

NBER WORKING PAPER SERIES

SPECULATIVE DYNAMICS OF PRICES AND VOLUME

Anthony A. DeFusco
Charles G. Nathanson
Eric Zwick

Working Paper 23449
<http://www.nber.org/papers/w23449>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2017, Revised August 2022

We thank Michael Fishman, Andreas Fuster, Stefano Giglio, Edward Glaeser, Adam Guren, Sam Hanson, Amir Kermani, Stijn Van Nieuwerburgh, Alp Simsek, Johannes Stroebel, Lawrence Summers, Richard Thaler, Rob Vishny, and Wei Xiong for helpful comments. William Cassidy, Jessica Henderson, Saul Ioffie, Harleen Kaur, Laurence O'Brien, 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Anthony A. DeFusco, Charles G. Nathanson, and Eric Zwick. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Speculative Dynamics of Prices and Volume

Anthony A. DeFusco, Charles G. Nathanson, and Eric Zwick

NBER Working Paper No. 23449

May 2017, Revised August 2022

JEL No. E32,G02,G12,R3

ABSTRACT

Using data on 50 million home sales from the last U.S. housing cycle, we document that much of the variation in volume came from the rise and fall in speculation. Cities with larger speculative booms have larger price booms, sharper increases in unsold listings as the market turns, and more severe busts. We present a model in which predictable price increases endogenously attract short-term buyers more than long-term buyers. Short-term buyers amplify volume by selling faster and destabilize prices through positive feedback. Our model matches key aggregate patterns, including the lead-lag price-volume relation and a sharp rise in inventories.

Anthony A. DeFusco
Kellogg School of Management
Northwestern University
2211 Campus Drive
Evanston, IL 60208
and NBER
anthony.defusco@kellogg.northwestern.edu

Eric Zwick
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
ericzwick@gmail.com

Charles G. Nathanson
Kellogg School of Management
Northwestern University
2211 Campus Drive, Room 4479
Evanston, IL 60208-0898
nathanson@kellogg.northwestern.edu

The housing market in the United States underwent a tumultuous cycle between 2000 and 2011. The rise and fall in house prices caused several problems for the U.S. economy. During the boom, a surge in housing investment drew resources into construction from other sectors (Charles et al., 2018) and contributed to a capital overhang that slowed the economic recovery (Rognlie et al., 2017). During the bust, millions of households lost their homes in foreclosure, and falling house prices led many others to cut consumption (Mayer et al., 2009; Mian et al., 2013, 2015; Guren and McQuade, 2020). Large real estate cycles are not unique to the U.S. (Mayer, 2011) or to this time period (Case, 2008; Glaeser, 2013). Given the economic costs of these recurring episodes, understanding their cause is critical.

This paper presents evidence that speculation was a key driver of this real estate cycle.¹ Three stylized facts from the cycle guide our analysis. First, prices and volume jointly rise and fall throughout the cycle. Second, volume falls before prices, resulting in a pronounced lead-lag relation between prices and volume. Third, the period during which prices continue to rise despite falling volume coincides with rapidly accumulating unsold listings. We refer to this period as the *quiet*, which is preceded by the *boom* and followed by the *bust*. These stylized facts hold on average across cities and are especially pronounced in cities with larger cycles. They suggest that focusing on who was most active during each phase of the cycle can provide insight on the underlying mechanisms.

We study the behavior of speculative homebuyers during each of these three phases of the housing cycle using transaction-level data from CoreLogic on 50 million home sales between 1995 and 2014. We measure speculative buying and selling across 115 metropolitan statistical areas (MSAs), which represent 48% of the U.S. housing stock. We pursue two complementary approaches to identify speculative activity. First, following Bayer et al. (2020), we classify transactions based on their realized holding periods, denoting those buyers who resell the property within three years as short-term buyers. Second, following Chincio and Mayer (2015), we classify transactions based on the inferred occupancy status of the property, denoting buyers who list a mailing address distinct from the property address as non-occupant buyers. We supplement our transaction data with a separate CoreLogic data set on homes listed for sale, sourced from a consortium of local multiple listing service (MLS) boards. We link these data to transaction records to study the role of speculative buyers for

¹Harrison and Kreps (1978, p. 323) define speculation as follows: “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.”

inventory dynamics across MSAs.

The data reveal a strong relation between the differential entry of speculative buyers and the size of the cycle. While overall volume increases substantially during the boom of 2000–2005, both short-term and non-occupant volume rise dramatically more. In an accounting sense, growth in speculative volume explains 40% to 50% of total volume growth. This relation is also strong in the cross-section, as speculative volume growth can account for 30% to 50% of total volume growth across MSAs. Cities with stronger speculative volume booms also experience larger house price booms: MSAs with a one standard deviation larger short-volume and non-occupant boom see 25 and 15 percentage point larger cumulative price increases, respectively.

As the volume boom ends, price growth slows but remains positive, and unsold listings accumulate. Across MSAs, these patterns are more pronounced in cities with larger speculative volume booms. Our linked listing-transaction data further reveal that short-term buyers disproportionately contribute to the surge in aggregate inventories. MSAs with larger speculative volume booms also see substantially larger price busts, volume busts, and total foreclosures in the final phase of the cycle. We find that speculative volume is larger when house price growth over the past year is greater, which suggests that extrapolation—the belief that prices continue to rise after recent gains—draws speculators into the housing market. Consistent with our interpretation of the data, a National Association of Realtors survey reveals wide variation in expected holding times, shorter expected holding times among investors, and increases in the short-term buyer share following recent price gains.

In the second part of the paper, we provide a quantitative model to match these novel facts about the housing market. Our approach adapts core insights from Cutler et al. (1990), De Long et al. (1990), and Hong and Stein (1999) to study the housing market. As in these papers, extrapolation causes a predictable boom and bust in prices after a positive demand shock. In contrast, we relax the assumption of Walrasian market clearing, so that homes listed for sale may not sell immediately. To do so, we microfound extrapolation using the approach in Glaeser and Nathanson (2017) and then extend their framework to a non-Walrasian setting.

In our model, a mover attempts to sell her house by posting a list price. A potential buyer arrives and decides whether to purchase the house at that price. Potential buyers differ in the benefits they derive from owning a house; non-occupants benefit less than occupants. Buyers

also differ in the expected amount of time until becoming a mover; short-term buyers have shorter horizons *ex ante*. The average flow benefit of potential buyers fluctuates randomly over time. Agents cannot observe this demand process, but they observe the history of price growth and the share of listings that sell each period. Using these market data, agents infer the current level and growth rate of the demand process and optimally make decisions in light of these beliefs—the choice of list price for movers, and whether or not to purchase for potential buyers. As in Glaeser and Nathanson (2017), agents mistakenly believe that potential buyers neglect time variation in the growth rate when deciding whether to buy.

We study how our housing market responds to a large, unexpected increase to the growth rate of the demand process. The model matches key facts from our empirical work, including the lead–lag relation between prices and volume, the excess growth of short-term and non-occupant volume during the boom, and a growth in listings during the quiet coming disproportionately from short-holding-period sales. In the model, the quiet occurs when agents overestimate demand and believe it continues to grow, which causes movers to increase their list prices despite falling transaction volume.

We then use this setting to evaluate the effect of speculation on the housing cycle. When we shut down speculation by imposing rational expectations, almost all of the salient aspects of the housing cycle disappear or become quantitatively insignificant. We find similar patterns when we remove short-term and non-occupant buyers from the model. Therefore, speculators amplify the effects of non-rational expectations on prices and quantities over the housing cycle. Motivated by this result, we study transaction taxes on non-occupant buyers as well as on all buyers, as governments have used such taxes in attempts to curb speculation (Chi et al., 2021). Taxing all buyers attenuates the housing cycle, but even a large 5% tax on non-occupants has only a small effect on the price boom, price bust, and volume boom.

Previous and contemporaneous empirical work examines short-term buyers (Adelino et al., 2016; Bayer et al., 2021) and non-occupant buyers (Haughwout et al., 2011; Bhutta, 2015; Chincó and Mayer, 2015) in the housing market, as well as the importance of speculation for volume or prices (Gao et al., 2020; Bayer et al., 2020; Mian and Sufi, 2022). Our paper is the first to focus on the joint dynamics of volumes, prices, and inventories, along with speculative activity. We present stylized facts that any model of this episode should be able to match. Our focus on joint dynamics emphasizes the connection between speculation and the lead-lag relationship between prices and volume, a pattern which receives less at-

tion and has not been linked to speculation in past work. Beyond this, our data expands on past work through including more MSAs, non-mortgage sales, new microdata on homes listed for sale linked to prior transactions, and multiple measures of speculation.

Three strands of the literature theoretically explain the comovement of prices and volume in housing and other markets. In the first, investors disagree about asset values due to overconfidence (Daniel et al., 1998, 2001; Scheinkman and Xiong, 2003). The second exploits features specific to the housing market, such as credit constraints (Stein, 1995; Ortalo-Magné and Rady, 2006) or search and matching frictions (Wheaton, 1990; Díaz and Jerez, 2013; Head et al., 2014; Hedlund, 2016; Ngai and Sheedy, 2020; Anenberg and Bayer, 2020). The final strand incorporates psychology into models with extrapolative expectations to generate trade (Barberis et al., 2018; Liao and Peng, 2018). Some papers straddle multiple categories (Guren, 2014; Piazzesi and Schneider, 2009; Burnside et al., 2016). Relative to these studies, our model’s contribution is to simultaneously generate three key patterns from our empirical work: the existence of the quiet, the disproportionate growth in short-term volume during the boom and quiet, and the excess growth in non-occupant purchases during the boom. In addition, our model illustrates a mechanism for how speculation amplifies the housing cycle, allows us to disentangle the relative importance of short-term and non-occupant buyers, and provides a framework to evaluate the effects of transaction taxes on the housing market.

1 Data

In this section, we describe the data we use to establish the core motivating facts for our model and how we identify speculative buyers in that data. Further information regarding the data is in Online Appendix A.

1.1 Data sources and sample selection

Our main data come from CoreLogic, a private vendor that collects and standardizes publicly available tax assessments and deeds records from across the U.S., and include observations from 115 MSAs. In analyses that require us to identify an owner’s occupancy status, we use a subset of 102 MSAs for which we can be sure that there were no major changes in the way that mailing addresses were coded during our sample period. In Online Appendix A, we describe how we select these MSAs. Our analysis of the housing cycle covers the time

period 2000 through 2011 because measuring realized holding periods requires observing consecutive transactions.

We include all arms-length transactions of single-family homes, condos, or duplexes that occur at a non-zero price. We then drop a small number of duplicate transactions where the same property is observed selling 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. In Online Appendix A, we give the steps we follow to arrive at a final sample of 51,580,408 transactions. Given the geographic coverage of these data and their source in administrative records, our sample serves as a proxy for the population of transactions in the U.S. during the sample period.

Our listings data on individual homeowners is also provided by CoreLogic and is sourced from a consortium of local MLS boards throughout the country. For each listing, we observe the date the home was originally offered for sale, an indicator for whether the listing ever sold, and the date of sale for those that did. We link these data to the deeds data using the assessor’s parcel number (APN) for the property. When analyzing listings, we focus our attention on a subset of the 115 MSAs for which we can be relatively certain that the listings data are representative of the majority of owner-occupied home sales in the area. In Online Appendix A, we describe the approach we use to select these MSAs, leaving us with a final sample of 57 MSAs for our listings analysis.

We supplement these transaction- and listing-level data with national and MSA-level housing stock counts from the U.S. Census, national counts of sales and listings of existing homes from the National Association of Realtors (NAR), and national and MSA-level nominal house-price indices from CoreLogic. We also use survey data to study heterogeneity in expected holding horizons in the cross-section and over time. Each March, as part of the Investment and Vacation Home Buyers Survey, the 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 between 2008 and 2015.

1.2 Identifying Speculators

We identify speculators in our transaction-level data using two complementary approaches, each of which has been used in prior work. In the first approach, we categorize transactions based on their realized holding periods. We denote transactions held for less than three years as “short-term” sales and track the evolution of these sales over time. This approach follows Bayer et al. (2020), who classify speculators as those likely holding homes for short time periods for investment purposes. We similarly denote listings as short-term when the homeowner lists the house less than three years after buying it.

In the second approach, we classify homebuyers based on their occupancy status. Those who purchase a home without the intent to occupy it immediately are more speculative in the sense that a larger portion of their overall expected return is derived from capital gains rather than from the consumption value of living in the home. To identify these buyers, we follow Chinco and Mayer (2015) and mark buyers as non-occupants when the transaction lists the buyer’s mailing address as distinct from the property address. While this proxy may misclassify some non-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 of the level of non-occupant purchases.

One advantage of both methods is that they are based on the full sample of housing transactions. Other work has identified speculators based on the presence of multiple first-lien mortgage records in credit reporting data or self-reported occupancy status on loan applications (Haughwout et al., 2011; Gao et al., 2020; Mian and Sufi, 2022). While based on similar ideas, such approaches may omit a substantial fraction of speculative activity.

2 Dynamics of prices, volume, and inventory

In this section, we document the three phases of the housing cycle we mention above: boom, quiet, and bust. In Panel A of Figure 1, we plot aggregate trends in prices and volume between 2000 and 2011. In Panels B–E, we plot analogous series for four cities that represent regions with the largest boom–bust cycles during this time: Phoenix, AZ; Las Vegas, NV; Orlando, FL; and Bakersfield, CA. During the housing cycle, volume peaks before prices, and there is a sustained period during which volume is falling rapidly on high prices. This dynamic holds consistently across regions that experienced large price cycles. At the aggre-

gate level, volume rises to 150% of its level in 2000 and then falls back to this level before prices fall. In the four cities in Panels B–E, volume more than doubles during the boom. Prices subsequently peak between 200% and 300% of their 2000 levels.

Figure 2 shows that this lead–lag relation between prices and volume also holds on average across all MSAs in our sample from 2000 to 2011. We estimate correlations between prices and lagged volume by running regressions of the form:

$$p_{i,t} = \beta_k v_{i,t-k} + \eta_{i,t}, \quad (1)$$

where p is log price demeaned at the MSA level, v is volume normalized by the MSA’s 2000 housing stock and demeaned at the MSA–calendar month level, i indexes MSAs, and time is measured in months. Figure 2 plots the correlations implied by each β_k coefficient for up to four years of lags ($k = 48$) and one year of leads ($k = -12$). The correlation is positive at most leads and lags but reaches its maximum at a positive lag of 24 months. Thus, changes in volume generally lead changes in prices by about two years.

In Panel A of Figure 3, we plot aggregate trends in prices and inventories of homes listed for sale between 2000 and 2011. In Panels B–E, we plot analogous series for four cities that represent the same regions as in Figure 1. Because Las Vegas and Orlando are not in our listings data, we replace them with the nearby MSAs of Reno and Daytona Beach. During the period when the relation between volume and prices reverses, aggregate inventories rise dramatically to nearly double their level from earlier in the cycle. This pattern also characterizes the joint dynamic of prices and inventories across cities in Panels B–E. In Phoenix, Reno, and Bakersfield, inventories rise during the quiet to between double and triple their earlier levels. In Daytona Beach, inventories rise to 450% of their pre-quiet levels.²

These stylized facts suggest that focusing on the dynamic of quantities—both volume and inventories—can provide insight on the drivers of the cycle. In particular, determining who was most heavily participating in the housing market during each phase may help us differentiate between various explanations for that cycle.

²We repeat the analyses for Figures 1–3 for MSAs outside the sand states. The results in Figures IA1, IA2, and IA3 of the online appendix reveal that the patterns we document are not exclusive to these states.

3 Speculators during the cycle

This section explores the role of speculators throughout the housing cycle and their correlation with the aggregate dynamics of prices, volume, and inventory.

3.1 Quantities and prices in the boom

Figure 4 presents a simple illustration of the quantitative importance of speculation during the cycle. The figure shows monthly aggregate time series for total transaction volume (with and without new construction), short-holding-period volume, and non-occupant volume calculated using our deeds data. Each series is normalized relative to its average value in 2000 and seasonally adjusted by removing calendar-month fixed effects. For reference, we also report the raw counts of each type of transaction in 2000, 2005, and 2010. To abstract from the effect of foreclosures on speculative volume during the bust, we drop lender acquisitions and dispositions of foreclosed properties when constructing the series in this figure.

While overall volume increased by 40% during the boom years of 2000–2005, speculative volume increased dramatically more. Both short-term sales and purchases by non-occupants approximately doubled between 2000 and 2005. Not only did these speculative components of volume increase more rapidly, but their increase also accounted for a non-trivial portion of the overall increase in volume. For example, total volume increased from 2.73 million transactions in 2000 to 3.82 million in 2005. During the same time period, short-holding-period volume increased from 510 to 940 thousand transactions, which implies that volume growth in this category alone can account for 39% of the total volume increase during the boom.³ A similar calculation for non-occupant volume (in the 102 MSAs with reliable non-occupant data) implies that this measure of speculative activity can account for 53% of the volume increase during the boom. If we exclude new construction from the total volume statistics—because short-term sales can only involve homes previously sold—short-term volume accounts for 57% of the aggregate increase in existing home sales. These calculations illustrate that speculators were, in an accounting sense, a key driver of the

³Part of the increase in short-term volume during the boom happens mechanically because total volume is increasing. In Appendix B.1, we use conditional selling hazards by buyer cohort to quantify the contribution of an overall increase in total volume to the share of late-boom volume coming from short-term sales. Approximately 90% of the rise in short-term volume comes from the changing composition of buyers, rather than mechanical forces.

volume boom.

The shift in the composition of volume toward speculative buyers also correlates highly with changes in total volume across local markets. This correlation can be seen in the top two panels of Figure 5. Panel A presents scatter plots of the percentage change in total volume at the MSA level from 2000 to 2005 versus the percentage change in volume for short holding periods and long holding periods separately. Not only does the growth in volume of short-holding-period transactions correlate strongly with the increase in total volume across MSAs, but the magnitude of this relation is also much stronger for short holding periods relative to long holding periods.⁴ A similar conclusion arises from Panel B, which presents analogous scatter plots grouping transactions according to the occupancy status of the buyer rather than the holding period of the seller. The relation between total volume growth and non-occupant volume growth across MSAs is strong, positive, and larger in magnitude than the corresponding relation with growth in sales to owner-occupants.

Panels C and D further show that cross-MSA differences in speculative volume growth explain much of the differences in total volume growth. For each MSA, we plot the change in either short-holding-period volume (Panel C) or non-occupant volume (Panel D) divided by initial total volume on the y-axis against the percentage change in total volume on the x-axis. The slope provides an estimate of how much of a given increase in total volume during this period came in the form of short-holding-period or non-occupant volume. For short-holding-period volume, the answer is 30% (or 36% excluding new construction). For non-occupant volume, the slope is even larger and implies that, for the average MSA in our sample, 54% of the increase in total volume between 2000 and 2005 came from non-occupant purchases. Thus, shifts in the composition of volume toward speculative buyers are a major determinant of changes in total volume during the boom.

Table 2 shows how speculative volume relates to the size of the price and quantity cycles in the cross-section of MSAs (Table 1 shows summary statistics).⁵ We estimate the correla-

⁴One concern with our short-term speculation measure is that it is based on realized rather than expected holding periods. This way of measuring short-term speculation may complicate the interpretation of our results if buyers' intended holding periods endogenously respond to changes in economic conditions during the boom. Appendix B.2 presents instrumental variable regressions that predict short-term volume using pre-cycle demographics. The change in realized short-term volume is quantitatively important for overall volume growth and the size of the price cycle, even when using only the portion of short-term volume growth predicted by ex-ante buyer characteristics.

⁵We focus our empirical analysis on MSA-level outcomes for two reasons. First, the variation across cities is likely more informative for the aggregate housing cycle. Second, and related to the first, spatial correlation across ZIP Codes within cities hinders interpretation of cross-sectional results for some housing

tion between growth in each speculative measure and various housing market outcomes by separately regressing these outcomes on each measure of speculation. To aid interpretation, we scale the change in outcomes for all quantity measures relative to total volume in 2003.

In Panel A, the first two columns show that house price booms are strongly related to the size of speculative volume booms across cities. Cities with a one standard deviation larger short-volume boom (12.9%) see a 24.9 percentage point larger cumulative price increase during the boom. Cities with a one standard deviation larger non-occupant boom (27.1%) see a 15.4 percentage point larger cumulative price increase during the boom. On average across cities, prices rise by 97% during the boom and quiet. Thus, the relation between speculative volume and prices is economically large in the cross-section of MSAs.

Consistent with the aggregate evidence in Figure 3, which shows a modest increase in listings during the boom, we find a small, statistically insignificant relation across MSAs between speculative booms and the change in listings during the boom (Panel B, columns 1–2). Given the strong relation between the short-term and total volume booms, this suggests that the increase in demand during the boom was sufficient to absorb the rising flow of listings from short-term buyers.

3.2 Quantities and prices in the quiet and bust

As discussed in Section 2, there is a quiet period in the housing cycle during which prices rise, transaction volumes rapidly fall, and there is a large increase in unsold listings. In Panel B of Table 2, columns 3 and 4 show that the rise in listings during the quiet correlates strongly with the run-up of speculative volume during the boom across MSAs. Cities with a one standard deviation larger short-volume boom (12.9%) see a larger cumulative increase in listings during the quiet of 76.9 percentage points relative to the total volume in 2003. Cities with a one standard deviation larger non-occupant boom (27.1%) see a cumulative increase in listings during the quiet of 71.7 percentage points relative to the total volume in 2003. Across cities, the mean increase in inventories during the quiet is 178% of 2003 total volume with a standard deviation of 144%.

Thus, the relation between speculative booms and the market outcomes. For example, MSA fixed effects account for 86% of the variation in house price booms across ZIP Codes. This fact is likely due to data limitations in house price index estimation, with local price indices often derived from spatial interpolation, and helps explain differences in results in cross-MSA analyses, as in our paper, and cross-ZIP Code, within-MSA analyses, as in Griffin et al. (2020).

rise of listings is quantitatively important in accounting for the cross-section of inventories.⁶

In Figure 6, we supplement this cross-MSA evidence by showing that short-term listings account for the majority of the increase in new listings from 2003 to 2007. We plot monthly series for total and short-term new listings, normalizing each series relative to its 2003 average and seasonally adjusting by removing calendar-month effects. These data only include a home listed for sale the first time it appears during a listing spell to avoid double-counting unsold listings. While total new listings rise to 150% of their 2003 average at the quiet's peak, short-term listings rise to 250% of their 2003 average and remain above 200% well into the bust. Short-term listings rise from 280 to 590 thousand, accounting for 55% of the rise in total new listings from 1.17 million to 1.73 million. In later stages of the bust, short-holding-period listings fall well below the 2003 level, consistent with the idea that purchases in the quiet and early bust are more likely to include fundamental buyers and longer-term investors.⁷ This evidence suggests that attempted sales by speculators who bought during the boom explain much of the increase in listings during the quiet, and that the reduced entry of speculators during the quiet contributes to the eventual decline in total volume.

Larger speculative booms also predict stronger contractions in total volume and prices during the end of the cycle. Panel C of Table 2 shows that cities with a one standard deviation larger short-volume boom and non-occupant boom respectively see cumulative declines in total volume (relative to 2003 volume) that are 13.5 and 13.9 percentage points larger. The analogous results for prices, shown in columns 3 and 4 of Panel A, imply 7.4 and 4.5 percentage point larger declines during the bust. Thus, speculative booms explain much of the 63% average decline in volume during the quiet and bust (relative to 2003 volume) and 28% decline in prices during the bust. These cross-MSA results are consistent with the aggregate pattern in Figure 4, in which speculative volume declines more sharply during the

⁶Table 2 reports the change in the inventory of unsold listings. In the online appendix, Table IA6 reports analogous results using the change in the flow of new listings and shows qualitatively similar results. The rise in unsold listings during the quiet is driven both by an increase in the rate at which homes were listed for sale and a reduction in the probability of sale conditional on listing. In the online appendix, we repeat the analysis for Tables 2 and IA6, while including an indicator for whether the MSA is in a sand state. The results in Tables IA7 and IA8 are similar, though somewhat weaker for the non-occupant volume boom.

⁷This evidence complements Genesove and Mayer (1997, 2001), who document the role of home equity and loss aversion, respectively, in preventing list prices from adjusting downward during a market downturn in Boston. Short-holding-period buyers are more likely to maintain high list prices because—in the home equity view—they will have paid down less of their mortgages when they turn to sell and because—in the loss aversion view—they will have paid higher initial prices than long-holding-period buyers. In our model, extrapolation creates another force causing recent buyers to set overly optimistic list prices, the same force that helps explain their initial entry into the market.

quiet and bust than does total volume. Turning points in both short-holding-period and non-occupant volume exactly coincide with the turning point in aggregate volume, the sharp rise in listings during the quiet, and the decline in price growth before its reversal.

Finally, we find that cities with larger short-term speculative booms experienced more severe foreclosure crises. The estimate in column 3 of Panel C implies that a one standard deviation increase in the short-volume boom is associated with 11.5 percentage points more foreclosures (relative to 2003 volume) in the bust, equal to 370 thousand more foreclosures. This effect is large relative to the 2.68 million foreclosures across the 115 MSAs in our data. In contrast, the relation between foreclosures and the non-occupant boom is insignificant (column 4 of Panel C).

3.3 Summary of main empirical results

Our results show strong relations between speculative purchases during the boom and the amplitude of the housing cycle. Across cities, a larger speculative boom predicts sharper increases in prices and volume during the boom, a greater boom and bust in prices, a larger surge in listings during the quiet, and a more pronounced fall in volume during the quiet and bust. Time series evidence also indicates that speculation accounts for much of the increase in volume during the boom and listings during the quiet.

These results suggest the following narrative linking short-term speculators to the housing cycle. As prices increase in the boom, short-term speculators buy houses in anticipation of capital gains, and this buying activity pushes up prices further. As price growth eventually slows, speculative volume slows, contributing disproportionately to the decrease in total volume. At the same time, speculative buyers from the recent past—who are now looking to sell—continue to generate a new flow of listings. Because smaller expected capital gains attract fewer new speculative buyers to the market, many of these new listings fail to sell. Prices rise as volume falls, which suggests sellers are still posting higher prices. The result is a quiet period with falling volume, rising inventories, and slowing price growth. Accumulating inventories and falling demand eventually result in negative price growth, which creates a lead-lag pattern between the drops in volume and prices. The goal of our model is to illustrate this causal narrative theoretically.

4 Characterizing speculative buyers

In this section, we use our microdata and other data to provide additional insight on speculative purchases. These facts motivate how we model speculation.

4.1 Extrapolation among speculators

Using multiple measures of speculation, we examine whether house price growth can predict subsequent speculative purchases and beliefs in the housing market. Our first measures use our deeds dataset. For each MSA and year from 2000 to 2011, we count total non-occupant purchases and divide by the equivalent count from 1999 as a normalization. We do the same for short-term purchases, defined here as those for which we observe another sale on the same property in the next three years. Panels A and B of Figure 7 present binned scatter plots of normalized speculative purchases against house price growth in the past year. Both non-occupant and short-term purchases are much higher in the years and MSAs that witness higher house price appreciation in the last year.⁸

The second measure of speculation uses responses from the NAR’s Investment and Vacation Home Buyers Survey. For each year of the survey, we calculate the fraction of respondents (except those reporting “don’t know”) who report an expected holding time of less than three years or had already sold their home by the time of the survey. This measure captures the intention of buyers at the time of purchase. Thus, it complements our transaction-based metric that relies on realizations of short horizons after the fact. In Panel C of Figure 7, we plot this measure of speculation against annual house price growth at the national level. A gain of 10% in house prices over the past year is associated with an 8.2 percentage point larger short-term buyer share.

