Journal of Urban Economics 108 (2018) 72–84

Contents lists available at ScienceDirect

Journal of Urban Economics

journal homepage: www.elsevier.com/locate/jue

The role of price spillovers in the American housing boom[☆]

Anthony DeFusco[a], Wenjie Ding[b], Fernando Ferreira[c], Joseph Gyourko[c]

[a] Kellogg School of Management, Northwestern University, United States
[b] The Wharton School, University of Pennsylvania, United States
[c] The Wharton School, University of Pennsylvania and NBER, United States

A B S T R A C T

One of the striking features of the last U.S. housing boom was the heterogeneity in the timing of its onset across local markets. In this paper, we exploit this heterogeneity to estimate the extent to which the boom was spread via spatial spillovers from one market to another. Our analysis focuses on spillovers that occur around the time that a local market enters its boom, which we identify using sharp structural breaks in house price growth rates. On the extensive margin, there is evidence that the likelihood of a market booming increases substantially if nearby neighbors boom. On the intensive margin, we also find statistically significant but economically modest effects of the size of a neighbor's boom on subsequent price growth in nearby markets. These affects appear to be unrelated to local market fundamentals, suggesting a potential role for non-rational factors.

1. Introduction

Given the housing market's central role in fomenting the global financial crisis, the last U.S. housing cycle was arguably one of the defining economic events of the last half century. In this paper, we test for whether spillovers across markets materially influenced how the housing boom began and spread both spatially and over time. This is an important issue for a variety of reasons. Shiller (2005, 2006) has argued that a type of psychological contagion may have led to an irrational exuberance that could have a spatial dimension. Other recent research hints at the potential for significant geographic spillovers in the housing market. For example, Bailey et al., (2016) show that recent house price experiences within an individual's geographically distant social network can directly affect that individual's own expectations and housing market behavior in her local market. Others have provided evidence that out-of-town speculators could also be a potential source of cross-market spillovers. Both Haughwout et al. (2011) and DeFusco et al., (2017) report that investment purchases constituted a large share of the transaction volume during the run-up to the housing bust, and Chinco and Mayer (2014) suggest that a significant fraction of these purchases were made by out-of-town buyers.[1]

Motivated by this evidence, we ask whether spatial spillovers were an important contributing factor to the spread of the housing boom across markets. To answer this question, we focus on the pattern of local market house price changes around the time that neighboring markets enter their housing booms. In doing so, we investigate both extensive and intensive margin spillovers. On the extensive margin, we ask whether the probability of a boom starting in a given focal market is materially influenced by whether a boom has recently begun in a nearby neighboring market. To the best of our knowledge, this type of extensive margin spillover has not been considered in prior work. On the intensive margin, we investigate the magnitude of focal market price changes around the time that neighboring markets enter their housing booms.

The nature of the housing market and the richness of our data allow us to systematically identify the beginning of local housing booms using an empirical approach that exploits sharp changes in house price growth rates. In particular, we define the beginning of a housing boom to be the quarter in which each market experienced a positive and statistically significant structural break in its house price growth series. This approach has been used to identify house price shocks in previous empirical work (Ferreira and Gyourko, 2011; Charles, Hurst, and Notowidigdo, 2018; Dokko et al., 2015), and we use these events here as focal points for studying the effect of geographic spillovers on both the extensive and intensive margins.

A key descriptive finding from our analysis is that the timing of local housing booms was highly spatially correlated. This can be seen in Fig. 1, which documents the geography of the time line of the start of local housing booms across the 94 metropolitan areas in our

☆ The authors thank the Research Sponsors Program of the Zell/Lurie Real Estate Center at Wharton for financial support. We also appreciate the comments of participants in presentations at the NBER Conference on Housing and Financial Crisis, the University of Miami, and the University of California-Berkeley.

E-mail addresses: anthony.defusco@kellogg.northwestern.edu (A. DeFusco), fferreir@wharton.upenn.edu (F. Ferreira), gyourko@wharton.upenn.edu (J. Gyourko).

[1] Within a metropolitan area, Bayer et al. (2011) also find that homeowners are more likely to engage in speculative activity after having observed a recently successful house "flip" in their local neighborhood.

https://doi.org/10.1016/j.jue.2018.10.001
Received 14 April 2017; Received in revised form 5 September 2018
Available online 12 October 2018

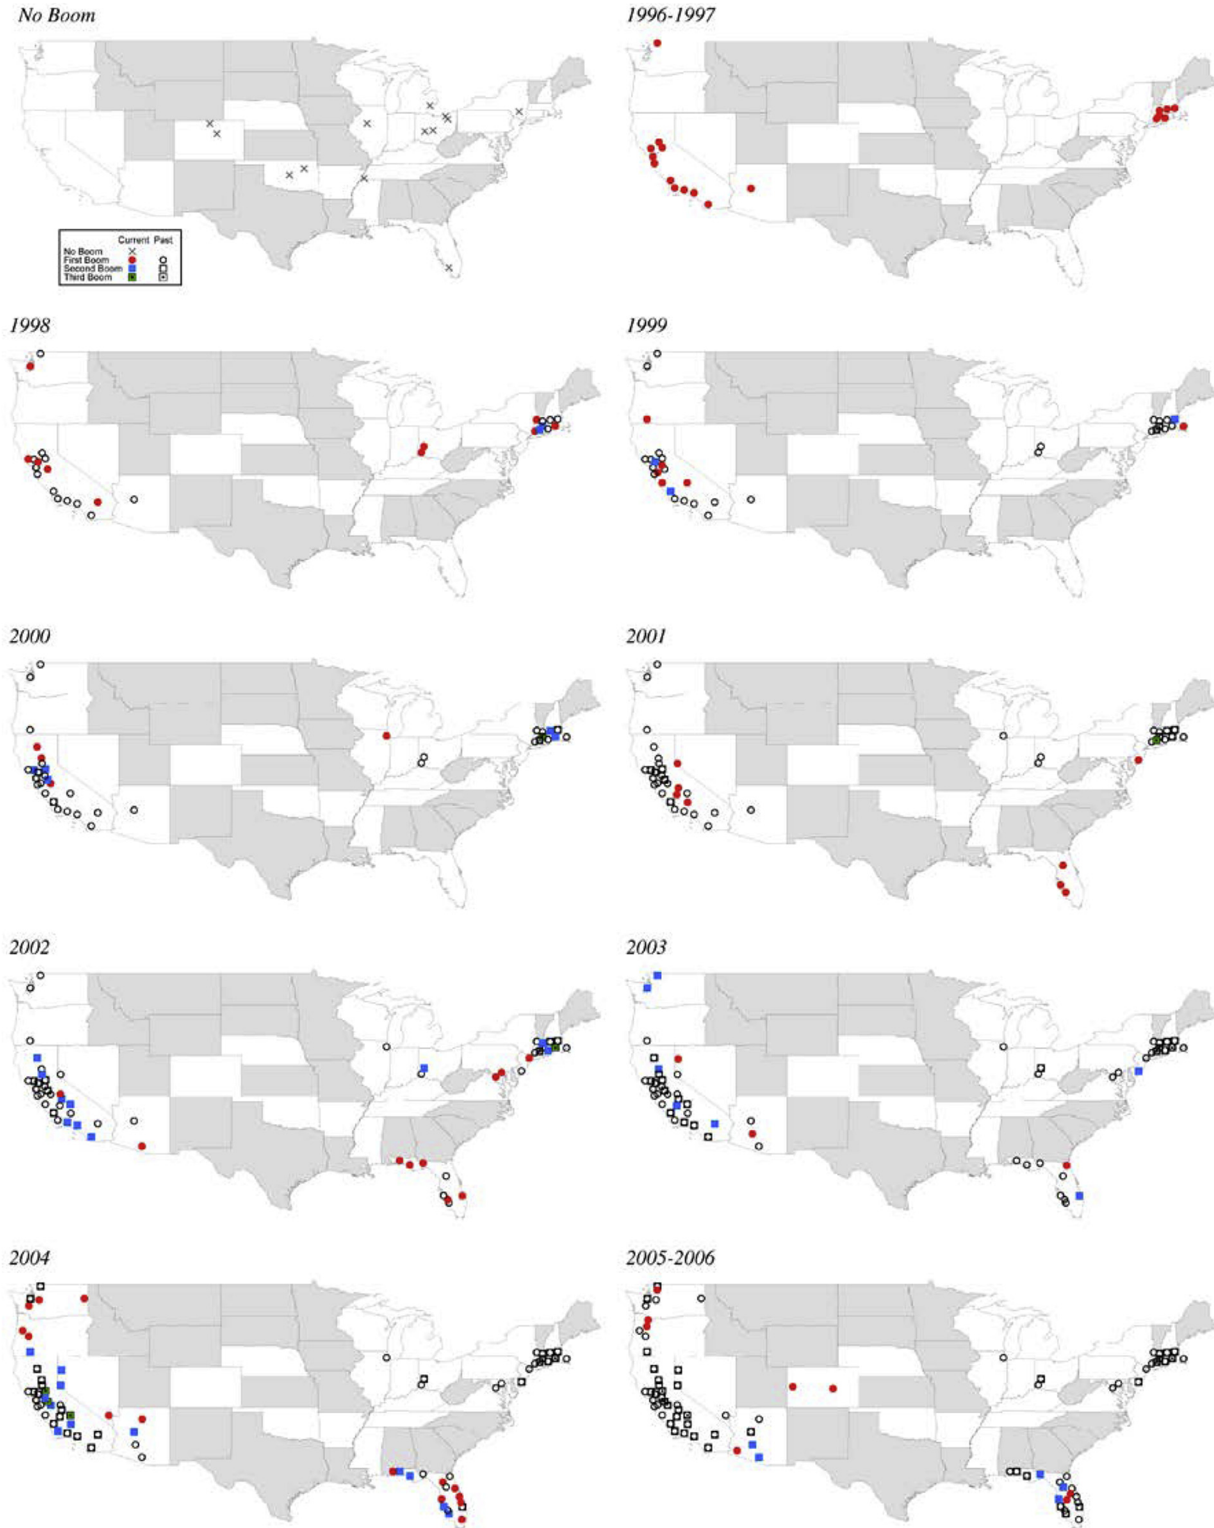

Fig. 1. Timing and geography of housing booms by MSA.

