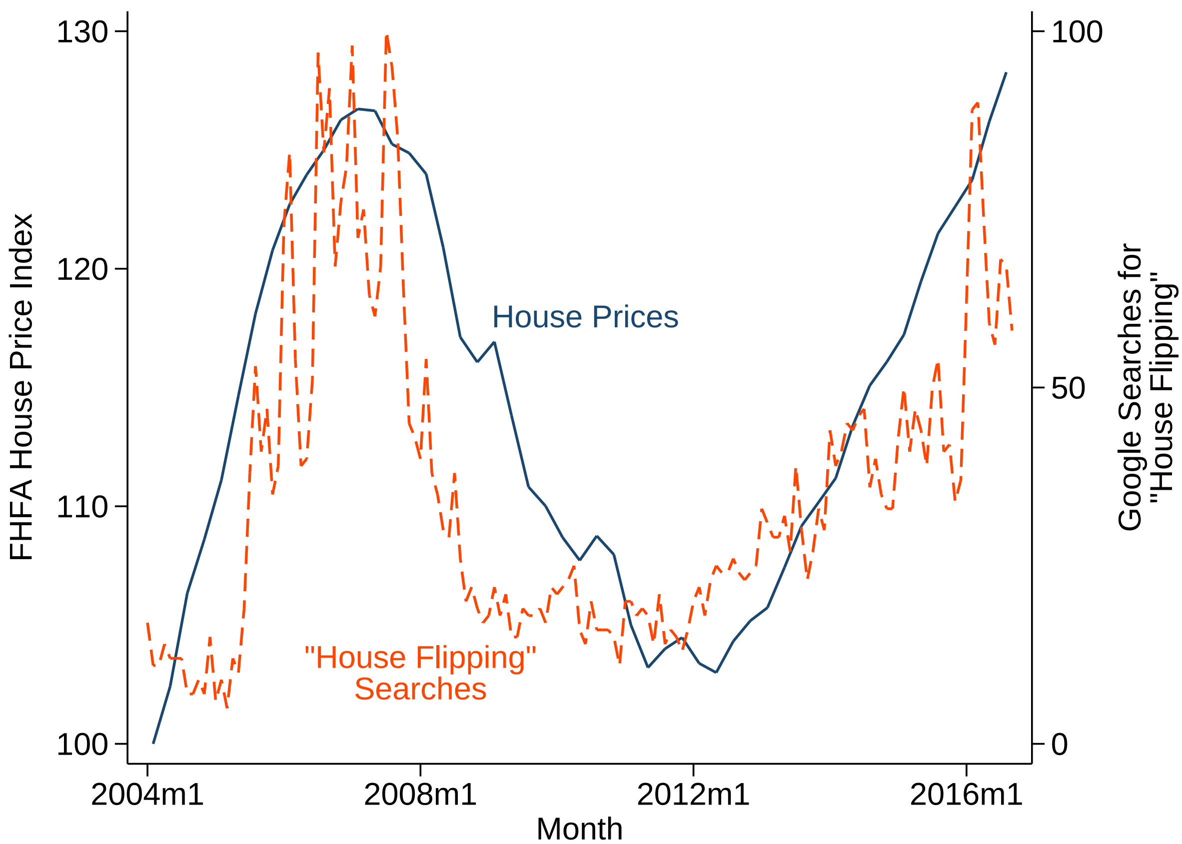

Notes: Search data were downloaded from Google Trends on January 31, 2017. House price index is nominal and quarterly. We normalize it to 100 in 2004q1 and map quarters to months using the midpoints.

