

Time-varying housing supply elasticities and US housing cycles*

Knut Are Aastveit[†]
Norges Bank

Bruno Albuquerque[‡]
Ghent University

André Anundsen[§]
Norges Bank

January 31, 2018

Abstract

In a panel of 254 US metro areas, we estimate housing supply elasticities to be substantially lower today than what they were during the housing boom in the 2000s. An implication of this finding is that a given change in demand will have a stronger effect on house prices and therefore, *ceteris paribus*, contribute to greater house price volatility. Our estimates suggest that an exogenous monetary policy shock that lowers the interest rate by one percentage point leads in the current environment to a higher increase in house prices of almost three percentage points after three years. We also show that housing supply elasticities declined the most in areas that experienced the sharpest drop in house prices during the housing bust in the latter part of the 2000s. The fear of a new bust may have paved the way for a housing boom that so far looks remarkably similar to the 1996-2006 boom in terms of house price developments.

Keywords: *House prices; Heterogeneity; Time-varying housing supply elasticities;*

JEL classification: *C23, E32, E52, R31*

*This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank.

[†]E-mail: knut-are.aastveit@norges-bank.no

[‡]E-mail: bruno.albuquerque@ugent.be

[§]E-mail: andre-kallak.anundsen@norges-bank.no

1 Introduction

Has there been a change in housing supply elasticities in US metropolitan areas over time? The answer to this question is important, since any change in housing supply elasticities will affect the transmission of demand shocks to house prices, and thereby impact house price volatility. There is a vast number of papers showing that local differences in housing supply elasticities have been a central factor in explaining the cross sectional variation in US house price developments (see e.g., (Green et al. 2005, Gyourko et al. 2008, Saiz 2010, Huang and Tang 2012, Glaeser et al. 2014, Anundsen and Heebøll 2016, Aastveit and Anundsen 2017)). The common assumption in these papers is to consider supply elasticities as constant over time. In this paper, we aim to fill this gap by constructing time-varying and area-specific housing supply elasticities. After documenting a sizeable and economically important change in elasticities *within* MSAs over time, we make a novel contribution to the literature by showing that the time-varying elasticities have important implications for house price volatility.

We ask the following questions: (i) Have housing supply elasticities changed over the US housing cycle since the 1990's?; (ii) Do time-varying elasticities impact house price volatility and the transmission of demand shocks?; and (iii) What factors have contributed to the changes in housing supply elasticities over time?

To answer these questions, we consider a panel of 254 US Metropolitan Statistical Areas over 1987Q1–2016Q4. We tackle the first question by estimating a panel data model for different sub-samples: the first cycle from 1986-1996, the 1996-2006 boom, the 2006-2012 bust, and the recent recovery over 2012-2017. We allow supply elasticities within each period to depend on local variation in topography and regulatory conditions.

Theory suggests that there should be differences in housing supply elasticities depending on whether the market is in a boom or a bust. In particular, due to the durability of housing, it is reasonable to expect that the housing supply curve is nearly vertical when house prices are falling (Glaeser and Gyourko 2005). Moreover, since this should hold in all markets, local-specific factors such as differences in topography and housing market regulation should not matter for the responsiveness of housing supply during a bust. Our results support this view. We find the estimated supply elasticities during the 2006-2012 bust to be very low and to have negligible variation in the cross section. Theory also suggests that local differences in topography (Saiz 2010) and regulation (Gyourko et al. 2008) should impact housing supply elasticities in a

booming housing market. We find support for this during both the 1996-2006 boom and the more recent recovery. In addition, and more interesting, we provide evidence for a substantial and synchronised decline in housing supply elasticities in the recent boom.

This result implies an increased price responsiveness to a given aggregate demand shock. To investigate this conjecture, we follow Aastveit and Anundsen (2017) and estimate the effect of an exogenous monetary policy shock on house prices. Nevertheless, and differently from them, we allow the effect of an expansionary shock to vary over time. In particular, we interact the monetary policy shock with the estimated time-varying housing supply elasticities. This exercise reveals that a similarly sized monetary policy shock has a substantially greater impact on house prices during the current recovery than during the previous boom – estimated to be twice as large. For an area with a median housing supply elasticity, house prices are estimated to increase by nearly 6% after three years following a monetary policy shock that reduces the policy rate by one percentage point. For the same shock, we estimate house prices to increase by around 3% for the 1996-2006 boom.

We also explore the factors that have contributed to the decline in supply elasticities between the two booms. Our results suggest that the areas that experienced the largest decline in house prices at the end of the previous decade were also those areas that have seen the largest decline in supply elasticities. We interpret this as evidence that the fear of a new bust have led developers to be less price responsive than previously. The irony of this behaviour is that it may have paved the way for a new housing boom where house prices are more responsive to fluctuations in demand.

Some research argues that the current housing boom appears to be less vulnerable than the boom that ended in 2006, as it has not been fuelled by credit (Glick et al. 2015). The current rise in house prices is, however, surrounded by important vulnerabilities: (i) the current level of household debt is significantly higher than before; (ii) the underlying fundamentals are weaker this time around; and (iii) while the rise in house prices might benefit existing homeowners (through the equity home channel), inequality between this group and first-time homeowners or renters can be magnified in a context where supply does not accompany the strong rise in house prices. Our paper thus suggests that, contrary to common belief, the current housing boom appears to be more vulnerable than before, since aggregate demand shocks exert a much larger effect on house prices than previously.

The rest of the paper proceeds as follows. In the next section, we compare the similarities

and differences between the housing boom of the 2000's and the current boom. In Section 3 we consider a simple demand-supply model with durable housing to illustrate the implications of time-varying housing supply elasticities. In Section 4 we describe the data and some stylised facts about the US housing cycle over the last 30 years. We present our econometric approach and the estimates of the housing supply elasticities in Section 5. We then discuss the impact of the time-varying elasticities on the dynamics of the housing market in Section 6. In Section 7 we explore the factors that have led to the decline in elasticities between the two housing booms. Section 8 concludes the paper.

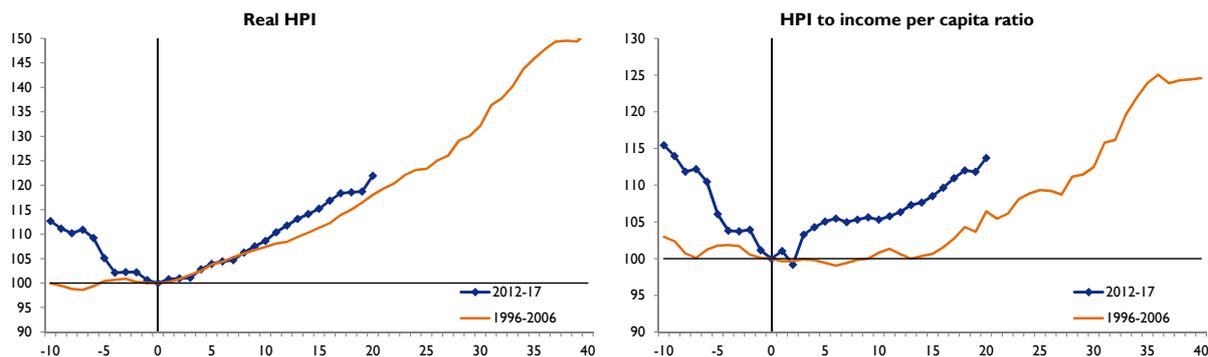
2 The two house price booms: 1996-2006 vs 2012-17

Real house prices at the national level have grown by almost 25% since the beginning of the current housing recovery in the second quarter of 2012. This pace has been in line with the previous house price boom over 1996-2006, a boom recognised by many to have been the root cause of the Great Recession (left panel of Figure 1). What is more remarkable is that the current boom in house prices looks stronger when scaling house prices by income per capita (right panel of Figure 1). Since the current economic recovery has been characterised by subdued income and consumption growth, weaker underlying aggregate demand makes the adjusted house prices look somewhat stronger compared with the 1996-2006 boom.

The previous boom was characterised by strong increases in both house prices and housing supply, in a context where the common belief was that house prices 'can only go up' (Piazzesi and Schneider 2009). Nevertheless, construction activity during the current boom has been growing considerably below what the rise in house prices would suggest. In fact, while building permits and housing starts had increased by a cumulative amount of almost 60% in the first five years of the previous boom, this time around permits and starts have responded quite sluggishly to house prices, with only a cumulative increase of roughly 15% (top panel of Figure 2).

Why is supply progressing at such a slow pace, while house prices have risen in line or above the previous boom? One potential explanation for the sluggish supply response is weaker aggregate demand, reflecting subdued wage and income growth. In addition, although credit standards have loosened considerably, and interest rates are persistently low, credit growth to households has been weak and it has likely been dampening housing demand as well – households have been deleveraging from excessive household debt, particularly mortgage debt, accumulated

Figure 1: House price developments across booms



Notes: The figure tracks the evolution of real house prices at a quarterly frequency (left panel) and house prices scaled by income per capita (right panel) during the two house price booms. The zero on the x-axis marks the beginning of each housing boom. We scale the series to 100 at that point. We measure real house prices with the FHFA house price index, a weighted, repeat-sales index, deflated by CPI. The solid orange line refers to the boom between 1996q4 and 2006q4, while the blue line is from 2012q2 to 2017q2.

in the run-up to the Great Recession. Nevertheless, weak aggregate demand cannot explain why there is a breakdown in the relationship between housing supply and house prices. A second theory posits that the strong rise in construction activity during the 1996-2006 boom, followed by a severe economic downturn during the following bust, might have led to an oversupply of houses during the current bust, which may explain why housing supply has been expanding slowly. This theory, however, is not validated by data. For instance, housing stock per capita has been trending consistently downwards, in contrast with the dynamics during the 1996-2006 boom, and the housing vacancy rate is at levels similar to the previous boom (bottom panel of Figure 2). The same picture is portrayed by the months' supply of houses, another indicator usually taken into account to assess the strength of the housing market.¹ These indicators clearly suggest that there is no evidence pointing to an oversupply of houses in the market; if anything, they signal a lack of houses to satisfy the existing housing demand.

A third potential explanation centres on the long shadow that the Great Recession might have cast on builders' expectations about future demand, making them more reluctant to invest significant resources in construction projects. While there is not much available data that can shed some light on this conjecture, existing survey data do not support the theory that builders have become more pessimistic about the housing market going forward. In fact, builders' sentiment published by the NAHB shows that their expectations for sales over the next 6 months has been evolving in line with the previous boom (show figures in appendix).

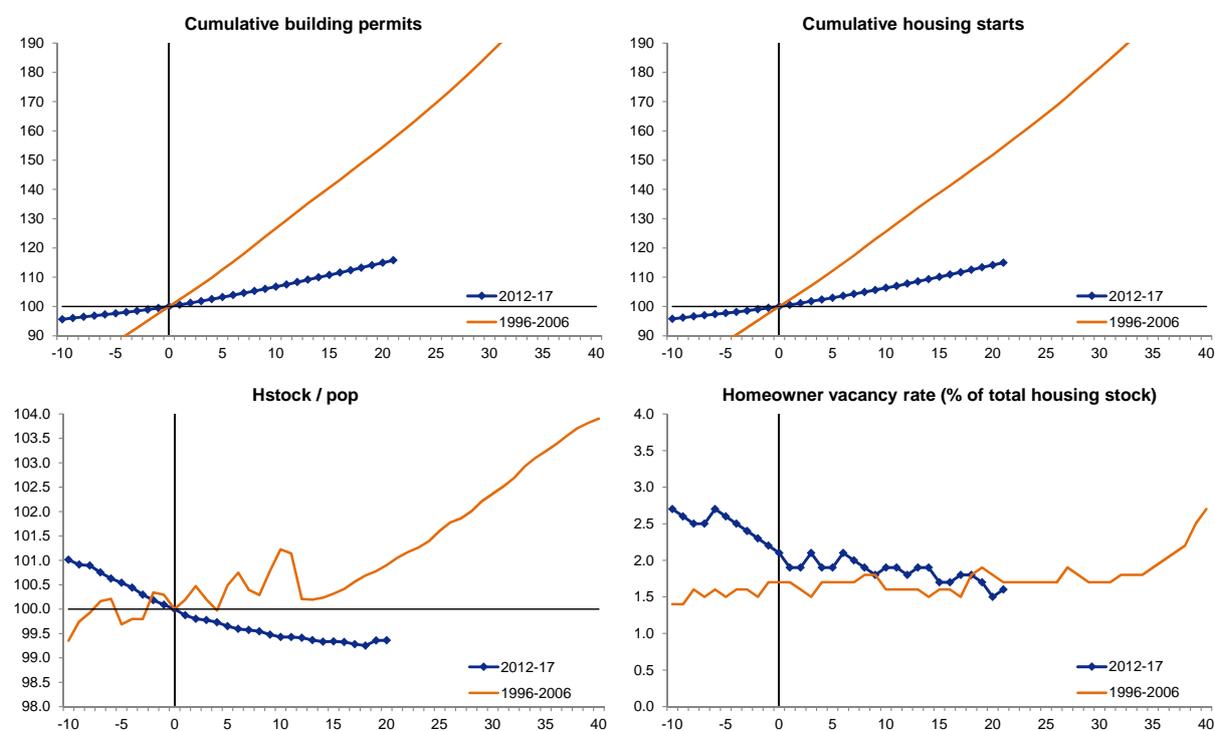
Finally, the relationship between housing supply and prices might have changed due to

¹The months' supply is the ratio of houses for sale to houses sold. It indicates how long the current for sale inventory would last given the current sales rate if no additional new houses were built.

tighter regulation over time, which makes it difficult to build and might result in higher house prices. There is reason to believe that regulation may change over a reasonable long period of time. For instance, recent literature has found that production costs in construction have been relatively stable compared to the early 80s, but productivity has been declining, suggesting that this decline in productivity in construction should be explained by increased costs stemming from tighter regulation (Davis and Palumbo 2008, Glaeser and Gyourko 2018).