Note: Figure plots the distribution of housing booms over time and across MSAs. Colored markers represent booms beginning as of the indicated time period. Hollow markers represent previous booms. The shape and color of each marker denotes whether the boom is a first, second, or third boom. See Section 2 for details on how the beginning of each boom was estimated. States shaded in grey are those for which there is no DataQuick data.

sample.[2] The top left panel plots the 13 primarily rust belt and interior markets that never boomed. The other panels show that the remaining 81 markets in our sample boomed at very different times over the ten-year period from 1996 to 2006. The housing boom spread from what were initially highly concentrated areas on the two coasts, with the earliest booms beginning between 1996 and 1999 in California and the mid-New England region. On the west coast, housing booms eventually spread inland towards central California and to neighboring states to the east and north. On the east coast, housing booms spread to other markets in New England and then to neighboring regions, eventually reaching the majority of Florida markets between 2004 and 2006. This timing appears to be non-random and the patterns are suggestive of spatial spillover effects that disseminate positive housing price shocks from one market to another.

To more formally analyze this pattern, we estimate whether the probability that a given market enters its housing boom in a particular period is related to the timing of recent booms in neighboring markets. As explained more fully later in the paper, we use the timeline of a neighbor's boom as our source of variation in the data to identify this type of extensive margin spillover effect. Our baseline specification involves regressing an indicator for whether a focal market enters a housing boom in a given quarter on a series of indicators reflecting whether it's neighboring markets are booming and how proximate a given period is in time to the start of those market's booms.

Our results reveal a large impact on the extensive margin and generally confirm the visual impression given by Fig. 1. Unconditionally, the probability of a focal market entering a housing boom in a given quarter roughly doubles if its nearby neighbors have recently begun to boom. As would be expected, the magnitude of this effect falls and standard errors increase once additional controls are included. However, even in our preferred specification, which includes both regional trends and a detailed set of local market controls, the economic magnitude of the effect is quite large and implies a roughly 50% increase in the likelihood of booming.

We also find a statistically significant impact of spillovers on the intensive margin. To investigate this possibility, we ask whether there is any evidence of changes in price levels around the timing of neighboring market booms. We find robust evidence that there was. For the average MSA in our sample, prices jump by roughly 0.6 to 1% in the year that nearby neighboring markets enter their housing boom. While statistically significant, these effects imply a relatively modest elasticity of focal market price with respect to neighboring market prices of roughly 0.1 to 0.25 in the period immediately following the neighbors' boom.

Having documented the existence of both intensive and extensive margin spillovers, we next investigate the importance of several plausible mechanisms that could be driving these effects. These mechanisms include the impact of neighboring market housing booms on the average income of potential buyers in the focal market, the behavior of lenders in both sets of markets (e.g., whether subprime share rises), migration patterns across markets, and speculative activity in the focal market. We find that these measures of focal market fundamentals have little effect on our main estimates, which raises the possibility that the price spillovers we document may be due to forces unrelated to fundamentals.

In addition to these results, our research also makes a number of data and methodological contributions that help to address potential biases that could lead one to mistakenly overestimate the extent of spillovers or contagion across space. To estimate the structural break points that we use to demarcate the beginning of local booms, we use a data set containing over 23 million observations on individual home sales in 94 metropolitan areas dating back to the early 1990 s in most cases. Importantly, this very large micro-level data set enables us to address specification search bias of the type identified by Leamer (1983), which arises when the same sample is employed to identify both the timing of a shock and the magnitude of the volatility during that period. Most studies of contagion in other asset markets are not able to deal with this issue because they typically only have access to a single aggregate price index for each market.[3] In contrast, our empirical strategy leverages the availability of transaction-level micro data for the housing market to generate randomly split samples that we use to separately identify the timing of booms and the magnitude of price changes during those periods.[4]

The substantial variation in the timing of booms across markets documented in Fig. 1 also allows us to address several sources of more standard omitted variable bias. Most importantly, the added degrees of freedom afforded by the multiple, non-contemporaneous booms we observe allow us to control for omitted factors that might reflect changes in aggregate economic conditions. For example, many of our specifications include a full set of census division-by quarter fixed effects, which means that our estimates will pick up only the changes in focal market conditions explained by the neighbors' boom that are over and above the regional average trends. The fact that we find evidence of spillovers even conditional on these controls underscores the potential role that such forces may have played in the development of the last housing boom.

The remainder of the paper proceeds as follows. The next section discusses our method for dating the beginning of local housing booms. Section 3 describes our data sources and sample selection criteria. Section 4 discusses our empirical framework, presents our estimates of the spillover effects and explores potential mechanisms that might explain them. Section 5 concludes.

2. Identifying the timeline of local housing booms

Any analysis of spillovers during the recent housing boom first requires knowledge of the timing of the beginning of that boom in different markets. Our approach to identifying local booms follows that of Ferreira and Gyourko (2011). Specifically, we estimate the existence and timing of local booms at the MSA level based on whether and when there was a structural break in each area's price appreciation rate series.

This strategy is motivated by implications of the dynamic urban spatial equilibrium model developed in Glaeser et al., (2014). Their framework implies that, in steady state, each local market will exhibit constant and continuous growth paths for house prices, new construction and population.[5] Empirically, this suggests that house prices in a given

[2] These figures are based on the estimation of structural breaks in local house price growth rates that is described in detail in the next section.

[3] See Forbes (2013) for an excellent recent review of empirical work on contagion in financial markets and Dungey et al. (2005) for a technical analysis of the challenges involved in convincingly estimating contagion effects in these markets. Other early empirical work on financial market contagion includes studies of the 1987 U.S. stock market crash (King and Wadhwani, 1990; Lee and Kim, 1993), the 1994 Mexican peso crisis (Calvo and Reinhart, 1996), and the Hong Kong stock market and Asian currency crisis of 1997 (Corsetti et al. 2005).

[4] This approach decreases the likelihood of falsely concluding that there are more and bigger booms than truly exist and is the same strategy followed by Card et al. (2008) in their study of tipping points in residential segregation models.

[5] Glaeser et al. (2014) introduce dynamics into Rosen's (1979) and Roback's (1982) classic static model of spatial equilibrium. In this compensating differential framework, house prices (P_i) are the entry fee paid to access the wages (W_i, which reflect productivity) and amenities (A_i) of labor market area *i*. Their model is closed with an assumption that there is some elastically supplied reference market area which is always open to another household. The utility level available in the reference market is given by U^*, and establishes the lower bound on utility provided in any market. In the long run, perfect mobility ensures that U^* is achieved in all markets, so that in equilibrium, no one has an incentive to move to another place which offers higher utility. A simple, linear version of this framework would imply that $U^* = W_i + A_i - P_i$, so that $dP_i = dW_i + dA_i$ in equilibrium. The steady state rate of price appreciation need

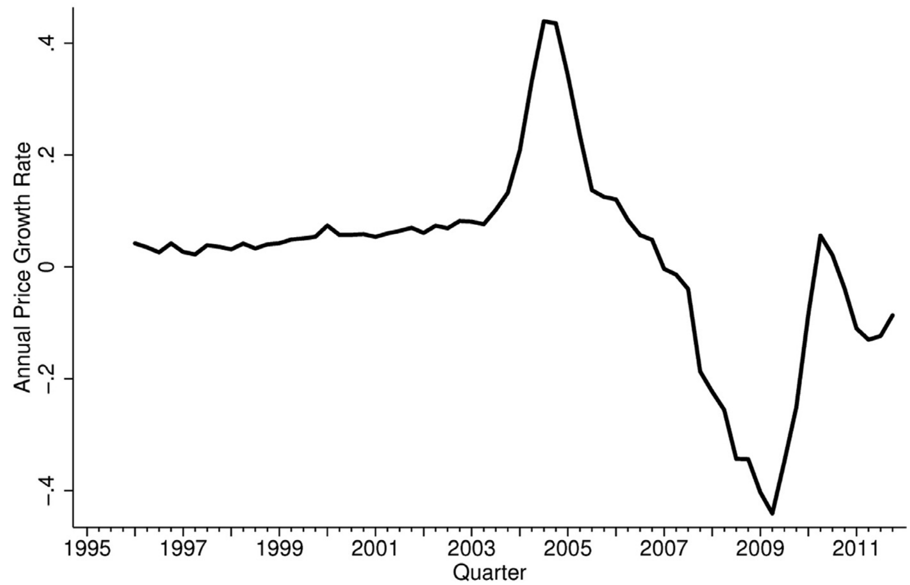

Fig. 2. Las Vegas house price growth rate.
Note: Figure plots year-over-year house piece growth rates at a quarterly frequency for the Las Vegas-Henderson-Paradise, NV MSA. The price growth series is constructed from a constant quality hedonic price index estimated using the DataQuick transactions data as described in Section 3.

market will grow at a (roughly) constant rate unless there is a shock to local productivity, amenities or expectations, in which case we would then observe a discrete jump in the appreciation rate for that market. The data are generally consistent with this predicted pattern. As an example, Fig. 2 plots year-over-year house price appreciation rates at a quarterly frequency for the Las Vegas market. House prices in this market were appreciating at a high, and roughly constant rate for many years before beginning to increase sharply starting in early 2004. Informally, our approach defines the beginning of the housing boom in a local market as the point at which house price growth rates exhibit this type of sharp change.[6]

To formalize this idea, we start with the following reduced form model of house price growth in MSA *i* at time *t:*

$$PG_{i,t} = d_{i,t} + \epsilon_{i,t}, \qquad t = 1, \dots, T, \tag{1}$$

where $PG_{i,\,t}$ represents year-over-year price growth in MSA *i* measured in quarter *t*. Glaeser et al., (2014) implies that $d_{i,t} = d_{i,0}$ for all *t* if the market is on its steady-state growth path. However, if there is a positive shock at time *t* then the price growth rate will exhibit a discrete jump in that period. The beginning of a local housing boom can thus be identified by testing for the existence of one or more structural breaks in the parameter $d_{i,\,t}$. To carry out this test we follow established methods in the time series literature for estimating such breaks.