Our final measures of speculation use responses from the 2014–2017 waves of the Federal Reserve Bank of New York’s Survey of Consumer Expectations.⁹ This survey asks respondents’ views on housing as an investment as well as their probability of buying a non-primary home in the next three years. Thus, the survey directly queries non-occupant housing de-

⁸In Appendix B.3, we estimate higher-frequency panel VAR specifications of speculative volume and lagged house price appreciation, in the style of Chinco and Mayer (2015). The positive relation between prices and speculative purchases continues to hold.

⁹The data come from the replication files of Armona et al. (2019). We thank Andreas Fuster for sharing this evidence with us.

mand, complementing the measure of non-occupant purchases in our deeds data. Panels D and E of Figure 7 present binned scatter plots of the survey measures against appreciation in the Zillow house price index over the past five years in the respondent’s ZIP Code. The share of respondents saying that housing is a very good investment rises with local house price appreciation; the opposite is true for those calling housing a bad or very bad investment. The reported probability of buying a non-primary home also rises with lagged house price growth.

In summary, house price growth predicts increased speculative purchases in three different datasets. These results complement survey evidence showing that expected future house price growth rises with realized past house price growth (Case et al., 2012; Armona et al., 2019). We incorporate extrapolative beliefs into our model in such a way that speculative purchases and posted list prices respond strongly to recent price growth. This modeling choice builds on prior studies that use extrapolative expectations to understand other aspects of the housing market (Glaeser et al., 2008; Guren, 2014; Glaeser and Nathanson, 2017).

4.2 Overlap between short-term and non-occupant buyers

In this section, we examine overlap between short-term and non-occupant buyers. Data from the NAR’s Investor and Vacation Home Buyers Survey report expected holding times separately for investor and non-investor buyers. As Figure 8 shows, about 20% of investor buyers report expected holding periods of under three years, larger than the corresponding share among non-investor buyers. Therefore, these data provide direct evidence of overlap between short-term and non-occupant buyers.

To focus on speculators who entered during the 2000–2005 boom, we also measure this overlap in our CoreLogic data. We find that 27% of 2000–2005 short-term volume came from non-occupant buyers, while 41% of the increase in short-term volume over this time came from non-occupants (see Online Appendix C.1 for details). Therefore, non-occupants account for an excess share of the growth in short-term buyers.

The evidence in this section indicates that there is substantial overlap between short-term and non-occupant buyers. In light of this evidence, we allow for such overlap in our model.

4.3 Credit utilization

To examine the role credit plays in enabling speculative volume, we present in Table 3 summary statistics on the proportion of all-cash purchases in our data. Column 1 shows that 29% of short-term buyers and 38% of non-occupant buyers do not use a mortgage. These shares exceed the all-cash share among all buyers, which is 20%. The remaining columns of the table report averages at the MSA-by-month level and show that all-cash transactions among speculators remain high at all points of the housing cycle. Thus, while credit may have enabled speculation, there is a disproportionately large group of speculators who do not use credit at all. The behavior of these buyers goes unobserved in any analysis of speculative activity based on mortgage data alone.¹⁰

In Online Appendix C.2, we study the relation between leverage and short-term volume growth. We find that short-term sales increase most strongly among sellers whose LTV when purchasing the home was between 60% and 85%. This evidence is consistent with prior work documenting credit growth among speculators during the boom (Haughwout et al., 2011; Bhutta, 2015; Mian and Sufi, 2022). However, it also suggests that very high credit utilization ($LTV \geq 85\%$) does not account for most of the rise in speculative buying.

Motivated by these findings, we omit credit constraints from our model of housing market speculation. We stress that, although we omit credit from the model, our findings are compatible with stories in which credit enables speculative entry during the cycle.

4.4 Buyer scale and experience

Next, we examine whether short-term buyers are individuals buying a few houses or firms buying many houses. In Online Appendix C.3, we present a methodology for classifying buyers as real estate developers, experienced investors holding three or more homes, or inexperienced buyers owning one or two homes. Of the short-term sales in 2000–2005, 15% of the initial purchases are from developers, 24% are from experienced investors, and 61% are from inexperienced buyers. This evidence is consistent with Bayer et al. (2020, 2021) who also find an important role for inexperienced short-term investors during this episode.

¹⁰The correlations between the speculative booms in Table 2 and their analogous counterparts that exclude cash transactions are approximately 0.9. Thus, while excluding all-cash transactions would understate the importance of speculators in the aggregate, the cross-sectional relationships in Table 2 are robust to excluding these transactions.

In light of the large share of inexperienced buyers among short-term sellers, we allow buyers to own only one house in our model.

Finally, we explore whether short-term sellers remain within the MSA by buying another house nearby. We link transactions within MSA in our data by comparing names of buyers and sellers. As we describe in Online Appendix C.3, 69% of short-term sellers do not buy in the MSA within a quarter of the sale. To match the high share of such sellers, we assume in our model that homeowners exit the local housing market upon selling their house.

5 The model

The goal of our model is to match the joint dynamics of prices, volume, and listings. Additionally, the model should explain the disproportionate role of non-occupants and short-term sales in generating these dynamics.¹¹ In doing so, the model complements our empirical analysis by permitting stronger causal statements about the role of speculation and allowing us to conduct counterfactual explorations of model assumptions and policy design.

5.1 Environment and preferences

We present a discrete-time model of a city with a fixed amount of perfectly durable housing, normalized to have measure 1. There are three types of agents in the model: movers, stayers, and potential buyers. Movers are city homeowners who are trying to sell their homes. Stayers are city homeowners who do not list their homes for sale. Potential buyers are people from outside the city who get a one-time chance to buy a house from a mover. In Figure 9, we illustrate how agents transition between these three types.

All agents are risk-neutral and can borrow or lend across periods at an interest rate of r . They maximize their expectation of the discounted present value of their per-period utility, which is the sum of two components: housing utility and non-housing consumption, whose price we normalize to 1.

Each period, a mover lists her house for sale by posting a list price, P . She then matches randomly to a potential buyer from outside the city, who decides whether to purchase at the listed price.¹² In the event of a sale, the mover exits the market and consumes her terminal

¹¹In Online Appendix D, we discuss the relation between our model and prior work in detail.

¹²In other models, some movers fail to match to a potential buyer due to search frictions (Head et al., 2014; Guren, 2018). We abstract from this possibility.

wealth. Movers who fail to sell remain movers next period. We denote the share of listings that sell at time t by π_t . Movers receive 0 housing utility while listing their homes. They are impatient and discount time at rate $r_m \geq r$.

Potential buyers who decide to buy become stayers at the beginning of the next period. Those who do not buy exit the market and consume their terminal wealth. Stayers receive housing utility e^δ at the beginning of each period, but cannot sell their house. With probability λ each period, a stayer transitions to being a mover, at which point she lists her home for sale. Housing utility e^δ and the mover hazard λ remain constant for a given stayer over time but may vary across stayers. All stayers discount time at rate r .¹³

At time t , each potential buyer knows the housing utility she would receive while being a stayer if she chooses to purchase and the probability λ that she would transition into becoming a mover each period. For each potential buyer within a given cohort, the log of her housing utility, δ , is the sum of a time-varying aggregate demand shifter, d_t , and an idiosyncratic term, a , that varies across potential buyers at a point in time:

$$\delta = d_t + a. \tag{2}$$

Potential buyers observe their own value of δ but do not separately observe d_t and a . That is, they cannot determine what fraction of their personal valuation is common to all potential buyers in their cohort.

The demand shifter d_t affects the distribution of housing utility across different cohorts of potential buyers over time. We model it as a difference-stationary process with a persistent growth rate:

$$\begin{aligned} d_t &= d_{t-1} + g_t + \epsilon_t^d \\ g_t &= (1 - \rho)\mu_g + \rho g_{t-1} + \epsilon_t^g, \end{aligned}$$

where $0 \leq \rho < 1$, and ϵ_t^d and ϵ_t^g are mean-zero independent normals. We denote $\sigma_d^2 = \text{Var}(\Delta d_t)$ and $\gamma = \text{Var}(g_t)/\text{Var}(\Delta d_t)$, which implies that the variances of ϵ_t^d and ϵ_t^g are $(1 - \gamma)\sigma_d^2$ and $\gamma(1 - \rho^2)\sigma_d^2$, respectively. As with d_t , the growth rate g_t is unobservable to all agents in the model.

¹³We assume that r is large enough to rule out rational bubbles and provide the precise condition for this in Appendix E.1.

The idiosyncratic term a generates within-cohort heterogeneity in housing utility. We assume that there are two types of potential buyers, indexed by n : non-occupants ($n = 0$) and occupants ($n = 1$). To capture the idea that non-occupants generally receive smaller flow benefits from their homes than occupants, we allow the distribution of a to vary across these two groups. Specifically, the distribution of a across potential buyers of type n at each time t is $\mathcal{N}(\mu_n, \sigma_a^2)$. Each potential buyer knows whether she is a non-occupant or an occupant.

Finally, to capture heterogeneity in expected holding periods, we allow λ to vary across potential buyers within each cohort. We assume that λ follows a discrete distribution with possible values $\lambda \in \{\lambda_1, \dots, \lambda_J\}$ and denote the joint probability that a potential buyer is of occupancy-type n and has mover hazard λ_j to be $\beta_{n,j}$. Thus, the distribution of expected investment horizons can also differ across non-occupants and occupants.

5.2 Inference about demand

To forecast the price at which they will eventually sell their house, agents must estimate the current level of the demand shifter, d_t , and its growth rate, g_t . Agents use historical data on city house prices to estimate these latent variables. We focus on equilibria in which all movers at a given time post the same list price, which we denote P_t (conditions for this outcome are below). Agents at time t observe the full history of price changes, $P_{t'}/P_{t'-1}$ for $t' < t$. They deduce any past price level, $P_{t'}$, by inflating the list price they observed as a potential buyer by cumulative price growth between the time of their purchase and t' . Agents also observe the history of the shares of listings that sell, $\pi_{t'}$ for $t' < t$.

To infer d_t and g_t from historical market data correctly, an agent needs to know how past potential buyers used market data to decide whether to buy a house. Following Glaeser and Nathanson (2017), we depart from rationality and propose that agents instead adopt a simplified model of how other agents decide to buy a house. Specifically, agents believe that other agents decide to buy a house if and only if:

$$e^\delta \geq \bar{\kappa}P, \tag{3}$$

where P is the list price of the house and $\bar{\kappa}$ is a time-invariant constant that is common across all potential buyers. As we discuss in Section 6.3, this is the key behavioral assumption that

generates positive feedback and bubble-like dynamics within our theoretical framework. In employing this mental model, agents neglect the fact that the beliefs, and therefore the decision rule, of potential buyers could vary over time based on the changing history of market data.¹⁴ However, conditional on the beliefs implied by this simplified model, agents make decisions optimally.

Given Eq. (2) and the decision rule in (3), agents believe that other agents buy if and only if:

$$a \geq \log P + \log \bar{\kappa} - d_t.$$

Therefore, according to agents' simplified model, the share of potential buyers at time t who would purchase at list price P is:

$$1 - F(\log P + \log \bar{\kappa} - d_t) \equiv \tilde{\pi}(P, d_t), \quad (4)$$

where $F(a) = \sum_{n=0}^1 \sum_{j=1}^J \beta_{n,j} \Phi(a - \mu_n)$ is the CDF of a across both non-occupants and occupants, and $\Phi(\cdot)$ is the CDF of a normal random variable with mean 0 and variance σ_a^2 .

Given market data on historical prices $P_{t'}$ and sales shares $\pi_{t'}$, agents at time t use Eq. (4) to infer past values of the demand shifter. In particular, by equating $\pi_{t'}$ to $\tilde{\pi}(P_{t'}, d_{t'})$, they infer that:

$$\tilde{d}_{t'} = \log P_{t'} - F^{-1}(1 - \pi_{t'}) + \log \bar{\kappa}, \quad (5)$$

where $\tilde{d}_{t'}$ denotes an agent's belief about the true value of the demand shifter $d_{t'}$. Given this inferred history of the demand shifter, agents employ a standard Kalman filter to arrive at posterior estimates of its current value, d_t , and its growth rate, g_t . Lemma 1 characterizes these posteriors (all proofs are in Online Appendix E).

Lemma 1. *Conditional on house prices and sale probabilities before t , the posterior distributions of d_t and g_t are $\mathcal{N}(\hat{d}_t, \hat{\sigma}_d^2)$ and $\mathcal{N}(\hat{g}_t, \hat{\sigma}_g^2)$, where:*

$$\begin{aligned} \hat{d}_t &= \tilde{d}_{t-1} + \hat{g}_t \\ \hat{g}_t &= \mu_g + (1 - \alpha)\rho \sum_{k=1}^{\infty} (\alpha\rho)^{k-1} \left(\Delta \tilde{d}_{t-k} - \mu_g \right), \end{aligned}$$

¹⁴This simplified model of other agents' willingness to pay is the same as the "cap rate error" that Glaeser and Nathanson (2017) introduce. That paper motivates this error by showing that common knowledge of rationality is not robust to small mistakes and involves unintuitive decision rules as a function of past prices.

and $\hat{\sigma}_d$, $\hat{\sigma}_g$, and $\alpha \in (0, 1)$ are constants depending on σ_d , γ , and ρ .

Together with Eq. (5), Lemma 1 shows that agents estimate the current level of the demand shifter, d_t , and its growth rate, g_t , from historical market data in a straightforward manner. In particular, differencing Eq. (5) yields:

$$\Delta \tilde{d}_{t-k} = \Delta \log P_{t-k} - \Delta F^{-1}(1 - \pi_{t-k}),$$

which implies that the expected growth rate, \hat{g}_t , is a weighted average of past price growth adjusted downward each period to reflect any increase in the share of unsold listings. The expected demand shifter, \hat{d}_t , equals this expected growth rate plus agents' belief about last period's demand shifter.

5.3 Mover problem

The mover's problem is to select a list price that maximizes the expected present value of utility conditional on beliefs about the demand shifter and its growth rate. We write the problem recursively as:

$$V^m(\hat{d}_t, \hat{g}_t) = \sup_P E \left(\tilde{\pi}(P, d_t)P + (1 + r_m)^{-1}(1 - \tilde{\pi}(P, d_t))V^m(\hat{d}_{t+1}, \hat{g}_{t+1}) \right), \quad (6)$$

where the expectation is over $d_t \sim \mathcal{N}(\hat{d}_t, \hat{\sigma}_d^2)$. If the potential buyer who matches to the mover buys, the mover receives P and exits the city. The first term, $\tilde{\pi}(P, d_t)P$, gives the mover's perceived probability of this event times the payoff. The second term gives the discounted value of continuing as a mover next period times the probability of that event.

All movers at time t post the same list price when a unique P maximizes the right side of Eq. (6). We verify the existence of such a price at each point of the state space in our quantitative exercise. Lemma 2 clarifies how this price depends on mover beliefs, \hat{d}_t and \hat{g}_t .

Lemma 2. *The optimal list price takes the form $P_t = e^{\hat{d}_t} p(\hat{g}_t)$ for some function $p(\cdot)$.*

The log list price scales one-for-one with the current belief about the level of the demand shifter, \hat{d}_t . It also depends on the belief about the growth rate, \hat{g}_t , because the option of selling next period becomes more valuable when movers expect faster demand growth.

Because \hat{d}_t and \hat{g}_t depend on historical market data, we can also characterize price posting

as a function of past prices and sales shares. To provide intuition about price posting, Lemma 3 shows that when r_m is large, movers set prices in a simple extrapolative fashion.

Lemma 3. *In the limit as $r_m \rightarrow \infty$, agents' expectation of house price growth over the next period conditional on house prices and sale probabilities before t is:*

$$E\Delta \log P_{t+1} = \mu_g + (1 - \alpha)\rho \sum_{k=1}^{\infty} \left(\frac{\rho}{1 + (1 - \alpha)\rho} \right)^k (\Delta \log P_{t-k} - \mu_g).$$

Given this expectation, movers at time $t + 1$ set prices according to the rule:

$$\Delta \log P_{t+1} = E\Delta \log P_{t+1} + (1 + (1 - \alpha)\rho) (\log(\bar{\kappa} \bar{p}) - F^{-1}(1 - \pi_t)),$$

for some constant \bar{p} .

In this limit, price growth expectations are a simple weighted average of past price changes, as in the reduced form extrapolation formulas that Barberis et al. (2015, 2018) and Liao and Peng (2018) assume. Similarly, price setting closely resembles the “backward-looking rule of thumb” that Guren (2018) assumes, except that movers here decrease list prices when they observe a high share of unsold listings in the prior period. Therefore, the bounded rationality of movers in our model endogenously leads to extrapolative expectations and price posting when movers are impatient.

5.4 Potential buyer problem

The potential buyer's problem is to decide whether to purchase or not, taking as given the price that movers post. At the end of time t , the expected utility for a potential buyer from purchasing a house is:

$$V^b(\hat{d}_t, \hat{g}_t; \lambda, \delta, n) = (1 + r)^{-1} E \left(e^\delta + \lambda V^m(\hat{d}_{t+1}, \hat{g}_{t+1}) + (1 - \lambda) V^s(\hat{d}_{t+1}, \hat{g}_{t+1}; \lambda, \delta) \right), \quad (7)$$

where the expectation is over $d_t \sim \mathcal{N} \left(\frac{\sigma_a^2 \hat{d}_t + \hat{\sigma}_d^2 (\delta - \mu_n)}{\sigma_a^2 + \hat{\sigma}_d^2}, \frac{\sigma_a^2 \hat{\sigma}_d^2}{\sigma_a^2 + \hat{\sigma}_d^2} \right)$.¹⁵ A potential buyer who purchases becomes a stayer and receives housing utility e^δ at the beginning of the next period.

¹⁵The posterior on d_t is different for potential buyers than for movers and stayers. A potential buyer's log housing utility, δ , conveys information about the current demand shifter, d_t , due to Eq. (2). Therefore, her posterior on d_t combines the posterior based on housing data, $\mathcal{N}(\hat{d}_t, \hat{\sigma}_d^2)$, with her prior based on her housing utility, $\mathcal{N}(\delta - \mu_n, \sigma_a^2)$. Movers and stayers, however, do not use their own δ to estimate d_t because it is a noisy observation of a past value of the shifter, $d_{t'}$, which they believe they infer directly as $\tilde{d}_{t'}$.

With probability λ , she then becomes a mover, the value of which is equal to $V^m(\hat{d}_{t+1}, \hat{g}_{t+1})$ and given by Eq. (6). With probability $1 - \lambda$, she continues on as a stayer, the value of which we denote by $V^s(\hat{d}_{t+1}, \hat{g}_{t+1}; \lambda, \delta)$. At any time t , the stayer value function can be written recursively as:

$$V^s(\hat{d}_t, \hat{g}_t; \lambda, \delta) = (1 + r)^{-1} E \left(e^\delta + \lambda V^m(\hat{d}_{t+1}, \hat{g}_{t+1}) + (1 - \lambda) V^s(\hat{d}_{t+1}, \hat{g}_{t+1}; \lambda, \delta) \right), \quad (8)$$

where the expectation is over $d_t \sim \mathcal{N}(\hat{d}_t, \hat{\sigma}_d^2)$.

A potential buyer decides to buy when the value of doing so is at least as large as the price: $V^b(\hat{d}_t, \hat{g}_t; \lambda, \delta, n) \geq P$. Lemma 4 recasts this decision rule in terms of the minimum housing utility at which a potential buyer decides to buy.

Lemma 4. *A potential buyer at time t with housing utility e^δ and occupancy type n and for whom $\lambda = \lambda_j$ decides to purchase a home with list price P if and only if:*

$$e^\delta \geq \kappa_{n,j}(\hat{g}_t)P,$$

for some function $\kappa_{n,j}(\cdot)$.

The potential buyer's decision rule is similar to the one in Eq. (3) that other agents believe she is using. She purchases if the per-period housing utility she would receive exceeds some fraction of the list price. The key distinction is that the fraction she actually uses depends on both the history of market data she observes and her type. In particular, because the potential buyer anticipates selling in the future, this fraction depends on \hat{g}_t , the expected growth rate of the demand shifter, and on λ , which determines the amount of time she expects until becoming a mover.

The cutoff rule in Lemma 4 determines both the share of listings that sell and the fraction of all purchases made by buyers of each of the $2J$ types. Specifically, a purchase occurs when:

$$a \geq \log P + \log \kappa_{n,j}(\hat{g}_t) - d_t,$$

which implies that the share of potential buyers of type n and λ_j who buy at time t is $1 - \Phi(\log P + \log \kappa_{n,j}(\hat{g}_t) - d_t - \mu_n)$. Substituting the expression for list prices from Lemma 2 and averaging these shares over all potential buyer types gives the share of all listings that

sell:

$$\pi_t = 1 - \sum_{n=0}^1 \sum_{j=1}^J \beta_{n,j} \Phi \left(\log p(\hat{g}_t) + \log \kappa_{n,j}(\hat{g}_t) + \hat{d}_t - d_t - \mu_n \right). \quad (9)$$

The share of sales at time t going to buyers of type n and λ_j , which we denote $b_{n,j,t}$, equals:

$$b_{n,j,t} = \pi_t^{-1} \beta_{n,j} \left(1 - \Phi \left(\log p(\hat{g}_t) + \log \kappa_{n,j}(\hat{g}_t) + \hat{d}_t - d_t - \mu_n \right) \right). \quad (10)$$

The share of listings that sell, π_t , and the share of sales going to each of the $2J$ types, $b_{n,j,t}$, determine the dynamics of all the aggregate quantity variables in the model.

5.5 Quantities

The model has three aggregate quantities of interest: transaction volume, Q_t , inventory available for sale, I_t , and new listings, L_t . The following accounting identities characterize the evolution of these aggregates as a function of sales probabilities, π_t , and the composition of buyers, $b_{n,j,t}$:

$$\begin{aligned} Q_t &= \pi_t I_t, \\ I_t &= (1 - \pi_{t-1}) I_{t-1} + L_t, \\ L_t &= \sum_{j=1}^J \lambda_j S_{j,t-1}, \end{aligned}$$

where $S_{j,t}$ measures the share of housing owned by stayers of type $\lambda = \lambda_j$ at the end of time t . This share evolves according to the following law of motion:

$$S_{j,t} = (1 - \lambda_j) S_{j,t-1} + (b_{0,j,t} + b_{1,j,t}) Q_t.$$

As these equations make clear, the current composition of buyers affects the composition of stayers, thereby altering future listings and volume. Volume rises when there are more listings or when the selling probability is higher.

In addition to these aggregates, the model generates dynamic patterns in quantities that vary across both realized holding periods and buyer occupancy types. For instance, one variable we track in the data is new listings of homes purchased within the last three years.

In the model, new listings at time t of homes purchased within the last K periods equals:

$$L_t^K = \sum_{k=1}^K \sum_{j=1}^J \lambda_j (1 - \lambda_j)^{k-1} (b_{0,j,t-k} + b_{1,j,t-k}) Q_{t-k}.$$

Similarly, our empirical analysis decomposes volume according to the occupancy type of the buyer and the realized holding period of the seller. In the model, the decomposition by occupancy is straightforward: volume to buyers of occupancy type n equals $\sum_{j=1}^J b_{n,j,t} Q_t$. Decomposing volume by realized holding period is more complicated. The sales volume at time t of houses purchased within the last K periods equals $\sum_{k=1}^K \pi_t I_t^k$, where I_t^k denotes the inventory of listings at time t of homes purchased at time $t - k$. This quantity satisfies the recursion:

$$I_t^k = (1 - \pi_{t-1}) I_{t-1}^{k-1} + \sum_{j=1}^J \lambda_j (1 - \lambda_j)^{k-1} (b_{0,j,t-k} + b_{1,j,t-k}) Q_{t-k}$$

for $k > 0$, with initial condition $I_t^0 = 0$.

6 Model results

6.1 Simulation and calibration methodology

We perform a series of simulations to analyze the baseline properties of our model and to study impulse responses to a shock. Each simulation corresponds to 148 sequential realizations of the two stochastic shocks, ϵ_t^d and ϵ_t^g . The first 100 periods burn in the simulation, leaving 48 analysis periods. Each period represents a quarter, so our analysis spans 12 years. We draw a control sample of 1,000 independent simulations to analyze the model's baseline properties. To analyze the impulse response to a shock, we draw a treatment sample of 1,000 additional simulations identical to the control except in periods 101–104 during which the growth rate shocks ϵ_t^g are two standard deviations higher, representing a large but plausible increase in demand. Impulse responses are average differences between treatment and control outcomes.

Solving the model at any point in time requires evaluating both the function that movers use to set prices, $p(\hat{g}_t)$, and the function that potential buyers use to decide whether to

purchase, $\kappa_{n,j}(\hat{g}_t)$. To do so, we discretize \hat{g}_t using the Rouwenhorst (1995) method and then calculate the function values at these discrete points. To evaluate the functions outside these points, we use cubic splines between mesh points and linear splines beyond the boundaries.

We set $r = 0.012$ and $\rho = 0.880$, corresponding to annual values of 5% and 0.51 in Guren (2018) and Glaeser and Nathanson (2017), respectively. We normalize $\mu_0 = 0$, so that μ_1 gives the average log difference in housing utility between occupants and non-occupants. We set $\mu_g = -\sigma_d^2/2$, which implies that the unconditional expected growth rate of e^{d_t} is 0, so that the average growth rate of housing utility across cohorts of potential buyers is the same as that for stayers already living in the city. We choose $\bar{\kappa}$ so that the average value of $d_t - \hat{d}_t$ in the control simulations equals 0. This choice ensures that agents' simplified model in Eq. (3) leads to inferences about the level of the demand shifter that are correct on average.

We select values of the remaining parameters so that moments from our simulation match the empirical counterparts in Table 4. The composition of buyers and the volatility of demand growth determine $\beta_{n,j}$ and σ_d , respectively, and the selling hazard disciplines r_m , as more patient movers take longer to sell by setting higher prices. We target three features of the national U.S. housing cycle: the ratio of price boom to bust, the volume boom relative to the price boom, and the degree to which the non-occupant volume boom exceeds the occupant boom. Intuitively, these moments determine γ , σ_a , and μ_1 through quantifying extrapolation, the elasticity of demand, and the excess sensitivity of non-occupants.

6.2 Parameter estimates

Table 5 reports parameter values that match the moments in Panels B and C of Table 4. Non-occupant housing utility is 0.9% less than occupant housing utility on average, corresponding to less than a standard deviation in each group's distribution. The mover discount rate is 14%. To map this number into a flow cost of moving, we calculate how much higher the mover value function would be if the mover discount rate were equal to r for a single period. The average difference is 3.7% of the list price, in line with the typical costs of selling a house (Han and Strange, 2015) and smaller than the estimate in Guren (2018) of 2.1% per month.

Relative to occupant potential buyers, a much larger fraction of non-occupant potential buyers have short horizons. According to the estimates for $\beta_{n,j}$, over half of non-occupant potential buyers expect to become movers six months after buying a house; the equivalent share of occupant potential buyers is 25%. These estimates come from targeting the data in

Figure 8, which show that a relatively large share of buyers of investment properties intend to own for less than one year. They imply significant overlap between non-occupant and short-term potential buyers within the model.