Overall, while it remains challenging to pin point exactly the reasons why housing supply has reacted less to a given change in house prices, compared with the previous boom, the picture painted above does suggest that housing supply elasticities vary over time.

Figure 2: Housing supply indicators across booms



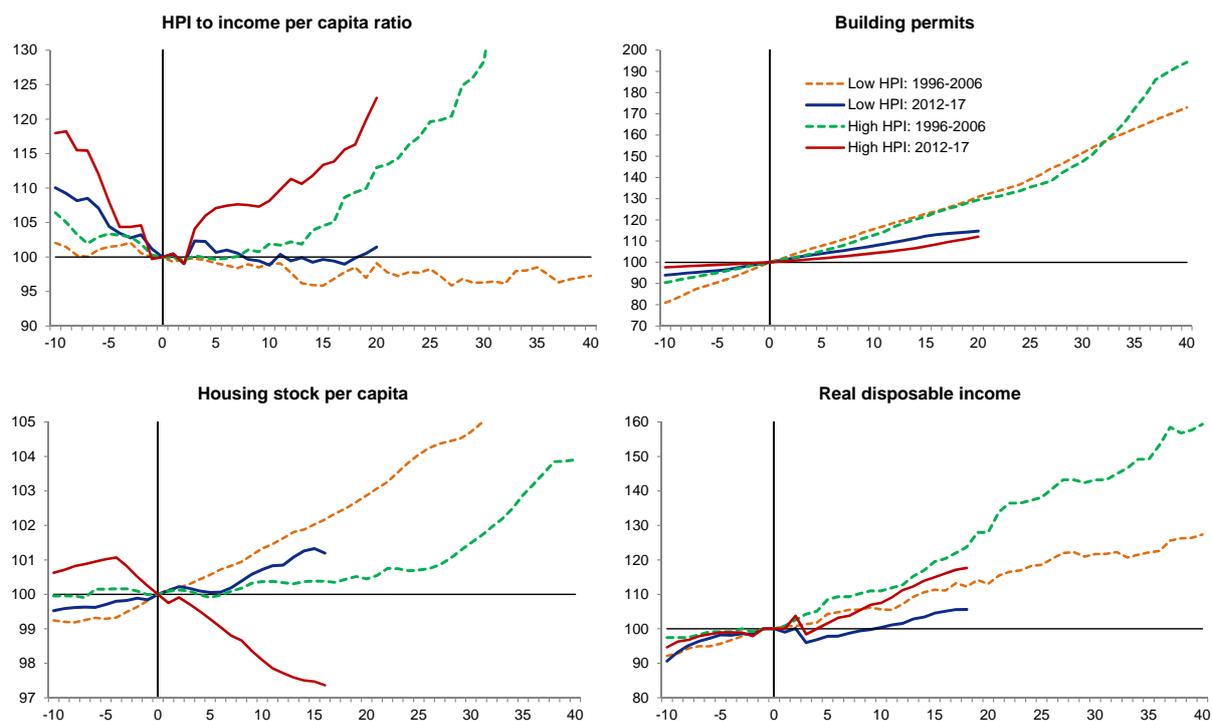
Notes: The figure tracks the evolution of selected housing indicators at a quarterly frequency during the two house price booms. The zero on the x-axis marks the beginning of each housing boom. We scale the series to 100 at that point, but not for the bottom right-hand panel; the homeowner vacancy rate is displayed in % of total housing stock. The solid orange line refers to the boom between 1996q4 and 2006q4, while the blue line is from 2012q2 to 2017q2.

So far we have been focusing our analysis on the US national. But the aggregate picture masks significant heterogeneity at the local level. To illustrate this point, we break the sample into quartiles of the cumulative house price change between 1996 and 2006. We define ‘Low HPI’ MSAs as those areas that recorded the smallest cumulative growth in house prices during that period, as measured by the first quartile, while ‘High HPI’ refers to the top quartile (Figure 3). When we look at the house price-to-income per capita ratio, and although this indicator has been stronger during the current boom for both groups, the gap between the two booms has

been larger for the ‘High HPI’ MSAs. In addition, the dynamics in housing stock per capita has been markedly different between the two groups in the current recovery. While the housing stock has been expanding slowly for the ‘Low HPI’ MSAs, in line with the previous boom, it has, however, been trending downwards for the ‘High HPI’ MSAs.

Overall, these two indicators highlight the importance of taking the heterogeneity in local housing markets into account. Only by looking at disaggregated data, can we be more confident in extracting valuable lessons and in providing policy implications for the US economy as a whole. The use of more disaggregated data follows the spirit of the most recent housing market literature which tends to look at the housing market as a collection of several markets that differ not only by geography but also by other attributes, as surveyed in Piazzesi and Schneider (2016).

Figure 3: Housing indicators and disposable income for MSA groups across housing booms



Notes: The figure tracks the evolution of selected housing indicators and disposable income at a quarterly frequency during the two house price booms across two groups of MSAs. The zero on the x-axis marks the beginning of each housing boom. We scale the series to 100 at that point. ‘Low HPI’ refers to MSAs belonging to the first quartile of the cumulative house price change between 1996 and 2006, while ‘High HPI’ refers to the top quartile. Dashed lines correspond to the first housing boom, 1996-2006, and solid lines to the on-going boom, 2012-2017.

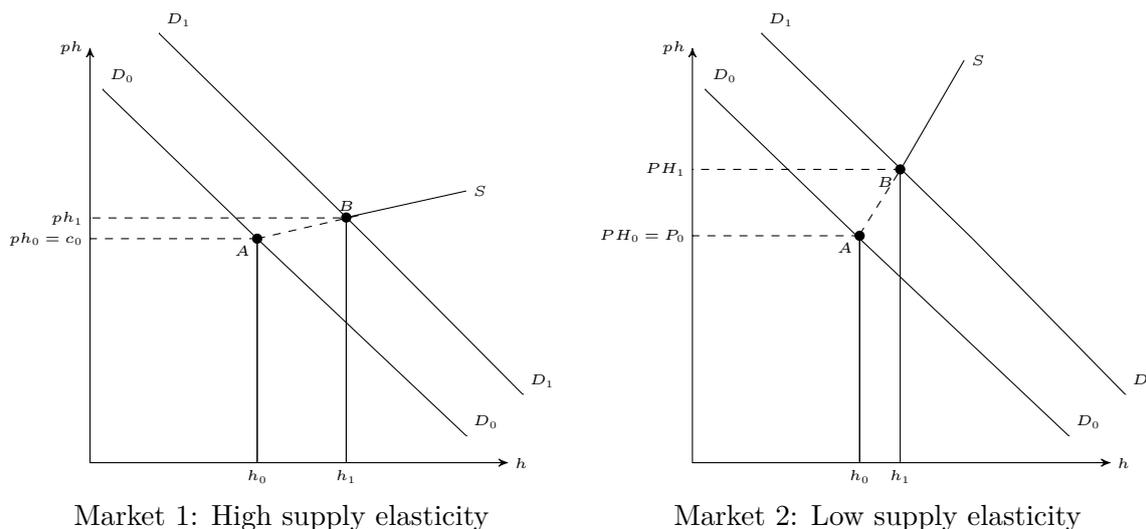
3 Theoretical background

[VERY INCOMPLETE SECTION]

The housing supply response might vary over time, and not only at the cross-section. As we

have argued before, several factors might change the slope of the supply curve, i.e. for the same % change in house prices, supply responds differently. Changes in regulatory and geographical conditions, demographic changes, financing conditions, homebuilders' expectations about future house prices and housing demand; which is related to the perceived profitability of the housing sector, precautionary reasons, and the business cycle where the supply response is arguably asymmetric in booms versus busts.

Figure 4: Positive demand shock in high vs low supply elasticity markets



Note: D_0 is the original demand curve, S the supply curve, and point A is the initial equilibrium with house prices ph_0 and quantity h_0 . The demand curve shifts to D_1 after a positive demand shock, and the new equilibrium to point B , with house prices at ph_1 and quantity at h_1 . The dotted part of the housing supply curve illustrates that housing supply is rigid downwards, so that the supply curve kinks at A before the shock and at B after the shock.

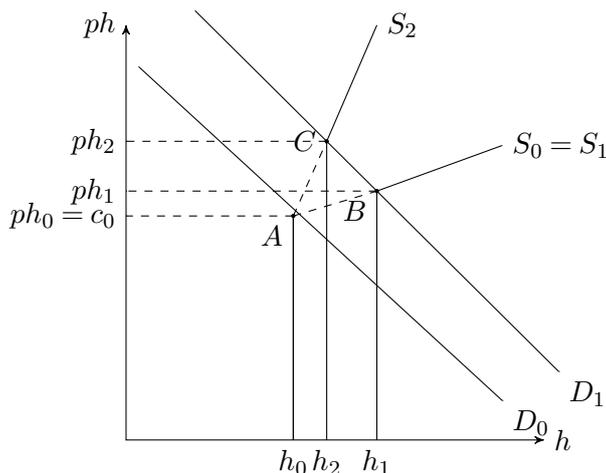
A more vertical supply curve implies that a demand shock can act as an amplification mechanism for house prices in the presence of positive demand shocks (Figure 5). This makes it interesting to analyse episodes of housing booms, not busts, as the downward rigidity of housing stock implies that the impact of a given negative demand shock should be independent of the elasticity (Aastveit and Anundsen 2017). In a boom, the lower the elasticity, the larger the house price response.

4 Data and housing market cycles

4.1 Data

To explore the time variation in housing supply elasticities and the heterogeneity across local housing markets in the United States, we use quarterly data for a panel of 254 MSA between

Figure 5: Positive demand shock when the supply elasticity is time varying



Note: D_0 is the original demand curve, S_0 the supply curve, and point A is the initial equilibrium with house prices ph_0 and quantity h_0 . The demand curve shifts to D_1 after a positive demand shock, and the new equilibrium is reached at point B, with house prices at ph_1 and quantity at h_1 , conditional on a time-invariant supply elasticity ($S_0 = S_1$). Alternatively, if there is a decline in the elasticity over time, whereby the supply curve becomes steeper (from S_1 to S_2), the equilibrium moves to point C, with higher house prices ph_2 and lower quantity h_2 .