Borrowing heavily from Estrella's (2003) notation, the null hypothesis is that $d_{i,\,t}$ is constant for the entire sample period:

$$H_0 : d_{i,t} = d_{i,0}, \qquad t = 1, \dots T.$$

The alternative is that $d_{i,\,t}$ changes at some proportion, $0 < \pi_i < 1$, of the sample which marks the beginning of a housing boom in market *i*. Specifically the alternative hypothesis is

$$H_1 : d_{i,t} = \begin{cases} d_{1,i}(\pi_i), & t = 1, \dots, \pi_i T \\ d_{2,i}(\pi_i), & t = \pi_i T + 1, \dots T. \end{cases}$$

For any given π_i, it is straightforward to carry out this hypothesis test. However, it is slightly more complicated when π_i is unknown and the determination of its value is the primary interest. To see how we estimate the value of π_i and assess its statistical significance, let $\Pi_i = [\pi_{i,1}, \pi_{i,2}]$ be a closed interval in (0, 1) and let S_i be the set of all observations from $t = int(\pi_{i,1}T)$ to $t = int(\pi_{i,2}T)$, where $int(\cdot)$ denotes rounding to the nearest integer. The estimated break point is the value t^* from the set S_i that maximizes the likelihood ratio statistic from a test of H_1 against H_0.[7] That is, for every $t \in S_i$ we construct the likelihood ratio statistic corresponding to a test of H_1 against H_0 for that value of *t*, and we take the *t* that produces the largest test-statistic as our estimated break point for MSA *i*.

Assessing the statistical significance of this breakpoint estimate requires knowing the distribution of the supremum of the likelihood ratio statistic as calculated from among the values in S_i. Let $\xi_i = sup_{S_i} LR$ denote this supremum. Andrews (1993) shows that this distribution can be written as

$$P(\xi_i > c) = P\left(sup_{\pi_i \in \Pi_i} Q_1(\pi_i) > c\right) = P\left(sup_{1<s<\lambda_i} \frac{B_1(s)}{s^{1/2}} > c^{1/2}\right), \tag{2}$$

where $B_1(s)$ is the Bessel process of order 1, $\lambda_i = \pi_{i,2}(1-\pi_{i,1})/\pi_{i,1}(1-\pi_{i,2})$, and

$$Q_1(\pi_i) = \frac{(B_1(\pi_i) - \pi_i B_1(1))'(B_1(\pi_i) - \pi_i B_1(1))}{\pi_i(1-\pi_i)}.$$

Direct calculation of the probability in (2) is non-trivial and prior research has relied on approximations that typically are based on simulation or curve-fitting methods (Andrews, 1993; Hansen, 1997). However, Estrella (2003) provides a numerical procedure for calculating exact *p*-values that does not rely on these types of approximations. We use this method to calculate *p*-values for the estimated break point, π_i, for each MSA in the sample.

not be zero. Secular trends in house prices can come from an underlying trend in housing demand as long as the market is not in perfectly elastic supply. It can also arise from trends in physical construction costs under certain conditions.

[6] Note that we are not testing for housing bubble, but are simply using the time series methods to estimate the beginning of local housing *booms*. As such, we are not interested in testing whether price growth patterns are explosive, but only in whether a given housing market has switched from one growth rate regime to another. See Holly, Pesaran, and Yamagata (2010) for a housing bubble test using state-level price data in the U.S. and Giglio et al. (2016) for related tests in Singapore and the UK.

[7] We use the terms supremum and maximum interchangeably in this exposition. Technically, all of the results are in terms of the supremum of the likelihood ratio statistic.

Note that this method does not provide an unbiased estimate of the magnitude of the change in price growth rates at the breakpoint, $d_{i,2}$. Under the null hypothesis that there is no break point, the estimate of $d_{i,2}$ has a nonstandard distribution and OLS estimates of its magnitude will be upwardly biased in absolute value. This can lead to an increased chance of falsely concluding that $d_{i,2} \neq 0$ and is a form of specification search bias arising from the fact that the same data is being used to estimate both the timing and the magnitude of the structural break (Leamer, 1983).

Several approaches for adjusting the estimate of the magnitude of structural break have been suggested and are typically based on simulations of the distribution of $d_{i,2}$ under the null hypothesis of no break point (Andrews, 1993; Hansen, 2000a,b). Our approach to correcting the estimates of $d_{i,t}$ follows the method used by Card et al., (2008) of randomly splitting the underlying sample of housing transactions into two and using one sample to estimate the timing of the boom and the other to estimate the magnitude of price changes around that time. The idea is that if the two subsamples are independent, then estimates of $d_{i,2}$ from the second sample, which was not used to estimate the location of the break point, will have a standard distribution even under the null hypothesis of no structural break in the first sample. In practice, we randomly split our sample of unique houses in two and create separate price growth series for each sample of houses. The first price series is used to estimate the timing of the boom following the method just discussed, while the second is used to analyze the magnitude of price changes following housing booms in neighboring markets.

A strength of the approach described above is that it yields an estimate of a single date for the structural change, which allows us to set up our empirical model in the spirit of an event-study around that date. In fact, it generates a breakpoint estimate regardless of whether the structural break represents a positive or negative change in the price growth rate. In the cases where the estimated break point is either insignificant or implies a negative change in growth rates, we conclude that the market did not have a boom. That is the case for the 13 interior markets shown in the first panel of Fig. 1. In Appendix A1, we show that our estimates of the timing of local booms are robust to an alternative Markov-switching model (Hamilton, 2016) that allows for the estimation of random, as opposed to deterministic, changes in regimes.

Allowing for only one potential breakpoint per MSA could lead to estimation errors for MSAs where two or more breaks are actually present. For all locations where we do find evidence of a statistically significant and positive break point, we also test for the existence of two breaks against the null hypothesis of only one. To do so, we closely follow Bai (1999) and Bai and Perron (1998) and we refer the reader to Appendix A2 for the details of this procedure. Similarly, if we can reject the null hypothesis of one break against the alternative of two, we also estimate and test for the significance of three breaks relative to two.[8]

About half of the MSAs were found to have experienced more than one structural break. However, for many of those cases, the secondary breaks either implied negative changes in growth rates or were positive but economically small. The estimation of a secondary break generally does not displace the location of the break point that is estimated when only allowing for one. Moreover, comparison of histograms of timing of local booms based on one-break or two-break methods lead to similar distributions of local booms over time. A small number of markets were found to have three structural breaks.[9]

Table 1
Summary statistics.

	Mean	Std. Dev.	25th percentile	75th percentile
Price index	134.41	47.92	98.94	162.30
Average income ($1000′s)	88.90	32.21	65.84	103.03
Percent minority	20.99	13.07	10.59	29.89
Percent speculators	5.53	4.25	2.56	6.95
Percent government insured	14.79	12.59	3.24	23.71
Percent subprime	9.79	8.21	3.42	14.35
Average LTV	0.66	0.11	0.59	0.74
Average square footage (1000′s)	1.65	0.21	1.56	1.74
Unemployment rate	6.86	3.65	4.38	8.51
Net migration	85.94	2991.06	−213.00	404.00
Number of observations	6225			

Note: This table presents descriptive statistics for the primary analysis sample. The level of observation is the MSA-quarter and the data run from the first quarter available for each MSA until the fourth quarter of 2011.

We use all estimated breaks in our empirical analysis below and always distinguish between positive and non-positive break points. Fig. 3 plots the distribution of all positive and statistically significant break points across MSAs and echoes the conclusion from Fig. 1. While there is an increased concentration of booms beginning in the mid-2000′s approaching the peak in national aggregate prices, there still is a great degree of heterogeneity in the timing of the beginning of local housing booms, which we use to estimate spillover effects below.

3. Data

Our house price data come from DataQuick, a private data vendor that collects the universe of housing transactions from county recorder's offices in markets across the country. The sample used is for 94 metropolitan areas, with information on over 23 million individual transactions ranging from the first quarter of 1993 through the last quarter of 2011. We randomly split the sample into two, and in each subsample, we create a constant quality quarterly price index for each MSA.[10] One of these indices is used to estimate the timing of the boom in each market and the other is used to assess how prices change following the beginning of the boom in neighboring markets. The mean, standard deviation, and interquartile range for the price index we use to measure price changes following neighboring market booms are reported in the first row of Table 1.

We also create a number of variables to measure fundamentals that may contribute to local housing market spillovers. These are reported in subsequent rows of Table 1. We consider three types of fundamentals: (1) demand shifters, such as the average income of mortgage applicants, MSA-level unemployment rates, and net migration flows; (2) buyer characteristics and property traits, including the percentage of speculators, the percentage of minority buyers and the average square footage of transacted housing units; and (3) credit market conditions, measured by the average loan-to-value ratio of home purchases, the percentage of mortgages originated by subprime lenders and those insured by the FHA or VA.

[8] As noted in Appendix A2, data limitations prevent us from being able to test for multiple breaks in some markets. In cases where we do not have enough data to test for the existence of two breaks, we use the estimates from the single break procedure. Similarly, in cases where we can reject the null of one break relative to two but do not have enough data to test for the existence of three breaks we use the estimates from the two-break procedure.

[9] We also experimented with versions of the Markov-switching model that allow for more regimes, in the same spirit of the multiple breakpoints described above. Overall, estimates seem to line up with our multiple breaks, but the results are noisier given that our sample sizes for each MSA are quite small. Those results are available upon request.