Lemma 3 shows that when $r_m \rightarrow \infty$, price growth expectations are a weighted average of past price changes. Here, r_m is finite, but nonetheless large enough to generate extrapolation. To measure extrapolation, we follow Armona et al. (2019) by focusing on the relation between realized price growth over the last year and expectations of annualized price growth over the next 1 and 2–5 years. We measure this relation by regressing movers’ 1- and 2–5-year expectations in period 105 of the control simulations against price growth in the prior four periods. The coefficients from these regressions of 0.127 and 0.042 are similar to though somewhat smaller than the corresponding values of 0.226 and 0.047 that Armona et al. (2019) find in survey data (see their Table 5).

6.3 Buyer cutoff rules

Agents in the model are fully rational except that they ignore the influence of historical market data on the home purchasing decisions of other agents. The effect of this departure from rationality on the model’s dynamics depends on the extent to which the cutoffs that agents actually use when deciding to buy, $\kappa_{n,j}(\hat{g}_t)$, differ from the constant cutoff other agents assume they use, $\bar{\kappa}$. In Figure 10, we plot these cutoffs. Four features of this figure are relevant for understanding the dynamics of our model.

First, the true buyer cutoffs, $\kappa_{n,j}(\hat{g}_t)$, decrease in the expected growth rate of the demand shifter, \hat{g}_t . Intuitively, potential buyers expect larger capital gains when the expected growth rate is high and are therefore willing to purchase at higher prices. Therefore, the expected growth rate of the demand shifter, \hat{g}_t , along with the demand shifter itself, d_t , both increase housing demand.

Second, when \hat{g}_t is high, the cutoffs buyers actually use are less than the constant cutoff that other agents believe they use. This error causes agents in the next period to misattribute the speculative behavior of this period’s buyers—who are purchasing due to high anticipated growth—to an increase in the level of the demand shifter, d_t , instead. As a result, when expected growth is high at time t , subsequent agents overestimate what the level of demand must have been at that time, i.e., $\tilde{d}_t > d_t$. Because the demand process is persistent, this error raises the expectations of next period’s agents about the demand shifter, \hat{d}_{t+1} , and its growth

rate, \hat{g}_{t+1} , leading movers to list their homes at a higher price. Thus, speculative buying raises subsequent house prices, causes overestimation of the demand shifter, and ignites positive feedback by raising the expected growth rate of next period’s potential buyers.

Third, the slopes of the buyer cutoff functions, $\kappa_{n,j}(\hat{g}_t)$, are steeper for higher values of λ_j . Intuitively, potential buyers with shorter horizons expect to sell sooner, so their demand is more sensitive to expected capital gains. As a result, short-term buyers disproportionately drive the positive feedback through which speculative buying today stimulates such buying next period.

Finally, the buyer cutoffs for a given mover hazard are nearly identical for occupants and non-occupants.¹⁶ Quantitatively, the threshold of housing utility at which a purchase occurs does not depend on occupancy status. The only difference in housing demand between occupants and non-occupants with the same horizon is that the distribution of housing utility for the non-occupants is shifted to the left of that of the occupants. As a result, because the non-occupants use the same cutoff as the occupants, a smaller share of them end up buying a house. For a given mover hazard, non-occupants’ demand is therefore more elastic than occupants’ with respect to the demand shifter, d_t , and its expected growth rate, \hat{g}_t .

6.4 Impulse responses

In Figure 11, we plot the impulse responses. As with the national U.S. cycle in Figures 1 and 3, the cycle in the model progresses through a boom, quiet, and bust (Panels A and B).¹⁷ We use grey shading to mark the transition points between these phases, defined as the peaks of volume and prices. The quiet lasts eight quarters, close to the duration in Figure 1 and the correlation-maximizing lag in Figure 2.

In the boom, demand rises because the demand shifter, d_t , is higher and because the expected growth rate, \hat{g}_t , rises in response to price growth. Together, these channels differentially stimulate buying from potential buyers with higher λ (Panel C) and non-occupants

¹⁶The cutoffs depend on occupancy type only because a potential buyer’s housing utility, δ , conveys information about the contemporaneous demand shifter, d_t . Quantitatively, this channel is irrelevant because $\sigma_a = 0.066$ is much larger than $\hat{\sigma}_d = 0.011$.

¹⁷The price boom in our model is smaller than the national boom shown in Figure 1. Potentially, the shocks that generated the national boom are stronger than the one year of two standard deviation shocks we feed into our model. Another possibility is that our assumed value of 0.023 for the annual volatility of demand growth (see Table 4) is too low. Finally, new construction and credit, which our model omits, may have amplified the national boom (Favilukis et al., 2017; Nathanson and Zwick, 2018). To ease comparison with the national cycle, we analyze outcomes in our model relative to the price boom it generates.

(Panel D). The overall increase in housing demand pushes up the share of listings that sell, π_t (Panel E). Short-term buyers re-list their houses quickly, increasing the flow of listings during the boom (Panel F). Prices and volume increase as a result. Tempering the volume boom is the decline in inventory (Panel B), which occurs as the stock of unsold listings diminishes.

The qualitative behavior of volume, inventories, and sale probabilities during the boom is similar in search and matching models, such as Guren (2014). The key difference is the increasing flow of listings coming differentially from short-term buyers (Panel F). This flow limits the decline in inventories to 1.5 log points, amplifying and sustaining the rise in volume. Relative to the price boom, this decline in inventories is an order of magnitude smaller than in Guren (2014). Furthermore, the differential flow of short-term listings leads to the short-term volume boom shown in Panel C, which matches Figure 4. The disproportionate increase in demand from non-occupants, together with the overall rise in volume, produces the strong non-occupant volume boom shown in Panel D that also matches Figure 4.

In the quiet, demand begins to fall because the price level has risen so high. Because they neglect time-variation in the cutoff rule that other potential buyers are using, agents misattribute demand growth during the boom entirely to d_t , though much of it comes from \hat{g}_t , the expected capital gains channel. Eventually, agents over-estimate the demand level so much and post prices that are so high that sale probabilities start to fall (Panel E). Nonetheless, movers increase their list prices throughout the quiet because they continue to revise upward their estimate of the demand shifter for two reasons. First, because of past price growth, the expected growth rate, \hat{g}_t , remains high, which mechanically causes upward revisions to the expected level of demand. Second, the sale probability, π_t , remains high even though it is falling, and these high realizations constitute positive surprises about demand that cause movers to increase their beliefs. Eventually, π_t falls below its pre-shock average, ending these upward revisions and the concomitant increase in list prices.

One of the distinguishing features of the quiet in both the model and the data is the sharp rise in unsold inventories. At their peak, unsold listings are 1.4% above their pre-shock level. The two causes of the excess inventories are the fall in selling probabilities (Panel E) and the elevated flow of short-term listings continuing throughout the quiet (Panel F), which matches the data in Figure 6. This second cause is novel to our model and may explain why inventories rise above their pre-shock level here whereas they fail to do so in models lacking

this channel, such as Guren (2014).¹⁸

The bust begins as movers cut list prices. Agents revise down their expectations of the growth rate, which further depresses demand and sale probabilities. However, because they continue to believe that potential buyer demand is independent of the expected growth rate, movers do not cut prices enough to restore demand, and the bust continues over several periods. Volume falls below its pre-shock level, as in Figure 1. The decline in \hat{g}_t leads to a smaller share of short-term buyers, depressing the flow of new listings (Panel F), which allows inventories to recover (Panel B).

The model generates a second boom in prices, volume, and listings in the last five years of the simulation. This second boom occurs because prices overshoot on the way down, as is common in models with extrapolative expectations (Hong and Stein, 1999; Glaeser and Nathanson, 2017). Underpricing occurs when agents think that demand is lower than its true value. In this case, sale probabilities rise, and volume increases. This increase in demand disproportionately affects short-term buyers, so short-term volume and listings also rise during the second boom.

6.5 Counterfactuals

Many features of the impulse responses discussed above closely match the patterns observed in the data. However, the fact that our model matches these patterns does not directly speak to the role that speculation plays in generating those patterns. To quantify the contribution of speculation to the housing cycle, we rerun the simulation under three counterfactuals, each of which shuts down a different aspect of our baseline model. Impulse responses corresponding to Panels A–D of Figure 11 are in Figure 12; those corresponding to Panels E and F of Figure 11 are in Figure IA4 of the online appendix.

6.5.1 Rational expectations

In the fully rational counterfactual, agents no longer use the simplified model for potential buyer behavior in Eq. (3). Instead, they correctly understand the problem that potential buyers are solving. As a result, they believe that the share of potential buyers at time t who

¹⁸Our model understates the rise in listings during the quiet because of our simplifying assumption that each mover matches to a potential buyer regardless of the number of contemporaneous movers. With a more realistic matching function, such as the one in Guren (2014), our model might also hit the peak of listings (relative to price growth) that appears in Figure 3.

would purchase at list price P is:

$$1 - \sum_{n=0}^1 \sum_{j=1}^J \beta_{n,j} \Phi(\log P + \log \kappa_{n,j}(\hat{g}_t) - d_t - \mu_n) \equiv \pi(P, d_t, \hat{g}_t).$$

Using this function, agents at time t correctly infer the past values of the demand shifter by equating $\pi_{t'}$ to $\pi(P_{t'}, d_{t'}, \hat{g}_{t'})$ and solving for $d_{t'}$. They calculate \hat{d}_t and \hat{g}_t using the Kalman filter in Lemma 1. The mover value function becomes:

$$V^m(\hat{d}_t, \hat{g}_t) = \sup_P E \left(\pi(P, d_t, \hat{g}_t) P + (1 + r_m)^{-1} (1 - \pi(P, d_t, \hat{g}_t)) V^m(\hat{d}_{t+1}, \hat{g}_{t+1}) \right),$$

where the expectation is over $d_t \sim \mathcal{N}(\hat{d}_t, \hat{\sigma}_d^2)$. By an argument analogous to the proofs of Lemmas 2 and 4, the optimal price takes the form $e^{\hat{d}_t} p(\hat{g}_t)$, and a potential buyer buys when $e^{\delta} \geq \kappa_{n,j}(\hat{g}_t) P$, although $p(\cdot)$ and $\kappa_{n,j}(\cdot)$ may differ from the corresponding functions in those lemmas.

We compute impulse responses using the same parameters and sequence of shocks in the baseline model. Results appear in Panels A–D of Figure 12. When expectations are rational, prices no longer overshoot, inventories never rise above their pre-shock value, and the volume boom lasts only four quarters and is only about one quarter of its size in the baseline model. The short- and long-horizon volume booms are nearly identical in size. In contrast, non-occupant volume continues to rise much more than occupant volume, because non-occupant demand is more elastic with respect to the demand shifter, d_t . Therefore, even when potential buyers have rational expectations, non-occupants react more strongly to the demand shock underlying the impulse response, but this reaction does not generate any positive feedback.

In summary, the price bust and the rise in listings above their initial value—two salient features of the data in Figure 3—depend on departing from rational expectations. These features appear in the baseline model but not the rational version. Quantitatively, a large volume boom, and one that is disproportionately short-term, likewise depend on departing from rationality. An excess non-occupant volume boom does not.

6.5.2 Walrasian market clearing

In Online Appendix F.1, we solve a Walrasian version of our model in which a mechanism selects a price P_t each period so that the number of potential buyers willing to buy at that price equals the number of movers willing to sell. We also describe technical changes to the model setup and parameters that aid comparison to the baseline model.

We find that the equilibrium price is $P_t = e^{d_t} p(\hat{g}_t)$, where $p(\cdot)$ is a function. In contrast to the baseline model, the demand shifter, not its expected value, directly affects prices. Here, demand from buyers directly pins down the price; in the baseline model, movers choose the price and demand pins down the share of listings that sell. As a result, prices incorporate changes to demand more quickly with Walrasian market clearing. In the Walrasian model, agents believe that the equilibrium house price is $P_t = e^{d_t} \tilde{p}$, where \tilde{p} is a constant. Therefore, when \hat{g}_t is high, equilibrium prices exceed what agents expect, which leads them to think mistakenly that d_t is high. This force in turn pushes up \hat{g}_{t+1} , which increases P_{t+1} . This positive feedback mechanism is similar to the one in the baseline model.

The results are in Panels E–H of Figure 12. Prices and volume both go through a large boom and bust cycle in the Walrasian model, as in the baseline model. However, volume now peaks after prices, so there is no longer a quiet. The price boom is faster, with prices reaching their peak nine quarters after the shock instead of 15. Under Walrasian market clearing, prices react more quickly to new information, explaining the absence of the quiet and the shorter duration of the price boom. Listings rise in the Walrasian model, but listings and volume coincide due to Walrasian market clearing, so these two variables never diverge as in the baseline model. Finally, short-term and non-occupant volume continue to rise in a large and disproportionate fashion in the Walrasian model.

In summary, many of the features of the baseline impulse response do not require departing from Walrasian market clearing, as they continue to appear in the Walrasian extension. These features include large price and volume cycles, high levels of listings while prices fall, and disproportionate volume booms from short-term sales and non-occupant purchases. However, the existence of the quiet—a period right after the boom in which volume falls while prices and listings rise—does require departing from Walrasian market clearing.

6.5.3 Absence of speculative buyers

The last counterfactual shuts down speculation by adjusting the distribution of potential buyer types while leaving the framework of the model unchanged. In particular, we set $\beta_{n,j} = 0$ for all n and j except for $n = 1$ and the j for which $\lambda_j = 0.03$. All potential buyers are occupants with a horizon of about eight years, which is close to the average horizon among potential buyers in the baseline model. By assigning all potential buyers the same (low) value of λ , this counterfactual removes both short-term buyers and the heterogeneity in holding periods that generates variation in the composition of buyers. We update $\bar{\kappa}$ so that the demand error is still zero and keep other parameters unchanged.

Panels I–L of Figure 12 display the results. Prices and volume still go through a cycle, but the volume boom is three times smaller, and the price overshoot almost disappears. Listings fall 7%, much more than the decline of 1.5% in the baseline model. There is a quiet during which listings rise, but they reach a smaller value of 0.4% (versus the 1.4% in the baseline model) at the end of this period. Short-term volume rises slightly more than long-term volume because of the mechanical channel discussed in Appendix B.1, but by far less than the 7.8-fold relative increase in the baseline model. Finally, non-occupant volume equals zero by assumption.

In Appendix F.2, we explore the distinct roles of short-term and non-occupant potential buyers in amplifying the housing cycle. Removing either group attenuates the housing cycle, but there is substantial overlap between the two groups. If we eliminate short-term buyers while holding constant the share of non-occupants, the housing cycle becomes small, but if we eliminate non-occupants while keeping constant the share of short-term buyers, the housing cycle remains strong. These results suggest that short horizons are the key amplifying force in the model, as opposed to non-occupancy.

While these counterfactuals suggest that removing short-term potential buyers dramatically reduces the magnitude of the cycle, they may overstate this effect because we conduct the counterfactuals using parameter values calibrated in the baseline model under the assumption of exogenous trading horizons. During the 2000–2005 housing boom, it is possible that homeowners who originally expected to stay in their homes for many years decided instead to sell early to exploit rising house prices. We rule out this possibility in our model by assuming that homeowners only list their homes after receiving an exogenous moving shock. To match the 2000–2005 volume boom, our calibration compensates for this omission

by assigning excess weight to the shares of potential buyers with high values of λ . Therefore, removing this large group of short-term buyers from the model may have an outsized effect relative to removing the likely smaller group of such buyers who exist in reality. Nonetheless, our counterfactual demonstrates that removing speculators qualitatively attenuates the price bust and volume cycle and amplifies the decline in inventories during the boom.

6.6 Transaction taxes

In this section, we use our model to study an ad valorem tax that buyers must pay at the time of purchase. The tax rate can depend on the buyer’s occupancy type n , so that a buyer pays a tax $\tau_n P$ when purchasing a home at price P . We denote the vector of tax rates by $\tau = (\tau_0, \tau_1)$. Analyzing capital gains taxes would complicate our model significantly, because contemporaneous movers who bought at different past prices would face different optimality problems and hence choose different list prices, so we leave that analysis to future work.

Holding prices constant, the share of potential buyers who complete a purchase is lower in the presence of this tax. As a result, $\bar{\kappa}$ must go up, as we select this constant so that the average value of $d_t - \hat{d}_t$ equals zero. Intuitively, the threshold $\bar{\kappa}$ rises to reflect the decrease in housing demand from the new tax. We denote this new value $\bar{\kappa}^\tau$. By analyzing the mover value function, it is straightforward to show that the new optimal price is $P_t = e^{\hat{d}_t} p(\hat{g}_t) \bar{\kappa} / \bar{\kappa}^\tau$, where $p(\cdot)$ is the same function that is in Lemma 2. That is, prices scale down by a constant amount that reflects the reduced demand due to the tax.

The reduction in housing demand operates through the cutoff functions, $\kappa_{n,j}(\cdot)$. Due to the proportional nature of the tax, Lemma 4 continues to hold, but now these cutoff functions depend on the tax. We denote them as $\kappa_{n,j}^\tau(\hat{g}_t)$. A potential buyer of occupancy type n and for whom $\lambda = \lambda_j$ buys at time t if:

$$a \geq \log p(\hat{g}_t) + \log \left(\frac{\bar{\kappa} \kappa_{n,j}^\tau(\hat{g}_t)}{\bar{\kappa}^\tau} \right) + \hat{d}_t - d_t.$$

We explore a tax that binds equally on all buyers, so that $\tau_0 = \tau_1$, and a tax that affects only non-occupant buyers, so that $\tau_1 = 0$. We consider taxes of 0.5%, 1%, and 5%, which span the tax rates in many large cities (Chi et al., 2021).

Table 6 reports a 5% tax on all buyers significantly attenuates the price cycle, reducing the bust from 8.2% to 1.1%. It also reduces the volume boom, but this reduction is smaller

than the corresponding one for prices. Smaller taxes of 0.5% and 1% also reduce the cycle amplitude, but these effects are much smaller.

The last three columns of Table 6 report results for the tax on non-occupant buyers. This tax is a weaker instrument for attenuating the house price cycle: the 5% tax reduces the price bust only to 5.8%, and the lower taxes have a smaller effect. The 5% tax nearly eliminates the non-occupant volume boom, reducing it to 0.1% from 12.3%. Therefore, targeting the tax to non-occupants limits its efficacy in reducing the house price cycle, as even a tax that nearly eliminates the non-occupant volume boom still leaves much of the house price cycle.

To understand the mechanism behind these results, in Figure IA5 in the online appendix, we plot the adjusted buying cutoffs, $\bar{\kappa}\kappa_{n,j}^\tau(\hat{g}_t)/\bar{\kappa}^\tau$, for both 5% tax scenarios. Comparing this figure to Figure 10 shows how each tax changes housing demand. The 5% tax on all buyers raises the cutoffs for the $\lambda = 0.5$ group by about half a standard deviation (σ_a), which makes the $\lambda = 0.17$ group more marginal than before. Therefore, the tax effectively skews the composition of buyers towards those with longer horizons. The tax on non-occupants similarly raises the cutoffs, but only for non-occupants. As a result, both the $\lambda = 0.5$ occupants and the $\lambda = 0.17$ non-occupants are marginal. Therefore, many of the buyers with the shortest horizons are still active in the market, which provides an explanation for why this tax has a weaker effect.

7 Conclusion

In this paper, we present evidence that speculators in general and short-term speculators in particular play a crucial role in the housing cycle. This evidence raises additional lines of inquiry.

First, do the expansions in credit that typically accompany housing booms appeal disproportionately to short-term investors? Barlevy and Fisher (2011) document a strong correlation across U.S. 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 to buyers who expect to resell quickly. The targeting of credit expansions to short-term buyers might explain the amplification effects of credit availability on real estate booms documented by Favara and Imbs (2015), Di Maggio and Kermani (2017), and Rajan and Ramcharan (2015). Mian and

Sufi (2022) explore this channel in contemporaneous work.

A second line of inquiry concerns tax policy. While we analyze a fixed transactions tax in this paper, in the spirit of Tobin (1978), Stiglitz (1989), Summers and Summers (1989), and Dávila (2015), natural alternatives such as a short-term capital gains tax might discourage housing speculation by lowering expected after-tax capital gains. However, such taxes discourage productive residential investment as well. Is this tax optimal, and if not, what type of tax policy would be better? It is also unclear empirically whether transaction and capital gains taxes would particularly discourage short-term investors, given that the incidence of this tax might fall more on buyers than sellers.

A third research question involves new construction, which is absent from our model. In a static model, Nathanson and Zwick (2018) predict that undeveloped land amplifies house price booms by facilitating speculation by developers. Developers have short investment horizons because the time from land purchase to home sale ranges from a few months to a few years. Moreover, because developers do not receive housing utility, their payoffs resemble those of the non-occupants in our model. Adding construction to the model in this paper might further clarify the role of land markets and new construction in housing cycles.

References


- Adelino, Manuel, Schoar, Antoinette and Severino, Felipe.** (2016). ‘Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class’, *Review of Financial Studies* 29(7), 1635–1670.
- Anenberg, Elliot and Bayer, Patrick.** (2020). ‘Endogenous Sources of Volatility in Housing Markets: The Joint Buyer-Seller Problem’. Forthcoming, *International Economic Review*.
- Armona, Luis, Fuster, Andreas and Zafar, Basit.** (2019). ‘Home Price Expectations and Behavior: Evidence from a Randomized Information Experiment’, *Review of Economic Studies* 86(4), 1371–1410.
- 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.** (2018). ‘Extrapolation and bubbles’, *Journal of Financial Economics* 129(2), 203–227.
- 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.** (2020). ‘Speculators and Middlemen: The Strategy and Performance of Investors in the Housing Market’, *Review of Financial Studies* 33(11), 5212–5247.
- Bayer, Patrick, Mangum, Kyle and Roberts, James W.** (2021). ‘Speculative Fever: Investor Contagion in the Housing Bubble’, *American Economic Review* 111(2), 609–51.
- Bhutta, Neil.** (2015). ‘The Ins and Outs of Mortgage Debt during the Housing Boom and Bust’, *Journal of Monetary Economics* 76, 284–298.
- Burnside, Craig, Eichenbaum, Martin and Rebelo, Sergio.** (2016). ‘Understanding Booms and Busts in Housing Markets’, *Journal of Political Economy* 124(4), 1088–1147.
- Case, Karl E.** (2008). ‘The Central Role of Home Prices in the Current Financial Crisis: How Will the Market Clear?’, *Brookings Papers on Economic Activity* 2, 161–193.
- 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.
- Charles, Kerwin Kofi, Hurst, Erik and Notowidigdo, Matthew J.** (2018). ‘Housing Booms and Busts, Labor Market Opportunities, and College Attendance’, *American Economic Review* 108(10), 2947–94.
- Cheng, Ing-Haw, Raina, Sahil and Xiong, Wei.** (2014). ‘Wall Street and the Housing Bubble’, *American Economic Review* 104(9), 2797–2829.

- Chi, Chun-Che, LaPoint, Cameron and Lin, Ming-Jen.** (2021), Flip or Flop? Tobin Taxes in the Real Estate Market. Working Paper.
- Chinco, Alex and Mayer, Christopher.** (2015). ‘Misinformed Speculators and Mispricing in the Housing Market’, *Review of Financial Studies* 29(2), 486–522.
- Choi, Hyun-Soo, Hong, Harrison G. and Scheinkman, José.** (2014). ‘Speculating on Home Improvements’, *Journal of Financial Economics* 111(3), 609–624.
- 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.
- Daniel, Kent, Hirshleifer, David and Subrahmanyam, Avaniidhar.** (1998). ‘Investor Psychology and Security Market Under- and Overreactions’, *Journal of Finance* 53(6), 1839–1885.
- Daniel, Kent, Hirshleifer, David and Subrahmanyam, Avaniidhar.** (2001). ‘Overconfidence, Arbitrage, and Equilibrium Asset Pricing’, *Journal of Finance* 56(3), 921–965.
- 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.
- 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.** (2017). ‘Credit-Induced Boom and Bust’, *The Review of Financial Studies* 30(11), 3711–3758.
- 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.** (2017). ‘The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium’, *Journal of Political Economy* 125(1), 140–223.
- Foote, Christopher L., Gerardi, Kristopher S. and Willen, Paul S.** (2012). ‘Why Did So Many People Make So Many *Ex Post* Bad Decisions? The Causes of the Foreclosure Crisis’, *Federal Reserve Bank of Boston Public Policy Discussion Papers* 12(2).
- Gao, Zhenyu, Sockin, Michael and Xiong, Wei.** (2020). ‘Economic Consequences of Housing Speculation’, *Review of Financial Studies* 33(11), 5248–5287.
- Genesove, David and Mayer, Christopher.** (2001). ‘Loss Aversion and Seller Behavior: Evidence from the Housing Market’, *Quarterly Journal of Economics* 116(4), 1233–1260.
- Genesove, David and Mayer, Christopher J.** (1997). ‘Equity and Time to Sale in the Real Estate Market’, *American Economic Review* pp. 255–269.
- Glaeser, Edward L.** (2013). ‘A Nation of Gamblers: Real Estate Speculation and American History’, *American Economic Review* 103(3), 1–42.


- Glaeser, Edward L., Gyourko, Joseph and Saiz, Albert.** (2008). ‘Housing Supply and Housing Bubbles’, *Journal of Urban Economics* 64(2), 198–217.
- Glaeser, Edward L. and Nathanson, Charles G.** (2017). ‘An extrapolative model of house price dynamics’, *Journal of Financial Economics* 126(1), 147 – 170.
- Griffin, John M, Kruger, Samuel and Maturana, Gonzalo.** (2020). ‘What drove the 2003–2006 house price boom and subsequent collapse? Disentangling competing explanations’, *Journal of Financial Economics* .
- Guren, Adam.** (2014), The Causes and Consequences of House Price Momentum. Working paper, Harvard University.
- Guren, Adam.** (2018). ‘House Price Momentum and Strategic Complementarity’, *Journal of Political Economy* 126(3), 1172–1218.
- Guren, Adam M and McQuade, Timothy J.** (2020). ‘How Do Foreclosures Exacerbate Housing Downturns?’, *Review of Economic Studies* 87(3), 1331–1364.
- 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.
- 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.
- Hedlund, Aaron.** (2016). ‘The Cyclical Dynamics of Illiquid Housing, Debt, and Foreclosures’, *Quantitative Economics* 7, 289–328.
- 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.
- Liao, Jingchi and Peng, Cameron.** (2018), Price and Volume Dynamics in Bubbles. Working Paper.
- Mayer, Christopher.** (2011). ‘Housing Bubbles: A Survey’, *Annual Review of Economics* 3, 559–577.
- Mayer, Christopher, Pence, Karen and Sherlund, Shane M.** (2009). ‘The Rise in Mortgage Defaults’, *Journal of Economic Perspectives* 23(1), 27–50.
- Mian, Atif, Rao, Kamallesh and Sufi, Amir.** (2013). ‘Household Balance Sheets, Consumption, and the Economic Slump’, *Quarterly Journal of Economics* 128(4), 1687–1726.