1986 and 2016. This sample is fairly representative, covering more than 80% of US income and population. The definition of the MSAs we follow are in accordance with the new delineations issued by the Office of Management and Budget (OMB), based on the 2010 Census.²

The MSA data on housing supply encompass building permits, housing starts, housing completions and housing stock. In addition, we have data on house prices, consumer price index (CPI), and controls for macroeconomic (disposable income, unemployment rate, non-farm employment), financial (mortgage originations), and socio-demographic conditions (population, population density and fraction of population with bachelor degree or higher). We deflate all nominal macroeconomic series with the CPI. The MSA data have been provided by Moody's Analytics, with the original sources of the data coming mainly from the Census Bureau, Bureau of Economic Analysis (BEA), Bureau of Labor Statistics (BLS), and the Federal Housing Finance Agency (FHFA) – see Appendix A for a full list of variables and sources, and Table B.1 for summary statistics.

We also consider two indicators that measure, in a different way, how much housing supply is restricted, commonly used in the related literature. These indicators are nevertheless only available at the cross-section and are thus time-invariant. The first one, Saiz (2010)'s land unavailability index, captures the geographical restrictions of putting up a construction project.

²It is important we follow consistently the same MSA definitions across variables and time, given that major changes may take place every 10 years, when population counts and commuting patterns are revised following the decennial census.

In particular, Saiz (2010) uses topographic maps over 1970-2000 to measure the proportion of land in a 50 km radius of the city centre that is lost to steep slopes and water bodies, such as oceans, rivers, lakes and wetlands. He finds that metropolitan areas that are more inelastic – the supply response to a given change in house prices is more limited – are typically more land constrained. The second indicator, Gyourko et al. (2008)’s Wharton Land Use Regulatory Index, refers to the degree of stringency of local zoning laws, i.e. the time and financial cost of acquiring building permits and constructing a new home. The summary indicator is constructed based on the results of a nation-wide survey in 2005, and on a separate study of state executive, legislative and court activity.

Finally, the data that are available only at the state level include wages and salaries in construction, which acts as a proxy for the builders’ costs, and the deregulation index computed in a similar spirit as in Mian et al. (2017), which measures the level of inter- and intra-state deregulation in the US banking sector.³

4.2 Identifying housing market cycles

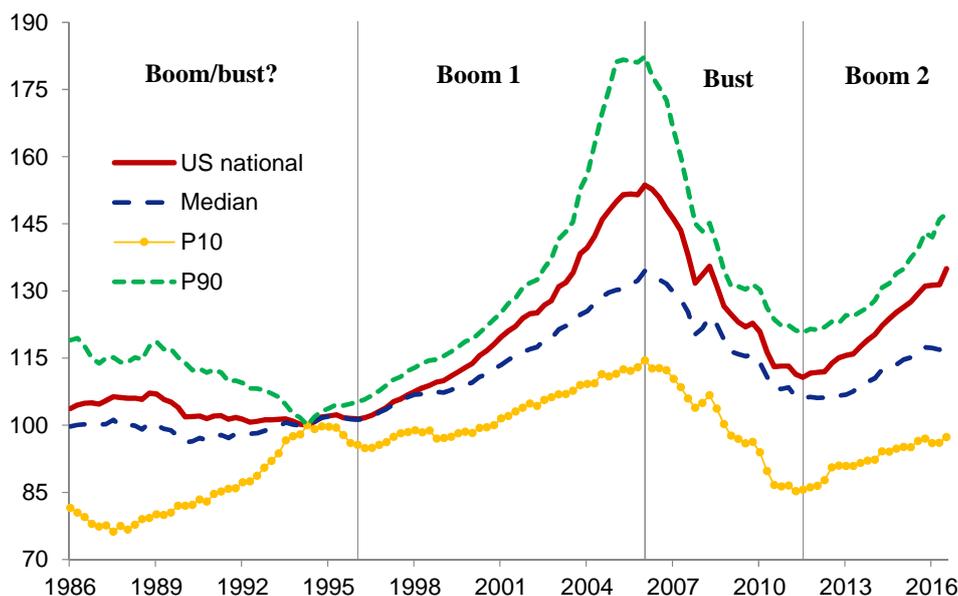
House prices have gone through several boom-bust cycles over the last 30/40 years. Although house prices tend to vary more at the local level, and less so when looking at the aggregate national picture, it is possible to identify some periods in the local (MSA) data when there is sufficient co-movement across areas to recognise a pattern in housing cycles at the national level. By analysing the peaks and troughs for average house prices for a sample of 79 MSA, Glaeser et al. (2008) identify a national boom over 1982-89, a subsequent bust until 1996, and a strong boom between 1996 and 2006.

When we use our significantly larger sample of 254 MSAs, three housing cycles appear to show up quite clearly: a strong boom from 1996 until 2006, which gives rise to a severe bust until 2012, and the on-going recovery that has started since then. For illustration, we plot the national house price index, together with the median, the 10th and 90th percentiles of the house price distribution at the MSA level (Figure 6). With our dataset we cannot, however, identify clearly neither a boom from 1982 to 1989, nor a boom over 1989-1996, as in Glaeser et al. (2008). Instead, we observe that there is significant heterogeneity across the MSAs over the 1986-96

³Differently from Mian et al. (2017), who compute the index for 1989 at the cross-section, we construct it for each state and year according to the following formula: $0.5 * \sum_{j \in (intra, inter)} \max(\text{year} - \text{Dereg_year}_{j,s}, 0)$, where $\text{Dereg_year}_{j,s}$ refers to the year when banking deregulation j took place in state s .

period, without any clear generalised co-movement in house prices across areas.⁴ In fact, the MSAs at the bottom of the distribution recorded a steady increase in house prices over 1986-96, while the MSAs at the top saw the opposite dynamics. This heterogeneity dissipates once we look at the median or at the aggregate index, with real house prices remaining relatively stable over that ten-year period (Figure B.1 in Appendix B complements the information provided in Figure 6, by showing the Kernel densities of cumulative changes in house prices for each one of the cycles.).⁵

Figure 6: Real house prices (1995q1=100)



Notes: We measure real house prices with the FHFA house price index, which is a weighted, repeat-sales index, deflated by CPI. The index assumes the value of 100 in 1995q1. The solid red line represents the US aggregate index, the long-dashed blue line, the yellow line with markers, and the dashed green line refer respectively to the median, 10th and 90th percentiles of the house price distribution for a sample of 254 MSA. The vertical lines divide the sample period by housing cycles, as defined in the main text.

To be more specific, we analyse the peaks and troughs in real house prices at the median to determine the housing market cycles. Given that our sample includes some MSAs with extremely large variation in prices over time, we look at the median, and not the average as in Glaeser et al. (2008), to minimise the effect that outliers have on determining the housing cycles. The first cycle, going from 1987q1 to 1996q4, is characterised by significant heterogeneity across MSAs. A bit less than 50% of the MSAs in our sample recorded cumulative increases in house prices over this period, whilst house prices at the median or national level were relatively stable (Table 1). The second housing market cycle that we identify at the local level is a housing boom

⁴We start our sample in 1987q1 as the coverage of house prices at the MSA level is relatively limited before that period.

⁵We get different housing market cycles from Glaeser et al. (2008) for the 1982-89 and 89-96 periods since we cover a substantially larger sample of MSAs. Only when we restrict our sample to the largest 79 MSAs, do we find a clear common housing bust across MSA over 1982-89 and a subsequent boom until 1996.

that runs from 1997q1 to 2006q4. All of the MSAs in our sample reach the end of the cycle with house prices above the previous trough in 1996q4, but the dispersion around house prices increases is high; while house prices increased on average by 17% for the MSAs belonging to the first quartile, they increased by around 86% for the top quartile. The bust in house prices that ensued lasted until 2012q2, a period when all MSAs but one saw house prices in real terms decline substantially. We define the last housing market cycle as the on-going boom, which began in 2012q3. By the end of 2016, house prices in almost 90% of the MSAs in our sample had already recovered from the losses during the crisis.

Table 1: Local house price cycles

	US	median	Q1	Q2	Q3	Q4	N	>0
1986-1996	-2.2	-1.8	-18.6	-5.5	4.7	16.9	220	103
1996-2006	51.5	32.8	17.2	27.7	41.0	85.9	254	254
2006-2012	-28.0	-21.0	-42.4	-25.5	-16.7	-10.5	254	1
2012-2017	22.0	9.2	0.3	6.7	13.1	37.1	254	223

Notes: The figures, apart from the last two columns, show cumulative changes in real house prices for each housing cycle. The first column refers to the US aggregate, while the following five columns show the median and the quartiles of the distribution of real house prices. N is the number of MSA in the sample, while >0 counts the number of MSAs that recorded cumulative house price increases over each cycle.

An alternative to our method of identifying local house prices would be to define the booms and busts by modelling MSA-specific housing cycles. For instance, Ferreira and Gyourko (2011) define the beginning of a boom in house prices as the global breakpoint in the rate of house price appreciation in each MSA. The idea is that each local housing market should exhibit constant and continuous growth paths for house prices in the steady state. Using micro data on housing transactions for a sample of 94 MSAs over 1993-2009, they find significant heterogeneity in the timing of the start of the house price boom across areas; the authors estimate the earliest booms to have begun between 1997-1999, primarily in coastal areas, and then spread afterwards to other MSAs. In turn, Hernández-Murillo et al. (2017) estimate a Markov-switching model of housing cycles that allow housing markets in clusters of MSAs to deviate from the national housing cycle. Using monthly data on 135 MSAs from 1989 to 2012, they find evidence of three clusters of cities that experience their own idiosyncratic contractions in housing markets, either before, after or independent of the national housing downturn.

We do not define MSA-specific housing market cycles, as proposed by the two papers above, for the following reasons. Firstly, after computing time-varying housing supply elasticities in the next section, we will be mostly interested in understanding why elasticities are so different

between the two pre-defined house price booms – the first over 1996-2006 and the second during 2012-2017 – and the potential macro implications of such a phenomenon, a topic first introduced in Section 2. The fact that the timing of the housing booms does not necessarily coincide for all MSA on the same quarter or year as the median should not constitute an issue for our analysis, given that we estimate the elasticities in the next section to have declined over time for all MSAs. In addition, it would have been impractical to analyse the two booms for each one of the 254 MSAs at potentially different points in time. The median MSA should instead portray a reasonable representation of the reality of the US housing market. And, secondly, Del Negro and Otrok (2007) make the point that the national housing cycle has been increasingly more important in affecting local housing markets, in particular surrounding the housing boom in the run-up to the Great Recession. This trend appears to have continued as our data show that housing markets have become more synchronised over time across booms and busts, despite an higher dispersion around the magnitude of house prices changes (Table 1). In other words, the data indicate that there has been an increasingly stronger co-movement in house prices across local markets, which alleviates some concerns about local housing markets having their own cycle completely independent of the national one.

5 Estimating time-varying housing supply elasticities

5.1 Baseline IV specification

We take Green et al. (2005)’s single-equation approach as the starting point to estimate housing supply elasticities over time. The authors estimate time-invariant housing supply elasticities for a sample of 45 MSA over 1979-1996 by regressing a proxy for annual growth in housing stock on lagged house price growth. We adjust their model in the following ways. First, we use building permits and not the housing stock as the main supply variable, given that permits lie at the beginning of the construction stage and are thus a better indicator to capture more adequately the reaction of builders to a change in house prices.⁶ Second, given the stationary properties of building permits, we adopt a level specification with housing supply being determined by the price-to-cost ratio (Tobin’s Q), which is standard in the housing literature (see, for instance, Glaeser et al. 2008). Due to data availability issues at the metropolitan level, we use wages

⁶Just to give an idea about the process of building a housing unit in the United States, in most cases builders need to apply for a permit to get their construction project approved (building permits), which can take a few months. After the approval is granted, the construction works start (housing starts). The process ends when the housing unit is occupied or available for occupancy (housing stock).

and salaries in the construction sector at the state level as a proxy for the construction costs faced by builders.⁷ Third, we account for geographical and land-use regulatory constraints in the response of housing supply to a change in house prices (Saiz 2010, Gyourko et al. 2008).