[10] We create a MSA-level constant quality house price series by quarter using hedonic regressions. Price, in logarithmic form, is modeled as a function of the square footage of the home entered in quadratic form, the number of bedrooms, the number of bathrooms, and the age of the home. We also created a version of the Case and Shiller (1987) repeat sales price index for 14 Case-Shiller markets that overlap with the DataQuick files, and found that the simple correlation of appreciation rates on the two different indexes based on DataQuick is usually higher than 0.9. We employ hedonic price indexes because their data requirements are much less onerous.

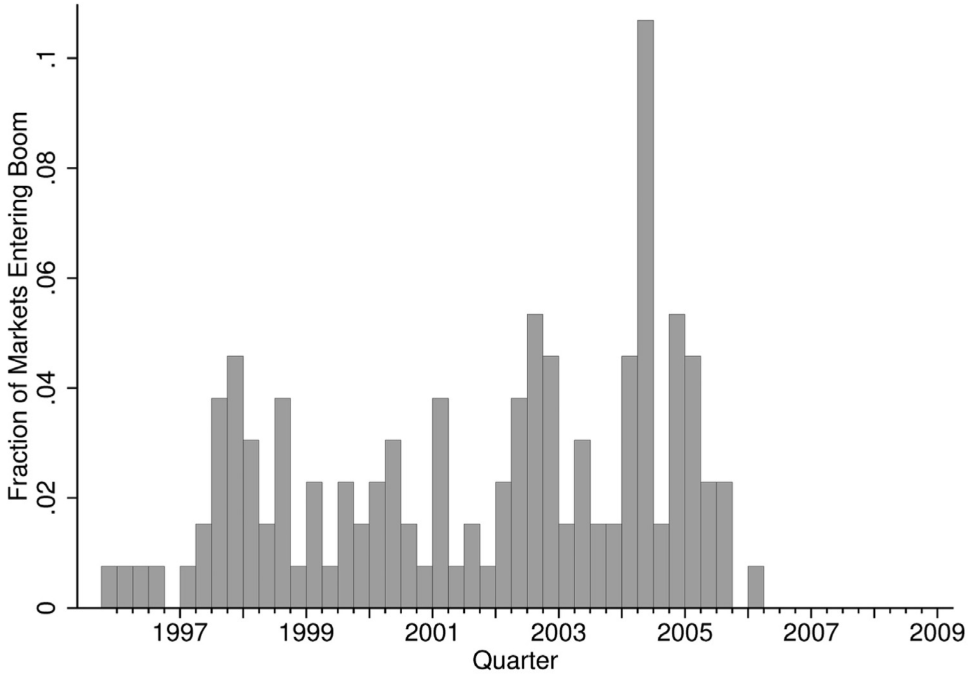

Fig. 3. Distribution of beginning of booms across MSAs.
Note: Figure plots the fraction of all markets in the analysis sample that enter a housing boom in the indicated quarter. A market is defined as entering a housing boom in a given quarter if we find a positive and statistically significant structural break in that market's annualized house price growth series for that quarter. Structural breaks are estimated as described in Section 2.

To construct many of the demographic measures of homebuyers, we merge the DataQuick files with Home Mortgage Disclosure Act (HMDA) data, which provide information on the income and race of all mortgage applicants. In each period, we calculate the average income of all local loan applicants as reported in HMDA.[11] Similarly, the "Percent Minority" variable reflects the fraction of African-American and Hispanic loan applicants as coded in the HMDA files. Because these measures reflect the characteristics of all mortgage applicants, and not only those who purchase a home, we take them to be an accurate description of the race and reported income of potential homebuyers in each market.

MSA-level unemployment rates come from the Bureau of Labor Statistics' Local Area Unemployment Statistics (LAUS) series, and net migration flows are calculated using data on county-to-county migration patterns provided on an annual basis by the Internal Revenue Service.

The variable "Percent Speculators" refers to the fraction of transactions involving a speculator on either the buyer or the seller side of the transaction. We identify speculators in one of two ways. First, we follow Chinco and Mayer (2014) who reasoned that since speculators would not be living in the purchased unit, they would have their tax bills sent to another address. We compare the precise street address of the housing unit with the address to which the tax bill is sent – the 'Tax Address' in the DataQuick files. Whenever the two are appreciably different, we call that purchaser a speculator.[12] The second way we identify whether a purchaser is a speculator is by whether the buyer has a name that is a business. This includes corporate or commercial names that include LLC or INC in them, homebuilders, or trusts (especially mortgage-backed securities trusts that are typically identified by a four-digit number in their names).[13]

Credit market variables include the average loan-to-value ratio (LTV) among homebuyers in DataQuick (including zeros for all-cash buyers), the fraction of FHA/VA-insured loans, and the fraction of subprime loans. To calculate the share of subprime loans, we compare the names of the underlying mortgage lenders from the DataQuick files to the list of subprime lenders compiled by the Department of Housing and Urban Development (HUD) as well annual lists of the top twenty subprime lenders from 1990-onward contained in a publication now called *Inside Mortgage Finance*.[14] When calculating the subprime share, we exclude borrowers who took out loans insured by the Federal Housing Administration (FHA) or Veterans Administration (VA), even if the lender is on one of the subprime lists. While FHA/VA-insured loans have many subprime-like traits, Ferreira and Gyourko (2015) document that their market shares over time are quite different.

4. Econometric model and estimates

4.1. The extensive margin: is there a timing effect?

To study the role of spatial spillovers during local housing booms, we begin by investigating whether the likelihood that a given market enters its housing boom in a particular period is related to the timing of recent booms in nearby neighboring markets. While a naive visual inspection of the results reported in Fig. 1 would suggest a potential role for such spillovers, it obviously is important to control for aggregate and regional

[11] HMDA income may not represent a precisely accurate measure of true homebuyer income during some parts of our sample period. To the extent that HMDA incomes are over-reported, changes in this variable may be better thought of as proxying for credit market conditions and changes in lending standards.

[12] By appreciably different, we generally mean that more than one number in the street address before the zip code differs.

[13] Other research has identified speculators by whether they 'flip' properties quickly (e.g., Bayer et al., 2011). We also investigated those cases, and found that many of them were already encompassed by our measures of tax address and names of business.

[14] This publication claims to capture up to 85 percent of all subprime originations in most years. Previously, it was named *B&C Mortgage Finance*. See Chomsisengphet and Pennington-Cross (2006) for more details on these lenders and lists.

trends that could be affecting all markets simultaneously. To do so, our approach leverages the heterogeneity documented above in the timing of those booms across local markets.

We begin by calculating pairwise straight-line distances between each of the 362 MSAs defined by the 2000 Census. For a given focal MSA, m, we then rank all other MSAs from nearest to farthest according to these distances and estimate variants of the following regression:

$$Boom_{m,d,t} = \gamma_{d,t} + X'_{m,t}\beta + \sum_{\rho=-5}^{5}\sum_{b\in\mathcal{B}} \theta^{b}_{\mathcal{N},\rho} \cdot 1\left\{\left\lfloor \frac{t - t^{*}_{n\in\mathcal{N}}}{4} \right\rfloor + 1 = \rho\right\} + \epsilon_{m,d,t}. \tag{3}$$

where $Boom_{m,\,d,\,t}$ is an indicator for whether MSA m located in census division d experienced a statistically significant and positive structural break in its house price growth series (i.e. entered a boom) in quarter t, $\gamma_{d,\,t}$ is a set of census division-by-quarter fixed effects, and $X_{m,\,t}$ is a set of possibly time-varying controls.

The third term on the right-hand side of Eq. (3) with the multiple summation signs contains the primary variables of interest. For a given set of neighboring markets $\mathcal{N}$ (e.g., the 5 closest neighbors), we construct a series of "relative year" indicator variables which identify whether the current quarter occurs ρ years before or after a quarter in which any of the neighbors in that set experienced a structural break in its house price growth series ($t^{*}_{n\in\mathcal{N}}$).[15] If the neighboring market had more than one break point, these indicators will mark all quarters that occur ρ years before or after *any* of that market's break points. To distinguish between neighboring markets that had a housing boom and those that did not, we do this separately for break points that are positive and statistically significant ($b = boom$), and those that are either negative or insignificant ($b = no\ boom$). The coefficients, $\theta^{boom}_{\mathcal{N},\rho}$, on the relative year indicator variables for positive and statistically significant break points describe how the likelihood that the focal market enters a housing boom evolves over the course of its neighboring markets' housing booms.[16]

The indicator variables are defined so that Relative Year 0 denotes the 12-month period prior to the estimated break point for the relevant neighbor.[17] Relative Year 1 then includes the quarter in which the break point occurred as well as the subsequent three quarters. Relative Year 0 is the omitted category in all specifications so that the coefficients should be interpreted as the difference between the probability that the focal market enters a housing boom in given year relative to the probability in the year prior to when the neighboring market entered its boom. We report results for the three years preceding the estimated break point and for four years after that time (Relative Years −2 through + 4).[18] This allows us to see whether there are pre-trends and to track the build-up of the neighbor's boom after it starts.

The results in Table 2 allow for spillover effects from the 5 closest neighbors. We group the 5 closest neighbors together in our analysis of the extensive margin largely for reasons of statistical power.[19] However, in unreported results, we find that the effects are qualitatively similar when we allow for spillovers from just the closest neighbor or consider larger groupings such as the closest 10.[20] Our analysis of the intensive margin below will also consider alternative groupings and yield similar results.

The first column of Table 2 reports results from an unconditional version of Eq. (3). Note that there is virtually no pre-trend, but that the probability of the focal market booming jumps markedly in the year that any of its five nearest neighbors enters their booms. The coefficient on Relative Year 1 indicates that the probability jumps by over 5 percentage points. This is very large economically, given that the overall probability of having a boom in any given quarter is only about 2%.