- Mian, Atif and Sufi, Amir.** (2022). ‘Credit supply and housing speculation’, *Review of Financial Studies* 35(2), 680–719.
- Mian, Atif, Sufi, Amir and Trebbi, Francesco.** (2015). ‘Foreclosures, House Prices, and the Real Economy’, *Journal of Finance* 70(6), 2587–2634.
- Nathanson, Charles G. and Zwick, Eric.** (2018). ‘Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market’, *Journal of Finance* 73(6), 2587–2633.
- Ngai, L. Rachel and Sheedy, Kevin D.** (2020). ‘The Decision to Move House and Aggregate Housing-Market Dynamics’, *Journal of the European Economic Association* 18(5), 2487–2531.
- 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.
- 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.
- 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.
- Rognlie, Matthew, Shleifer, Andrei and Simsek, Alp.** (2017). ‘Investment Hangover and the Great Recession’, *American Economic Journal: Macroeconomics* 10(2), 113–53.
- Rouwenhorst, K. Geert.** (1995), Asset Pricing Implications of Equilibrium Business Cycle Models, in **Thomas F. Cooley.**, ed., ‘Frontiers of Business Cycle Research’, Princeton University Press, Princeton, NJ, pp. 294–330.
- Scheinkman, Jose A. and Xiong, Wei.** (2003). ‘Overconfidence and Speculative Bubbles’, *Journal of Political Economy* 111(6), 1183–1220.
- 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.
- 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 Dynamics of Prices and Volume
Panel A. National




Panel B. Phoenix, AZ




Panel C. Las Vegas, NV



Panel D. Orlando, FL




Panel E. Bakersfield, CA




Notes: This figure displays the dynamic relation between prices and volume in the U.S. housing market between 2000 and 2011. Panel A shows monthly prices and sales volume at the aggregate level. Panels B–E show analogous series for a set of cities that represent regions with the largest boom–bust cycles during this time: Phoenix, AZ; Las Vegas, NV; Orlando, FL; and Bakersfield, CA. Monthly price index information comes from CoreLogic and monthly sales volume is based on aggregated transaction data from CoreLogic for 115 MSAs representing 48% of the U.S. housing stock. We apply a calendar-month seasonal adjustment for volume. Shaded regions denote the quiet, defined as the period between the peak of volume and the last peak of prices before their pronounced decline.

FIGURE 2
The Lead-Lag Relation between Prices and Volume




Notes: This figure shows that the correlation between prices and lagged volume is robust across MSAs and maximized at a positive lag of 24 months. We regress the demeaned log of prices on seasonally adjusted lagged volume divided by the 2000 housing stock following Eq. (1) 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. The implied correlation equals $\beta_k \text{std}(v_{i,t-k}) / \text{std}(p_{i,t})$, where $v_{i,t-k}$ and $p_{i,t}$ are the demeaned regressors.


FIGURE 3
The Dynamics of Prices and Inventories
Panel A. National




Panel B. Phoenix, AZ




Panel D. Daytona Beach, FL



Panel C. Reno, NV




Panel E. Bakersfield, CA



Notes: This figure displays the dynamic relation between prices and inventory in the U.S. housing market between 2000 and 2011. Panel A shows monthly prices and the inventory of listings at the aggregate level. Panels B–E show analogous series for a set of cities that represent regions with the largest boom–bust cycles during this time: Phoenix, AZ; Reno, NV; Daytona Beach, FL; and Bakersfield, CA. Aggregate inventory information comes from the National Association of Realtors, which are available starting in 2000. Our MSA-level inventory data are available for these cities starting in 2001. Monthly price index information comes from CoreLogic and monthly inventory by MSA is based on aggregated data from CoreLogic for 57 of the 115 MSAs in our main sample for which listings data are available. We apply a calendar-month seasonal adjustment for inventories. Shaded regions denote the quiet, defined as the period between the peak of volume and the last peak of prices before their pronounced decline.

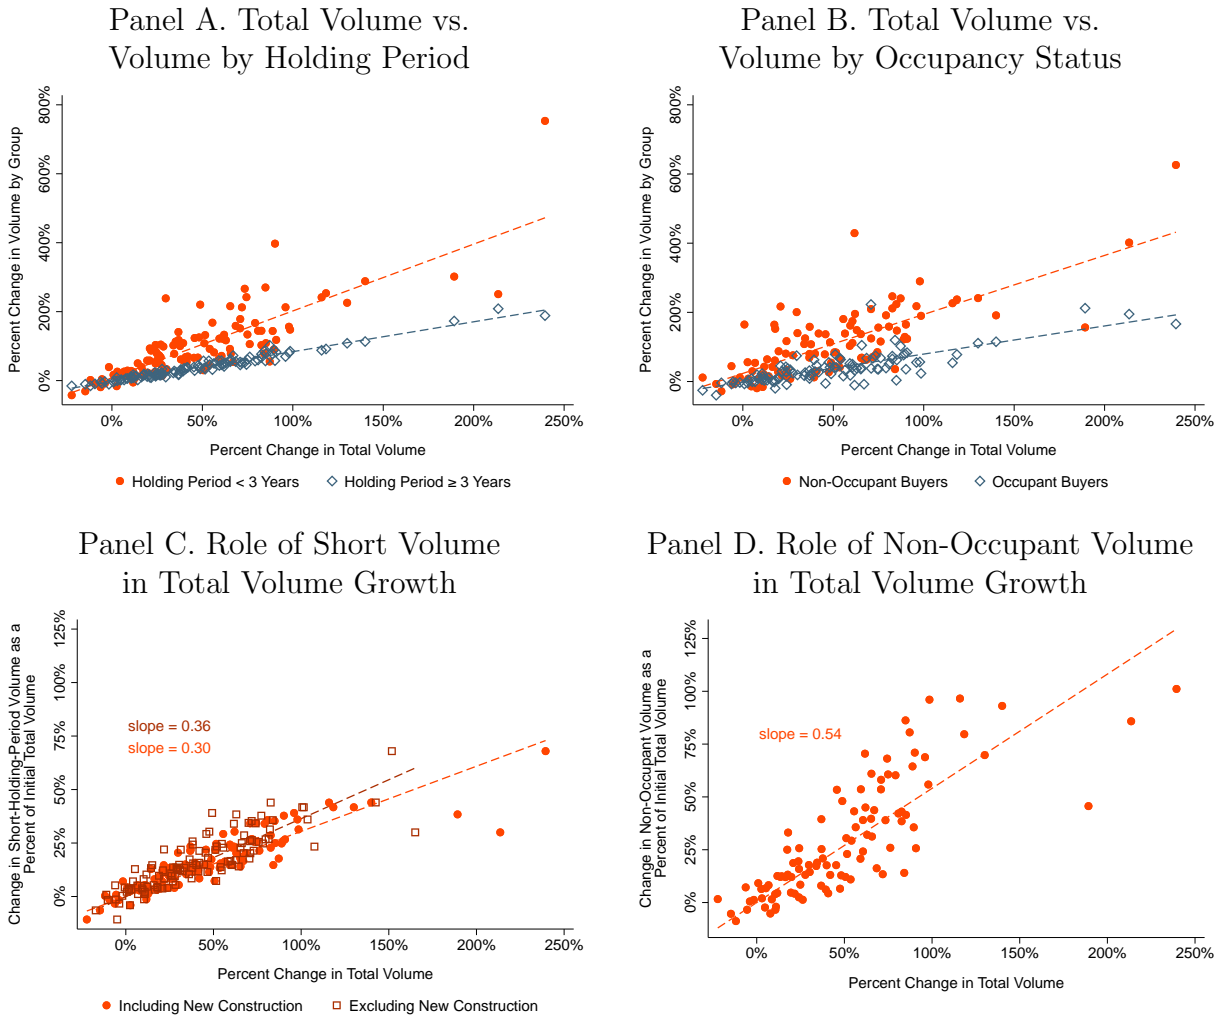
FIGURE 4
Normalized Aggregate Volume by Transaction Type



Notes: This figure shows monthly aggregate time series for total transaction volume (navy triangles), total volume excluding new construction (blue circles), short-holding-period volume (red squares), and non-occupant volume (orange diamonds) between 2000 and 2011. All series exclude lender acquisitions and dispositions of foreclosed properties to remove the mechanical increase in short-term spells driven by forced sales during the bust. The non-occupant volume series only includes observations from the 102 MSAs for which we can consistently identify these transactions; the other series include observations for all 115 MSAs. Each series is separately normalized relative to its average value in the year 2000 and seasonally adjusted by removing calendar-month fixed effects. The raw counts of each type of transaction in the years 2000, 2005, and 2010 are reported in the upper right corner of the figure. In the table, *S1* refers to the short-holding-period sample of 115 MSAs and *S2* refers to the non-occupant sample of 102 MSAs.

FIGURE 5

Short Holding Period, Non-Occupant, and Total Volume Growth Across MSAs



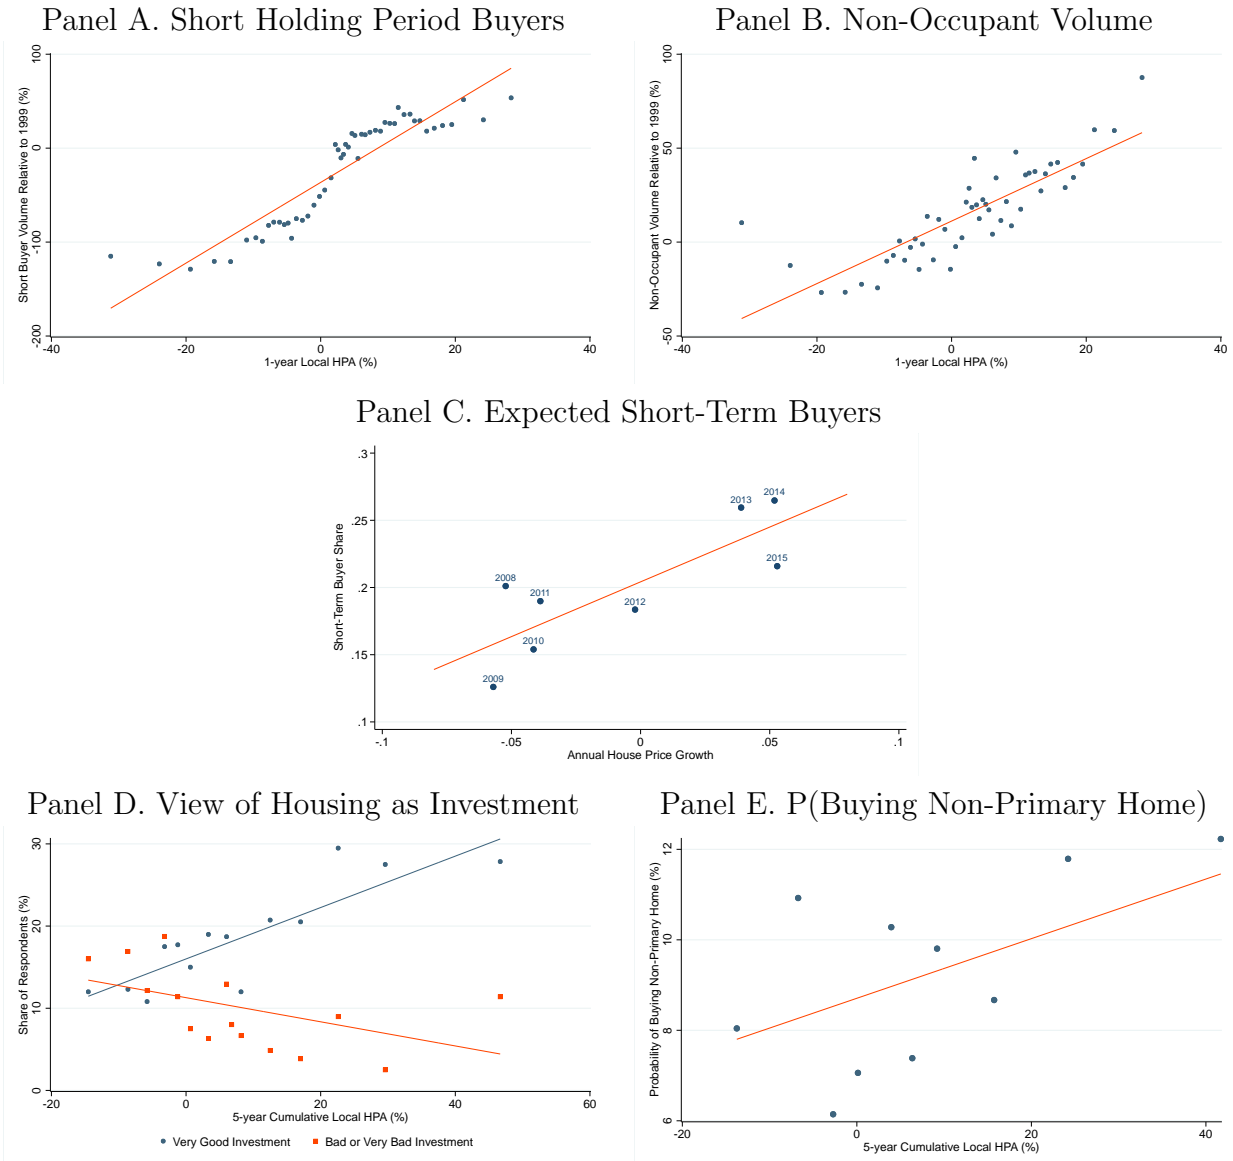
Notes: This figure illustrates the quantitative importance of short holding period and non-occupant volume in accounting for the increase in total volume across MSAs between 2000 and 2005. The top two panels present MSA-level scatter plots of the percentage change in total volume from 2000 to 2005 versus the percentage change in volume for short and long holding periods (Panel A) and the percent change in volume for occupant and non-occupant buyers (Panel B). The bottom two panels show that the growth in short-holding-period and non-occupant volume were quantitatively important components of the growth in total volume across MSAs. For each MSA, we plot the change in short-holding-period volume (Panel C) and non-occupant volume (Panel D) divided by initial total volume on the y-axis against the percentage change in total volume on the x-axis. Because short-holding-period volume is based on the holding period of the seller and therefore cannot, by definition, include sales of newly constructed homes, Panel C also includes a version of the scatter plot that excludes new construction from total volume.

FIGURE 6
The Flow of Listings for Short-Holding-Period Buyers



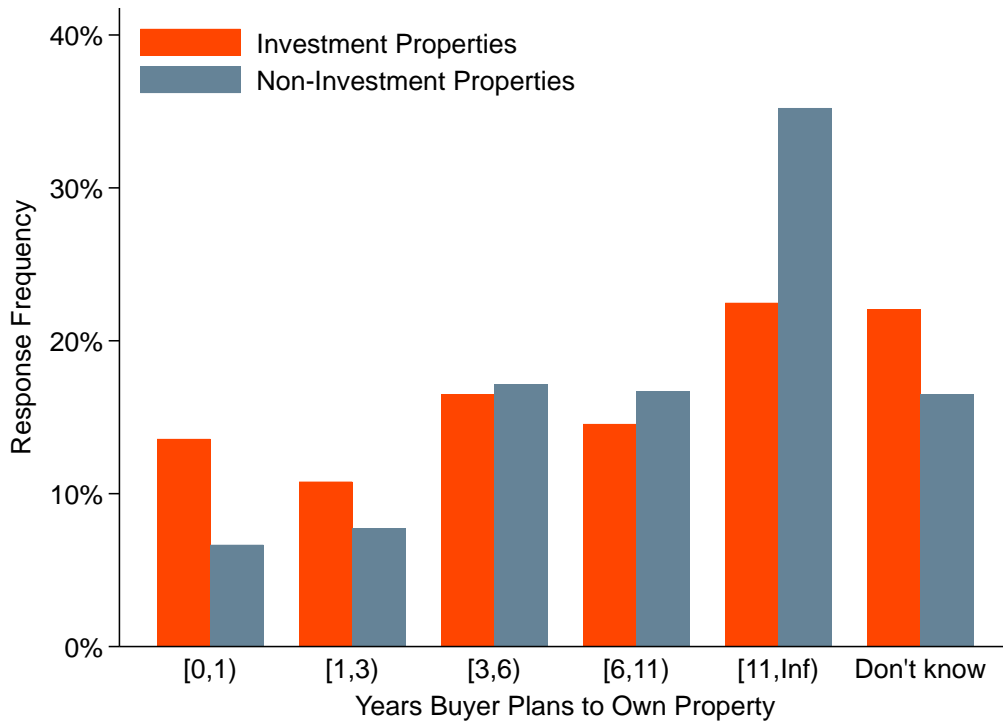
Notes: In this figure, we illustrate the time variation in propensities to list among recent buyers versus all buyers between 2000 and 2011 in the U.S. We link listings micro data to transaction data at the property level to identify short-holding-period listings. We plot monthly aggregate time series for total listings (blue circles) and short-holding-period listings (red squares), defined as a listing where the previous sale occurred within the past three years. The series include observations for the 57 MSAs in our listings sample. Each series is separately normalized relative to its average value in the year 2003 and seasonally adjusted by removing calendar-month fixed effects. The raw counts of each type of listing in the years 2003, 2007, and 2010 are also reported in the upper right corner of the figure.

FIGURE 7
Speculative Homebuying and Recent House Price Appreciation



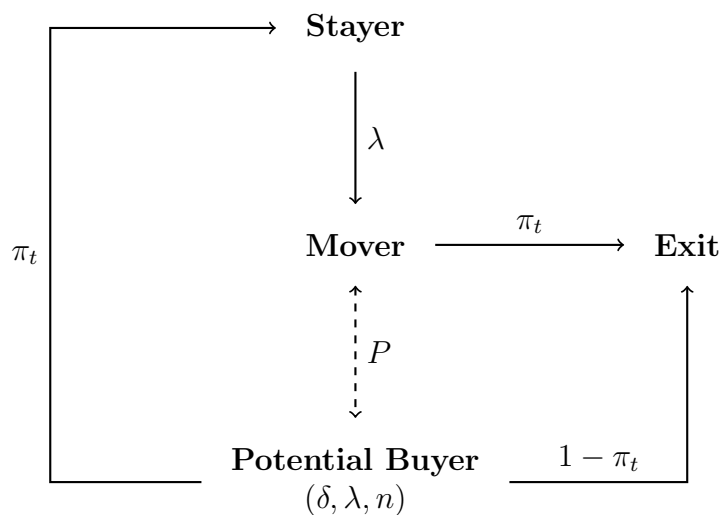
Notes: Panels A and B use CoreLogic data to show the relation between short holding period volume and non-occupant volume at the MSA level, respectively, and the past year’s house price appreciation. Volume measures are scaled relative to their level in 1999. Short-holding-period volume in Panel A is forward-looking, i.e., it is based on whether the buyer sells within three years. Panel C uses data from the NAR Investment and Vacation Home Buyers Survey; “annual house price growth” equals the average across that year’s four quarters of the log change in the all-transactions FHFA U.S. 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 an expected horizon of less than three years. We use the FHFA index here because it covers the 2015–2016 period. Panels D and E use data from the Federal Reserve Survey of Consumer Expectations and Armona et al. (2019) to study the relation between recent house price growth and the probability of buying a non-primary home. In these data, local house price appreciation is computed at the ZIP Code-level from Zillow.

FIGURE 8
 Expected Holding Times of Homebuyers, 2008–2015




Notes: This figure presents evidence on heterogeneity in expected holding times among recent homebuyers from the NAR Investment and Vacation Home Buyers Survey. We plot the response frequency averaged equally over each survey year from 2008 to 2015. We reclassify buyers who have already sold their properties by the time of the survey as having an expected holding time in [0,1).

FIGURE 9
Model Flowchart



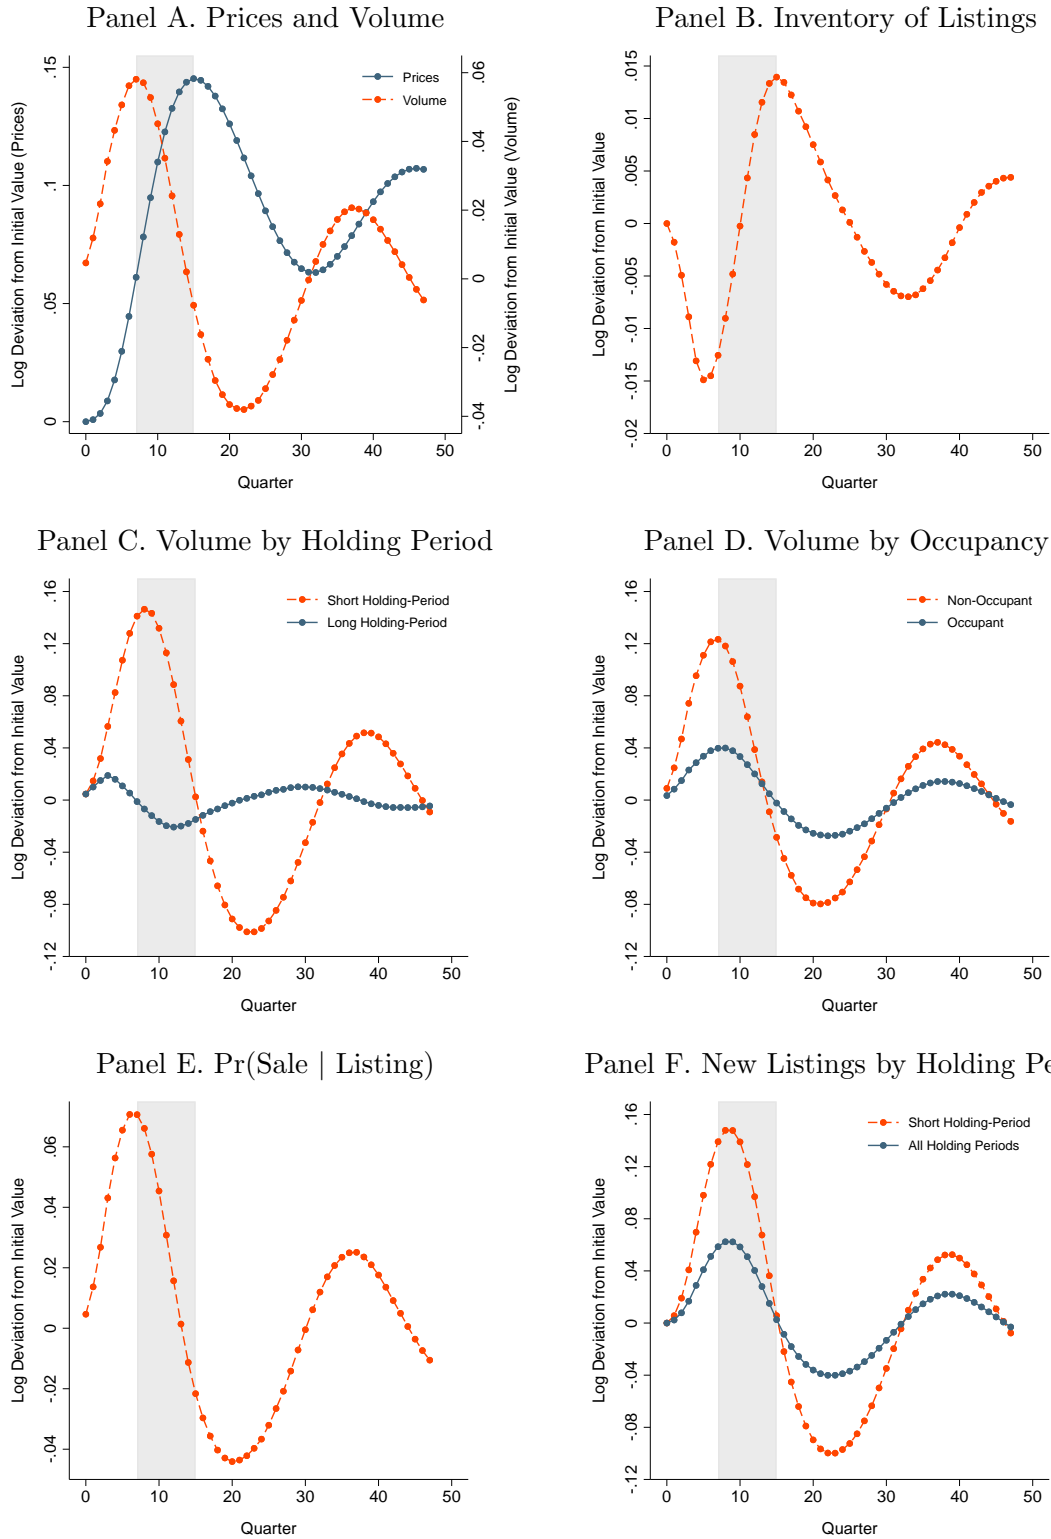
Notes: This figure illustrates how agents in the model transition between different types. At each time period, t , all movers list their homes for sale and are matched to potential buyers. Potential buyers decide whether to purchase at the mover's listed price P . If a potential buyer purchases, she becomes a stayer and receives a constant per-period housing utility e^δ during each period that she remains a stayer. If she does not purchase, she exits the market and consumes her terminal wealth. A stayer transitions into being mover with probability λ each period. The log of the housing utility, δ , that potential buyers who purchase will receive as stayers is the sum of a time-varying aggregate demand shifter, d_t , which is common to all potential buyers matched at time t , and an idiosyncratic term, a , which varies across potential buyers within a given cohort. The idiosyncratic term a is distributed $\mathcal{N}(\mu_n, \sigma_a)$ and depends on the potential buyer's type, n , which can be either occupant ($n = 1$) or non-occupant ($n = 0$). The mover hazard λ also differs across potential buyers and follows a discrete distribution given by $\Pr(\lambda = \lambda_j) = \beta_{n,j}$, for $j \in \{1, \dots, J\}$. Potential buyers are aware of both δ and λ at the time they decide to purchase. Movers receive zero housing flow utility during the time which they are attempting to sell. Those who do not sell remain movers in the next period and those who do sell exit the market and consume their terminal wealth. We denote the probability of a sale by π_t .

FIGURE 10
 Buying Cutoffs for Different Expected Growth Rates



Notes: The buying cutoff, $\kappa_{n,j}(\hat{g}_t)$, determines how large a potential buyer's housing utility must be relative to the price of a house for her to decide to buy. It depends on the potential buyer's occupancy type, n , her quarterly moving hazard, λ_j , and the current expected quarterly growth rate of the demand shifter, \hat{g}_t . We plot values of these functions for the λ values in our calibration, which appear in the legend. Solid lines correspond to occupants ($n = 1$); dashed lines correspond to non-occupants ($n = 0$). The horizontal grey dashed line gives $\bar{\kappa}$, which agents mistakenly believe is the time-invariant buying cutoff for other potential buyers.


FIGURE 11
Impulse Responses




Notes: Impulse responses are average differences between log outcomes in control simulations and treatment simulations, in which a 2-standard-deviation shock to ϵ_t^g (the demand growth innovation) occurs in quarters 0 through 3. The shaded grey area denotes the beginning and end of the quiet. A short holding period is defined as less than or equal to 12 quarters and a long holding period is defined as greater than 12 quarters.

FIGURE 12
Impulse Responses in Counterfactuals


Panel A. Prices and Volume, Rational




Panel B. Inventory of Listings, Rational




Panel C. Volume by Holding Period, Rational




Panel D. Volume by Occupancy, Rational




Panel E. Prices and Volume, Walrasian




Panel F. Inventory of Listings, Walrasian




Panel G. Volume by Holding Period, Walrasian




Panel H. Volume by Occupancy, Walrasian




Panel I. Prices and Volume, No Speculation




Panel J. Inventory of Listings, No Speculation



Panel K. Volume by Holding Period, No Speculation



Panel L. Volume by Occupancy, No Speculation



Notes: Impulse responses are average differences between log outcomes in control simulations and treatment simulations, in which a 2-standard-deviation shock to ϵ_t^d (the demand growth innovation) occurs in quarters 0 through 3. A short holding period is defined as less than or equal to 12 quarters and a long holding period is defined as greater than 12 quarters.