We focus on the last three housing cycles mentioned before, and leave out for the moment the 1986-96 period, as we have shown before that we can identify with more clarity the housing cycles since the end of 1996 for a large fraction of the MSAs. The first cycle is the 1996-2006 boom, followed by the subsequent 2006-2012 bust, and the final cycle is the 2012-2017 boom. Given that some of these cycles are as short as five years, we would lose too many degrees of freedom if we were to estimate the supply elasticities by running a regression for each MSA as in Green et al. (2005). What we do is to resort to panel data techniques with the advantage of preserving more degrees of freedom, the estimated parameters being more efficient, and our specification being less prone to identification problems compared to MSA-individual models. Specifically, we run the following model with Fixed Effects for 254 MSA on quarterly data over each of the three housing cycle:

$$\begin{aligned} \log(H_{i,t}) = & \eta_i + \beta \log(HPI_{i,t}) + \lambda [\log(HPI_{i,t}) * Unaval_i] + \delta [\log(HPI_{i,t}) * WRLURI_i] \\ & + \gamma X'_{i,t} + \zeta_{i,t} + \epsilon_{i,t} \end{aligned} \quad (1)$$

where the dependent variable is the logarithm of building permits, HPI is the logarithm of the real house price index, $Unaval$ is Saiz (2010)'s land unavailability index, $WRLURI$ is Gyourko et al. (2008)'s Wharton Land Use Regulatory Index, the control variables X'_i include the lagged dependent variable, the logarithm of real wages in the construction sector, and of its interaction with the supply restriction indices. We add η_i to account for MSA fixed effects, while the state-by-time fixed effects $\zeta_{i,t}$ capture time-varying unobserved common factors at the state level. The parameters λ and δ , referring to the interaction of house prices with the supply restrictions, allow housing supply to respond differently across MSAs to a given change in house prices conditional on land available for construction and regulation. We expect β to be positive, as builders typically apply for more building permits as a result of an increase in house prices, and λ and δ to be negative, i.e. for the same increase in house prices, tighter geographical and regulatory restrictions on housing supply should lead to a smaller expansion in building permits.

⁷The costs associated with delivering a housing unit encompass other components besides labour costs, such as the land value and the costs of equipment. Glaeser and Gyourko (2018) provide annual estimates of the total costs of delivering a housing unit (including a profit to compensate the home builder) for an unbalanced panel of 78 MSA over 1985-2013.

Estimating the regression above by OLS might not be a valid option due to endogeneity issues and a potential omitted variable bias that may plague our coefficients. There may be a third variable, such as income shocks or household expectations about future economic conditions, driving simultaneously house prices and housing supply, therefore invalidating our parameter estimates. We deal with this potential econometric issue by using Instrumental Variable (IV) estimation. The challenge, however, is in finding an instrument that can pass both the 1st stage requirement – correlated with local house prices – and the 2nd – not correlated with omitted supply factors and with the dependent variable. Even though the literature has typically used Saiz (2010)’s elasticity indicator as the instrument for house prices, such as Mian et al. (2013) and Kaplan et al. (2016) – although not free of criticism (Davidoff 2016) – it makes little sense to use it in our framework given that is precisely the supply elasticity that we want to compute at the end of the process.

We take as the starting point the works of Glaeser et al. (2012), and Glaeser and Gottlieb (2009) who use the mean January temperature as the main instrument for house prices.⁸ January temperatures capture the exogenous variation in amenities that lead house prices to change. For example, Glaeser and Gottlieb (2009) focus on agglomeration economies, particularly on understanding the rise of Sunbelt cities and the reasons behind a strong link between population density and income within the United States. The rationale for using mean January temperatures as the instrument for house prices is mentioned at the outset on page 984, ‘*No variable can better predict city growth over the past fifty years than January temperature, yet it is unclear a priori why warm places have grown so dramatically.*’ In turn, Glaeser et al. (2012) interact the January temperature, and the Wharton supply constraint index, with year dummies to create instruments for house prices, so as to study the effect of an exogenous change in loan approval rates on house prices at the metropolitan level. In their framework, the January temperature variable proxies for housing demand, while supply restrictions proxy for housing supply.

In our paper, we need an instrument for house prices that varies over time, otherwise our panel model drops it automatically (Eq. 1 above). Given that the mean January temperature variable only varies across MSAs and not over time – a single data point reflecting monthly average temperatures in January calculated over 1941-1970 – we interact it with quarterly dummies to transform it into a time-varying variable, in a similar fashion to Glaeser et al.

⁸In a study of the impact of immigration in housing rents, Saiz (2007) uses the same variable to control for amenities that influence the growth in rents across US cities.

(2012). To make it clear, we have three instruments in the model due to the presence of the interaction terms: the transformed January temperature variable acts as an instrument for house prices, and the same variable interacted with *unaval* and with *wrluri* as instruments for house prices interacted with the same supply restrictions. We assess the relevance and strength of the instruments based on the correlation of the mean January temperature with house prices (1st stage requirement of the IV approach), and on the exclusion restriction (the covariance between the instruments and the residuals from the 1st stage equation should be zero). Particularly for the former, we rely on the Cragg-Donald weak identification F-test together with the Stock and Yogo (2002) critical values for relative size and bias at the 5% level.

Our initial results for the 1st stage regression (not reported here) show that our instruments are relatively weak for the first housing cycle over 1996-2006 – they fail to pass the Stock and Yogo (2002) critical values. In this context, we add another instrument to our specification, the logarithm of real disposable income (and the interaction of it with the two supply restriction variables) to improve on the overall strength of the instruments. Income is highly correlated with house prices, while the exclusion restriction is supported by the fact that income is usually included in standard housing demand equations, but not in housing supply ones (which is precisely what we are estimating in this paper).

The new specification with mean January temperatures and income as instruments for house prices improves substantially the quality of the IV regression in terms of the strength of the instruments, compared with the previous approach where January temperatures were used alone: for all housing cycles the Cragg-Donald F-test exceeds the Stock and Yogo (2002) critical value of 12.2 for the 5% relative bias (Table 2).⁹ The coefficient on house prices is statistically significant at conventional levels, and positive, for all housing cycles. Moreover, the slow increase in building permits during the current cycle, shown in Section 2, is consistent with a weakening in the response of permits to a given change in house prices after the strong 1996-2006 boom. In particular, our estimates indicate that building permits increased by 4.4% over the long term for every 1% increase in house prices during the 1996-2006 boom, which is more than twice as large compared with the subsequent downturn and current housing recovery – a response of

⁹We have also run alternative IV regressions by complementing the average January temperatures with (i) a variable capturing the fraction of the population with bachelor degree or higher, and (ii) the unemployment rate. In addition, we have also interacted the mean January temperature with a common variable to the whole panel, particularly the national house price indicator, thus replacing the time dummies as the interaction term. Glaeser et al. (2008), for instance, show that local building permits do not correlate well with national house prices (see their footnote 19), which justifies that the exclusion restriction is respected. Although the strength of the overall instruments also improves compared with a specification that uses January temperatures as the sole instrument, adding disposable income instead supersedes the aforementioned alternative IV specifications.

roughly 2%.¹⁰ In addition, the interaction of house prices with the supply restriction variables yield the expected negative sign, i.e. the tighter the geographical and regulatory restrictions to housing supply for a given house price, the smaller the increase in building permits.

For the sake of completeness, we also show the coefficients from an OLS regression in columns (1) to (3). Although the results are relatively similar to the IV specification, the coefficients on house prices are a bit lower than those from the IV for the last two cycles, suggesting that the potential endogeneity between permits and house prices is biasing the coefficients down. Moreover, the supply restrictions appear to exert a stronger influence on permits for the IV approach, while also being more precisely estimated.

Table 2: Regression estimates by cycle

	OLS			IV		
	(1) 1996-2006	(2) 2006-12	(3) 2012-17	(4) 1996-2006	(5) 2006-12	(6) 2012-17
Log (HPI)	2.505*** (0.194)	0.483* (0.273)	0.932** (0.400)	2.256*** (0.214)	1.030** (0.455)	1.613*** (0.379)
Log (HPI)*Unaval	-0.086 (0.124)	-0.027 (0.354)	-0.719 (0.557)	-1.178*** (0.246)	-0.266 (0.651)	-1.272 (0.975)
Log (HPI)*Wrluri	-0.187*** (0.032)	-0.114 (0.103)	-0.256 (0.167)	-0.736*** (0.065)	-0.468** (0.235)	-1.141*** (0.333)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Number of MSA	254	254	254	243	243	243
Observations	10,160	5,588	4,572	9,720	5,346	4,374
Cragg-Donald F-test	-	-	-	137.6	50.29	240.6

Notes: Regression estimates of Eq. 1, where the dependent variable is the logarithm of building permits. Columns (4) to (6) show IV estimates in which we instrument house prices with disposable income, and with the mean January temperature interacted with quarterly dummies. The Cragg-Donald test for weak instruments assumes that under the null the excluded instruments are weakly correlated with the endogenous regressor (house prices). We compare the resulting F-test with the Stock and Yogo (2002) critical value of 12.2, referring to the 5% relative bias. Additional control variables are not reported. Robust heteroskedastic standard errors shown in parentheses. Asterisks, *, **, and ***, denote statistical significance at the 10, 5, and 1% levels.

5.2 Estimated elasticities

The supply elasticity for each MSA is computed as the derivative of Eq. 1 with respect to house prices: $\beta + \lambda * Unaval_i + \delta * WRLURI_i$. The interaction terms allow the elasticities to differ across MSAs, stemming from differences in geographical and regulatory restrictions.¹¹

¹⁰The long-term coefficient is the result of dividing its short-run coefficient by 1 minus the lagged coefficient on the dependent variable; for instance, for the 1996-2006 cycle: $4.4 = 2.256 / (1 - 0.483)$.

¹¹Our baseline estimated elasticities remain qualitatively robust to alternative methods, as measured by the rank correlation between the different specifications: (i) using housing starts instead of building permits as the housing supply measure; (ii) estimating panel dynamic OLS à la Kao and Chiang (2000), who extend the method

When we estimate the model above for the full sample (1987q1-2016q4), we find a strong and positive rank correlation of approximately 0.88 of our estimated elasticities with Saiz (2010)'s own elasticity measure.

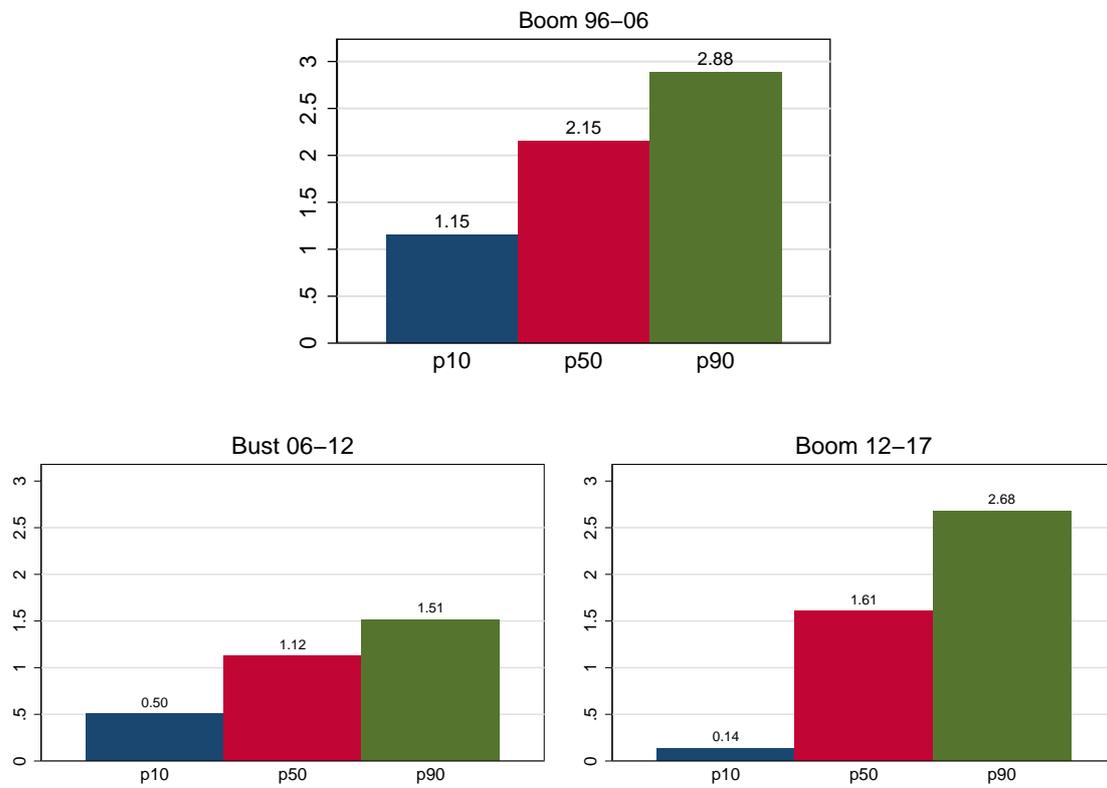
We plot in Figure 7 the estimated elasticities at the median, 10th and 90th percentiles for each housing cycle. First, we find considerable variation over time, with the 1996-2006 housing boom displaying the highest supply elasticities across the whole distribution. This is in line with the view that the increase in house prices during that period was accompanied by a strong supply response from builders. Second, we estimate an across-the-board decline in elasticities during the 2006-12 bust, as a consequence of the downward rigidity of housing stock. Third, despite the recovery in the elasticities during the current boom, they remain well below the 96-06 values. Finally, the dispersion in the supply elasticities has increased in the current cycle, suggesting that the supply response of the different MSAs has become more heterogeneous, and possibly driven more by idiosyncratic factors.