Adding time controls (quarter fixed effects in column 2) reduces the coefficient considerably, but it remains statistically and economically significant. Controlling for national trends does not eliminate the intuition arising from Fig. 1. However, finer geographic controls weaken the results further. Column 3 adds regional controls by interacting census division with quarter dummies, and column 4 adds the full set of focal market fundamentals described in Section 3.[21] The point estimates for Relative Year 1 fall by about two-thirds, and the coefficients are no longer statistically significant. However, the magnitude of the estimates are still economically quite large given the 2% overall probability of a boom beginning in any quarter. Moreover, there is now a marginally significant impact in Relative Year 2. These results suggest that it may take some time (an additional year) for booms in a nearby neighbor to influence the likelihood of the focal market itself starting to boom. This pattern is robust to alternative functional forms. This is shown in columns 5 and 6 of Table 2, which report marginal effects from probit and logit specifications that are directly analogous to the specification in column 4. The same time pattern holds and the spillover effects are, if anything, slightly higher in these specifications.

In sum, Table 2's results are more consistent than not with there being spatial spillovers on the extensive margin. The point estimates themselves are large, and imply that the probability of the focal market booming roughly doubles within the next one to two years if any of its 5 closest neighboring markets enters a boom this year. The fact that statistical significance weakens and becomes marginal as finer geographical controls are included is likely an issue of sample size. The spatial and temporal heterogeneity in the time lines of local market booms is much greater than exists in studies of contagion in other asset markets (e.g. stock market and currency crises), but we still are limited to only 94 individual housing markets.

4.2. The intensive margin: is there a price effect?

Having documented evidence of spillovers on the extensive margin, we now turn to the intensive margin and ask whether there are any

[15] The notation $\lfloor x \rfloor$ is used to represent the floor function, which is the largest integer smaller than x. We need to divide the number inside the bracket by 4 because the underlying data is by quarter, but we estimate the spillover effects in terms of relative years.

[16] While we construct our distance rankings using the entire set of 362 MSAs, our price data only covers 94 of those markets. In cases where neighbor n is not included among those 94 markets, we set all relative year dummies for that neighbor to zero and include a dummy variable in $X_{m,\,t}$ denoting that neighbor n is missing for focal market m. We are missing price data for the nearest neighbor for 18 of our 94 focal markets. Our results are qualitatively similar when we drop MSAs with a missing nearest neighbor and also when we calculate distances and construct the ranking of nearest neighbors using only the 94 MSAs for which we have price data.

[17] We work with 12-month periods because there is noise in the quarterly data that is not due solely to error in the estimation of the break point. For example, it is common for there to be at least a one quarter difference between the time that a transactions price is agreed upon and when the actual closing occurs. In addition, we know that prices in housing markets do not follow a random walk, but move slowly and are strongly positively correlated over short horizons (Case and Shiller, 1987, 1989).

[18] The MSA samples are almost equally balanced using that time span. The coefficients for relative years outside this window are based on a smaller number of MSAs since not all markets entered their booms at the same time and our data only go back to 1993. We estimate separate coefficients for up to five years preceding and following the estimated breakpoint. Reporting those coefficients does not introduce any net new relevant information.

[19] For example, given that the baseline probability of booming in any given quarter is only 2%, an analysis that considered only the closest neighbor would generate a severely underpowered regression that contained many zeros on both the left and right-hand side.

[20] For example, the estimates in the unconditional regressions reported in column 1 of Table 2 are nearly identical if we consider only the first neighbor or group neighbors 1–10 together.

[21] The full set of controls includes mortgage applicant income, migration into the focal market, subprime and FHA/VA lending market shares, percentage of speculative buyers, percentage of minority buyers, average LTV at origination, average square footage of purchased homes, and the local unemployment rate

Table 2
The impact of nearby neighbors' housing booms on the probability of the focal market entering a boom.

	(1) Unconditional	(2) National trend controls	(3) Regional trend controls	(4) Focal market Fundamental controls	(5) Probit	(6) Logit
Relative Year = −2	0.001	0.001	−0.011	−0.013	−0.031	−0.032
	(0.007)	(0.009)	(0.011)	(0.011)	(0.024)	(0.025)
Relative Year = −1	0.004	0.005	0.007	0.005	0.006	0.006
	(0.010)	(0.007)	(0.007)	(0.007)	(0.015)	(0.017)
Relative Year = 1	0.052***	0.036***	0.012	0.011	0.009	0.009
	(0.010)	(0.009)	(0.012)	(0.012)	(0.020)	(0.021)
Relative Year = 2	0.011	0.013*	0.019*	0.019*	0.036**	0.037**
	(0.007)	(0.008)	(0.010)	(0.010)	(0.017)	(0.018)
Relative Year = 3	0.004	0.007	0.004	0.004	0.015	0.017
	(0.007)	(0.007)	(0.009)	(0.009)	(0.025)	(0.026)
Relative Year = 4	−0.005	−0.004	−0.004	−0.003	−0.014	−0.016
	(0.006)	(0.007)	(0.010)	(0.010)	(0.027)	(0.027)
Quarter FEs		X				
Quarter-by-division FEs			X	X	X	X
Focal market fundamental controls				X	X	X
Number of observations	6225	6225	6225	6225	6225	6225

Note: Each cell reports the coefficient estimate on the dummy variable for the indicated relative year of the closest 5 geographic neighbors. Relative Year 0 denotes the 12-month period preceding the break point of the neighboring MSA and is the omitted category. Coefficients are reported only for relative years that are associated with positive and statistically significant break points. Standard errors are clustered at the census division by year level and are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively. See the discussion of Eq. (3) in the main text for exact details on the specification.

changes in price levels around the timing of neighboring market boom. Our primary interest is in gauging how prices in a given focal market, m, evolve as its neighboring markets enter their respective housing booms. To measure these effects, we estimate versions of Eq. (3) that use log focal market prices, $\log(P_{m,d,t})$, as the outcome rather than an indicator for whether the focal market enters a boom.

Table 3 reports our core results from this exercise. The top panel shows the coefficient estimates on the relative year dummies associated with positive and statistically significant break points. The bottom panel reports the analogous coefficients associated with statistically insignificant or negative break points. For the sake of comparison, we continue to group the five closest neighbors together in this table, but will consider alternative groupings below.

Column 1′s results are from an unconditional specification that simply regresses the focal market's (log) house price on the timeline of its nearest neighbors' booms without any other controls. These results highlight the strong trend growth in house prices during our time span, especially among markets whose near neighbors experienced a boom (top panel).

The second column removes the effect of this overall trend growth by including a series of quarter dummies, which soak up aggregate trends, and four lags of (log) focal market prices, which control for short-run persistence in price growth at the local level. We intentionally do not control for contemporaneous focal market fundamentals in this baseline specification because they could represent intermediate outcomes through which the spillover effect may operate. In Section 4.3 below, we will explore this possibility by testing whether their inclusion in Eq. (3) mitigates the estimated spillover effect.

Note that this baseline specification of Eq. (3) yields clear evidence of spillover effects that only manifest if the one of the nearest five neighboring markets experienced a statistically significant positive boom. In the top panel, the coefficients for the two years prior to the neighbor's boom (i.e., Relative Years −1 and −2) become very small economically and are statistically insignificant. However, prices in the focal market jump sharply by roughly 0.9% beginning immediately the year that one of the neighboring markets enters a boom. Prices then stay higher for another three years throughout our reported timeline. There is no such evidence of this pattern in the bottom panel, which is reassuring given that we should not expect to find evidence of spillovers if the neighboring market did not experience a housing boom.

This pattern survives in column 3, which also includes a full set of relative year fixed effects for the focal market itself. By including the own-market relative years in this specification, we are controlling for all average factors that could explain the price variation around a housing boom in the focal market itself.[22] Thus, the spillover effect we document appears to exist even beyond the average price path experienced over the course of a local boom. Column 4 controls for even more granular common aggregate trends across markets by including census division-by-quarter fixed effects. The coefficients are very similar to those in column 3, which suggests that unobserved common factors across markets within a region are unlikely to be driving the estimated spillover effect.

Thus far, we have focused on the combined spillover impact from the 5 closest neighbors. In Fig. 4 we report additional results from a more flexible specification that is analogous to Eq. (3) but which allows for separate spillover effects from each of the 10 closest neighbors independently. The first panel in the figure reports coefficient estimates and 95% confidence intervals for the relative year dummies associated with positive and statistically significant break points for the closest neighbor only. The remaining panels plot the analogous estimates for neighbors 2–10. Including separate relative year dummies for neighbors 2–10 not only allows us to see if there are meaningful spillover effects beyond the closest 5 neighbors, but also serves as a useful control for differential regional trends that are not entirely picked up by the census division-by-quarter fixed effects.

This specification is considerably more flexible than our main specification in Table 3, which causes us to lose some statistical power. However, several patterns are apparent. Neighbors 1, 2, 3 and 5 all have at least one statistically significant effect in the set of post-boom relative years, and their magnitudes are relatively similar. Meanwhile, we almost never find statistically significant coefficients for Neighbors 6 and above (the only exception is neighboring market 8). These patterns suggest that spillover effects arise primarily from very close neighbors. Moreover, the patterns for neighbors 1–5, which are generally indicative of positive spillover effects but somewhat noisy, justify our decision to pool these neighboring markets together in the main analysis.

[22] Controlling for the focal market's own cycle also helps account for the potentially higher volatility in prices when the boom starts in that market (Forbes and Rigobon (2002)).

Table 3
The impact of nearby neighbors' housing booms on log focal market price.