TABLE 1
Speculators and Housing Market Outcomes (Summary Statistics)

Panel A. Short-Volume Sample			
Variable	Mean	Standard Deviation	Observations
Short-Volume Boom	15.97	12.93	115
Price Boom	97.06	47.88	115
Price Bust	-27.9	13.64	115
Δ Volume Quiet + Bust	-62.96	18.87	115
Foreclosures Bust	82.84	55.96	115

Panel B. Non-Occupant Volume Sample			
Variable	Mean	Standard Deviation	Observations
Non-Occupant Volume Boom	29.29	27.05	102
Short-Volume Boom	16.88	13.36	102
Price Boom	100.57	49.27	102
Price Bust	-28.99	13.97	102
Δ Volume Quiet + Bust	-63.32	19.47	102
Foreclosures Bust	86.57	58.08	102

Panel C. Short-Volume Sample with Listings			
Variable	Mean	Standard Deviation	Observations
Short-Volume Boom	14.64	12.33	57
Δ Listings Boom	91.67	94.93	57
Δ Listings Quiet	178.39	143.86	57

Panel D. Non-Occupant Volume Sample with Listings			
Variable	Mean	Standard Deviation	Observations
Non-Occupant Volume Boom	27.81	27.32	48
Short-Volume Boom	15.84	12.88	48
Δ Listings Boom	82.11	93.67	48
Δ Listings Quiet	171.74	151.29	48

Notes: This table reports summary statistics for MSA-level variables in different samples of MSAs. Δ Volume Quiet + Bust is defined as the change in total volume from 2005 through 2011. Δ Listings Boom is defined as the change in total listings from 2003 through 2005. Δ Listings Quiet is defined as the change in total listings from 2005 through 2007. Foreclosures Bust is defined as total foreclosures from 2007 through 2011. Price Boom is defined as the change in prices from 2000 through 2006. Price Bust is defined as the change in prices from 2006 through 2011. To aid interpretation of these relations, we scale the change in outcomes for all quantity measures relative to total volume in 2003 and multiply by 100. Total volume in 2003 has mean 28,061 and standard deviation 43,708 in the Short Volume Sample and mean 25,167 and standard deviation 35,967 in the Short Volume Sample with Listings.

TABLE 2
Speculative Booms and Housing Market Outcomes

Panel A. MSA-Level Prices				
	Price Boom		Price Bust	
Short-Volume Boom	1.930*** (0.297)		-0.571*** (0.083)	
Non-Occupant Volume Boom		0.570*** (0.173)		-0.166*** (0.049)
Number of Observations	115	102	115	102
R-squared	0.272	0.098	0.293	0.103
Panel B. MSA-Level Inventories				
	Δ Listings Boom		Δ Listings Quiet	
Short-Volume Boom	-1.133 (1.027)		5.961*** (1.353)	
Non-Occupant Volume Boom		-0.070 (0.505)		2.645*** (0.718)
Number of Observations	57	48	57	48
R-squared	0.022	0.000	0.261	0.228
Panel C. MSA-Level Volume Quiet and Bust				
	Δ Volume Quiet + Bust		Foreclosures Bust	
Short-Volume Boom	-1.047*** (0.096)		0.895** (0.398)	
Non-Occupant Volume Boom		-0.512*** (0.051)		-0.060 (0.215)
Number of Observations	115	102	115	102
R-squared	0.515	0.505	0.043	0.001

Notes: This table reports estimates of the relation between speculative volume and housing cycle measures at the MSA level. Δ Volume Quiet + Bust is defined as the change in total volume from 2005 through 2011. Δ Listings Boom is defined as the change in total listings from 2003 through 2005. Δ Listings Quiet is defined as the change in total listings from 2005 through 2007. Foreclosures Bust is defined as total foreclosures from 2007 through 2011. Price Boom is defined as the change in prices from 2000 through 2006. Price Bust is defined as the change in prices from 2006 through 2011. To aid interpretation of these relations, we scale the change in outcomes for all quantity measures relative to total volume in 2003 and multiply by 100. Table 1 presents summary statistics for each sample. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively. Standard errors appear in parentheses.

TABLE 3
All-Cash Buyer Shares

	Transaction-Level	MSA-Level			
	All Months	All Months	Boom	Quiet	Bust
Short Buyers	0.29	0.38 (0.21)	0.29 (0.16)	0.28 (0.17)	0.52 (0.20)
Non-Occupant Buyers	0.38	0.41 (0.18)	0.36 (0.15)	0.32 (0.14)	0.50 (0.18)
All Buyers	0.20	0.25 (0.16)	0.22 (0.15)	0.20 (0.14)	0.30 (0.16)

Notes: This table presents statistics on the share of buyers of various types who purchased their homes without the use of a mortgage (“all-cash buyers”). In column 1, the all-cash buyer share is measured at the transaction level and includes all transactions recorded between January 2000 and December of 2011 from the CoreLogic deeds records described in Section 1.1. The first row includes only transactions by homebuyers who are observed to have sold the home within three years of purchase. The second row includes only non-occupant buyers. The third row includes all buyers. In columns 2–5, all-cash buyer shares are first calculated at the MSA-by-month level and then averaged across MSA-months within a given time period. The standard deviation of these MSA-month means is reported in parentheses for reference. Column 2 includes all MSA-months between January 2000 and December 2011. Column 3 includes only MSA-months between January 2000 and August 2005. Column 4 includes only MSA-months between August 2005 and December 2006. Column 5 includes only MSA-months between December 2006 and December 2011. All statistics are calculated in the full sample of 115 MSAs with the exception of those for non-occupants, which are calculated in the sample of 102 MSAs with valid non-occupancy data.

TABLE 4
Inputs into model calibration

Parameter or target	Value	Source
Panel A: Assumed parameters		
r (non-mover discount rate)	0.012	Guren (2018)
Potential λ values	{0.50, 0.17, 0.05, 0.03, 0.01}	Figure 8
ρ (demand growth persistence)	0.880	GN (2017)
Panel B: Steady-state targets		
Occupant buyer shares	(0.06, 0.07, 0.16, 0.16, 0.34)	Figure 8
Non-occupant buyer shares	(0.04, 0.03, 0.04, 0.04, 0.06)	Figure 8
Annual volatility of demand growth	0.023	GN (2017)
Quarterly selling hazard	0.75	Guren (2018)
Mean demand error	0	Model
Mean demand growth	0	Model
Panel C: Cycle targets		
Price overshoot	2.3	Figure 1
Volume boom/price boom	0.4	Figure 1
Non-occupant boom/occupant boom	3.1	Figure 4

Notes: This table reports parameters that we assume in the calibration, as well as targets we use to determine the remaining parameters. In the model, we target the mean buyer shares, quarterly selling hazard, and demand error across all analysis periods in control simulations. We theoretically derive the annual volatility of demand growth as well as the mean demand growth as functions of parameters. Price overshoot is the ratio of log price growth from the beginning to peak to log price growth from the beginning to the trough after the peak. Volume boom/price boom is the ratio of log existing volume growth from the beginning to the peak of volume (2000 to 2005, using numbers from Figure 4) to aforementioned log price growth. Non-occupant boom/occupant boom is the ratio of each category of log volume growth from 2000 to 2005 in the sample of MSAs we use for non-occupant analysis. In the model, we use quarterly minimums and maximums instead of aggregating at the year level. We match all targets to within rounding. GN (2017) denotes Glaeser and Nathanson (2017).

TABLE 5
Outputs from model calibration

Parameter or outcome	Interpretation	Value
Panel A: Derived parameters		
σ_a	Flow utility dispersion	0.066
μ_1	Occupant premium	0.009
γ	g variance share	0.070
$\bar{\kappa}$	Assumed buying cutoff	0.029
σ_d	Demand volatility	0.011
μ_g	Mean demand growth	-0.000
r_m	Mover discount rate	0.141
$\beta_{0,j}$	Non-occupant shares	(0.143, 0.022, 0.030, 0.030, 0.045)
$\beta_{1,j}$	Occupant shares	(0.185, 0.052, 0.119, 0.119, 0.254)
Panel B: Steady-state outcomes		
1-year extrapolation	–	0.127
2–5-year extrapolation	–	0.042

Notes: See text for definitions of parameters in Panel A. We find these values by searching for parameters such that moments from the model match targets in Table 4. Panel B reports regression coefficients of annualized price growth in the next year and between 2 and 5 years from now on last year’s price growth. We run these regressions across control simulations at the beginning of the analysis period.

TABLE 6
Outcomes for different tax regimes

Outcome	Base- line	Tax on all buyers			Tax on non-occupant buyers		
		0.5%	1%	5%	0.5%	1%	5%
Price boom	14.5	13.1	12.2	9.7	13.5	12.8	12.6
Price bust	-8.2	-6.4	-5.1	-1.1	-7.0	-6.0	-5.8
Volume boom	5.8	5.5	5.2	4.1	5.3	4.8	4.5
Listings, end of boom	-1.3	-1.2	-1.1	-1.0	-0.8	-0.8	-0.6
Listings, end of quiet	1.4	1.3	1.2	0.6	1.2	0.9	0.9
Short volume boom	14.1	13.7	13.2	10.3	13.5	11.8	11.4
Non-occupant volume boom	12.3	11.7	11.1	8.9	7.5	2.3	0.1
Sale probability boom	7.1	6.7	6.3	5.1	6.2	5.5	5.0

Notes: We report 100 times changes in log outcomes between treatment and control simulations. We define the end of the quiet as the first local maximum in the impulse response of log prices, and we measure the following outcomes at that time: price boom and listings end of quiet. We define the end of the boom as the first local maximum in the impulse response of log volume before the end of the quiet, and we measure the following outcomes at that time: volume boom, listings end of boom, short volume boom, non-occupant volume boom, and sale probability boom. The price bust is the change from the end of the quiet to the first local minimum of the impulse response of log prices after the end of the quiet. The tax is relative to the purchase price, payable at time of sale. We alter $\bar{\kappa}$ in each column to maintain a zero demand error while keeping other parameters the same. The baseline values correspond to Figure 11.

Internet Appendix for “Speculative Dynamics of Prices and Volume”

ANTHONY A. DEFUSCO
CHARLES G. NATHANSON
ERIC ZWICK

A Data

To conduct our empirical analysis we make use of a transaction-level data set containing detailed information on individual home sales taking place throughout the US between 1995 and 2014. The raw data was purchased from CoreLogic and is sourced from publicly available tax assessment and deeds records maintained by local county governments. In some analyses we supplement this transaction-level data with additional data on the listing behavior of individual homeowners. Our listings data is also provided by CoreLogic and is sourced from a consortium of local Multiple Listing Service (MLS) boards located throughout the country.

Selecting Geographies

To select our sample of transactions, we first focus on a set of counties that have consistent data coverage going back to 1995 and which, together, constitute a majority of the housing stock in their respective MSAs. In particular, to be included in our sample a county must have at least one “arms length” transaction with a non-negative price and non-missing date in each quarter from 1995q1 to 2014q4.¹ Starting with this subset of counties, we then further drop any MSA for which the counties in this list make up less than 75 percent of the total owner-occupied housing stock for the MSA as measured by the 2010 Census. This leaves us with a final set of 250 counties belonging to a total of 115 MSAs. These MSAs are listed below in Table IA1 along with the percentage of the housing stock that is represented by the 250 counties for which we have good coverage. Throughout the paper, when we refer to counts of transactions in an MSA we are referring to the portion of the MSA that is accounted for by these counties.

Selecting Transactions

Within this set of MSAs, we start with the full sample of all arms length transactions of single family, condo, or duplex properties and impose the following set of filters to ensure that our final set of transactions provides an accurate measure of aggregate transaction volume over the course of the sample period:

1. Drop transactions that are not uniquely identified using CoreLogic’s transaction ID.
2. Drop transactions with non-positive prices.

¹We rely on CoreLogic’s internal transaction-type categorization to determine whether a transaction occurred at arms length.

3. Drop transactions that appear to be clear duplicates, identified as follows:
 - (a) If a set of transactions has an identical buyer, seller, and transaction price but are recorded on different dates, keep only the earliest recorded transaction in the set.
 - (b) If the same property transacts multiple times on the same day at the same price keep only one transaction in the set.
4. If more than 10 transactions between the same buyer and seller at the same price are recorded on the same day, drop all such transactions. These transactions appear to be sales of large subdivided plots of vacant land where a separate transaction is recorded for each individual parcel but the recorded price represents the price of the entire subdivision.
5. Drop sales of vacant land parcels in MSAs where the CoreLogic data includes such sales.² We define a vacant land sale to be any transaction where the sale occurs a year or more before the property was built.

Table IA2 shows the number of transactions that are dropped from our sample at each stage of this process as well as the final number of transactions included in our full analysis sample.

Identifying Occupant and Non-Occupant Buyers

We identify non-occupant buyers using differences between the mailing addresses listed by the buyer on the purchase deed and the actual physical address of the property itself. In most cases, these differences are identified using the house numbers from each address. In particular, if both the mailing address and the property address have a non-missing house number then we tag any instance in which these numbers are not equal as a non-occupant purchase and any instance in which they are equal as occupant purchases. In cases where the mailing address property number is missing we also tag buyers as non-occupants if both the mailing address and property address street names are non-missing and differ from one another. Typically, this will pick up cases where the mailing address provided by the buyer is a PO Box. In all other cases, we tag the transaction as having an unknown occupancy status.

Restricting the Sample for the Non-Occupant Analysis

Our analysis of non-occupant buyers focuses on the growth of the number of purchases by these individuals between 2000 and 2005. To be sure that this growth is not due to changes in the way mailing addresses are coded by the counties comprising the MSAs in our sample, for the non-occupant buyer analysis we keep only MSAs for which we are confident such changes do not occur between 2000 and 2005. In particular, we first drop any MSA in which the share of transactions in any one year between 2000 and 2005 with unknown occupancy status exceeds 0.5. Of the remaining MSAs, we then drop those for which the increase in the number of non-occupant purchases between any year and the next exceeds 150%, with the

²MSAs are flagged as including vacant land sales if more than 5 percent of the sales in the MSA occur more than two years before the year in which the property was built.

possible base years being those between 2000 and 2005.³ The 102 MSAs that remain after these two filters are marked with an “x” in columns 3 and 7 of Table IA1.

Restricting the Sample for Listings Analysis

The geographic and time series coverage of the CoreLogic MLS data is not as comprehensive as the transaction-level data. As a result, our analysis of listings behavior is restricted to a subset of markets for which we can be relatively certain that the MLS data is representative of the majority of owner-occupied home sales in the area. We impose several filters to identify this subset of MSAs. First, starting with the full set of 115 MSAs contained in the transaction-level data, we drop any MSA for which there is not at least one new listing in every month and in every county subcomponent of the MSA between January 2000 and December 2014. Within the remaining set of MSAs we then drop any MSA for which the number of new listings between 2006 and 2008 is more than 2.5 times the number of new listings between 2003 and 2005. This filter eliminates MSAs that experience large jumps in coverage during the quiet. Finally, we also drop any MSA for which the number of sold listings (from the MLS data) is less than 25 percent of total sales volume (from the transaction data) over the period 2003-2012. This filter eliminates MSAs for which the listings data is likely to be unrepresentative of sales activity during our main sample period. This leaves a final sample of 57 MSAs for our listings analysis. These MSAs are marked with an “x” in columns 4 and 8 of Table IA1.

Identifying New Construction Sales

In several parts of our analysis we omit new construction sales from the calculation of total transaction volume. To identify sales of newly constructed homes, we start with the internal CoreLogic new construction flag and make several modifications to pick up transactions that may not be captured by this flag. CoreLogic identifies new construction sales primarily using the name of the seller on the transaction (e.g. “PULTE HOMES” or “ROCKPORT DEV CORP”), but it is unclear whether their list of home builders is updated dynamically or maintained consistently across local markets. To ensure consistency, we begin by pulling the complete list of all seller names that are ever identified with a new construction sale as defined by CoreLogic. Starting with this list of sellers, we tag any transaction for which the seller is in this list, the buyer is a human being, and the transaction is not coded as a foreclosure sale by CoreLogic as a new construction sale. We use the parsing of the buyer name field to distinguish between human and non-human buyers (e.g. LLCs or financial institutions). Human buyers have a fully parsed name that is separated into individual first and name fields whereas non-human buyer’s names are contained entirely within the first name field.

This approach will identify all new construction sales provided that the seller name is recognized by CoreLogic as the name of a homebuilder. However, many new construction sales may be hard to identify simply using the name of the seller. We therefore augment this definition using information on the date of the transaction and the year that the property was built. In particular, if a property was not already assigned a new construction sale using

³This step drops only Chicago-Naperville-Elgin, IL-IN-WI.

the builder name, then we search for sales of that property that occur within one year of the year that the property was built and record the earliest of such transactions as a new construction sale.

Finally, for properties that are not assigned a new construction sale using either of the two above methods, we also look to see if there were any construction loans recorded against the property in the deeds records. If so, we assign the earliest transaction to have occurred within three years of the earliest construction loan as a new construction sale. We use a three-year window to allow for a time lag between the origination of the construction loan and the actual date that the property was sold. Construction loans are identified using CoreLogic’s internal deed and mortgage type codes.

B Robustness

B.1 Mechanical Short-Term Volume

In Figure 4 we document a rise in the share of volume coming from short-term sales during the boom. Our interpretation of this pattern is that short-term volume rises due to a shift in the composition of buyers toward those with shorter intended holding periods. However, even in the absence of such a shift, any increase in total volume during the early part of the boom will generate a mechanical increase in the share of late-boom volume coming from short-term sales. The richness of our data allows us to quantify the contribution of this mechanical force relative to changes in the composition of buyers.

For each pair of distinct months between 1995 and 2005, we compute a conditional selling hazard $\pi_{t',t}$. This hazard is the share of homes purchased in month t' —and that have not yet sold by month t —that sell in month t . By focusing on selling hazards instead of total volume, we remove the mechanical force that comes from volume increasing over the cycle.

We estimate the following regression at the month-pair level:

$$\pi_{t',t} = \alpha_{y(t')}^{buy} + \alpha_{y(t)}^{sell} + \alpha_{t-t'}^{duration} + \epsilon_{t',t},$$

where $y(\cdot)$ gives the year of the month. The first set of fixed effects, $\alpha_{y(t')}^{buy}$, captures the average propensity of buyer cohorts from year $y(t')$ to sell in any future year. The second set of fixed effects, $\alpha_{y(t)}^{sell}$, captures the average propensity of all owners to sell in year $y(t)$. The third set of fixed effects, $\alpha_{t-t'}^{duration}$, measures time-invariant selling hazard profiles as a function of time elapsed since purchase $t - t'$. We interpret year-to-year movements in $\alpha_{y(t')}^{buy}$ as changes in the composition of buyers across those years, holding fixed both year-specific shocks to selling hazards that affect all cohorts equally and duration-specific drivers of selling hazards that do not vary over the cycle.

Table IA3 reports the buy-year fixed effects estimates for years 2000 to 2005 relative to 2000. The fixed effects are linear differences of a monthly selling hazard, so multiplying by 12 roughly gives the effect on the annualized selling probability. Therefore, buyers in 2005 have a 3.2 percentage point larger annual selling hazard than buyers in 2000 (12 times 0.0027 equals 0.0324).

We use these estimates to construct counterfactual growth of short-term volume from

2000 to 2005. For each $2000m1 \leq t' < t \leq 2005m12$, we construct the counterfactual selling hazard as

$$\pi_{t',t}^c = \pi_{t',t} - \left(\hat{\alpha}_{y(t')}^{buy} - \hat{\alpha}_{2000}^{buy} \right),$$

which subtracts away any increase due to the change in the composition of buyers from 2000 to the year of t' . We then compute the counterfactual of $v_{t',t}$, the volume of homes bought in t' and sold in t , using the following iterative procedure. Let $e_{t',t}$ count homes bought in t' that have not yet sold by t , and let c superscripts mark counterfactual values. We initialize counterfactuals with actuals: for each $1995m1 \leq t' < 2005m12$,

$$\begin{aligned} e_{t',t}^c &= e_{t',t} \\ v_{t',t}^c &= v_{t',t}. \end{aligned}$$

We then iteratively update the counterfactuals over t running from $t' + 1$ to $2005m12$:

$$\begin{aligned} e_{t',t}^c &= e_{t',t-1}^c - v_{t',t-1}^c \\ v_{t',t}^c &= \pi_{t',t}^c e_{t',t}^c. \end{aligned}$$

To compute short-term volume in year y , we sum $v_{t',t}$ across all subscripts for which $y(t) = y$ and $0 < t - t' < 36$; we sum $v_{t',t}^c$ across the same indices for counterfactual short-term volume.

The remaining columns of Table IA3 report the results. Between 2000 and 2005, total volume grows 36.7% and short-term volume grows 77.5% in the actual data. The disproportionate rise in short-term volume is the difference, 40.8%. Counterfactual short-term volume rises 41.5% between 2000 and 2005, giving a disproportionate rise of 4.8%. Therefore, $4.8\%/40.8\% = 11.8\%$ of the disproportionate rise in short-term volume remains in the counterfactual. We attribute the 88.2% that disappeared to the changing composition of buyers between 2000 and 2005.

B.2 Endogenous Holding Periods

The empirical evidence presented in Section 3 indicates that the differential entry of speculative buyers played a major role in driving the volume boom. However, the results for short-term volume growth are based on realized rather than expected holding periods. This way of measuring short-term speculation may complicate the interpretation of our results if buyers' intended holding periods endogenously respond to changes in economic conditions during the boom. The results on non-occupant buyers partially address this concern as they are based on a measure of speculative entry that does not suffer from the same issue. This section addresses this issue further using an instrumental variables strategy.

Our approach instruments for realized short-term volume growth using ex-ante demographic characteristics of an area that are likely to be correlated with intended short holding periods among potential homebuyers. We use the 2000 Census 5% microdata to calculate the share of recent homebuyers (within the last 5 years) in each MSA that were either younger than 35 or aged 65 and older at the time of questioning and include both shares as instruments for 2000–2005 short-term volume growth. This approach follows Edelstein and Qian (2014), who use data from the American Housing Survey to study demographic and mort-

gage characteristics as predictors of ex-ante investment horizon. Both older and younger buyers tend to have shorter horizons than middle-aged buyers, likely due to life cycle forces that affect the propensity to move, which gives the instrument its relevance.

The strength of this instrument is that it is predetermined relative to the realized holding periods for sellers in the boom and may therefore help purge our estimates of mechanical bias arising from endogenous changes in holding periods over the course of ownership spells. We stress this instrument does not remove the influence of age-specific shocks, so we do not interpret the IV regressions as demonstrating a causal relation. Rather our goal with this exercise is to mitigate potential mechanical feedback between total and short-term volume.

Table IA4 presents the results. Column 1 presents first stage regressions of the short-term volume boom on the old and young shares. The F-statistic of 39.95 indicates the IV regressions are well powered. Column 2 shows that an OLS regression of the 2000–2005 percent change in total volume on the 2000–2005 change in short-term volume divided by year-2000 total volume replicates the conclusion from Figure 5, Panel C. Because we are interested in instrumenting for short-volume growth, the left- and right-hand-side variables in this regression are swapped relative to their analogs in Figure 5. Thus, the coefficient estimate of 2.3 reported in Panel A is not directly comparable to the 0.3 number from Figure 5, Panel C. We rescale the coefficient using a variance decomposition, which indicates that 33 percent of the variation in total volume growth across MSAs can be explained by changes in short-term volume, thus matching the short-term volume result from Figure 5.

In Table IA4, column 3, the short-term volume coefficient does not change when we instrument using year-2000 homebuyer age. If a mechanical relation were driving this correlation, we would expect the IV coefficient to fall relative to the OLS. Columns 4 through 7 show that the relations between the price boom and bust and the short-term volume boom strengthen in the IV specifications. This result suggests a modest negative feedback between price growth and holding period, perhaps reflecting a disposition effect force in which price growth induces buyers to sell earlier than they otherwise would. Overall, the IV results present strong evidence that the change in realized short-term volume is quantitatively important for determining overall volume growth and the size of the price cycle, even when using only the portion of short-term volume growth predicted by ex-ante buyer characteristics.

B.3 High Frequency Analysis of Price Growth and Speculative Volume (pVAR)

To further investigate the link between house price changes and speculative entry, we examine higher frequency data. Speculative buyers may both cause and respond to house price changes. Because of the potential for this type of feedback mechanism, we do not attempt to directly identify the “causal” effect of speculators on house prices.⁴ Instead, we follow the approach in Chinco and Mayer (2015), who estimate predictive regressions that are flexible enough to allow for some types of feedback between speculative entry and prices. In particular, we estimate a series of panel vector auto-regressions (pVARs) that relate house

⁴Gao et al. (2020) exploit state capital gains tax changes as an instrument for speculation and use this variation to measure the consequences of housing speculation for the real economy.

price growth to the share of purchases made by non-occupant buyers and “short buyers” (i.e., those who will sell within three years of purchase) at a monthly frequency in each MSA between January 2000 and December 2006 (the year when prices peaked).

Table IA5 reports results from three different pVAR specifications. In column 1, we estimate a simple two-equation model that jointly links both month-over-month house price growth to the lagged share of transactions by short-buyers (top panel) and the contemporaneous short-buyer share to lagged house price appreciation (middle panel). Both equations also include lags of the relevant dependent variable (house price appreciation in the top panel and the short-buyer share in the middle panel).

The results indicate that a 1 percentage point increase in the fraction of purchases made by short-term buyers in a given month is associated with a 0.02 percentage point increase in the house-price appreciation rate in the following month. That is, short-buyer entry is predictive of subsequent house price growth, though we stress that these predictive regressions do not necessarily imply a causal relation. Interestingly, the results in the middle panel indicate that short-buyer entry can also be predicted by recent house price growth. A 1 percentage point increase in house price growth in the prior month is associated with a 0.16 percentage point increase in the short-buyer share of entrants.

In column 2, we estimate a similar model swapping out the short-buyer share for the non-occupant share of purchases. Unlike short-buyer entry, non-occupant entry does not appear to be predictive for house price growth. The coefficient on the lagged non-occupant share in the top panel is roughly half the magnitude of its short-buyer analog from column 1 and is not statistically significant. Non-occupants do, however, appear to respond similarly to past price growth. The estimate in the bottom panel indicates that a 1 percentage point increase in house price growth in the prior month is associated with a 0.12 percentage point increase in the non-occupant share of entrants. This estimate is qualitatively similar to and statistically indistinguishable from the analogous coefficient for short-term buyers.

Finally, in column 3 of the table we estimate a three-equation pVAR that allows for joint relations between all three variables of interest. The results from this specification are both qualitatively and quantitatively similar to those from columns 1 and 2. Short-buyer entry is strongly predictive of subsequent house price growth and predicted by recent past price growth, whereas non-occupant entry can be predicted by past price growth but is less informative for predicting subsequent prices.

These results are similar both qualitatively and quantitatively to those in Chinco and Mayer (2015) (see their Table 7). They find coefficients for lagged out-of-town second-house buyers versus house price growth of 0.02 percentage points, which matches our short-buyer share coefficient. They find that local second-house buyers do not predict future house price growth. Combining their two groups of second-house buyers would deliver an estimate identical to our non-occupant coefficient. Relative to their specification, we consider a sample of MSAs that is five times as large and focus on the distinction between short-term buyers and non-occupants rather than differences within the group of non-occupants.