The resulting estimated elasticities are qualitatively similar to those obtained from an OLS specification (Figure B.2 in Appendix B). Nonetheless, the dispersion across MSAs is larger with the IV approach, while the decline in elasticities during the 2006-12 bust is not as large as suggested by the OLS regression.

We shed more light on the heterogeneity between the MSAs by looking at the distribution of the elasticities within the US territory across the two housing booms (Figure 8). Coastal areas located in states such as California, Arizona, Florida, Oregon, and New York are characterised by the lowest elasticities in all housing cycles. This is not surprising given that geographical idiosyncrasies, such as steep ground and bodies of water, make it harder to build and limit the land available for construction in these areas compared to the rest of the country (Saiz 2010). In addition, land-use regulation which can also limit the expansion of housing supply also tends to be more stringent in these coastal areas (Gyourko et al. 2008). In fact, tighter regulation typically reduces the elasticity of housing supply, and thus raises house prices, reduces construction activity, and increases the volatility of house prices (see the review of the literature in Gyourko and Molloy 2015). In contrast, our estimated elasticities show that several areas in the Midwest have the highest elasticities, where builders face relatively fewer restrictions to

of Stock and Watson (1993) for time series to panel data with fixed effects; and (iii) estimating MSA-specific models that take into account other idiosyncratic factors, such as differences in population density, and the level of house prices, that may add additional heterogeneity in the estimated elasticities across metropolitan areas. Regarding the latter, we find that the elasticities from the individual models are in general less precisely estimated, and exhibit larger dispersion.

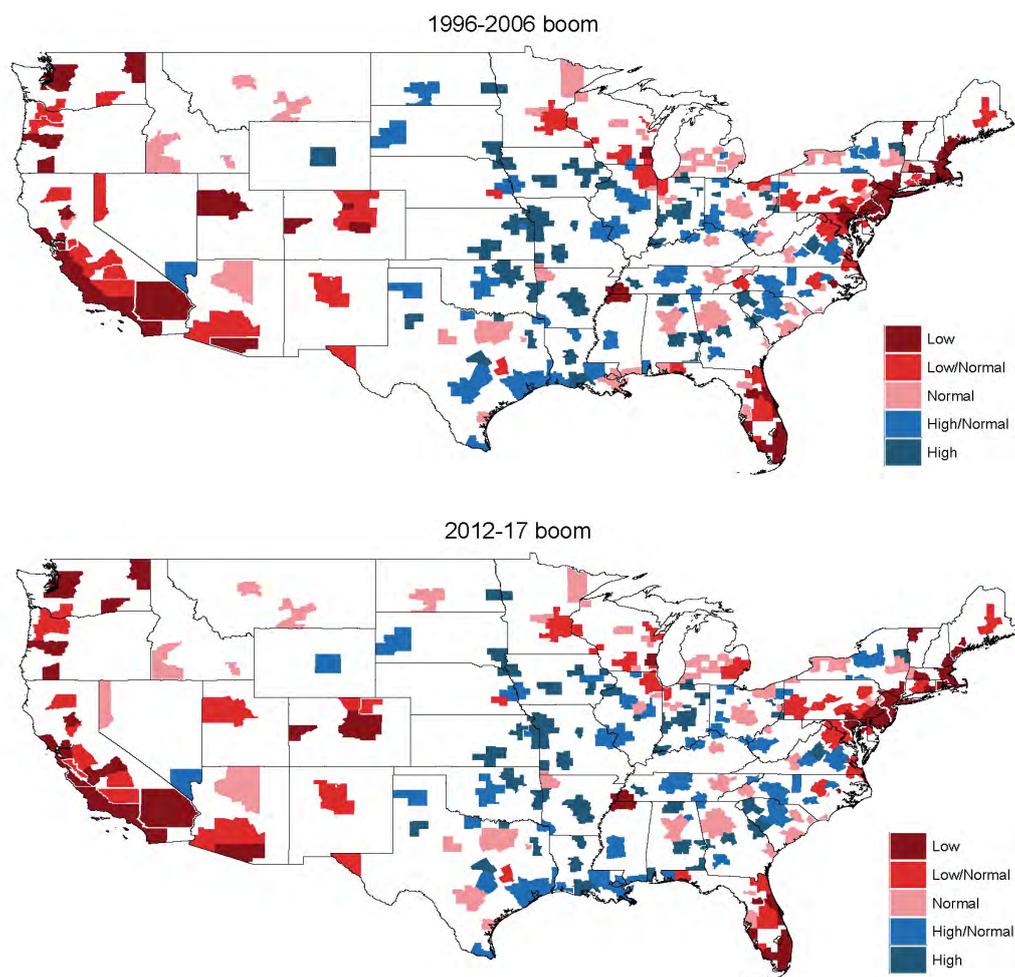
Figure 7: Estimated elasticities: IV specification



Notes: Estimated elasticities from Eq. 1 for the median, 10th and 90th percentiles for each housing cycle. The elasticities are computed from IV regressions using real personal disposable income, and average January temperatures interacted with quarterly dummies as instruments for house prices.

expand housing supply in response to a change in house prices.

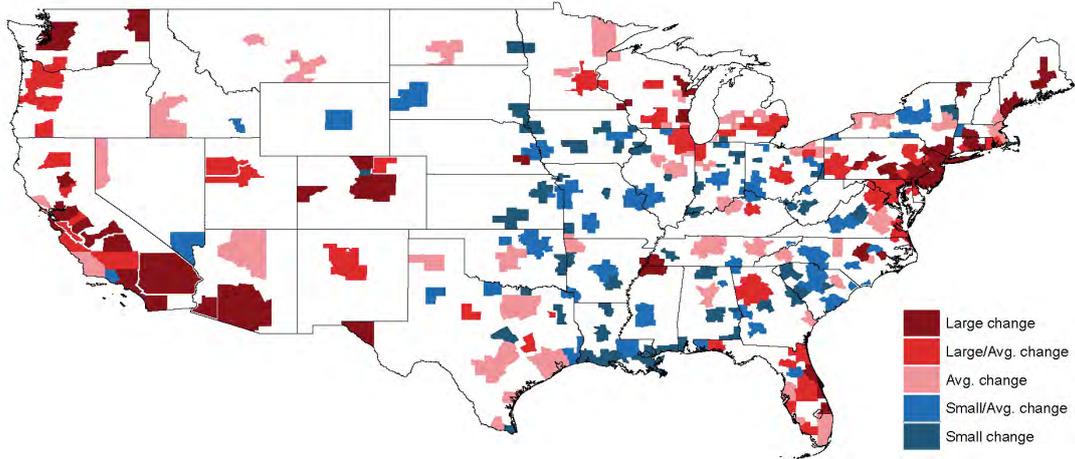
Figure 8: Estimated elasticities for the two housing booms



Notes: Estimated supply elasticities from Eq. 1 for the two housing booms. We split the elasticities for the MSAs into five groups, as represented by the different colours.

The maps also show that despite the elasticities displaying significant time variation, the rank order between the MSAs is relatively stable over time. In other words, even if the elasticities change for each MSA over time, these changes are not large enough to affect significantly the relative ranking of the MSA over the last 20 years. For instance, if an MSA has an estimated low supply elasticity that places it in the bottom part of the elasticity distribution in the first housing boom, chances are that it will remain in the same low-elasticity group over the following cycles, and vice-versa. Having said that, another interesting feature of our estimated supply elasticities is that we find the largest decline between the two booms to take place in those coastal areas that had the lowest elasticities during the first housing boom (Figure 9). We will study this phenomenon in more detail later in the paper.

Figure 9: Change in estimated elasticities between booms



Notes: The map figure shows the change in elasticities between the 1996-2006 boom and the 2012-17 boom. We split the elasticities for the MSAs into five groups, as represented by the different colours.

6 Results

6.1 House price dynamics and elasticities for the cross section

The estimated time-varying elasticities allows us to test empirically the theoretical conjecture made in Section 3 that during a housing boom, or a period of positive demand shocks, the lower the supply elasticities, the higher the volatility and cumulative growth in house prices. By contrast, the house price dynamics should be independent of the elasticities during housing busts, or a negative demand shock.

We test this by exploring the response of house prices for each cycle to the cross-sectional variation in the estimated elasticities:

$$\Delta \log(HPI_{i,c}) = \alpha + \beta Elast_{i,c} + \gamma \Delta Z_{i,c} + \epsilon_{i,c} \quad (2)$$

where the dependent variable is the cumulative change in real house prices for each housing cycle c , $Z_{i,c}$ includes a set of controls, measured in cumulative changes over each cycle, namely real construction wages, real disposable income, population, unemployment rate, non-farm employment, the level of initial house prices-to-income per capita ratio, and the fraction of population with a bachelor degree or higher. We focus on β , which measures the response of house prices to the estimated elasticities, after controlling for MSA-specific economic and social conditions.

We find that a one-standard deviation decline in the estimated elasticities during periods

of booms are associated with a highly statistically and significant increase in real house prices (Table 3). Our estimates also suggest that this association has become weaker in the last boom, with a 12% house price response over 1996-06 versus roughly only 5% over 2012-17. But when we put the last housing boom on a comparable basis, by splitting the first boom into two equally-sized periods, 1996-2001 and 2001-2006, we find that the the elasticities have played a stronger role in explaining the cross-sectional variation in house prices during the current boom, compared with the first five years of the 1966-2006 boom.

Moreover, and also as predicted by theory, the dynamics in house prices during the 2006-11 bust is independent of the supply elasticities, as shown in column (4). This point, however, has been disputed by Huang and Tang (2012), who find that supply restrictions, and by association supply elasticities, played a role in the 2006-09 housing downturn.

Table 3: Cumulative growth in real HPI

	(1)	(2)	(3)	(4)	(5)
	1996-2006	1996-2001	2001-2006	2006-2011	2011-2017
β	-12.039*** (1.425)	-1.532*** (0.528)	-9.119*** (1.249)	0.176 (0.448)	-5.076*** (0.795)
Controls	Yes	Yes	Yes	Yes	Yes
Observations	254	254	254	254	253
R-squared	0.582	0.337	0.584	0.840	0.669

Notes: Regression estimates of Eq. 2, where the dependent variable is the cumulative HPI growth for each housing cycle. For ease of interpretation, the coefficients have been multiplied by the standard deviation of the respective elasticities. Additional control variables are not reported. Robust heteroskedastic standard errors shown in parentheses. Asterisks, *, **, and ***, denote statistical significance at the 10, 5, and 1% levels.

6.2 Time-varying effects of aggregate demand shocks

We have shown that our estimated housing supply elasticities have declined considerably across the board between the two housing booms. What are the main macro implications of such a decline? While the literature does not provide any direct guidance on the matter, given that we are the first, to the best of our knowledge, to have focused on the time-varying nature of supply elasticities, studies based on cross sectional data have put forward the notion that areas with lower elasticities may experience non-negligible output costs. For instance, Hsieh and Moretti (2017) argue that spatial misallocation of labour across US cities related to stringent housing restrictions in highly-productive areas, such as New York and San Francisco Bay Area, has led to output costs over the last 50 years in the United States. In addition, Glaeser and Gyourko

(2018) posit that highly-regulated areas are characterised by higher house prices and smaller population growth relative to the level of demand. Regulation on housing creates an implicit tax on development that is higher in several areas than any reasonable externalities associated with new construction. They suggest that this phenomenon, which leads to overall welfare losses, are pretty much the result of revealed preferences of existing homeowners.