	(1) Unconditional	(2) Basic trend controls	(3) Focal market Timeline controls	(4) Local trend controls
	Panel A. Positive and Significant Break Points			
Relative Year = −2	−0.074***	−0.002	−0.001	−0.001
	(0.022)	(0.002)	(0.002)	(0.002)
Relative Year = −1	−0.081***	0.000	0.001	0.001
	(0.021)	(0.001)	(0.001)	(0.002)
Relative Year = 1	−0.025	0.009***	0.007***	0.006***
	(0.022)	(0.002)	(0.002)	(0.002)
Relative Year = 2	0.053*	0.008***	0.005**	0.005**
	(0.031)	(0.002)	(0.002)	(0.002)
Relative Year = 3	0.103***	0.009***	0.006***	0.009***
	(0.032)	(0.002)	(0.002)	(0.002)
Relative Year = 4	0.129***	0.005*	0.003	0.007***
	(0.029)	(0.003)	(0.002)	(0.003)
	Panel B. Insignificant or Negative Break Points			
Relative Year = −2	−0.082***	0.005	0.003	0.004
	(0.029)	(0.003)	(0.003)	(0.004)
Relative Year = −1	−0.087**	0.005	0.003	0.007
	(0.038)	(0.003)	(0.003)	(0.005)
Relative Year = 1	0.018	−0.003	−0.002	−0.000
	(0.041)	(0.004)	(0.004)	(0.004)
Relative Year = 2	0.082**	−0.003	−0.001	0.002
	(0.032)	(0.003)	(0.003)	(0.003)
Relative Year = 3	0.079**	−0.004	−0.001	−0.001
	(0.031)	(0.003)	(0.003)	(0.003)
Relative Year = 4	0.124***	−0.002	0.000	0.001
	(0.036)	(0.004)	(0.004)	(0.005)
Quarter FEs		X	X	
Four lags of focal market price		X	X	X
Focal market relative year FEs			X	X
Quarter-by-division FEs				X
Number of observations	6225	5849	5849	5849

Note: Each cell reports the coefficient estimate on the dummy variable for the indicated relative year of the closest 5 geographic neighbors. Relative Year 0 denotes the 12-month period preceding the break point of the neighboring MSA and is the omitted category. Coefficients are reported separately for relative years that are associated with positive and statistically significant break points (Panel A.) and those that are not (Panel B). Standard errors are clustered at the census division by year level and are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively. See the discussion of Eq. (3) in the main text for exact details on the specification.

Are these intensive margin spillover effects large or small? One way to gauge the economic magnitude of the effect we estimate is by calculating an elasticity of focal market housing price growth with respect to neighboring market price growth. The starting point of this exercise is to estimate a version of Eq. (3) that uses the log price of the nearest neighbor as the dependent variable to determine the magnitude of the change in the neighbor's price upon entering its own boom. Those results (which are available upon request) show that prices in the neighboring market jump discretely by 2.4% in the first year of the boom. By Relative Year 3, prices are 6.5% higher depending upon the specification.

An upper bound on the implied elasticity can be computed by using only the estimates of price changes in the focal market during first year of the neighbors' boom given in column 4 of Table 3. This yields an elasticity estimate of 0.25 (i.e., dividing 0.006 from the top panel of Table 3 by 0.024). A smaller elasticity of 0.09 results if we consider cumulative price changes through the third year after the neighbor's boom. However, the economic interpretation of the spillover estimates for the years after the beginning of the boom can be complicated because of potential feedback effects, which may be less of a concern in the first year of the boom.[23] Nonetheless, our preferred specification provides reasonable magnitudes for the elasticity ranging from 0.09 to 0.25.

4.3. Are the spillovers on price fundamentally based?

Table 4 reports estimates from specifications that include controls for various focal market fundamentals to see whether these fundamentals may help to explain the spillover effects that we estimate. In this table, we continue to pool neighbors 1–5 together, but the results are similar when we consider only the nearest neighbor. The first column's results are from a model that is directly analogous to that in column 4 of Table 3, but which also includes focal market fundamental controls on the right-hand side. The full set of controls includes mortgage applicant income, migration into the focal market, subprime and FHA/VA lending market shares, percentage of speculative buyers, percentage of minority buyers, average LTV at origination, average square footage of purchased homes, and the local unemployment rate. If these fundamentals can partially explain the spillover effect, then we should expect the estimates on the relative year dummies to exhibit less of a jump in the year that the nearest neighbor enters its boom. This is not the case, as these new estimates are very similar to the baseline results presented in column 4 of Table 3. This suggests that the spatial spillovers we have identified are not being transmitted via the fundamentals considered here.[24]

[23] These feedback effects may not even play a major role in subsequent years, especially if only 10% or less of the main spillover effect propagates across close neighbors. Nonetheless, the spillover estimates for Relative Years 2 and 3 are better thought of as reduced form estimates that include the impact of contemporaneous spillovers but that also embed a share of spillovers associated with the complete path of price appreciation since the beginning of the boom.

[24] In addition to simply controlling for fundamentals, we also directly investigated whether there were meaningful changes in several key focal market fundamental factors around the time of the neighbor's boom. Results from this analysis

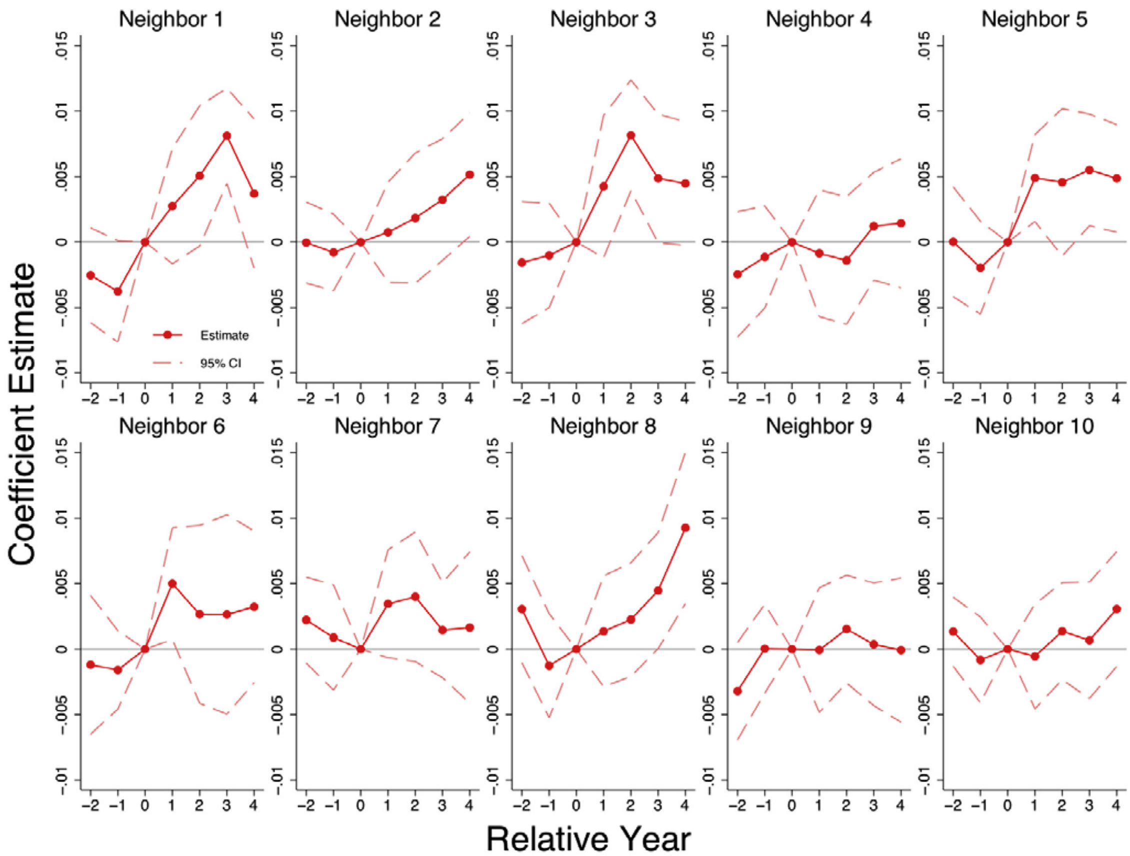

Fig. 4. Estimated price spillovers for neighbors 1–10.
Note: Figure plots the estimated coefficients and 95% confidence interval for the relative year dummies for the closest 10 neighbors. Estimates come from a version of Eq. (3) that allows for separate effects from each of the 10 closest neighbors. See text for a detailed description of the regression.

Table 4
The impact of nearest neighbor's housing booms on log focal market price controlling for focal market fundamentals.

	(1) Focal market Fundamental controls	(2) Focal market Income leads	(3) Focal market Fundamental leads
Relative Year = −2	−0.002	−0.001	−0.002
	(0.002)	(0.002)	(0.002)
Relative Year = −1	0.000	0.001	−0.000
	(0.002)	(0.002)	(0.002)
Relative Year = 1	0.005***	0.005***	0.005***
	(0.002)	(0.002)	(0.002)
Relative Year = 2	0.004*	0.004**	0.003
	(0.002)	(0.002)	(0.002)
Relative Year = 3	0.007***	0.007***	0.007***
	(0.002)	(0.002)	(0.002)
Relative Year = 4	0.006**	0.006**	0.005**
	(0.002)	(0.002)	(0.002)
Quarter-by-division FEs	X	X	X
Four lags of focal market price	X	X	X
Focal market relative year FEs	X	X	X
Focal market fundamental controls	X	X	X
Four leads of focal market mean income		X	X
Four leads of all fundamental controls			X
Number of observations	5849	5473	5473

Note: Each cell reports the coefficient estimate on the dummy variable for the indicated relative year of the closest 5 geographic neighbors. Relative Year 0 denotes the 12-month period preceding the break point of the neighboring MSA and is the omitted category. Coefficients are reported only for relative years that are associated with positive and statistically significant break points. Standard errors are clustered at the census division by year level and are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively. See the discussion of Eq. (3) in the main text for exact details on the specification.

Table 5
Heterogeneity in spillovers.