C Additional analysis of speculation

In this appendix, we provide details about the calculations using microdata in Section 4.

C.1 Overlap between short-term and non-occupant buyers

The statistics in the text focus on the non-occupant sample of 102 MSAs. Of 2000–2005 short-term volume, 790 thousand out of 2.93 million (27%) were non-occupant buyers (excluding developers). Short-term-non-occupant-buyer transactions increase over 2000–2005 from 90 thousand to 230 thousand, 41% of the overall growth in short-term transactions (370 thousand to 710 thousand, excluding developers). Therefore, non-occupants account for an excess share of the growth in short-term buyers.

In a related approach, we measure the share of 2000–2005 non-occupant purchases that later become short-term sales. These calculations afford a direct comparison to the 2000–2005 increase in non-occupant volume that we analyze in Section 3. However, they are not completely comparable to the ones above, because they look until 2008 to see if a purchase becomes a short-term sale. Of 2000–2005 non-occupant volume, 930 thousand out of 3.60 million (26%) become short-term sellers (excluding developers). Non-occupant purchases that become short-term sales increase over 2000–2005 from 110 thousand to 210 thousand, 23% of the overall growth in non-occupant transactions (440 thousand to 880 thousand, excluding developer buyers). These numbers imply there was not a shift in the composition of non-occupant buyers during the boom toward short-term behavior. However, it is difficult to measure short investment horizons of buyers at the end of the boom because many listings from 2006–2008 did not sell quickly. Another interpretation of these results is that there was secular growth in long-term non-occupants alongside the entry of short-term speculators during the boom.

C.2 Credit utilization

To further investigate the role of credit, we decompose the increase in short-term selling into groups of transactions based on how much leverage the buyer originally used. We focus on a low-leverage group (purchase loan-to-value (LTV) $< 60\%$), a medium-leverage group (purchase LTV $\in [60\%, 85\%)$), and a high-leverage group (purchase LTV $> 85\%$). Of the short-term sellers in 2000–2005, 31% were low-LTV buyers, 33% were medium-LTV buyers, and 36% were high-LTV buyers. In contrast, for the long-term sellers for whom we observe purchase LTVs (i.e., with initial purchase during or after 1995), the distribution skews more toward high-leverage buyers: 22% were low-LTV buyers, 30% were medium-LTV buyers, and 48% were high-LTV buyers. Between 2000 and 2005, low-LTV, medium-LTV, and high-LTV short-term-buyer transactions account for 15%, 58%, and 27% of the growth in short-term transactions, respectively.⁵ As in our analysis of cash transactions among speculators, these statistics reveal that short-term volume is associated with lower use of leverage in the cross-section relative to the general population.⁶ At the same time, the proportional growth in

⁵Of the short-term sellers in 2000–2005 with non-missing LTV, 1.24 million were low-LTV buyers, 1.33 million were medium-LTV buyers, and 1.46 million were high-LTV buyers. Between 2000 and 2005, the number of low-LTV, medium-LTV, and high-LTV short-term-buyer transactions increases from 210 to 270 thousand, from 140 to 380 thousand, and from 190 to 300 thousand, respectively.

⁶In Table IA9 of the Online Appendix, we extend Table 3 to look at average purchase LTVs for short-term and non-occupant buyers. Both speculative buyer types have lower average LTVs, which is exclusively driven by their higher all-cash shares.

short-term buying is stronger among medium- and high-LTV sellers, making a larger relative contribution to the overall growth in short-term volume.

C.3 Buyer scale and experience

Scale. We mark transactions as developer purchases when the buyer’s name is not parsed as a person by CoreLogic and contains strings reflecting developer names. We identify developer names using CoreLogic’s internal new construction flag, as described in Online Appendix A. Both this analysis and the analysis of inexperienced investors below exclude transactions with missing buyer names.

In our sample, these transactions account for 6% of total volume and 9% of the growth in volume between 2000 and 2005. Of the 3.95 million short-term sales in 2000–2005, the initial purchases for 580 thousand (15%) were from developer buyers. From 2000 to 2005, the number of short-term-buyer sales increases from 530 thousand to 930 thousand while the number of short-term-developer-buyer sales increases from 100 thousand to 130 thousand, or 8% of the growth in short-term volume. Though developers were active in the housing market, they did not contribute disproportionately to short-term volume growth in the boom. A possible reason is that developers were more likely to engage in speculation in the raw land market (Nathanson and Zwick, 2018).

Experience. To flag non-developers as experienced or inexperienced, we count the total number of transactions for each unique buyer name in an MSA. We classify buyers with one or two purchases as inexperienced and those with three or more as experienced. Of the 2000–2005 short-term sales, 2.42 million of 3.36 million (72%) were inexperienced buyers at the time of purchase (excluding developers). Thus, inexperienced buyers constitute 2.42 million of 3.95 million total short-term sales, or 61%. Between 2000 and 2005, the number of inexperienced short-term-buyer sales increases from 310 thousand to 560 thousand, or 66% of the growth in short-term sales (excluding developers). The quantitative relevance of inexperienced buyers for volume is consistent with the evidence in Bayer et al. (2020).

The patterns we document are consistent with speculative motives leading short-term buyers to enter and exit the market in response to expected capital gains. But some short-term sellers likely do not exit the market and instead choose to buy another house within the same MSA. Such a pattern may reflect move-up purchases enabled by higher home equity in the boom (Stein, 1995; Ortalo-Magné and Rady, 2006), or repeated buying and selling of homes within the same market by experienced “flippers” (Bayer et al., 2020; Choi et al., 2014).

To explore this alternative explanation, we follow the methodology of Anenberg and Bayer (2013) and construct a direct measure of repeated within-MSA purchases. We use the names of buyers and sellers to match transactions as being possibly linked in a joint buyer-seller event. For each sale transaction, we attempt to identify a purchase transaction in which the seller from the sale matches the buyer from the purchase. To allow the possibility that a purchase occurs before a sale or with a lag, we look for matches in a window of plus or minus one quarter around the quarter of the sale transaction. We only look for within-MSA matches, as purchases associated with cross-city moves are similar in spirit to our model.

Our match accounts for several anomalies that would lead a naive match strategy to understate the match rate.⁷ Our approach is likely to overstate the number of true matches, because it does not use address information to restrict matches, and it allows common names to match even if they represent different people. Because we find a low match rate even with this aggressive strategy, we do not make use of address information in our algorithm or otherwise attempt to refine matches.

We focus on transactions between 2002 and 2011 because the seller name fields are incomplete in prior years for several cities. We also restrict sales transactions to those with human sellers, as indicated by the name being parsed and separated into first and last name fields by CoreLogic. The sample includes 16.3 million sales transactions. Of these, we are able to match 3.9 million to a linked buyer transaction, or 24%. Thus, three-quarters of transactions do not appear to be associated with joint buyer-seller decisions. Among sellers who had bought within the last three years, the match rate is slightly higher, equal to 31%, consistent with move-up purchase or flipper behavior. In addition, the match rates peak in 2005 at 29% and 38% for all transactions and short-term transactions, respectively.⁸ These patterns confirm and extend the findings in Anenberg and Bayer (2013), who conduct a similar match for the Los Angeles metro area and show that internal moves account for a substantial share of the volatility of transaction volume in that city. However, the evidence supports the notion that sellers engaging in repeat purchases do not account for most of the short-term volume and its growth, even during the cycle’s peak.

D Relation of model to literature

As mentioned in the Introduction, existing theoretical papers explain the comovement of prices and volume. However, there are three additional results from our empirical work that no prior model seems able to explain simultaneously.

First, the increase in volume during the boom, and listings during the boom and quiet, come disproportionately from short-term sales (Figures 4 and 6). Search-and-matching models struggle to generate this pattern if the decision to list is independent of homeowner characteristics, as in Wheaton (1990), Piazzesi and Schneider (2009), Díaz and Jerez (2013), Guren (2014), Head et al. (2014), and Anenberg and Bayer (2020).⁹ These models cannot explain the result that homeowners who bought later in the boom were more likely to resell than homeowners who bought earlier. Overconfidence models, such as Daniel et al. (1998,

⁷These include: inconsistent use of nicknames (e.g., Charles versus Charlie), initials in place of first names, the presence or absence of middle initials, transitions from a couples buyer to a single buyer via divorce, transitions from a single buyer to a couples buyer via cohabitation, and reversal of order in couples purchases.

⁸The importance of internal volume varies across cities and years during the boom, with the internal move share of MSA-level short-volume growth ranging from 35% to 46% on average. On average across MSAs, growth in internal short-volume accounts for 35% of the growth in total short volume in 2005, the peak year in total volume.

⁹Two exceptions are Hedlund (2016) and Ngai and Sheedy (2020), who respectively focus on credit constraints and within-market moves. As we explain in Section 4.4 and Online Appendix C.3, short-term volume increases significantly among low-LTV sellers, and most short-term sellers do not relocate within the same MSA. Therefore, these two papers do not explain a substantial share of the disproportionate rise in short-term volume during the boom.

2001), generate speculative trading that accompanies booms and busts in asset prices. In these models, an initial increase in asset prices boosts the confidence of optimistic investors, leading them to push prices up further. However, these models are not designed to fit the rise in short-term volume that occurs during booms, because the same overconfident investors buy the asset in the early as well as the late stages of a boom. Other disagreement and extrapolation–psychology papers can generate a disproportionate short-term volume boom, as long as rising prices generate more disagreement or stronger psychological urges to both buy and sell.

Second, non-occupants constitute a disproportionate share of the increase in buying activity during the boom (Figure 4). Non-occupant purchasing is absent from many search-and-matching models, either because the owner-occupied and rental markets are separate (Guren, 2014), or because all non-occupant owners are previous occupants of the same house (Head et al., 2014; Burnside et al., 2016). The extrapolation–psychology papers also provide no role for non-occupants, as they model more general asset markets where all owners receive the same flow benefits from the asset. Nathanson and Zwick (2018) present a disagreement model in which non-occupants disproportionately buy housing during a boom, but their model is static and is therefore not suited to explain the dynamics at the heart of this paper.

The third result is the existence of the quiet, during which prices and volume diverge while listings accumulate (Figures 1 and 3). Disagreement papers and credit-constraint housing models predict a monotonic relation between prices and volume, and therefore do not explain a period when these outcomes move in opposite directions.¹⁰ Barberis et al. (2018) and Liao and Peng (2018) generate a divergence of prices and volume, but listings fall with volume because of Walrasian market clearing. A similar pattern of prices, volume, and listings appears in Burnside et al. (2016). In contrast, Guren (2014) matches all three variables. However, in that model, listings sharply decline during the boom (more than one-for-one with respect to prices), and they never rise above their pre-shock level in the impulse response. Empirically, listings modestly rise during the boom in the aggregate. The sharp rise in listings during the quiet, far above their 2000 level, is perhaps the most salient aspect of Figure 3.

E Proofs

E.1 Lemma 1

Agents at t believe that they observe $d_{t-k} = \tilde{d}_{t-k}$ for all $k > 0$. Let g_t^* denote the mean of the posterior on g_{t-1} from this information, and σ_t^2 its variance. We solve for these outcomes using standard Kalman filtering. Denote $\sigma_{\epsilon^d}^2 = (1 - \gamma)\sigma_d^2$ and $\sigma_{\epsilon^g}^2 = \gamma(1 - \rho^2)\sigma_d^2$.

We have $g_{t-1} = g_t^* + \zeta_t^g$, where $\zeta_t^g \sim \mathcal{N}(0, \sigma_t^2)$. Therefore, $g_t = (1 - \rho)\mu_g + \rho g_{t-1} + \epsilon_t^g =$

¹⁰An exogenous increase in overconfidence raises volume in Daniel et al. (2001) and Scheinkman and Xiong (2003); it raises conditional return volatility in Daniel et al. (2001) while raising the price level in Scheinkman and Xiong (2003). Disagreement accounts for some of the average prices and volume in the housing market (Bailey et al., 2016) and can generate dispersion in beliefs about house price growth over the period we are studying (Piazzesi and Schneider, 2009; Burnside et al., 2016). By definition, disagreement is less suited to explain the high average level of these beliefs (Case et al., 2012; Foote et al., 2012; Cheng et al., 2014).

$(1 - \rho)\mu_g + \rho g_t^* + \rho \zeta_t^g + \epsilon_t^g$. The prior on g_t at $t + 1$ is thus $\mathcal{N}((1 - \rho)\mu_g + \rho g_t^*, \rho^2 \sigma_l^2 + \sigma_{\epsilon_g}^2)$. The information is $\Delta \tilde{d}_t$, which according to agents equals $g_t + \epsilon_t^d$. Therefore, the new posterior variance satisfies $\sigma_l^2 = \sigma_{\epsilon_d}^2 (\rho^2 \sigma_l^2 + \sigma_{\epsilon_g}^2) (\sigma_{\epsilon_d}^2 + \rho^2 \sigma_l^2 + \sigma_{\epsilon_g}^2)^{-1}$. Solving yields

$$\sigma_l^2 = (2\rho^2)^{-1} \left(-(1 - \rho^2)\sigma_{\epsilon_d}^2 - \sigma_{\epsilon_g}^2 + \sqrt{((1 - \rho^2)\sigma_{\epsilon_d}^2 + \sigma_{\epsilon_g}^2)^2 + 4\rho^2 \sigma_{\epsilon_d}^2 \sigma_{\epsilon_g}^2} \right).$$

The new posterior mean satisfies $g_{t+1}^* = (1 - \alpha)\Delta \tilde{d}_t + \alpha((1 - \rho)\mu_g + \rho g_t^*)$, where $\alpha = \sigma_{\epsilon_d}^2 / (\sigma_{\epsilon_d}^2 + \rho^2 \sigma_l^2 + \sigma_{\epsilon_g}^2)$. Iterating (and then subtracting one from the time subscripts everywhere) gives

$$g_t^* = \mu_g + (1 - \alpha) \sum_{k=1}^{\infty} (\alpha\rho)^{k-1} \left(\Delta \tilde{d}_{t-k} - \mu_g \right).$$

Because $\hat{g}_t = (1 - \rho)\mu_g + \rho g_t^*$, we have proved the formula in the lemma for \hat{g}_t . We showed above that $\hat{\sigma}_g^2 = \rho^2 \sigma_l^2 + \sigma_{\epsilon_g}^2$. We have $d_t = d_{t-1} + g_t + \epsilon_t^d = (d_{t-1} - \tilde{d}_{t-1}) + \tilde{d}_{t-1} + (1 - \rho)\mu_g + \rho g_{t-1} + \epsilon_t^g + \epsilon_t^d = (d_{t-1} - \tilde{d}_{t-1}) + \tilde{d}_{t-1} + \hat{g}_t + \rho \zeta_t^g + \epsilon_t^g + \epsilon_t^d$, which immediately gives $\hat{d}_t = \tilde{d}_{t-1} + \hat{g}_t$, with $\hat{\sigma}_d^2 = \rho^2 \sigma_l^2 + \sigma_{\epsilon_g}^2 + \sigma_{\epsilon_d}^2$.

The bound we assume for r (see Section 5.1) is

$$r > e^{\mu_g + \frac{(1-\alpha\rho)^2 \hat{\sigma}_d^2}{2(1-\rho)^2}} - 1. \quad (\text{E1})$$

E.2 Lemma 2

By Lemma 1, $\Delta \tilde{d}_t = \hat{g}_t + (\tilde{d}_t - \hat{d}_t)$. Furthermore, $\hat{g}_{t+1} = \mu_g + (\alpha\rho)(\hat{g}_t - \mu_g) + (1 - \alpha)\rho(\Delta \tilde{d}_t - \mu_g) = (1 - \rho)\mu_g + \rho\hat{g}_t + (1 - \alpha)\rho(\tilde{d}_t - \hat{d}_t)$. Finally, $\hat{d}_{t+1} = \tilde{d}_t + \hat{g}_{t+1} = \hat{d}_t + (1 - \rho)\mu_g + \rho\hat{g}_t + (1 - \alpha)\rho(\tilde{d}_t - \hat{d}_t)$. From the point of view of agents, $\tilde{d}_t = d_t$. Therefore,

$$\hat{d}_{t+1} = \hat{d}_t + (1 - \rho)\mu_g + \rho\hat{g}_t + (1 + (1 - \alpha)\rho)\zeta_t \quad (\text{E2})$$

$$\hat{g}_{t+1} = (1 - \rho)\mu_g + \rho\hat{g}_t + (1 - \alpha)\rho\zeta_t, \quad (\text{E3})$$

where $\zeta_t \equiv d_t - \hat{d}_t$.

Write $V^m(\hat{d}_t, \hat{g}_t) = e^{\hat{d}_t} v^m(\hat{d}_t, \hat{g}_t)$ and $P = e^{\hat{d}_t} p$. Then $\tilde{\pi}(P, d_t) = 1 - F(\log p + \log \bar{\kappa} - \zeta_t)$, which we denote $\tilde{\pi}(p, \zeta_t)$ by abuse of notation. Substituting these expressions into (6) and using (E2) yield

$$v^m(\hat{d}_t, \hat{g}_t) = \sup_p E \left(\tilde{\pi}(p, \zeta_t) p + \frac{(1 - \tilde{\pi}(p, \zeta_t)) e^{(1-\rho)\mu_g + \rho\hat{g}_t + (1+\rho-\alpha\rho)\zeta_t} v^m(\hat{d}_{t+1}, \hat{g}_{t+1})}{1 + r_m} \right), \quad (\text{E4})$$

with the expectation over $\zeta_t \sim \mathcal{N}(0, \hat{\sigma}_d^2)$ and \hat{g}_{t+1} given by (E3). Because \hat{d}_t and \hat{d}_{t+1} appear only in the first argument of v^m , this function does not depend on \hat{d}_t , so

$$V^m(\hat{d}_t, \hat{g}_t) = e^{\hat{d}_t} v^m(\hat{g}_t). \quad (\text{E5})$$

It follows that the argmax of (E4) does not depend on \hat{d}_t . We denote it $p(\hat{g}_t)$.

E.3 Lemma 3

When $r_m \rightarrow \infty$, $p(\cdot)$ becomes constant, as is clear from (E4). In this case, the formula for \hat{d}_{t+1} at the beginning of the proof of Lemma 2 implies that $\Delta \log P_{t+1} = (1 - \rho)\mu_g + \rho\hat{g}_t + (1 + (1 - \alpha)\rho)(\tilde{d}_t - \hat{d}_t)$. Solving for $\tilde{d}_t - \hat{d}_t$ and substituting it into the formula for \hat{g}_{t+1} there yields $\hat{g}_{t+1} = (1 + (1 - \alpha)\rho)^{-1}((1 - \rho)\mu_g + \rho\hat{g}_t + (1 - \alpha)\rho\Delta \log P_{t+1})$. Iterating this formula backwards (and then subtracting 1 from the time subscripts) gives

$$\hat{g}_t = \mu_g + (1 - \alpha) \sum_{k=1}^{\infty} \left(\frac{\rho}{1 + (1 - \alpha)\rho} \right)^k (\Delta \log P_{t-k} - \mu_g).$$

Conditional on market data before t , agents at t believe that $E(\tilde{d}_t - \hat{d}_t) = 0$. Therefore, $E\Delta \log P_{t+1} = (1 - \rho)\mu_g + \rho\hat{g}_t$. Substituting in the expression just derived for \hat{g}_t gives the first equation in the lemma.

To derive the second equation, we let \bar{p} denote the constant value of $p(\cdot)$ that holds in the limit as $r_m \rightarrow \infty$. From (5), $\tilde{d}_t = \hat{d}_t + \log(\bar{\kappa}\bar{p}) - F^{-1}(1 - \pi_t)$. Therefore, the equation above for $\Delta \log P_{t+1}$ implies that $\Delta \log P_{t+1} = E\Delta \log P_{t+1} + (1 + (1 - \alpha)\rho)(\log(\bar{\kappa}\bar{p}) - F^{-1}(1 - \pi_t))$, as claimed.

E.4 Lemma 4

A potential buyer at t observes the history of price changes, P_t/P_{t-1} , but not past price levels. Therefore, her information set is different than the one in the statement of Lemma 1. Nonetheless, she still computes \hat{g}_t using the formula in Lemma 1, as that formula depends only on past price changes and not past price levels. However, the formula for \hat{d}_t does not work because it requires knowledge of P_{t-1} . Therefore, she imputes \hat{d}_t using her knowledge of equilibrium and the list price she observes. In particular, given Lemma 2, $P = e^{\hat{d}_t}p(\hat{g}_t)$, which implies that $\hat{d}_t = \log(P/p(\hat{g}_t))$. The potential buyer's decision rule is therefore

$$V^b \left(\log \left(\frac{P}{p(\hat{g}_t)} \right), \hat{g}_t; \lambda, \delta, n \right) \geq P. \quad (\text{E6})$$

The proof proceeds by showing that this inequality is equivalent to the one in Lemma 4 through suitable choice of $\kappa_{n,j}(\hat{g}_t)$.

Write $V^s(\hat{d}_t, \hat{g}_t; \lambda, \delta) = (r + \lambda)^{-1}e^\delta + e^{\hat{d}_t}v^s(\hat{d}_t, \hat{g}_t; \lambda, \delta)$. Substituting this equation, (E2), and (E5) into (8) yields

$$v^s(\hat{d}_t, \hat{g}_t; \lambda, \delta) = (1 + r)^{-1}E \left(e^{(1-\rho)\mu_g + \rho\hat{g}_t + (1+\rho-\alpha\rho)\zeta_t} (\lambda v^m(\hat{g}_{t+1}) + (1 - \lambda)v^s(\hat{d}_{t+1}, \hat{g}_{t+1}; \lambda, \delta)) \right), \quad (\text{E7})$$

with the expectation over $\zeta_t \sim \mathcal{N}(0, \hat{\sigma}_d^2)$ and \hat{g}_{t+1} given by (E3). Because \hat{d}_t , \hat{d}_{t+1} , and δ appear only in the arguments of v^s , that function does not depend on \hat{d}_t and δ , allowing us to write $V^s(\hat{d}_t, \hat{g}_t; \lambda, \delta) = (r + \lambda)^{-1}e^\delta + e^{\hat{d}_t}v^s(\hat{g}_t; \lambda)$. Substituting this equation, (E2), and

(E5) into (7) yields

$$V^b(\hat{d}_t, \hat{g}_t; \lambda, \delta, n) = \frac{e^\delta}{r + \lambda} + \frac{e^{\hat{d}_t}}{1 + r} E \left(e^{(1-\rho)\mu_g + \rho\hat{g}_t + (1+\rho-\alpha\rho)\zeta_t} (\lambda v^m(\hat{g}_{t+1}) + (1 - \lambda)v^s(\hat{g}_{t+1}; \lambda)) \right),$$

with the expectation over $\zeta_t \sim \mathcal{N} \left(\frac{\hat{\sigma}_d^2(\delta - \hat{d}_t - \mu_n)}{\sigma_a^2 + \hat{\sigma}_d^2}, \frac{\sigma_a^2 \hat{\sigma}_d^2}{\sigma_a^2 + \hat{\sigma}_d^2} \right)$ and \hat{g}_{t+1} given by (E3). Let $\Psi(\zeta_t, \hat{g}_t; \lambda)$ denote the argument inside the expectation. We can simplify the buying decision, (E6), to

$$\frac{e^\delta}{P} \geq (r + \lambda) \left(1 - \frac{E\Psi(\zeta_t, \hat{g}_t; \lambda)}{(1 + r)p(\hat{g}_t)} \right), \quad (\text{E8})$$

with the expectation over $\zeta_t \sim \mathcal{N} \left(\frac{\hat{\sigma}_d^2(\log(e^\delta/P) + \log p(\hat{g}_t) - \mu_n)}{\sigma_a^2 + \hat{\sigma}_d^2}, \frac{\sigma_a^2 \hat{\sigma}_d^2}{\sigma_a^2 + \hat{\sigma}_d^2} \right)$.

To proceed, we use the following lemma about $v^m(\hat{g}_t)$ and $v^s(\hat{g}_t; \lambda)$:

Lemma IA1. *For all $\lambda > 0$, $v^m(\hat{g}_t)$ and $v^s(\hat{g}_t, \lambda)$ are continuous and weakly increasing functions of \hat{g}_t .*

Proof. Appendix E.5. □

From IA1, it follows immediately that $\Psi(\zeta_t, \hat{g}_t; \lambda)$ is a continuous and weakly increasing function of ζ_t for any \hat{g}_t and $\lambda > 0$, which implies that the right side of (E8) continuously and weakly decreases in e^δ/P . The left side continuously and strictly increases in e^δ/P . Therefore, for n, j , and \hat{g}_t such that the right side does not limit to a positive number as $\delta \rightarrow \infty$ for $\lambda = \lambda_j$, then (E8) holds for all δ , meaning that Lemma 4 holds with $\kappa_{n,j}(\hat{g}_t) = 0$. If the right side limits to a positive number as $\delta \rightarrow -\infty$ when $\lambda = \lambda_j$, then by the Intermediate Value Theorem, there exists a unique $\kappa_{n,j}(\hat{g}_t)$ such that the inequality holds if and only if $e^\delta/P \geq \kappa_{n,j}(\hat{g}_t)$, which proves Lemma 4.

E.5 Value Function Monotonicity

This section establishes that the functions $v^m(\cdot)$ and $v^s(\cdot; \lambda)$, which we define in the proofs of Lemmas 2 and 4, weakly and continuously increase. We follow Stokey et al. (1989). To apply their results, we need to work with a one-point (Alexandroff) compactification of a subset of the real numbers. For a topological set X , the Alexandroff compactification is the set $X^* = X \cup \{\infty\}$, whose open sets are those of X together with sets whose complements are closed, compact subsets of X ; X^* is compact (Kelley, 1955).