We tackle the issue from a different perspective. By exploring both the cross-section and the time variation in our dataset, we argue that lower housing supply elasticities over time suggest that an aggregate demand shock of the same magnitude can exert stronger effects on house prices in the current housing boom. The rationale for this conjecture builds upon the mechanism of Figure 5 from Section 3, whereby a steepening of the supply curve moves the equilibrium to higher prices and lower quantity. To prove our conjecture, we study the sensitivity of house prices changes to several stylised aggregate demand shocks across the two housing booms: 1996-2006 versus 2012-2017.

The first one is a monetary policy shock computed through a structural VAR (SVAR) along the lines of Christiano et al. (1996) and Sims (1992), using Wu and Xia (2016)'s shadow rate as the policy instrument to address the zero lower bound (ZLB) and unconventional policy after 2007.¹² We identify the shocks through a Cholesky identification scheme with the following order in our 5-variable VAR: real GDP, GDP deflator, real global commodity prices, real house prices, and the shadow rate. As is standard in the VAR literature, the fed funds rate (or the shadow rate in our case) is ordered last, practically implying that monetary policy reacts contemporaneously to all the variables in the VAR, but monetary policy is only allowed to influence macro and financial conditions with a lag of one quarter. All variables are expressed in logs, apart from the shadow rate which is in percent.

Some may argue that the identification of monetary policy shocks from the SVAR may be particularly challenging during the ZLB after 2007, a period characterised by several unconventional monetary policy measures deployed by the Fed. Although we use the shadow rate to deal with that issue – the effective policy rate goes significantly below zero during the ZLB – we cross-check our results with the high-frequency policy news shock from Nakamura and Steinsson (2018). Their shock measure is based on the surprise changes in the fed funds rate and on changes in the path of future interest rates in a 30-minute window around scheduled

¹²Wu and Xia (2016)'s shadow rate measures the effective policy rate in the economy during ZLB periods by translating changes in the Fed's balance sheet into Fed funds rate equivalents. The shadow rate stayed in negative territory for more than six years, from 2009q2 to 2015q4, reaching a minimum of -2.9% in 2014q2.

FOMC announcements. The limitation of using this shock is that it goes until 2014q1, therefore reducing to some extent the number of available observations for the on-going housing boom.

Our third shock variable is an exogenous credit supply shock driven by branching deregulation that leads to an increase in housing demand, in the spirit of Favara and Imbs (2015). In their paper, Favara and Imbs (2015) show that an exogenous expansion in mortgage credit results in a significant increase in house prices, and that this increase is less pronounced in areas with elastic housing supply. The branching deregulation index is taken from Rice and Strahan (2010) and is available since 1994.

Since we are interested in studying how a given aggregate shock affects real house prices in a non-linear fashion, given the different supply elasticities, we resort to Jordà (2005)'s Local Projection method. The Jordà method is indeed more flexible and better suited to capture non-linearities than VARs. Specifically, we estimate a FE model over 1987q1-2016q4 for each horizon $h=1,2,\dots,20$ quarters:

$$\Delta_h HPI_{i,t+h} = \beta^h Shock_{i,t}^h + \gamma^h Shock_{i,t}^h * Elast_{i,t}^h + \sum_{j=1}^4 \lambda_j^h X_{i,t-j}^h + \eta_i^h + \epsilon_{i,t+h} \quad (3)$$

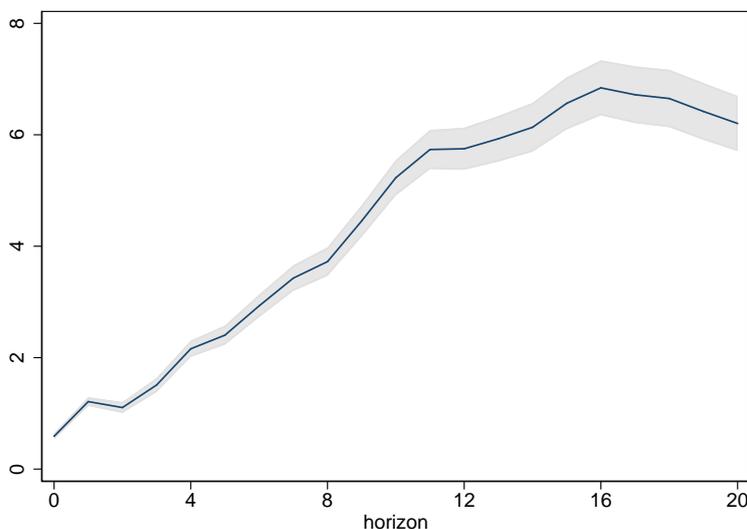
where the dependent variable is the cumulative % increase in real house prices from period t to horizon $t+h$, *Shock* is the aggregate demand shock which is also interacted with our previously estimated time-varying supply elasticities *Elast*, and the control variables $X_{i,t-j}^h$ refer to 4 lags of the change in the logarithm of house prices and of the shock variables. We add fixed effects η_i^h to control for time-invariant idiosyncratic MSA fixed characteristics. Our conjecture is that a positive aggregate demand shock should boost house prices ($\beta > 0$), but this effect should become smaller the higher the housing supply elasticity ($\gamma < 0$).

In line with our conjecture, and with Aastveit and Anundsen (2017), an expansionary monetary policy shock raises house prices over the short and medium term in a statistically significant way: for an exogenous decline of 100 basis points in the fed funds rate, real house prices rise by roughly 6% after 5 years for the mean elasticity MSA (Figure 10). What is novel in our results is that we find that, for a given positive shock, house prices rise by considerably more in the current housing boom, compared with the 1996-2006 boom. For the same 100 basis points decline in the fed funds rate, real house prices are 4 percentage points higher during the current boom (Figure 11). In addition, we show that there is considerable heterogeneity across MSAs. While the previous finding – a given positive aggregate shock hits house prices strongly in the current

boom – holds for all MSAs, we find that a typical low-elasticity MSA, such as San Francisco, is hit harder by a given shock that raises housing demand compared to a typical high-elasticity MSA, such as Kansas City (Figure 12). The main findings are qualitatively robust to using as an alternative the high-frequency policy news shock from Nakamura and Steinsson (2018).

Against this background, we find that lower supply elasticities during the current housing boom makes the housing market more susceptible than during the strong boom in 1996-2006 to stronger house price increases for a given increase in demand. The second finding centres on the effects of a given shock on regional asymmetries; given that we have shown before that low-supply elasticity metro areas experienced a larger decline in their elasticities between the two booms, the effect of a given shock on their housing market is larger in the current period than for high-supply elasticity areas. Accordingly, given the mechanism through which a positive demand shock transmits to prices in the housing market in both areas, the heterogeneity in the response of house prices between the two types of metro areas is even more pronounced than it was before.

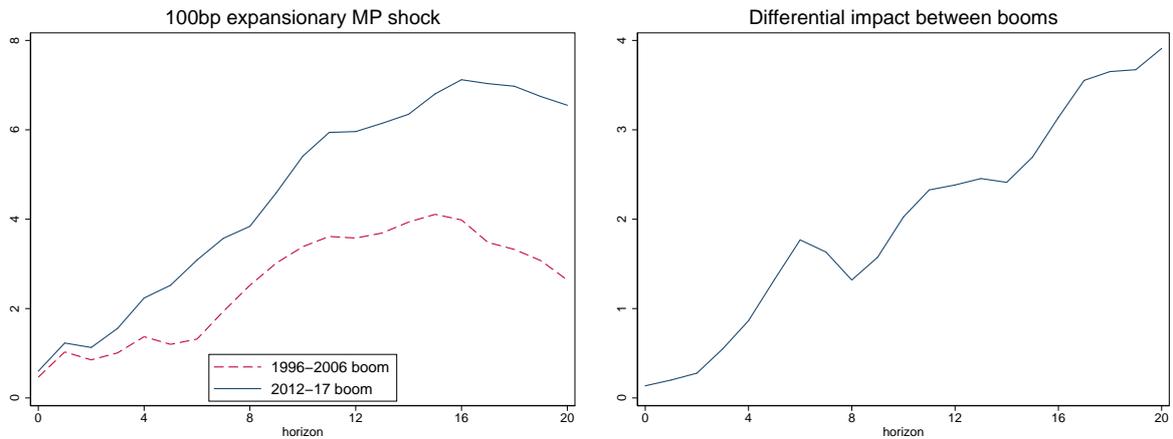
Figure 10: House prices response to an expansionary monetary policy shock



Notes: Cumulative impulse responses of real house prices to a 100 basis points decline in the fed funds rate, together with the associated 90th/10th confidence bands, assessed at the sample mean elasticity.

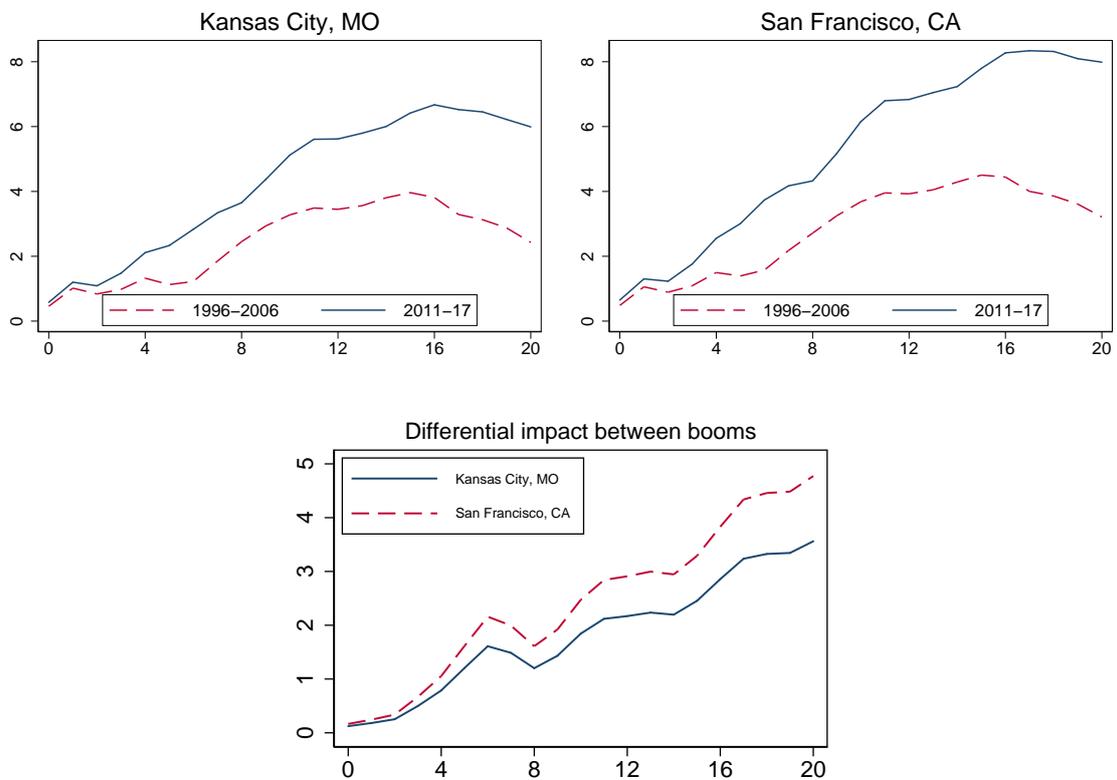
Our last shock is a credit supply shock driven by branching deregulation in the US banking system at the state level. Using the branching deregulation index developed by Rice and Strahan (2010), we find similar evidence as Favara and Imbs (2015), in that a decline in the index (more deregulation) leads to an increase in house prices: house prices rise by a cumulative 8% over 5 years following a one-standard deviation decline in the branching deregulation index (Table B.3 in Appendix B). Moreover, and in line with the findings shown before related to

Figure 11: House prices response to an expansionary monetary policy shock across booms



Notes: Cumulative impulse responses of real house prices to a 100 basis points decline in the fed funds rate, assessed at the sample mean elasticity for each housing boom period. The right-hand chart depicts the difference in the impact on house prices of a monetary policy shock between the 2012-17 and the 1996-2006 booms.