	Distance to nearest neighbor		Relative size (Population)		Nearest neighbor Boom size		Focal market Supply elasticity	
	(1) Nbr. Close	(2) Nbr. Far	(3) Nbr. Larger	(4) Focal Larger	(5) Small Boom	(6) Large Boom	(7) Inelastic	(8) Elastic
Relative Year = −2	−0.000	−0.001	0.000	−0.002	−0.002	−0.001	0.001	−0.003
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Relative Year = −1	0.000	−0.003	−0.001	−0.002	−0.003	−0.000	−0.003	−0.002
	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
Relative Year = 1	0.008**	0.004	0.003	0.008***	0.002	0.010***	0.005**	0.008***
	(0.003)	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)
Relative Year = 2	0.008***	0.006**	0.004	0.008***	0.005**	0.007**	0.004	0.008**
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)	(0.004)
Relative Year = 3	0.009***	0.010***	0.013***	0.007**	0.006***	0.012***	0.006***	0.009***
	(0.003)	(0.002)	(0.003)	(0.003)	(0.002)	(0.003)	(0.002)	(0.003)
Relative Year = 4	0.002	0.006*	0.003	0.003	0.001	0.005	0.008***	0.001
	(0.004)	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)	(0.005)
Quarter-by-division FEs	X		X		X			
Four lags of focal market price	X		X		X			
Focal market relative year FEs	X		X		X			
Number of observations	5849		5849		5849			

Note: Each cell reports the coefficient estimate on the dummy variable for the indicated relative year of the closest geographic neighbor. As described in detail in the main text, relative year dummies are further interacted with dummies for the dimension of heterogeneity indicated in the column headings. Relative Year 0 denotes the 12-month period preceding the break point of the neighboring MSA and is the omitted category. Coefficients are reported only for relative years that are associated with positive and statistically significant break points. Standard errors are clustered at the census division by year level and are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

Thus far, we have abstracted from expectations of future fundamental factors, effectively treating actors as myopic. The second column in Table 4 begins to address this issue by adding four leads of focal market income to the previous specification, essentially presuming that local residents can predict the path of local incomes over the next four quarters. Including this proxy for expectations does not change the estimated spillover effect. The third and final column reports results from adding four quarterly leads of all fundamentals included in Column 1, not just income. Once again, the magnitudes of the point estimates as well as the time pattern are unchanged, leading us to conclude that the intensive margin spillover effects we document do not appear to be related to measurable economic fundamentals.

4.4. Heterogeneity in the spillover effect

Our final tests look for heterogeneity in the average spillover effect along a number of dimensions. Given the relatively small number of degrees of freedom provided by our 94 metropolitan area sample, heterogeneity tests are not likely to have much power, but they still yield interesting insights as Table 5 shows.

For example, a natural extension of the result that spillover effects are due primarily to geographically close neighbors is to ask whether the strength of that impact weakens with distance. The first two columns of Table 5 indicate that the answer is ‘yes’ with respect to the impact of the timeline of the boom of the physically closest neighbor. Those results are the output from a regression like that in column 4 of Table 3, but which includes only the effect of the closest neighbor and further interacts that nearest neighbor’s relative year dummies with an indicator for whether that neighbor is more or less than the median distance of about 40 miles away from the focal market.[25] The estimates are not always terribly precise, but the point estimates suggest that proximity does matter, with physically closer neighbors having larger spillover effects just after booms begin.

It is also natural to ask whether spillover impacts depend upon the relative sizes of the focal and neighbor markets. To investigate this, we again estimated a regression for spillover effects from the closest neighbor only, but this time allowing the effect to vary based on whether the focal or neighboring market had a larger population (as of the 2000 Census). The estimates reported in the third and fourth columns of Table 5 are not consistent over time. Larger neighbors have appreciably larger impacts by Relative Year 3, but this is not the case in Relative Years 1 or 2. Hence, the evidence on size is not nearly robust enough to conclude that it is an important source of heterogeneity in the spillover effect.

A third dimension of heterogeneity investigated is the magnitude of the nearest neighbor’s boom. We classified the statistically significant positive booms as large or small based on whether the magnitude of the jump in price growth rates implied by the structural break point estimation procedure described in Section 2 was larger or smaller than the median implied change in growth rates (which was about 10%).[26] The point estimates reported in columns 5 and 6 of Table 5 indicate that larger booms are indeed associated with larger spillover effects, but the standard errors are too large to draw definitive conclusions on this margin.

The final two columns of Table 5 investigate whether there was any heterogeneity in the spillover effect by the degree of the focal market’s

are omitted here in the interest of space, but are available on request. In that work, we directly examined five fundamental factors, each of which has received prominent mention in previous academic research or by policy makers and the popular press. These factors were: focal market income, the percent of sales due to speculators, net migration flows into the focal market, the fraction of new mortgage originations by subprime lenders, and the fraction of new mortgage originations insured by the FHA or VA. The results show that these factors generally were not found to exhibit large increases (or decreases) around the time that neighboring markets enter their boom. For example, focal market income (which is defined as the average income reported by all mortgage applicants in that market and quarter) is higher in Relative Year 1, but it typically is not statistically distinguishable from its pre-boom level. Thus, there is no convincing evidence that spillovers operate via changes in focal market income. Similar conclusions pertain to each of the other four variables mentioned above.

[25] The interquartile range of distances between neighboring markets runs from 30 to 56 miles, so there is not much variation for much of the sample. We also experimented with alternative groupings such as dividing markets into whether their nearest neighbor was less than 30 miles away, from 30 to 60 miles away, and greater than 60 miles away. The results were not materially different from those reported here.

[26] For example, using the notation from Section 2 in the one-break case, the implied magnitude of the change in price growth rates is given by the difference $d_{2,i}(\pi_i) - d_{1,i}(\pi_i)$.

elasticity of housing supply. For this test, we split the focal MSAs into two groups according to Saiz's (2010) elasticity estimates.[27] Prices jump more in Relative Year 1 (and in the subsequent two years) in the elastically supplied markets, but once again the standard errors are large enough that we cannot conclude the point estimates are different at standard confidence levels. Even so, this result is interesting in light of the difficulty that the literature has faced in explaining large price movements in elastically-supplied markets (Nathanson and Zwick, 2018), and suggests that some of the disproportionately large swings in prices in these markets during the most recent cycle may have been due to a larger spillover effect.[28]

5. Conclusion

The temporal patterns by which housing booms began in different markets suggest a potential role for spatial spillovers in helping to foment the great American housing boom. We first analyzed the extensive margin to determine if a nearby neighbor beginning to boom raised the probability of a focal market itself booming. The point estimates are economically large—if your close neighbors start to boom, the probably that your market will also begin to boom roughly doubles. However, this impact is noisy and declines in magnitude when finer geographic and time fixed effects are included. With our limited number of markets, statistical power is low given the need to control for other common factors. While these results should be interpreted cautiously, we believe that they are nonetheless consistent with the presence of economically meaningful spatial spillover on the extensive margin.

We also provide evidence of spatial spillovers on the intensive margin. Price *levels* increase modestly around the time that neighboring markets enter their respective housing booms. Our rich data and the heterogeneity in the timing of housing booms allows us to estimate these price impacts with more confidence than is typical in analyses such as this. While the results are statistically significant, and indicate some role for spatial spillovers in the last cycle, they are not large enough economically to account for the bulk of the last boom.

The price effects we identify are not driven by shocks to income, migration patterns, or changes in lender behavior (i.e., they are not associated with contemporaneous changes in subprime mortgage activity). They also appear unrelated to fundamentals or expectations of fundamentals, which suggests some role for non-rational forces. This as an interesting and potentially important area for future research.

Appendix A1. Markov switching estimates

To be sure that our approach to estimate the timing of the begging of housing boom was robust, we also estimated a simple Markov Switching model that allows for probabilistic switching between two growth rate regimes, closely following Hamilton (2016). We modeled house price growth rates as a function of the constant (intercept that may change across regimes) and an error term. The underlying goal of this model is to estimate the transition probabilities between house price growth rate regimes at each point in time rather than taking a stand on any specific date.

In Appendix Fig. A.1 below, we show how the average estimated probabilities of being in state 2 change before and after the timing of the single breakpoint estimates that are used in our deterministic method (marked by the vertical line in the figure). To construct this figure, we estimate transition probabilities in each quarter separately by MSA and then average those transition probabilities within relative year across MSAs, where the relative years are determined using the timeline implied by our deterministic break point estimates.

Before year 0, average probabilities of being in regime 2 hover around 10%. They quickly and non-linearly jump to around 90% right after the beginning of a boom. Fig. A.2 also plots similar estimates for the MSAs that have non-statistically significant breaks. Here we find almost no changes in the average estimated probabilities before and after a break. This strong relationship between both sets of estimates give us confidence that our preferred method is robust to the assumptions of the Markov Switching model.

Appendix A2. Estimating multiple breakpoints

In estimating the break points, we allow for the possibility that a given market might experience more than one housing boom during the course of our sample period. Our method is recursive in that we first test for the existence of one break point against the null hypothesis of zero. Given the existence of at least one break point, we can then test the hypothesis of $m + 1$ break points against the null of m using the results from Bai (1999). Bai and Perron (1998) show that the test for one break is consistent in the presence of multiple breaks, which is what allows for this sequential estimation procedure.

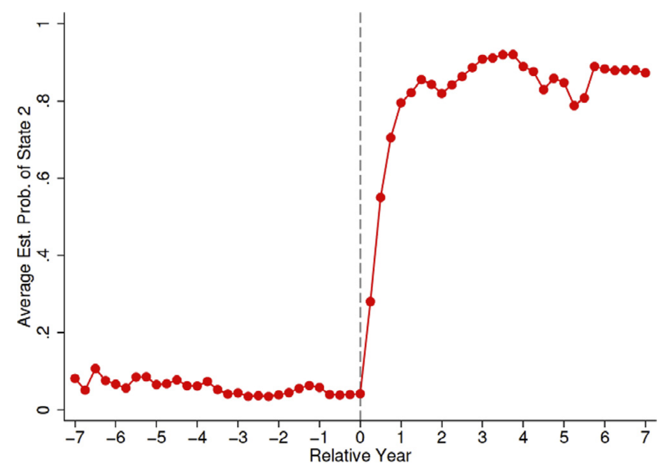

Fig. A.1. Estimated Markov switching probabilities by relative year around positive and statistically significant break points.