Lemma IA2. *Let $f : (0, \infty) \times \mathbb{R} \rightarrow \mathbb{R}$ be continuous. Suppose there exists functions $g_0 : \mathbb{R} \rightarrow \mathbb{R}$ and $g_\infty : \mathbb{R} \rightarrow \mathbb{R}$ such that $\lim_{x \rightarrow 0} f(x, y) = g_0(y)$ and $\lim_{x \rightarrow \infty} f(x, y) = g_\infty(y)$ uniformly. Define $\tilde{f} : [0, \infty]^* \times \mathbb{R} \rightarrow \mathbb{R}$ by $\tilde{f}(x, y) = f(x, y)$ for $x \in (0, \infty)$ and $\tilde{f}(x, y) = g_x(y)$ for $x \in \{0, \infty\}$. Then \tilde{f} is continuous.*

Proof. Let $Z \subset \mathbb{R}$ be open. We show that $\tilde{f}^{-1}(Z)$ is open by demonstrating that for each $(x, y) \in \tilde{f}^{-1}(Z)$, there exists an open set U such that $(x, y) \in U \subset \tilde{f}^{-1}(Z)$. If $x \in (0, \infty)$, then set $U = f^{-1}(Z)$, which is open by the continuity of f . Consider the case $x = 0$. Because Z is open, there exists $\epsilon > 0$ such that all z with $|z - g_0(y)| < \epsilon$ are in Z . By

uniform convergence, there exists $\delta > 0$ such that $|f(x', y') - g_0(y)| < \epsilon$ for all $x \in [0, \delta]$ and $y \in \mathbb{R}$. Therefore, $U = [0, \delta) \times \mathbb{R}$ suffices. Consider the case $x = \infty$. There likewise exists $\epsilon > 0$ such that all z with $|z - g_\infty(y)| < \epsilon$ are in Z . By uniform convergence, there exists $N > 0$ such that $|f(x', y') - g_\infty(y)| < \epsilon$ for all $x > N$ and $y \in \mathbb{R}$. Therefore, $U = (N, \infty) \times \mathbb{R}$ suffices. \square

We next establish the existence of a continuous solution $v^m(\cdot)$ to (E4). Let \mathcal{C} be the space of bounded continuous functions from \mathbb{R} to itself. Let $a > 0$ be a constant. For $v \in \mathcal{C}$, we define the operator T by $(Tv)(\hat{g}) = \sup_p f(p, \hat{g})$, where

$$f(p, \hat{g}) = \int_{-\infty}^{\infty} \left(\frac{\tilde{\pi}(p, \zeta)p}{a + e^{\frac{\rho\hat{g}}{1-\rho}}} + \frac{(1 - \tilde{\pi}(p, \zeta))e^{(1-\rho)\mu_g + \rho\hat{g} + (1+\rho-\alpha\rho)\zeta}}{1 + r_m} \times \right. \\ \left. \frac{\left(a + e^{\rho\mu_g + \frac{\rho^2\hat{g}}{1-\rho} + \frac{\rho^2(1-\alpha)\zeta}{1-\rho}} \right) v((1-\rho)\mu_g + \rho\hat{g} + \rho(1-\alpha)\zeta)}{a + e^{\frac{\rho\hat{g}}{1-\rho}}} \right) \phi(\zeta) d\zeta,$$

where $\phi(\cdot)$ is the probability density function of $\mathcal{N}(0, \hat{\sigma}_d^2)$. If v is a fixed point of T , then $v^m(\hat{g}) = (a + e^{\frac{\rho\hat{g}}{1-\rho}})v(\hat{g})$ solves (E4). We find a fixed point by demonstrating that $T : \mathcal{C} \rightarrow \mathcal{C}$ and then showing that for a sufficiently small value of a , T satisfies the Blackwell conditions and is hence a contraction mapping.

We first show that $Tv \in \mathcal{C}$. We have the bound

$$\|Tv\| \leq \sup_p \int_{-\infty}^{\infty} a^{-1} \tilde{\pi}(p, \zeta) p \phi(\zeta) d\zeta + \\ (1 + r_m)^{-1} e^{(1-\rho)\mu_g} \|v\| \sup_x \frac{ae^{\rho x + \frac{(1+\rho-\alpha\rho)^2 \hat{\sigma}_d^2}{2}} + e^{\rho\mu_g + \frac{\rho x}{1-\rho} + \frac{(1-\alpha\rho)^2 \hat{\sigma}_d^2}{2(1-\rho)^2}}}{a + e^{\frac{\rho x}{1-\rho}}},$$

so Tv is bounded.

Demonstrating continuity is much more complicated. We first apply Lemma 12.14 of Stokey et al. (1989) to establish the continuity of $f(\cdot, \cdot)$.

In their terminology, $X = (0, \infty)$, $Z = \mathbb{R}^2$, their y corresponds to our p , their z corresponds to our (\hat{g}, ζ) , and the transition function Q puts mass $\phi(\zeta')$ on (\hat{g}, ζ') and mass 0 on other elements of Z . To apply their lemma, we must show that Q has the Feller property, which means (see their page 375) that $\int h(z')Q(z, z')dz'$ is continuous in z as long as h is continuous and bounded.¹¹ Given our specification of Q , this integral reduces to $\int_{-\infty}^{\infty} h(\hat{g}, \zeta')\phi(\zeta')d\zeta'$, which is trivially continuous in ζ . To demonstrate continuity in \hat{g} , we closely follow the proof of their Lemma 9.5. Choose a sequence \hat{g}_n converging to \hat{g} . Then $|\int_{-\infty}^{\infty} h(\hat{g}_n, \zeta')\phi(\zeta')d\zeta' - \int_{-\infty}^{\infty} h(\hat{g}, \zeta')\phi(\zeta')d\zeta'| \leq \int_{-\infty}^{\infty} |h(\hat{g}_n, \zeta') - h(\hat{g}, \zeta')|\phi(\zeta')d\zeta'$. Each function $\zeta' \mapsto |h(\hat{g}_n, \zeta') - h(\hat{g}, \zeta')|$ converges pointwise to the zero function (by the continuity of h), so by the Lebesgue Dominated Convergence Theorem (their Theorem 7.10), this integral

¹¹Their lemma also requires that the term inside the integral defining $f(\cdot, \cdot)$, other than $\phi(\zeta)d\zeta$, is bounded in p , \hat{g} , and ζ . This boundedness holds because v is bounded, because $\lim_{p \rightarrow \infty} \tilde{p}(\zeta, p)p = 0$, and because $\lim_{\zeta \rightarrow \infty} (1 - \tilde{\pi}(p, \zeta))e^{c\zeta} = 0$ for any $c > 0$.

limits to zero. Therefore, $\hat{g} \mapsto \int_{-\infty}^{\infty} h(\hat{g}, \zeta') \phi(\zeta') d\zeta'$ is continuous in \hat{g} , and Q has the Feller property. As a result, $f(\cdot, \cdot)$ is continuous on $(0, \infty) \times \mathbb{R}$.

The next step is to invoke our Lemma IA2. To do so, we must show uniform convergence of $f(p, \hat{g})$ for $p \rightarrow 0$ and $p \rightarrow \infty$. In the first limit, $f(p, \hat{g}) \rightarrow 0$, and this convergence is uniform because terms with \hat{g} multiplying the terms with p are uniformly bounded in \hat{g} . In the second limit, the convergence is to the integral in which $\tilde{\pi} = 0$, and the convergence is uniform for the same reason. Hence, Lemma IA2 applies, and the induced \tilde{f} is continuous.

The final step is to show that $(Tv)(\hat{g})$ is continuous. This statement follows immediately from Berge's Maximum Theorem on general topological spaces (see, for instance, page 570 of Aliprantis and Border (2006)) because $\sup_{p \in (0, \infty)} f(p, \hat{g}) = \sup_{p \in [0, \infty)^*} \tilde{f}(p, \hat{g})$ and because $[0, \infty)^*$ is compact. Therefore, $Tv \in \mathcal{C}$.

We next verify the Blackwell conditions for T (Theorem 3.3 in Stokey et al. (1989)). Monotonicity is trivial. Given the bound above, discounting holds as long as

$$(1 + r_m)^{-1} e^{(1-\rho)\mu_g} \sup_x \frac{a e^{\rho x + \frac{(1+\rho-\alpha\rho)^2 \hat{\sigma}_d^2}{2}} + e^{\rho\mu_g + \frac{\rho x}{1-\rho} + \frac{(1-\alpha\rho)^2 \hat{\sigma}_d^2}{2(1-\rho)^2}}}{a + e^{\frac{\rho x}{1-\rho}}} < 1.$$

We are free to choose any positive value of a . By considering the limit as $a \rightarrow 0$, we find that we can choose such an a to satisfy this inequality as long as

$$(1 + r_m)^{-1} e^{\mu_g + \frac{(1-\alpha\rho)^2 \hat{\sigma}_d^2}{2(1-\rho)^2}} < 1.$$

This inequality holds because $r_m \geq r$ and we assume that (E1) holds. Therefore, by Theorem 3.3 of Stokey et al. (1989), T is a contraction mapping. By the Contraction Mapping Theorem (their Theorems 3.1 and 3.2), T has a unique fixed point in \mathcal{C} , as desired. Call this function v^* . As mentioned above, $v^m(\hat{g}) = v^*(\hat{g})(a + e^{\frac{\rho\hat{g}}{1-\rho}})$ then solves (E4); this function clearly inherits the continuity of v^* .

Finally, we show that v^m is weakly increasing. Let $\mathcal{C}' \subset \mathcal{C}$ be the set of v such that $v(\hat{g})(a + e^{\frac{\rho\hat{g}}{1-\rho}})$ weakly increases. We claim that \mathcal{C}' is closed. Let $\{v_n\}$ be in \mathcal{C}' converging in \mathcal{C} to v . For any $\hat{g}_0 < \hat{g}_1$, $v_n(\hat{g}_1)(a + e^{\frac{\rho\hat{g}_1}{1-\rho}}) - v_n(\hat{g}_0)(a + e^{\frac{\rho\hat{g}_0}{1-\rho}}) \geq 0$. Because v_n converges pointwise to v , we must have $v(\hat{g}_1)(a + e^{\frac{\rho\hat{g}_1}{1-\rho}}) - v(\hat{g}_0)(a + e^{\frac{\rho\hat{g}_0}{1-\rho}}) \geq 0$ as well. Therefore, Corollary 1 to Theorem 3.2 of Stokey et al. (1989) shows that $v^m \in \mathcal{C}'$ as long as $T : \mathcal{C}' \rightarrow \mathcal{C}'$, which is immediate from (E4).

The task remaining for this appendix is to show that each $v^s(\cdot; \lambda)$ weakly and continuously increases. The argument proceeds as with $v^m(\cdot)$, but we use (E7), and we can skip the steps involving a supremum. Define the map T on \mathcal{C} by

$$(Tv)(\hat{g}) = (1 + r)^{-1} \int_{-\infty}^{\infty} \left(\frac{a e^{(1-\rho)\mu_g + \rho\hat{g} + (1+\rho-\alpha\rho)\zeta}}{a + e^{\frac{\rho\hat{g}}{1-\rho}}} + \frac{e^{\mu_g + \frac{\rho\hat{g}}{1-\rho} + \frac{(1-\alpha\rho)\zeta}{1-\rho}}}{a + e^{\frac{\rho\hat{g}}{1-\rho}}} \right) ((1 - \lambda)v(g') + \lambda v^*(g')) \phi(\zeta) d\zeta,$$

where $g' = (1 - \rho)\mu_g + \rho\hat{g} + \rho(1 - \alpha)\zeta$, and $a > 0$ is a constant to be specified later. If v is a fixed point of T , then $v^s(\hat{g}; \lambda) = (a + e^{\frac{\rho\hat{g}}{1-\rho}})v(\hat{g})$ solves (E7). Clearly, Tv is bounded. To prove continuity, we again apply Lemma 12.14 of Stokey et al. (1989), this time with $X = Z = \mathbb{R}$, our \hat{g} corresponding to their y , and our ζ corresponding to their z . In order to apply their lemma, we have to absorb the ζ terms into the Q transition function so that their f is bounded. Using the identity $e^{-z^2/(2\sigma^2)+bz} = e^{\sigma^2 b^2/2} e^{-(z-\sigma^2 b)^2/(2\sigma^2)}$, we have

$$e^{(1+\rho-\alpha\rho)\zeta}\phi(\zeta) = e^{\hat{\sigma}_d^2(1+\rho-\alpha\rho)^2/2}\phi(\zeta - \hat{\sigma}_d^2(1 + \rho - \alpha\rho))$$

and

$$e^{\frac{(1-\alpha\rho)\zeta}{1-\rho}}\phi(\zeta) = e^{\frac{\hat{\sigma}_d^2(1-\alpha\rho)^2}{2(1-\rho)^2}}\phi\left(\zeta - \frac{\hat{\sigma}_d^2(1-\alpha\rho)}{1-\rho}\right).$$

These functions serve as constants times a valid transition function (we showed above that the normal distribution with 0 mean and variance $\hat{\sigma}_d^2$ has the Feller property), and the remainder of the integrand is bounded in both \hat{g} and ζ . Thus, Lemma 12.14 applies and Tv is continuous. As a result, $T : \mathcal{C} \rightarrow \mathcal{C}$.

Next we verify the aforementioned Blackwell conditions for T . Monotonicity again is trivial. Discounting holds if

$$\frac{1-\lambda}{1+r} \sup_{\hat{g}} \frac{ae^{(1-\rho)\mu_g + \rho\hat{g} + \frac{(1+\rho-\alpha\rho)^2\hat{\sigma}_d^2}{2}} + e^{\mu_g + \frac{\rho\hat{g}}{1-\rho} + \frac{(1-\alpha\rho)^2\hat{\sigma}_d^2}{2(1-\rho)^2}}}{a + e^{\frac{\rho\hat{g}}{1-\rho}}} < 1.$$

Because we are free to pick any $a > 0$, the inequality holds for some such a if

$$(1-\lambda)e^{\mu_g + \frac{(1-\alpha\rho)^2\hat{\sigma}_d^2}{2(1-\rho)^2}} < 1+r,$$

which always holds because $\lambda \in [0, 1]$ and we assume that (E1) holds. Therefore, T satisfies the Blackwell conditions and is a contraction mapping. As a result, it has a unique fixed point in \mathcal{C} . Call it v^{**} . Then $v^s(\hat{g}; \lambda) = (a + e^{\frac{\rho\hat{g}}{1-\rho}})v^{**}(\hat{g})$ solves (E7).

Finally, we show that $v^s(\cdot; \lambda)$ weakly and continuously increases. Continuity follows from the continuity of v^{**} . As argued above, weak monotonicity holds as long as $T : \mathcal{C}' \rightarrow \mathcal{C}'$, where this set is defined as above. That T maps \mathcal{C}' into itself is immediate from (E7) and the fact that v^m weakly increases. QED

F Details on Counterfactuals

F.1 Walrasian extension

In the Walrasian version of our model, a mechanism selects a price each period so that the number of potential buyers willing to buy at that price equals the number of movers willing to sell. The main model assumes that each mover matches to a potential buyer with probability one, which implicitly assumes that the potential buyer population moves in proportion to the mover population. To maintain comparability with the main model, we make an analogous

assumption in the Walrasian variant that the number of potential buyers at time t is NI_t , where $N > 1$ is a constant.

Here, we describe equilibrium in which all movers sell. In this case, (3) implies:

$$I_t = NI_t(1 - F(\log \bar{\kappa} + \log P_t - d_t)).$$

Solving for P_t yields what agents believe is the equilibrium pricing function:

$$\tilde{P}(d_t) = \kappa^{-1} e^{F^{-1}(1-N^{-1})} e^{d_t} = \tilde{p} e^{d_t}.$$

In equilibrium, movers must weakly prefer selling at this price versus waiting to sell next period. Therefore, we must have $e^{d_t} \geq (1 + r_m)^{-1} E_t e^{d_{t+1}}$, where E_t denotes the mover expectation that we now specify. By observing the current and past prices, movers believe that they observe the history of demand as $\tilde{d}_{t-j} = \log(\tilde{p}^{-1} P_{t-j})$ for $j \geq 0$. By a Kalman filtering argument similar to the proof of Lemma 1, the mover posterior on g_t at t has mean

$$\hat{g}_t^m = \mu_g + (1 - \alpha) \sum_{j=0}^{\infty} (\alpha \rho)^j (\Delta \tilde{d}_{t-j} - \mu_g) = \mu_g + (1 - \alpha) \sum_{j=0}^{\infty} (\alpha \rho)^j (\Delta \log P_{t-j} - \mu_g)$$

and variance σ_l^2 . We have $d_{t+1} = d_t + g_{t+1} + \epsilon_{t+1}^d = d_t + (1 - \rho)\mu_g + \rho g_t + \epsilon_{t+1}^g + \epsilon_{t+1}^d = d_t + (1 - \rho)\mu_g + \rho \hat{g}_t^m + \rho \zeta_t^g + \epsilon_{t+1}^g + \epsilon_{t+1}^d$. Therefore,

$$E_t e^{d_{t+1}} = e^{d_t} e^{(1-\rho)\mu_g + \rho \hat{g}_t^m} e^{(\rho^2 \sigma_l^2 + \sigma^2)/2}.$$

Mover optimality therefore requires that

$$\hat{g}_t^m \leq \rho^{-1} (\log(1 + r_m) - (1 - \rho)\mu_g - (\rho^2 \sigma_l^2 + \sigma^2)/2).$$

This inequality cannot hold at all times because \hat{g}_t^m is unbounded. Therefore, when the expected growth rate is sufficiently high, some movers will refrain from selling their homes at the Walrasian equilibrium price. However, we check that the inequality holds for all \hat{g}_t^m in the discrete mesh and also for all realized values in the simulations. For our parameters, the right side equals 0.15, which is much larger than the maximal realized value of 0.03. Therefore, in our simulations, we assume the approximation that the equilibrium always features full sale by all movers at all times.

We now solve for the optimal potential buyer decision, which determines the true pricing function. For $j \geq 1$, potential buyers set $\Delta \tilde{d}_{t-j} = \Delta \log P_{t-j}$. They face the same filtering problem on g_t as potential buyers in the main model, so their posterior mean \hat{g}_t follows the formula in Lemma 1. Because they sell immediately in the approximate equilibrium we consider, the mover value is just the price, $V^m = \tilde{p} e^{\tilde{d}_t}$. (In fact, even in the exact equilibrium, the mover value coincides with the price because movers are indifferent between selling and not.) The remainder of the derivation follows the proof of Lemma 4 closely, so we omit it. That is, there exist functions $\kappa_j(\hat{g}_t)$ such that a potential buyer purchases a house if and only if $e^\delta \geq \kappa_j(\hat{g}_t) P_t$. The functions no longer depend on n because the private flow utility δ is uninformative about d_t , as potential buyers believe that they observe d_t perfectly via

$\tilde{d}_t = \log(\tilde{p}^{-1}P_t)$. The actual equilibrium price must satisfy

$$I_t = NI_t \left(1 - \sum_{j=1}^J (\beta_{0,j} + \beta_{1,j}) F(\log \kappa_j(\hat{g}_t) + \log P_t - d_t) \right),$$

for which it is clear that a unique solution always exists of the form $P_t = p(\hat{g}_t)e^{d_t}$. We discretize the \hat{g}_t space and solve for the pricing function $p(\cdot)$ and the $\kappa_j(\cdot)$ functions at these values, interpolating/extrapolating in between and beyond the mesh.

To maintain comparability with the main model, we decrease γ to 0.042 so that the price overshoot is the same in the Walrasian model as in the main model, and we update $\bar{\kappa}$ so that the demand error is still zero on average. Under the baseline parameters, the price paths in the Walrasian model seem to be explosive. We believe that prices explode because they adjust more quickly with Walrasian market clearing. Choosing a lower γ leads to more stable price paths as in the baseline model. Other parameters remain the same.

F.2 Comparing Short-term and Non-Occupant Buyers

To study the role of short-term buyers, we re-run the simulations setting $\beta_{n,j} = 0$ for all values of j except that for which $\lambda_j = 0.03$. Unlike the counterfactual in Section 6.5.3, we keep a positive mass of non-occupant potential buyers, and we do so in two ways. In the first, the share of non-occupants among potential buyers with $\lambda = 0.03$ equals its baseline. In the second, we change this ratio to the non-occupant share in the whole baseline population. The second version controls for the non-occupant share as we alter the λ distribution.

We perform a similar pair of counterfactual exercises to measure the effect of removing non-occupant buyers. The first counterfactual sets the non-occupant shares, $\beta_{0,j}$, to zero, and then scales up the occupant shares, $\beta_{1,j}$, so that they sum to one. This method skews the λ distribution toward long-term potential buyers because occupants have longer horizons than non-occupants. Therefore, we explore a second counterfactual in which we maintain the original λ distribution while eliminating non-occupants. We continue to set each $\beta_{0,j}$ to zero, but now we update $\beta_{1,j}$ to the baseline share of all potential buyers for whom $\lambda = \lambda_j$.

Table IA10 reports key outcomes from the impulse responses under the baseline and each of these four counterfactuals. In the counterfactuals with only long-term buyers, the price boom falls to 8.7% from its baseline of 14.5%, meaning that short-term buyers amplify the price boom by 67%. Furthermore, in the counterfactuals, the price bust nearly disappears, the volume boom is half its baseline size, and sale probabilities rise less. Inventories fall more during the boom and attain a smaller level at the end of the quiet.¹² Therefore, eliminating short-term buyers prevents the model from matching key aggregate facts (Figures 1 and 3).¹³

We obtain similar results in the first counterfactual with only occupants: the price bust,

¹²These counterfactuals do a better job matching inventory levels during the bust, which reach 1.6% above the initial level, a higher peak than the baseline. In the baseline, new listings fall sharply during the bust because short-term buyers exit the market (Panel F of Figure 11). Thus, the baseline does a better job matching listing behavior in the boom and quiet than in the bust.

¹³The occupant adjustment does not affect the cycle because agents in the model correctly understand the distribution of housing utility, meaning that changing the housing utility distribution does not destabilize prices.

volume boom, rise in sale probabilities, and end-of-quiet listings become significantly smaller. However, when we adjust the λ distribution in the last counterfactual, eliminating non-occupants fails to attenuate the cycle. In fact, the cycle outcomes grow in this scenario. Evidently, non-occupants amplify the housing cycle, but only because many of them have short horizons. Long-term non-occupants fail to amplify the cycle and may even dampen it.

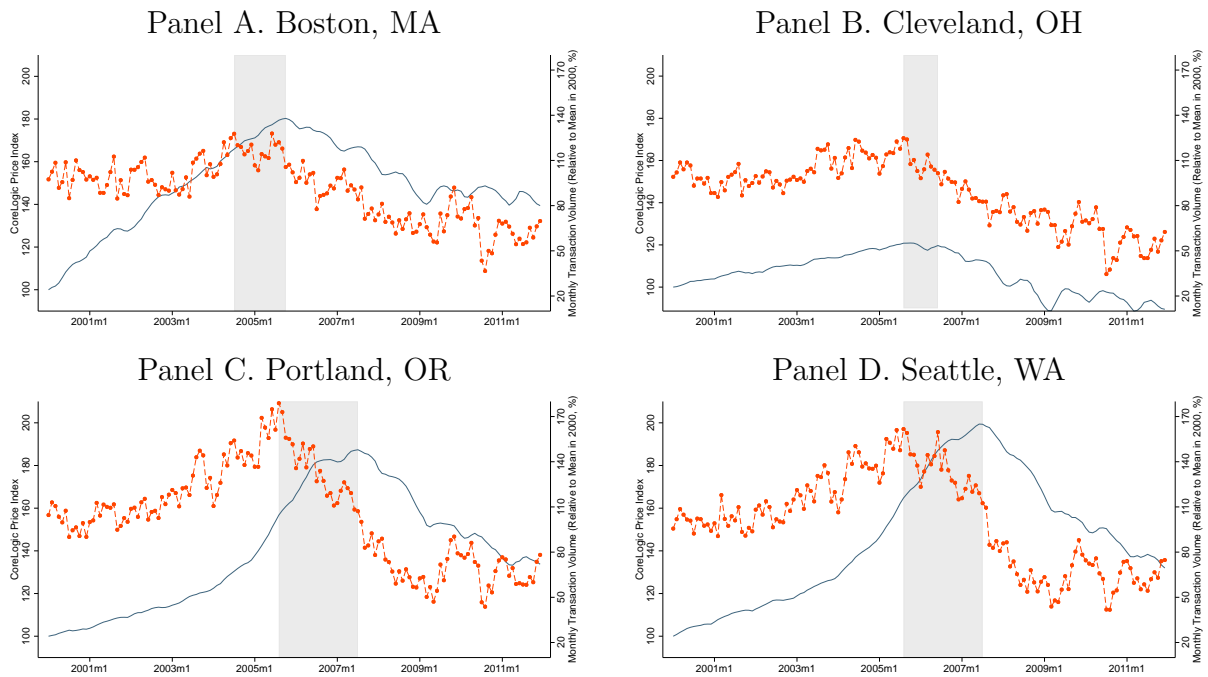
One concern is that the occupant premium, μ_1 , is about 7 times smaller than the standard deviation of flow utility, σ_a . Therefore, non-occupants may play a small role in amplifying the cycle solely because of parameter values in which non-occupants closely resemble occupants. To investigate this possibility, Table IA11 regenerates the first, third, and fifth columns of Table IA10 under the larger values of $\mu_1 = 0.033$ and $\mu_1 = 0.066$, corresponding to 50% and 100% of the baseline value of σ_a . We continue to find significant attenuation of the cycle with all long-term buyers if we adjust for the occupant distribution, but not with all occupant buyers if we adjust for the λ distribution.

These results speak to the finding in Table 2 that a short-volume boom more robustly predicts price booms and busts than does a non-occupant boom. Our findings are consistent with Gao et al. (2020), who find that non-occupants amplify the housing bust, as that paper does not look separately at long-term versus short-term non-occupants. Chinco and Mayer (2015) find a stronger effect of out-of-town than local non-occupant buyers on subsequent price growth. This finding is consistent with our results if out-of-town buyers have shorter horizons than local ones. Finally, our results echo Nathanson and Zwick (2018), who theoretically predict larger house price booms in cities with a greater share of non-occupant buyers when those buyers disagree about future prices and the housing stock is fixed. Static disagreement in that model functions similarly to how, in this model, variation in horizons interacts with extrapolative expectations to generate heterogeneous expected returns.

References


- Aliprantis, Charalambos D. and Border, Kim C.** (2006), *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Springer.
- Anenberg, Elliot and Bayer, Patrick.** (2013), Endogeneous Sources of Volatility in Housing Markets: The Joint Buyer-Seller Problem, Technical report, National Bureau of Economic Research.
- Bayer, Patrick, Geissler, Christopher, Mangum, Kyle and Roberts, James W.** (2020). 'Speculators and Middlemen: The Strategy and Performance of Investors in the Housing Market', *Review of Financial Studies* 33(11), 5212–5247.
- Chinco, Alex and Mayer, Christopher.** (2015). 'Misinformed Speculators and Mispricing in the Housing Market', *Review of Financial Studies* 29(2), 486–522.
- Choi, Hyun-Soo, Hong, Harrison G. and Scheinkman, José.** (2014). 'Speculating on Home Improvements', *Journal of Financial Economics* 111(3), 609–624.
- Edelstein, Robert and Qian, Wenlan.** (2014). 'Short-Term Buyers and Housing Market Dynamics', *Journal of Real Estate Finance and Economics* 49, 654–689.
- Gao, Zhenyu, Sockin, Michael and Xiong, Wei.** (2019). 'Economic Consequences of Housing Speculation'. NBER Working Paper 26457.
- Kelley, John L.** (1955), *General Topology*, Springer-Verlag.
- Nathanson, Charles G. and Zwick, Eric.** (2018). 'Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market', *Journal of Finance* 73(6), 2587–2633.
- 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.
- 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.
- Stokey, Nancy L., Lucas, Jr., Robert E. and Prescott, Edward C.** (1989), *Recursive Methods in Economic Dynamics*, Harvard University Press.