Figure 12: House prices response to an expansionary monetary policy shock for selected MSAs



Notes: Cumulative impulse responses of real house prices to a 100 basis points decline in the fed funds rate for selected MSAs for each housing boom. Kansas City represents a high-supply elasticity MSA, while New York a low-supply elasticity MSA.

an expansionary monetary policy shock, for the same decline in the branching indicator, house prices react significantly by more in the current boom period than it was the case during the previous housing boom (Table B.4 in Appendix B). Finally, we can also square our results with Favara and Imbs (2015)’s findings that an expansion in mortgage credit driven by deregulation leads to a smaller increase in house prices in areas with more elastic housing supply (Table B.5 in Appendix B).

7 Why have elasticities declined?

[VERY INCOMPLETE SECTION]

One of the questions that remains to be answered is related to understanding the factors that have been behind the generalised decline in supply elasticities over the last 20 years. We explore this question more formally by focusing on the decline in the estimated elasticities between two comparable periods in terms of the business cycles dynamics, more specifically the two housing booms. We estimate the following model in differences for the cross section:

$$\begin{aligned} Elast_i^{B2} - Elast_i^{B1} = & \alpha_i + \beta_1(\Delta Pop_i^{B2} - \Delta Pop_i^{B1}) + \beta_2(\Delta Wage_i^{B2} - \Delta Wage_i^{B1}) + \\ & + \beta_3 \Delta HPI_i^{06-12} + \beta_4 Z_i t + \epsilon_i \end{aligned} \quad (4)$$

We regress the change in the estimated elasticities between the two booms on the change in population growth (Pop) and wage growth in the construction sector ($Wage$) between the two booms, on the cumulative change in house price growth during the 2006-12 bust (HPI^{06-12}) and on other controls $Z_i t$, such as the level of house prices and of population. The superscripts $B1$ and $B2$ refer respectively to the first (1996-2006) and second (2012-17) housing booms.

Our results provide strong statistical evidence in favour of population growth and the decline in house prices experienced during the bust as being behind the heterogeneity in the decline in elasticities over time across metro areas. In particular, we find that areas where population growth increased by the most between the two booms, and that experienced the largest decline in house prices at the end of the previous decade, were the same areas where supply elasticities have declined by more (Table 4). We interpret this as evidence that the fear of a new bust have led developers to be less price responsive than previously. The irony of this behaviour is that it may have paved the way for a new housing boom where house prices are more responsive to fluctuations in demand.

Table 4: Δ Elasticity between booms

	(1)	(2)	(3)
$\Delta Pop^{B2} - \Delta Pop^{B1}$	-0.002 (0.002)	-0.003*** (0.001)	-0.003*** (0.002)
$\Delta Wage^{B2} - \Delta Wage^{B1}$	0.001** (0.001)	0.001 (0.001)	0.000 (0.001)
ΔHPI^{06-12}		0.009*** (0.001)	0.009*** (0.001)
Hpinc_pc		-0.007*** (0.001)	-0.007*** (0.001)
Constant	-0.891*** (0.020)	-0.200*** (0.073)	-0.259*** (0.082)
Observations	254	254	254
R-squared	0.028	0.527	0.528

Notes: Regression estimates of Eq. 4, where the dependent variable is the change in the estimated supply elasticities between the 2012-17 and the 1996-2006 housing booms. Robust heteroskedastic standard errors shown in parentheses. Asterisks, *, **, and ***, denote statistical significance at the 10, 5, and 1% levels.

8 Conclusion

In this paper we have provided evidence for a substantial and synchronised decline in local housing supply elasticities from the strong 1996-2006 housing boom to the on-going boom that started in mid-2012. In the current environment of very low supply elasticities, the immediate implication of our finding at the aggregate level is that the house price responsiveness to a given demand shock should be higher today.

We have investigated this conjecture by estimating the effect of an exogenous monetary policy shock on house prices, conditional on different supply elasticities over time. We have found that a similarly sized monetary policy shock has a substantially greater impact on house prices during the current recovery than during the previous boom – estimated to be twice as large. For an area with a median housing supply elasticity, house prices are estimated to increase by nearly 6% after three years following a monetary policy shock that reduces the policy rate by one percentage point. For the same shock, we estimate house prices to increase by around 3% for the 1996-2006 boom.

We have also shed some light on the factors that may have contributed to the decline in supply elasticities between the two booms. Our main findings suggest that areas where population growth increased by the most between the two booms, and that experienced the

largest decline in house prices at the end of the previous decade, were the same areas where supply elasticities have declined by more. Our interpretation is that the fear of a new bust have led builders to be less price responsive than previously. The irony of this behaviour is that it may have paved the way for a new housing boom where house prices are more responsive to fluctuations in demand.

Our findings cast some doubts about the commonly held view that the on-going housing market looks more ‘healthy’ and sustainable compared to the previous boom. In addition, since house prices have become more responsive to an aggregate demand shock, monetary policy makers face additional challenges related to the trade-off between supporting the economic recovery – therefore reinforcing the wealth and home equity channels – and financial stability concerns of ‘creating’ a bubble in the housing market.

References

- Aastveit, K. A. and Anundsen, A. (2017), Asymmetric effects of monetary policy in regional housing markets, Working Paper 2017/25, Norges Bank.
- Anundsen, A. K. and Heebøll, C. (2016), ‘Supply restrictions, subprime lending and regional US house prices’, *Journal of Housing Economics* **31**(C), 54–72.
- Christiano, L. J., Eichenbaum, M. and Evans, C. (1996), ‘The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds’, *The Review of Economics and Statistics* **78**(1), 16–34.
- Davidoff, T. (2016), ‘Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated With Many Demand Factors’, *Critical Finance Review* **5**(2), 177–206.
- Davis, M. A. and Palumbo, M. G. (2008), ‘The price of residential land in large US cities’, *Journal of Urban Economics* **63**(1), 352–384.
- Del Negro, M. and Otrok, C. (2007), ‘99 Luftballons: Monetary policy and the house price boom across U.S. states’, *Journal of Monetary Economics* **54**(7), 1962–1985.
- Favara, G. and Imbs, J. (2015), ‘Credit Supply and the Price of Housing’, *American Economic Review* **105**(3), 958–992.
- Ferreira, F. and Gyourko, J. (2011), Anatomy of the Beginning of the Housing Boom: U.S. Neighborhoods and Metropolitan Areas, 1993-2009, NBER Working Papers 17374, National Bureau of Economic Research, Inc.
- Glaeser, E. and Gyourko, J. (2018), ‘The Economic Implications of Housing Supply’, *Journal of Economic Perspectives* **32**(1), 3–30.
- Glaeser, E. L. and Gottlieb, J. D. (2009), ‘The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States’, *Journal of Economic Literature* **47**(4), 983–1028.
- Glaeser, E. L., Gottlieb, J. D. and Gyourko, J. (2012), Can Cheap Credit Explain the Housing Boom?, in ‘Housing and the Financial Crisis’, NBER Chapters, National Bureau of Economic Research, Inc, pp. 301–359.
- Glaeser, E. L. and Gyourko, J. (2005), ‘Urban Decline and Durable Housing’, *Journal of Political Economy* **113**(2), 345–375.

- Glaeser, E. L., Gyourko, J., Morales, E. and Nathanson, C. G. (2014), ‘Housing dynamics: An urban approach’, *Journal of Urban Economics* **81**(C), 45–56.
- Glaeser, E. L., Gyourko, J. and Saiz, A. (2008), ‘Housing supply and housing bubbles’, *Journal of Urban Economics* **64**(2), 198–217.
- Glick, R., Lansing, K. J. and Molitor, D. (2015), ‘What’s different about the latest housing boom?’, *FRBSF Economic Letter* .
- Green, R. K., Malpezzi, S. and Mayo, S. K. (2005), ‘Metropolitan-Specific Estimates of the Price Elasticity of Supply of Housing, and Their Sources’, *American Economic Review* **95**(2), 334–339.
- Gyourko, J. and Molloy, R. (2015), *Regulation and Housing Supply*, Vol. 5 of *Handbook of Regional and Urban Economics*, Amsterdam: Elsevier, pp. 1289–1337.
- Gyourko, J., Saiz, A. and Summers, A. (2008), ‘A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index’, *Urban Studies* **45**(3), 693–729.
- Hernández-Murillo, R., Owyang, M. T. and Rubio, M. (2017), ‘Clustered Housing Cycles’, *Regional Science and Urban Economics* **66**, 185–197.
- Hsieh, C.-T. and Moretti, E. (2017), *Housing Constraints and Spatial Misallocation*, NBER Working Papers 21154, National Bureau of Economic Research, Inc.
- Huang, H. and Tang, Y. (2012), ‘Residential land use regulation and the US housing price cycle between 2000 and 2009’, *Journal of Urban Economics* **71**(1), 93–99.
- Jordà, O. (2005), ‘Estimation and Inference of Impulse Responses by Local Projections’, *The American Economic Review* **95**(1), 161–182.
- Kao, C. and Chiang, M.-H. (2000), *On the Estimation and Inference of a Cointegrated Regression in Panel Data*, Vol. 15 of *Advances in Econometrics: Nonstationary Panels, Panel Cointegration and Dynamic Panels*, pp. 179–222.
- Kaplan, G., Mitman, K. and Violante, G. L. (2016), *Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data*, NBER Working Papers 22232, National Bureau of Economic Research, Inc.

- Mian, A., Rao, K. and Sufi, A. (2013), ‘Household Balance Sheets, Consumption, and the Economic Slump’, *The Quarterly Journal of Economics* **128**(4), 1687–1726.
- Mian, A., Sufi, A. and Verner, E. (2017), How do Credit Supply Shocks Affect the Real Economy? Evidence from the United States in the 1980s, NBER Working Papers 23802, National Bureau of Economic Research, Inc.
- Nakamura, E. and Steinsson, J. (2018), ‘High Frequency Identification of Monetary Non-Neutrality: The Information Effect’, *Quarterly Journal of Economics* (forthcoming).
- Piazzesi, M. and Schneider, M. (2009), ‘Momentum Traders in the Housing Market: Survey Evidence and a Search Model’, *American Economic Review* **99**(2), 406–411.
- Piazzesi, M. and Schneider, M. (2016), *Housing and Macroeconomics*, Vol. 2 of *Handbook of Macroeconomics*, Elsevier, pp. 1547–1640.
- Rice, T. and Strahan, P. E. (2010), ‘Does Credit Competition Affect Small-Firm Finance?’, *Journal of Finance* **65**(3), 861–889.
- Saiz, A. (2007), ‘Immigration and housing rents in American cities’, *Journal of Urban Economics* **61**(2), 345–371.
- Saiz, A. (2010), ‘The Geographic Determinants of Housing Supply’, *The Quarterly Journal of Economics* **125**(3), 1253–1296.
- Sims, C. A. (1992), ‘Interpreting the macroeconomic time series facts : The effects of monetary policy’, *European Economic Review* **36**(5), 975–1000.
- Stock, J. H. and Watson, M. W. (1993), ‘A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems’, *Econometrica* **61**(4), 783–820.
- Stock, J. H. and Yogo, M. (2002), Testing for Weak Instruments in Linear IV Regression, NBER Technical Working Papers 0284, National Bureau of Economic Research, Inc.
- Wu, J. C. and Xia, F. D. (2016), ‘Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound’, *Journal of Money, Credit and Banking* **48**(2-3), 253–291.

A Data

MSA-level data

Building permits: number of permits issued by a local jurisdiction to proceed on a construction project, i.e. includes housing units issued in local permit-issuing jurisdictions by a building or zoning permit. Source: Census Bureau, and Moody's Analytics.

Housing starts: number of housing units in which construction work has started. The start of construction is when excavation begins for the footings or foundation of a building. Source: Census Bureau, and Moody's Analytics.

Housing completions: number of completed housing units. A one-unit structure is defined as completed when all finished flooring has been installed or when it has been occupied before all construction is finished. In buildings with two or more housing units, the building is counted as completed when 50% or more of the units are occupied or available for occupancy. Source: Census Bureau, and Moody's Analytics.

Housing stock: a house, apartment, mobile home or trailer, a group of rooms, or a single room that is occupied or available for occupancy. Updated from 2010q3 onwards with housing completions. Source: Census Bureau, and Moody's Analytics.

FHFA house price index: weighted, repeat-sales index, measuring average price changes in repeat sales or refinancings on the same single-family properties whose mortgages have been purchased or securitised by Fannie Mae or Freddie Mac. Source: Federal Housing Finance Agency, and Moody's Analytics.