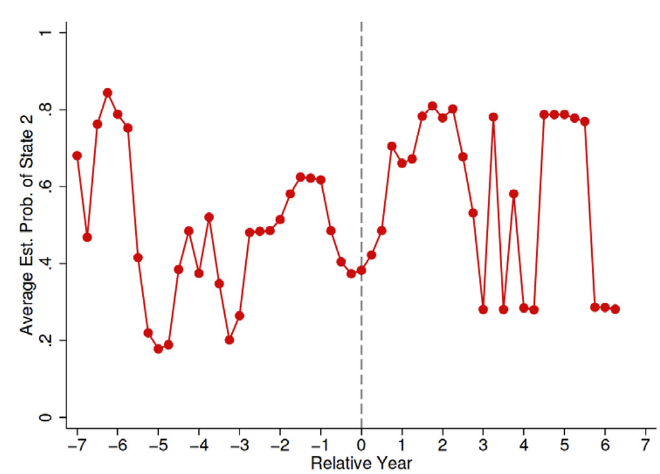

Fig. A.2. Estimated Markov switching probabilities by relative year around statistically insignificant break points.

[27] Saiz's (2010) supply elasticity estimates are available for only 78 of our metropolitan areas, so we start with a smaller sample for this particular analysis.

[28] Of course, this is not the only possible explanation and caution is in order against over-interpreting this particular result. Some of the small effect for inelastically supplied markets may arise from the fact that a significant share of them is comprised of larger coastal metropolitan areas with relatively small neighbors. So, at least some of the variation we document here could have been driven by the size results just discussed.

More specifically, let $0 < \varphi_{i,1} < \cdots < \varphi_{i,m} < 1$ mark the proportions of the sample generated by the m break points estimated under the null hypothesis for MSA i. For technical reasons, we require that $\varphi_{i,j} - \varphi_{i,j-1} > \pi_{i,0}$ for some small $\pi_{i,0}$ where we define $\varphi_{i,0} = 0$, $\varphi_{i,m+1} = 1$. Further, let $\eta_{i,j} = \frac{\pi_{i,0}}{\varphi_{i,j} - \varphi_{i,j-1}}$, $j = 1, \ldots, m+1$. The likelihood ratio test compares the maximum of the likelihood ratio obtained when allowing for $m+1$ breaks to that from only allowing for m. The distribution of this likelihood ratio statistic is given by

$$(A1.1)\ P(LR > c) = 1 - \prod_{i=1}^{m+1} \left(1 - P\left(sup_{\pi_i \in [\eta_{i,j}, 1-\eta_{i,j}]} Q_1(\pi_i) > c\right)\right),$$

which we calculate by recursive application of the method provided in Estrella (2003).

We apply this procedure to test for the existence of two break points against the null of one as well as three against the null of only two among those MSAs for which we find at least two statistically significant break points. There are some noteworthy practical issues involved with carrying out this procedure. We have not until this point said where the sample proportions $\pi_{i,0}$, $\pi_{i,1}$, $\pi_{i,2}$ come from. In practice, we restrict the full sample period for each MSA to lie between the first quarter in the data and the peak of price growth. We then do not allow any break points to lie in either the first or last two quarters of this sample for each MSA. This determines the fractions $\pi_{i,1}$ and $\pi_{i,2}$ which, because different MSAs have a different number of quarters, will vary across areas.

When estimating multiple break points, we further require that any two break points be at least four quarters apart. This determines the fraction $\pi_{i,0}$ which, again, will vary across areas due to differing sample sizes. Because of these restrictions, we are not able calculate p-values for many MSAs in the case of multiple breaks. The reason for this can be seen from the expression in (A1.1). Because this expression requires that $\eta_j < 0.5$, we must require that $\frac{\pi_{i,0}}{\varphi_{i,j} - \varphi_{i,j-1}} > 0.5$ for all j. This implies that we will not be able to calculate p-values for the two-break case in MSAs (neighborhoods) where the first break is less than $\pi_{i,0}/0.5$ from the beginning of the sample period. Naturally, this restriction is more burdensome when trying to calculate p-values in the three break case.

Reference

Andrews, D., 1993. Tests for parameter instability and structural change with unknown change point. Econometrica 61 (1), 821–856.

Bai, J., Perron, P., 1998. Estimating and testing linear models with multiple structural changes. Econometrica 66 (1), 47–78.

Bai, J., 1999. Likelihood ratio tests for multiple structural changes. J. Econom. 91 (2), 299–323.

Bailey, M., Cao, R., Kuchler, T. and Stroebel J. “Social networks and housing markets”, NBER Working Paper No. 22258, May 2016.

Bayer, P., Geissler, C., Magnum, K. and Roberts, J. “Speculators and middlemen: the role of flippers in the housing market”, NBER Working Paper No. 16784, February 2011.

Calvo, Sara, and Carmen M. Reinhart, “Capital Flows to Latin America: Is There Evidence of Contagion Effect?”, World Bank Policy Research Working Paper No. 1619 (1996).

Card, D., Mas, A., Rothstein, J., 2008. Tipping and the dynamics of segregation. Q. J. Econ. 123 (1), 177–218.

Case, K., Shiller, R., 1987. Prices of single family homes since 1970: new indexes for four cities. N. Engl. Econ. Rev. 45–56.

Case, K., Shiller, R., 1989. The efficiency of the market for single family homes. Am. Econ. Rev. 79 (1), 125–137.

Charles, K.K., Hurst, E., Notowidigdo, M.J., 2018. Housing booms and busts, labor market opportunities, and college attendance. Am. Econ. Rev. 108 (10), 2947–2994.

Chinco, A. and Mayer, C. “Misinformed speculators and mispricing in the housing market,” NBER Working Paper No. 19817, January 2014.

Chomsisengphet, S., Pennington-Cross, A., 2006. The evolution of the subprime mortgage market. Federal Reserve Bank St. Louis *Rev.* 88 (1), 31–56.

Corsetti, G., Pericoli, M., Sbracia, M., 2005. Some contagion, some interdependence: more pitfalls in tests of financial contagion. J. Int. Money Financ. 24, 1177–1199.

DeFusco, A.A., Nathanson, C.G. and Zwick, E. “Speculative dynamics of prices and volume”, 2017, Working Paper.

Dokko, J.K., Keys, B.J. and Relihan, L.E. “Affordability, financial innovation, and the start of the housing boom”, 2015, Working Paper.

Dungey, M., Renee, F., Brenda, G.-H., Vance, M., 2005. Empirical Modeling of Contagion: A Review of Methodologies. Quant. Financ. 5 (1), 1–16.

Estrella, A., 2003. Critical values and *P* values of bessel process distributions: computation and application to structural break tests. Econom. Theory 19 (6), 1128–1143.

Ferreira, F. and Gyourko, J. “Anatomy of the beginning of the housing boom: U.S. neighborhoods and metropolitan areas, 1993–2009”, NBER Working Paper No. 17374, August 2011 (latest version is May 22, 2012).

Ferreira, F. and Gyourko, J. “A new look at the U.S. foreclosure crisis: panel data evidence of prime and subprime borrowers from 1997 to 2012”, NBER Working Paper No. 21261, June 2015.

Forbes, K. “The “big C”: identifying and mitigating contagion”. The Changing Policy Landscape. 2012 Jackson Hole Symposium hosted by the Federal Reserve Bank of Kansas City, pps. 23–87. 2013.

Forbes, K., Rigobon, R., 2002. No contagion, only interdependence: measuring stock market comovements. J. Financ. LVII (5), 2223–2261.

Giglio, S., Maggiori, M., Stroebel, J., 2016. No-bubble condition: model-free tests in housing markets. Econometrica 84 (3), 1047–1091.

Glaeser, E., Gyourko, J., Morales, E., Nathanson, C., 2014. Housing dynamics: an urban approach. J. Urban Econ. 81, 45–56.

Hamilton, J.D., 2016. Macroeconomic Regimes and Regime Shifts. In: Uhlig, H., Taylor, J. (Eds.), Handbook of Macroeconomics, 2A. Elsevier, Amsterdam, pp. 163–201.

Hansen, B.E., 1997. Approximate asymptotic P values for structural-change tests. J. Bus. Econ. Stat. 15, 60–67.

Hansen, B.E., 2000a. Testing for structural change in conditional models. J. Econom. 97 (1), 93–115.

Hansen, B.E, 2000b. Sampling splitting and threshold estimation. Econometrica 68 (3), 575–603.

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

Holly, S., Hashem Pesaran, M., Yamagata, T., 2010. A spatio-temporal model of house prices in the USA. J. Econom. 158, 160–173.

King, M.A., Wadhwani, S., 1990. Transmission of volatility between stock markets. Rev. Financ. Stud. 3 (1), 5–33.

Leamer, E., 1983. Let’s take the con out of econometrics. Am. Econ. Rev. 73 (1), 31–43.

Lee, S.B., Kim, K.J., 1993. Does the October 1987 crash strengthen the comovements among national stock markets? Rev. Financ. Econ. 3, 89–102.

Nathanson, C.G., Zwick, E., 2018. Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market. J. Financ. Forthcoming.

Roback, J., 1982. Wages, rents, and the quality of life. J. Polit. Econ. 90 (4), 1257–1278.

Rosen, S., 1979. Wage-based indexes of urban quality of life. In: Mieszkowski, P., Straszheim, M. (Eds.), Current Issues in Urban Economics. Johns Hopkins Univerity Press, Baltimore.

Saiz, A., 2010. The geographic determinants of housing supply. Q. J. Econ. 125 (3), 1253–1296.

Shiller, R., 2005. Irrational Exuberance, 2nd edition Princeton University Press, Princeton, NJ.