FIGURE IA1
The Dynamics of Prices and Volume (Non-Sand-State Cities)



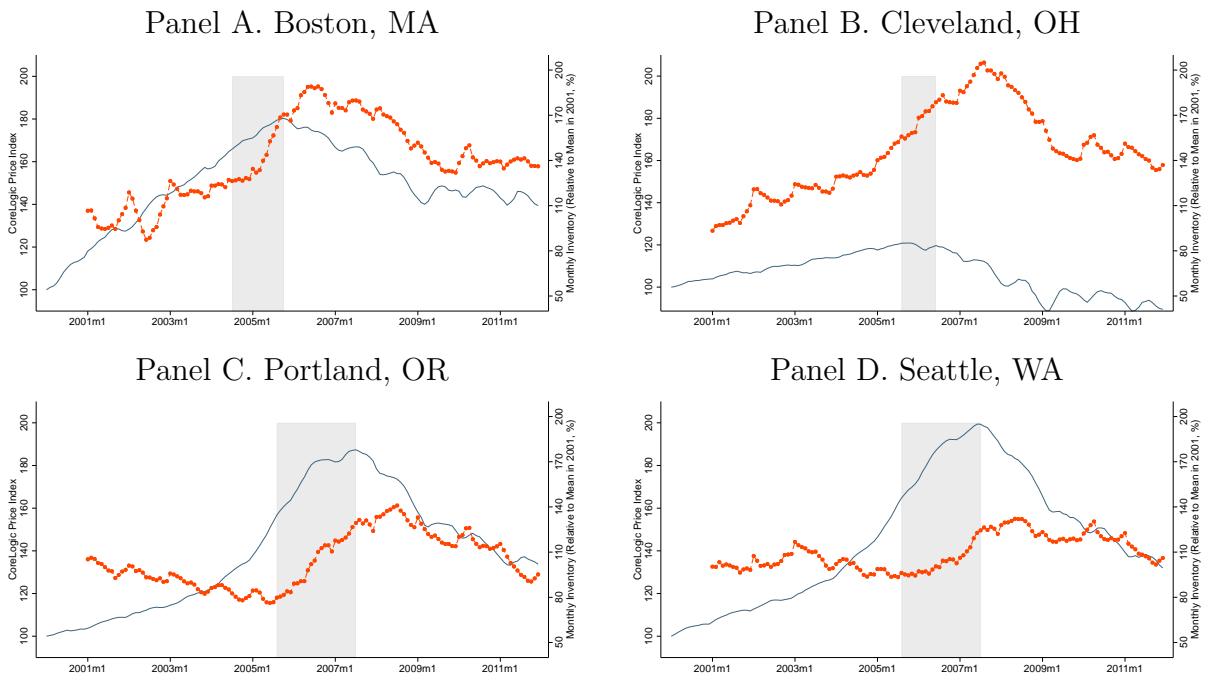
Notes: This figure displays the dynamic relation between prices and volume in the U.S. housing market between 2000 and 2011. In Figure 1, we focus on cities that represent the largest boom–bust cycles. Here, we focus on the largest cities outside of the sand states for which we have both volume and listings data. Variables are defined as in Figure 1. Shaded regions denote the *quiet*, defined as the period between the peak of volume and the last peak of prices before their pronounced decline.

FIGURE IA2
The Lead-Lag Relationship between Prices and Volume (No Sand States)



Notes: This figure shows that the correlation between prices and lagged volume is robust across MSAs. The figure is constructed as in Figure 2 but excludes MSAs in Arizona, California, Florida, and Nevada.


FIGURE IA3
The Dynamics of Prices and Inventories (Non-Sand-State Cities)




Notes: This figure displays the dynamic relation between prices and inventory in the U.S. housing market between 2000 and 2011. In Figure 3, we focus on cities that represent the largest boom–bust cycles. Here, we focus on the largest cities outside of the sand states for which we have both volume and listings data. Variables are defined as in Figure 3. Shaded regions denote the *quiet*, defined as the period between the peak of volume and the last peak of prices before their pronounced decline.

FIGURE IA4
 Additional Impulse Responses in Counterfactuals


Panel A. Pr(Sale | Listing),
 Rational




Panel B. New Listings by Holding Period,
 Rational




Panel C. Pr(Sale | Listing),
 Walrasian




Panel D. New Listings by Holding Period,
 Walrasian



Panel E. Pr(Sale | Listing),
 No Speculation

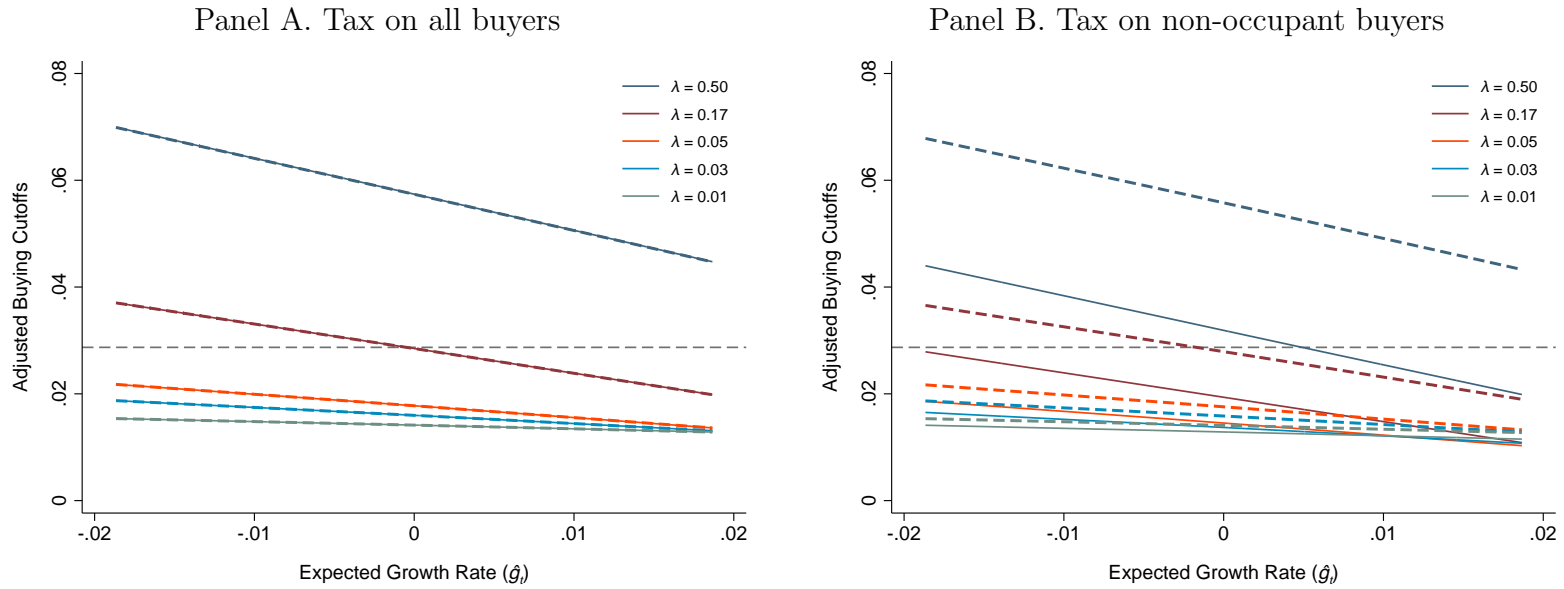


Panel F. New Listings by Holding Period,
 No Speculation



Notes: Impulse responses are average differences between log outcomes in control simulations and treatment simulations, in which a 2-standard-deviation shock to ϵ_t^g (the demand growth innovation) occurs in quarters 0 through 3. A short holding period is defined as less than or equal to 12 quarters.

FIGURE IA5
Adjusted Buying Cutoffs for Different Expected Growth Rates



Notes: The adjusted buying cutoff for occupancy type n and horizon type λ_j is $\bar{\kappa}\kappa_{n,j}^\tau(\hat{g})/\bar{\kappa}^\tau$, where $\tau = (\tau_0, \tau_1)$ is the vector of tax rates. In Panel A, we explore a 5% on all buyers, so that $\tau = (0.05, 0.05)$. In Panel B, we explore a tax that binds only on non-occupants, so that $\tau = (0.05, 0)$. Solid lines correspond to occupants ($n = 1$); dashed lines correspond to non-occupants ($n = 0$). The horizontal grey dashed line gives $\bar{\kappa}$.

TABLE IA1
List of Metropolitan Statistical Areas Included in the Analysis Sample

Metropolitan Statistical Area	Share of Housing Stock Represented	Included in Non-Occupant Analysis	Included in Listings Analysis	Metropolitan Statistical Area	Share of Housing Stock Represented	Included in Non-Occupant Analysis	Included in Listings Analysis
Akron, OH	1.00	x	x	New York-Newark-Jersey City, NY-NJ-PA	0.97	x	
Ann Arbor, MI	1.00	x	x	North Port-Sarasota-Bradenton, FL	1.00	x	
Atlanta-Sandy Springs-Roswell, GA	0.80			Norwich-New London, CT	1.00		x
Atlantic City-Hammonton, NJ	1.00	x	x	Ocala, FL	1.00	x	x
Bakersfield, CA	1.00	x	x	Ocean City, NJ	1.00	x	x
Baltimore-Columbia-Towson, MD	1.00	x		Olympia-Tumwater, WA	1.00	x	x
Barnstable Town, MA	1.00		x	Orlando-Kissimmee-Sanford, FL	1.00	x	
Bellingham, WA	1.00	x	x	Oxnard-Thousand Oaks-Ventura, CA	1.00	x	x
Bend-Redmond, OR	1.00	x		Palm Bay-Melbourne-Titusville, FL	1.00	x	
Boston-Cambridge-Newton, MA-NH	0.89		x	Pensacola-Ferry Pass-Brent, FL	1.00	x	
Boulder, CO	1.00	x	x	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	1.00	x	
Bremerton-Silverdale, WA	1.00	x	x	Phoenix-Mesa-Scottsdale, AZ	1.00	x	x
Bridgeport-Stamford-Norwalk, CT	1.00		x	Pittsfield, MA	1.00		
Buffalo-Cheektowaga-Niagara Falls, NY	0.80	x	x	Portland-Vancouver-Hillsboro, OR-WA	0.97	x	x
California-Lexington Park, MD	1.00	x		Port St. Lucie, FL	1.00	x	x
Canton-Massillon, OH	0.92	x	x	Prescott, AZ	1.00	x	x
Cape Coral-Fort Myers, FL	1.00	x	x	Providence-Warwick, RI-MA	0.78		x
Champaign-Urbana, IL	0.82	x		Punta Gorda, FL	1.00	x	
Charleston-North Charleston, SC	0.79	x		Raleigh, NC	0.78	x	
Chicago-Naperville-Elgin, IL-IN-WI	0.90			Reading, PA	1.00	x	
Chico, CA	1.00	x		Redding, CA	1.00	x	
Cincinnati, OH-KY-IN	0.78	x	x	Reno, NV	0.99	x	x
Cleveland-Elyria, OH	1.00	x	x	Riverside-San Bernardino-Ontario, CA	1.00	x	x
Colorado Springs, CO	0.95	x		Rockford, IL	0.84	x	
Crestview-Fort Walton Beach-Destin, FL	1.00	x		Sacramento-Roseville-Arden-Arcade, CA	1.00	x	x
Dallas-Fort Worth-Arlington, TX	0.85	x		Salem, OR	0.79	x	
Dayton, OH	0.86	x		Salinas, CA	1.00	x	
Deltona-Daytona Beach-Ormond Beach, FL	1.00	x	x	San Diego-Carlsbad, CA	1.00	x	x
Denver-Aurora-Lakewood, CO	0.95	x		San Francisco-Oakland-Hayward, CA	1.00	x	x
El Centro, CA	1.00	x		San Jose-Sunnyvale-Santa Clara, CA	1.00	x	
El Paso, TX	0.99	x	x	Santa Cruz-Watsonville, CA	1.00	x	
Elmira, NY	1.00	x		San Luis Obispo-Paso Robles-Arroyo Grande, CA	1.00	x	x
Erie, PA	1.00	x		Santa Maria-Santa Barbara, CA	1.00	x	
Eugene, OR	1.00	x	x	Santa Rosa, CA	1.00	x	
Flagstaff, AZ	1.00	x		Seattle-Tacoma-Bellevue, WA	1.00	x	x
Fort Collins, CO	1.00	x	x	Sebastian-Vero Beach, FL	1.00	x	
Fresno, CA	1.00	x		Sebring, FL	1.00	x	
Gainesville, FL	0.91	x		Sierra Vista-Douglas, AZ	1.00	x	
Gainesville, GA	1.00			Spokane-Spokane Valley, WA	0.87	x	
Hanford-Corcoran, CA	1.00	x		Springfield, IL	0.93	x	
Hartford-West Hartford-East Hartford, CT	1.00		x	Springfield, MA	1.00		x
Homosassa Springs, FL	1.00	x	x	Springfield, OH	1.00	x	
Ithaca, NY	1.00	x	x	Stockton-Lodi, CA	1.00	x	x
Jacksonville, FL	0.98	x		Tampa-St. Petersburg-Clearwater, FL	1.00	x	
Kahului-Wailuku-Lahaina, HI	1.00	x	x	The Villages, FL	1.00	x	
Kingston, NY	1.00	x	x	Toledo, OH	0.92	x	x
Lake Havasu City-Kingman, AZ	1.00	x	x	Trenton, NJ	1.00	x	
Lakeland-Winter Haven, FL	1.00	x		Tucson, AZ	1.00	x	x
Lancaster, PA	1.00	x	x	Urban Honolulu, HI	1.00	x	x
Las Vegas-Henderson-Paradise, NV	1.00	x		Vallejo-Fairfield, CA	1.00	x	
Los Angeles-Long Beach-Anaheim, CA	1.00	x	x	Vineland-Bridgeton, NJ	1.00	x	x
Madera, CA	1.00	x		Visalia-Porterville, CA	1.00	x	
Merced, CA	1.00	x	x	Washington-Arlington-Alexandria, DC-VA-MD-WV	0.95	x	
Miami-Fort Lauderdale-West Palm Beach, FL	1.00	x		Worcester, MA-CT	1.00		x
Modesto, CA	1.00	x	x	Youngstown-Warren-Boardman, OH-PA	0.80	x	x
Napa, CA	1.00	x		Yuba City, CA	1.00	x	
Naples-Immokalee-Marco Island, FL	1.00	x	x	Yuma, AZ	1.00	x	
New Haven-Milford, CT	1.00		x				

Notes: This table lists the Metropolitan Statistical Areas that are included in the final analysis sample along with the share of the total 2010 owner-occupied housing stock for each MSA that is represented by the subset of counties for which CoreLogic has consistent data coverage back to 1995.

TABLE IA2
Number of Transactions Dropped During Sample Selection

Original number of Transactions	57,668,026
Dropped: Non-unique CoreLogic ID	50
Dropped: Non-positive price	3,309,100
Dropped: Duplicate transaction	618,129
Dropped: Subdivision sale	1,321,261
Dropped: Vacant lot	839,078
Final Number of Transactions	51,580,408

Notes: This table shows the number of transactions dropped at each stage of our sample-selection procedure.

TABLE IA3
Mechanical Short-Term Volume Estimates

Year	$\hat{\alpha}_y^{buy} - \hat{\alpha}_{2000}^{buy}$	Total Volume	Actual Short-Term Volume	Counterfactual Short-Term Volume
2000	0	2821596	512787	512787
2001	0.0003	2757954	499643	494741
2002	0.0008	2985550	556987	534342
2003	0.0014	3226968	614429	557701
2004	0.0023	3667997	772708	659111
2005	0.0027	3857236	909976	725847
2000–2005 growth	–	36.7%	77.5%	41.5%

Notes: Total Volume gives annual transaction counts in our analysis sample. Actual Short-Term Volume are sales of properties for which the previous purchased occurred less than 36 months in the past. We estimate α_y^{buy} , a fixed effect for the propensity to sell a house having bought it in year y , using the regression equation in Section B.1. In the counterfactual, we assume that α_y^{buy} remains constant at its level in $y = 2000$ for $y \in \{2001, 2002, 2003, 2004, 2005\}$.

TABLE IA4
Instrumental Variables Estimation of the Role of Short-Term Volume

	First Stage	Volume Boom		Price Boom		Price Bust	
		OLS	IV	OLS	IV	OLS	IV
Short-Volume Boom		2.28*** (0.12)	2.28*** (0.18)	2.18*** (0.38)	2.78*** (0.57)	-0.77*** (0.09)	-1.05*** (0.13)
Old Share	1.69*** (0.26)						
Young Share	0.66** (0.32)						
Number of Observations	102	102	102	102	102	102	102
R-squared	0.45	0.79	0.79	0.25	0.23	0.45	0.39
F-Statistic	39.95						

Notes: This table presents OLS and IV regressions at the MSA level of price and volume housing cycle measures on the change in short-holding-period volume from 2000 to 2005 relative to total volume in 2000. In the IV regressions, Short-Volume Boom is instrumented with demographic data from the 2000 Census 5% microdata. The instruments are the share of recent buyers under 35 and the share of recent buyers aged 65 or older. Census microdata was not available for 13 MSAs in our sample, hence the lower sample count in this table. The first column presents the first-stage regression and F-statistic.

TABLE IA5
House Price Appreciation and Speculative Buyer Shares (Monthly Panel VAR)

	House Price Appreciation Rate		
Lagged Price Appreciation	0.375*** (0.026)	0.387*** (0.027)	0.372*** (0.026)
Lagged Short-Buyer Share	0.021*** (0.005)		0.023*** (0.005)
Lagged Non-Occupant Share		0.009 (0.008)	0.006 (0.006)
	Short-Buyer Share		
Lagged Price Appreciation	0.163*** (0.048)		0.162*** (0.048)
Lagged Short-Buyer Share	0.780*** (0.024)		0.781*** (0.023)
Lagged Non-Occupant Share			0.001 0.017
	Non-Occupant Share		
Lagged Price Appreciation		0.124*** (0.044)	0.172*** (0.045)
Lagged Short-Buyer Share			-0.071*** (0.016)
Lagged Non-Occupant Share		0.892*** (0.025)	0.900*** (0.021)

Notes: This table presents estimates from MSA-by-month panel vector autoregressions (pVARs) describing the relation between house price growth and the share of purchases made by non-occupant buyers and “short buyers,” defined as buyers who will sell within three years of purchase. The left-hand-side variables are house price appreciation from $t - 1$ to t , the short-buyer share of total volume in t , and the non-occupant share of total volume in t . The right-hand-side variables are lagged versions of these variables. The sample includes 8,568 observations for 102 MSAs for which we can consistently identify non-occupant buyers. House price appreciation has a mean of 0.84% and a standard deviation of 1.32%. Short-buyer share has a mean of 21.0% and a standard deviation of 5.5%. Non-occupant share has a mean of 32.8% and a standard deviation of 18.9%. Column (1) includes only house price appreciation and the short-buyer share. Column (2) includes only house price appreciation and the non-occupant share. Column (3) includes both speculative volume measures. The sample period includes the boom and quiet, which runs from January 2000 through December 2006. Regressions include MSA and month fixed effects and thus report the average autoregressive relations within MSAs over time. We seasonally adjust house prices by removing MSA-by-calendar-month fixed effects before computing house price growth. Standard errors are clustered at the MSA level.

TABLE IA6
Speculators and Housing Market Outcomes (Extra Listing Outcomes)

Panel A. Propensity to List				
	Δ New Listings Boom		Δ New Listings Quiet	
Short-Volume Boom	0.270		0.649***	
	(0.182)		(0.160)	
Non-Occupant Volume Boom		0.115		0.308***
		(0.092)		(0.080)
Number of Observations	57	48	57	48
R-squared	0.038	0.033	0.229	0.243

Panel B. Sale Probability				
	Δ P(Sale) Boom		Δ P(Sale) Quiet	
Short-Volume Boom	0.142***		-0.163***	
	(0.032)		(0.031)	
Non-Occupant Volume Boom		0.058***		-0.047**
		(0.017)		(0.018)
Number of Observations	57	48	57	48
R-squared	0.268	0.206	0.332	0.122

Notes: This table reports estimates of the relation between speculative volume and housing cycle measures at the MSA level. Short-Volume Boom has a mean of 16.0% and a standard deviation of 12.9%. Non-Occupant Volume Boom has a mean of 29.3% and a standard deviation of 27.1%. Δ New Listings Boom is defined as the change in the flow of listings from 2003 through 2005. Δ New Listings Quiet is defined as the change in the flow of listings from 2005 through 2007. These outcomes correspond to listing propensities among existing homeowners. Δ P(Sale) Boom is defined as the change in the probability of sale among the observed stock of listings from 2003 through 2005. Δ P(Sale) Quiet is defined as the change in the probability of sale among the observed stock of listings from 2005 through 2007. To aid interpretation of these relations, we scale the change in outcomes for all quantity measures relative to total volume in 2003. We do not scale the sale probability. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE IA7
Speculative Booms and Housing Market Outcomes (Sand State Control)

Panel A. MSA-Level Prices

	Price Boom		Price Bust	
Short-Volume Boom	1.022*** (0.272)		-0.237*** (0.061)	
Non-Occupant Volume Boom		0.228 (0.142)		-0.044 (0.032)
Number of Observations	115	102	115	102
R-squared	0.514	0.453	0.696	0.662

Panel B. MSA-Level Inventories

	Δ Listings Boom		Δ Listings Quiet	
Short-Volume Boom	-1.581 (1.163)		4.276*** (1.461)	
Non-Occupant Volume Boom		-0.206 (0.525)		1.930*** (0.642)
Number of Observations	57	48	57	48
R-squared	0.034	0.020	0.337	0.440

Panel C. MSA-Level Volume Quiet and Bust

	Δ Volume Quiet + Bust		Foreclosures Bust	
Short-Volume Boom	-1.145*** (0.105)		-0.233 (0.377)	
Non-Occupant Volume Boom		-0.516*** (0.053)		-0.451** (0.185)
Number of Observations	115	102	115	102
R-squared	0.533	0.505	0.317	0.333

Notes: This table reports estimates of the relation between speculative volume and housing cycle measures at the MSA level. The table follows Table 2 while adding a control for “Sand States,” which is an indicator for MSAs in Arizona, California, Florida, and Nevada.

TABLE IA8

Speculators and Housing Market Outcomes (Extra Listing Outcomes, Sand State Control)

Panel A. Propensity to List

	Δ New Listings Boom		Δ New Listings Quiet	
Short-Volume Boom	0.050 (0.198)		0.431** (0.171)	
Non-Occupant Volume Boom		0.040 (0.087)		0.228*** (0.072)
Number of Observations	57	48	57	48
R-squared	0.131	0.213	0.323	0.451

Panel B. Sale Probability

	Δ P(Sale) Boom		Δ P(Sale) Quiet	
Short-Volume Boom	0.146*** (0.036)		-0.086*** (0.028)	
Non-Occupant Volume Boom		0.058*** (0.018)		-0.021 (0.013)
Number of Observations	57	48	57	48
R-squared	0.268	0.206	0.598	0.607

Notes: This table reports estimates of the relation between speculative volume and housing cycle measures at the MSA level. The table follows Table IA6 while adding a control for “Sand States,” which is an indicator for MSAs in Arizona, California, Florida, and Nevada.

TABLE IA9
All-Cash Buyer Shares and Mean LTV by Buyer Type

	Transaction-Level	MSA-Level			
	All Months	All Months	Boom	Quiet	Bust
All-Cash Buyer Share					
Short Buyers	0.29	0.38 (0.21)	0.29 (0.16)	0.28 (0.17)	0.52 (0.20)
Non-Occupant Buyers	0.38	0.41 (0.18)	0.36 (0.15)	0.32 (0.14)	0.50 (0.18)
All Buyers	0.20	0.25 (0.16)	0.22 (0.15)	0.20 (0.14)	0.30 (0.16)
Mean LTV					
Short Buyers	0.59 (0.40)	0.52 (0.18)	0.60 (0.13)	0.59 (0.13)	0.41 (0.17)
Non-Occupant Buyers	0.50 (0.41)	0.48 (0.14)	0.52 (0.12)	0.54 (0.11)	0.41 (0.15)
All Buyers	0.65 (0.36)	0.62 (0.13)	0.64 (0.12)	0.64 (0.11)	0.59 (0.14)
Mean LTV LTV > 0					
Short Buyers	0.84 (0.16)	0.85 (0.06)	0.84 (0.05)	0.82 (0.04)	0.85 (0.07)
Non-Occupant Buyers	0.81 (0.17)	0.82 (0.06)	0.82 (0.06)	0.80 (0.05)	0.82 (0.06)
All Buyers	0.82 (0.16)	0.83 (0.05)	0.82 (0.04)	0.80 (0.04)	0.85 (0.05)

Notes: This table presents statistics on LTV ratios and the share of buyers of various types who purchased their homes without the use of a mortgage. In column 1, statistics are measured at the transaction level and includes all transactions recorded between January 2000 and December 2011 from the CoreLogic deeds records described in Section 1.1. The first row of each panel includes only transactions by homebuyers who are observed to have sold the home within three years of purchase. The second row of each panel includes only non-occupant buyers. The third row of each panel includes all buyers. In columns 2–5, means are first calculated at the MSA-by-month level and then averaged across MSA-months within a given time period. The standard deviation of these MSA-month means is reported in parentheses. Column 2 includes all MSA-months between January 2000 and December 2011. Column 3 includes only MSA-months between January 2000 and August 2005. Column 4 includes only MSA-months between August 2005 and December 2006. Column 5 includes only MSA-months between December 2006 and December 2011. All statistics are calculated in the full sample of 115 MSAs with the exception of those for non-occupants, which are calculated in the sample of 102 MSAs with valid non-occupancy data.

TABLE IA10
Model counterfactuals

Outcome	Baseline	All long-term buyers		All occupants	
		No occupant adjustment	Occupant adjustment	No short-term adjustment	Short-term adjustment
Price boom	14.5	8.7	8.7	9.4	14.6
Price bust	-8.2	-0.4	-0.4	-0.6	-8.3
Volume boom	5.8	2.9	2.9	2.1	5.8
Listings, end of boom	-1.3	-3.1	-3.1	-0.2	-1.3
Listings, end of quiet	1.4	0.4	0.4	0.0	1.4
Short volume boom	14.1	3.4	3.4	6.4	14.1
Non-occupant volume boom	12.3	3.6	3.6	-	-
Sale probability boom	7.1	6.0	6.0	2.3	7.1

Notes: We report 100 times changes in log outcomes between treatment and control simulations. See notes to Table 6 for outcome definitions. A two-sided minimum for prices does not occur in the 48 analysis periods in the fourth column, so we extend the analysis 60 additional periods to find such a minimum in order to measure the price bust. The counterfactuals involve different values of the underlying distribution of potential buyers, $\beta_{n,j}$, that the text describes. We alter κ in each counterfactual to maintain a zero demand error while keeping other parameters the same. The baseline values correspond to Figure 11.

TABLE IA11
Robustness to larger occupant premium (μ_1)

Outcome	$\mu_1 = 0.033$			$\mu_1 = 0.066$		
	Baseline	All long-term buyers	All occupants	Baseline	All long-term buyers	All occupants
Price boom	14.0	8.7	14.6	12.8	8.6	14.6
Price bust	-7.6	-0.4	-8.3	-5.6	-0.3	-8.3
Volume boom	5.6	2.9	5.8	4.9	2.9	5.8
Listings, end of boom	-1.2	-3.1	-1.3	-1.0	-3.0	-1.3
Listings, end of quiet	1.3	0.4	1.4	0.9	0.3	1.4
Short volume boom	13.8	3.4	3.4	12.5	3.4	14.1
Non-occupant volume boom	11.3	5.6	-	7.5	9.1	-
Sale probability boom	6.8	6.0	7.1	6.0	5.9	7.1

Notes: We report 100 times changes in log outcomes between treatment and control simulations. See notes to Table 6 for outcome definitions. For each value of μ_1 , we re-choose the other parameters in Table 5 by matching the targets in Table 4 other than non-occupant boom/occupant boom. The Baseline column reports outcomes under each new set of parameters. In the All long-term buyers column, we further change the $\beta_{n,j}$ distribution to put all weight on $\lambda = 0.03$ while keeping the occupancy distribution unchanged, corresponding to the third column of results in Table IA10. In the All occupants column, we further change the $\beta_{n,j}$ distribution to put all weight on $n = 1$ while keeping the λ distribution unchanged, corresponding to the fifth column of results in Table IA10.