FIGURE A2
Phase Diagram

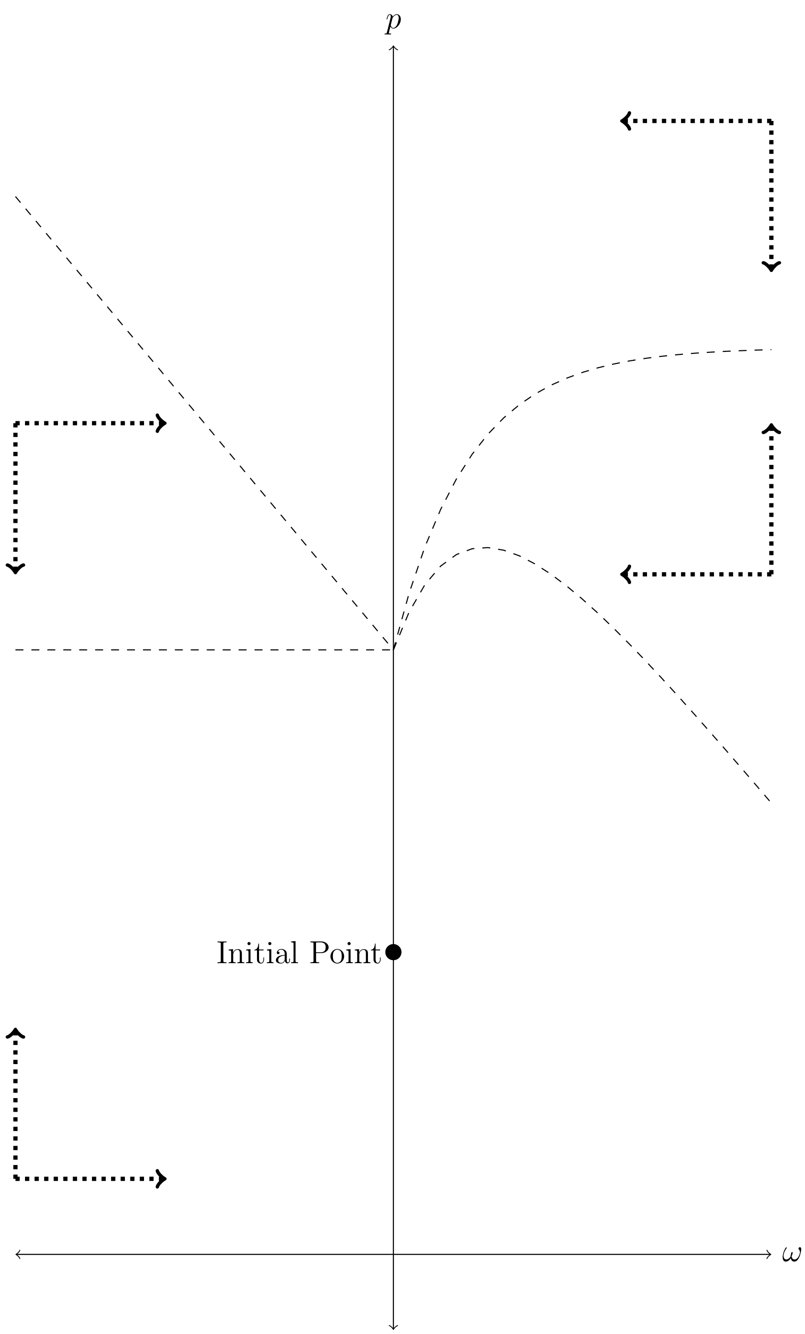

Notes: This figure illustrates the phase diagram for the (p, ω) system specified by equations (A2) and (A3); p denotes the log house price, and ω denotes the historical average log price change given by equation (10). The dashed loci indicate points at which either $\dot{p} = 0$ or $\dot{\omega} = 0$. The dotted arrows indicate the directions p and ω move in each of the four areas demarcated by the dashed loci. The system begins at the marked point on the p-axis.

TABLE A1
Sensitivity of S&P 500 Return Forecasts to
Historical Returns, 2000Q3 - 2011Q4

	Lagged Annual Return History		75-Year Weighted Average	
	One-Year Forecast	Ten-Year Forecast (Annualized)	One-Year Forecast	Ten-Year Forecast (Annualized)
Historical Return	0.029**	−0.013*	0.045***	0.011
	(0.014)	(0.007)	(0.010)	(0.015)
R^2	0.12	0.06	0.17	0.03
Observations	44	44	44	44

Notes: Return forecasts come from the Duke CFO Global Business Outlook, a quarterly survey of chief financial officers of U.S. firms. Historical returns on the S&P 500 come from CRSP's daily dividend-inclusive value-weighted return series `vwretd`. The historical return equals $P_t/P_{t-1} - 1$ in the first two columns and $\mu(1 - e^{-\mu T})^{-1} \int_0^T e^{-\mu\tau} \dot{p}_{t-\tau} d\tau$ in the latter two columns, with $p = \log P$, $\mu = 0.5$, and $T = 75$. The sample period is chosen to match that used by Greenwood and Shleifer (2014). Observations for 2001Q3 and 2002Q3 are dropped due to errors or gaps in the Duke CFO Global Business Outlook. Newey-West standard errors are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.