CPI: consumer price index for all urban consumers. Source: BLS, and Moody's Analytics.

UNAVAIL: the land unavailability index captures housing supply geographical constraints. It is constructed using topographic maps measuring the proportion of land in a 50 km radius of the city centre that is lost to steep slopes and water bodies, such as oceans, rivers, lakes and wetlands. Source: Saiz (2010).

WRLURI: the Wharton local land-use regulatory index captures regulatory restrictions in the housing market, i.e. measures the time and financial cost of acquiring building permits and constructing a new home. It refers to the principal component of 11 survey-based measures which is interpreted as the degree of stringency of local zoning laws. Source: Gyourko et al. (2008).

Average January temperatures: Monthly average temperatures in January over 1941-1970, measured in Fahrenheit degrees. We aggregate the original county level data to the MSA level. Source: United States Department of Agriculture (USDA) Economic Research Service Natural Amenities Scale Database.

Disposable personal income: The income available to persons for spending or saving. It is equal to personal income less personal current taxes. Source: BEA, and Moody's Analytics.

Unemployment rate: the number of unemployed as a % of total labour force. Data available since 1990q1. Source: BLS, and Moody's Analytics.

Non-farm employment: total number of people employed, excluding farm workers, private household employees, or non-profit organization employees. Data available since 1990q1. Source: BLS, and Moody's Analytics.

Mortgage originations: dollar amount of new mortgage loans approved by the mortgage broker or loan officer. Data available since 1990q1. Source: Home Mortgage Disclosure Act (HMDA), and Moody's Analytics.

Population: resident population in each MSA. Data available up to 2016q2; assumed to be constant until 2016q4. Source: Census Bureau, and Moody's Analytics.

Population density: population per square mile. Annual data since 2000. Source: Census Bureau, and Moody's Analytics.

Bachelor: fraction of population with bachelor degree or higher. Annual data since 1988. Source: Census Bureau, and Moody's Analytics.

State-level data

Construction wages: wages and salaries in the construction sector. Data available since 1990q1. The original quarterly series has been adjusted for seasonality using X-12-ARIMA from the Census Bureau. Source: BEA.

Deregulation index: measure of state-level banking deregulation that is based on the number of years since deregulation began in a specific state. It is computed as the average of inter-state and intra-state deregulation. A higher measure indicates more deregulation. The index is not available for South Dakota and Delaware. Source: Mian et al. (2017).

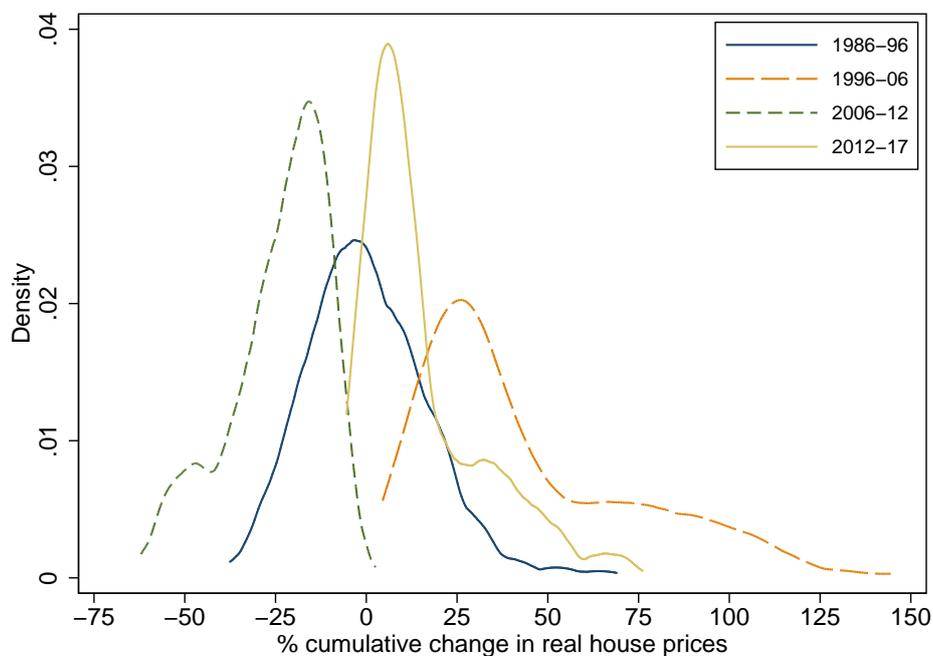
B Tables and figures

Table B.1: Descriptive statistics

	Obs	Mean	Std. Dev.	Min	Max
Real HPI (logs)	30,033	4.7	0.2	4.1	5.5
Building permits (logs)	30,480	7.3	1.4	1.0	12.1
Housing starts (logs)	30,480	7.4	1.4	0.8	11.6
Housing stock (logs)	30,480	5.1	1.1	3.2	9.0
Real personal income (logs)	30,480	16.3	1.2	14.0	20.8
UNAVAL	30,480	0.3	0.2	0.0	0.9
WRLURI	30,480	-0.1	0.8	-1.8	4.3
Deregulation index	18,768	2.2	1.5	0.0	4.0
Mean Jan. temperature (°F)	29,160	35.6	12.5	5.9	65.2
Δ CPI (%)	30,480	1.1	1.1	-7.0	8.2
Δ Real HPI (%)	29,992	0.1	1.9	-27.9	22.4
Δ Real personal income (%)	30,480	0.7	1.2	-8.9	11.9
Δ Unemployment rate	27,178	0.0	0.4	-8.4	8.1
Δ Non-farm employment (%)	27,999	0.3	0.8	-25.4	7.7
Δ Real construction wages (%)	27,178	0.5	3.2	-19.2	17.1
Δ Bachelor (yoy %)	28,444	0.4	1.4	-9.4	8.0
Δ Population (%)	29,972	0.3	0.4	-44.3	10.2
Δ Pop. density (yoy %)	16,216	0.9	1.1	-28.7	6.1

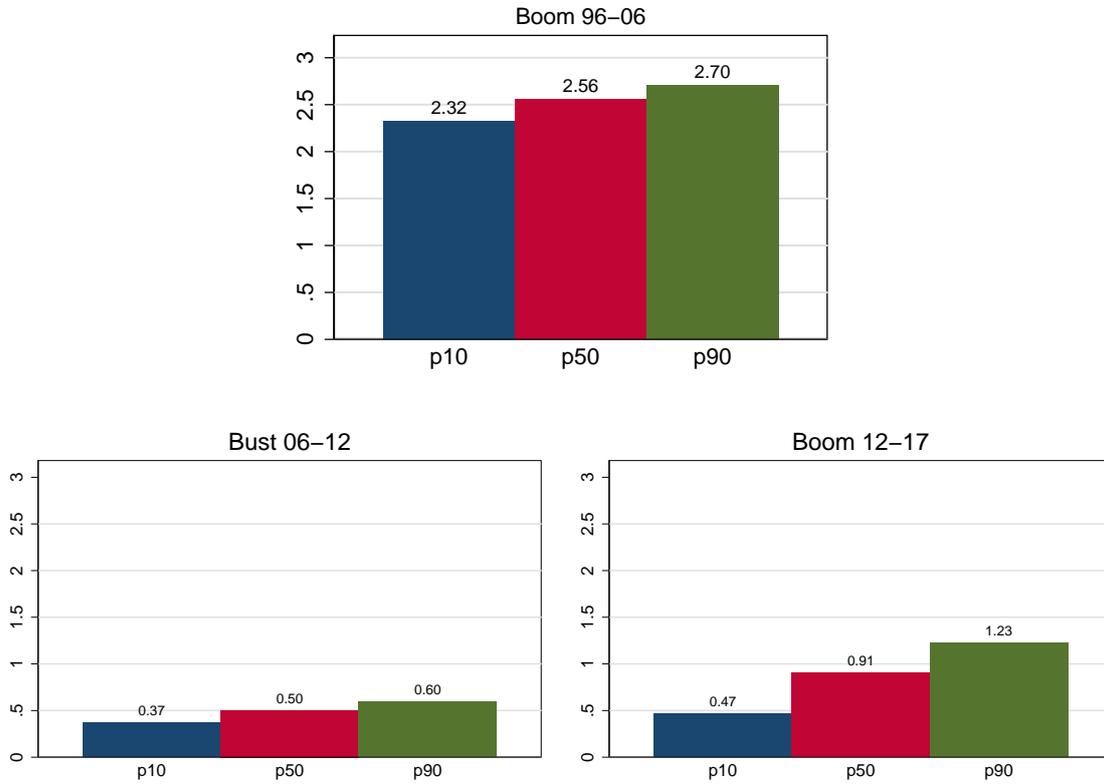
Sources: Bureau of Economic Analysis, Bureau of Labor Statistics, Census Bureau, Federal Housing Finance Agency, Gyourko et al. (2008), Moody's Analytics, Saiz (2010), Rice and Strahan (2010), and United States Department of Agriculture,.

Figure B.1: Kernel densities



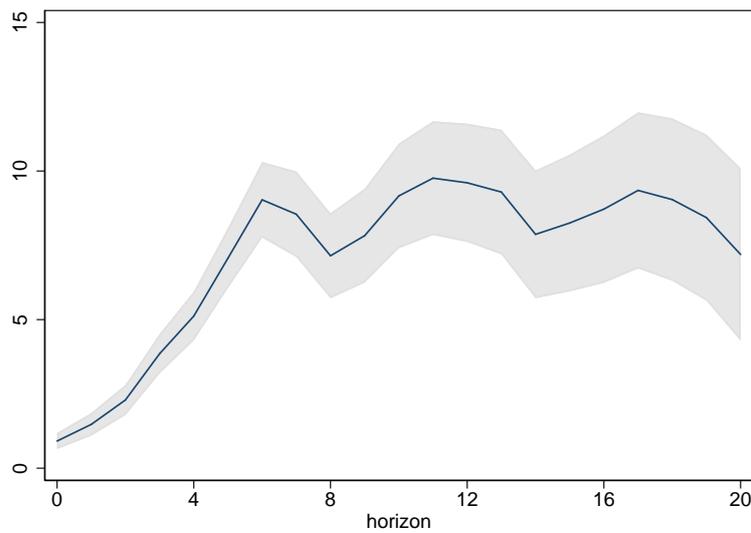
Notes: Kernel densities for the % cumulative change in real house prices for the local housing market cycles.

Figure B.2: Estimated elasticities: OLS specification



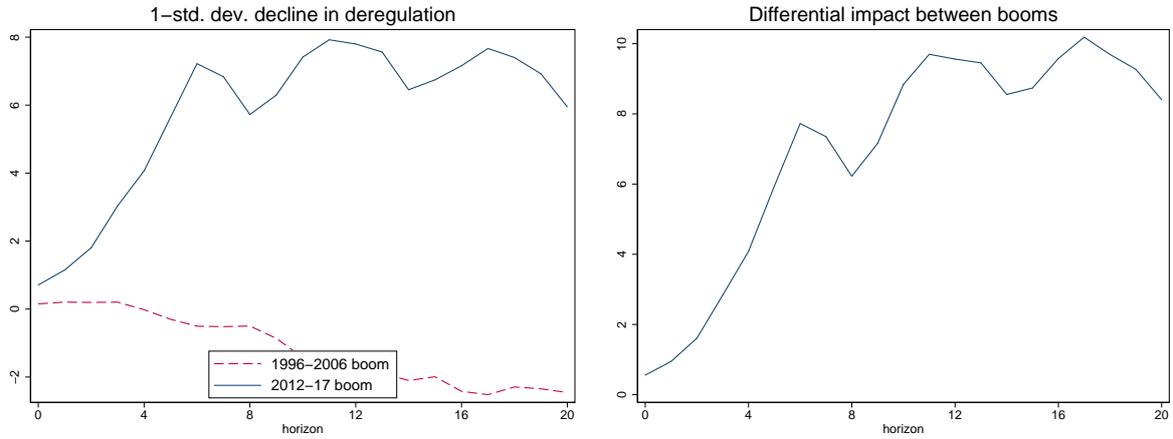
Notes: Estimated elasticities from Eq. 1 for the median, 10th and 90th percentiles for each housing cycle.

Figure B.3: House prices response to a credit supply shock



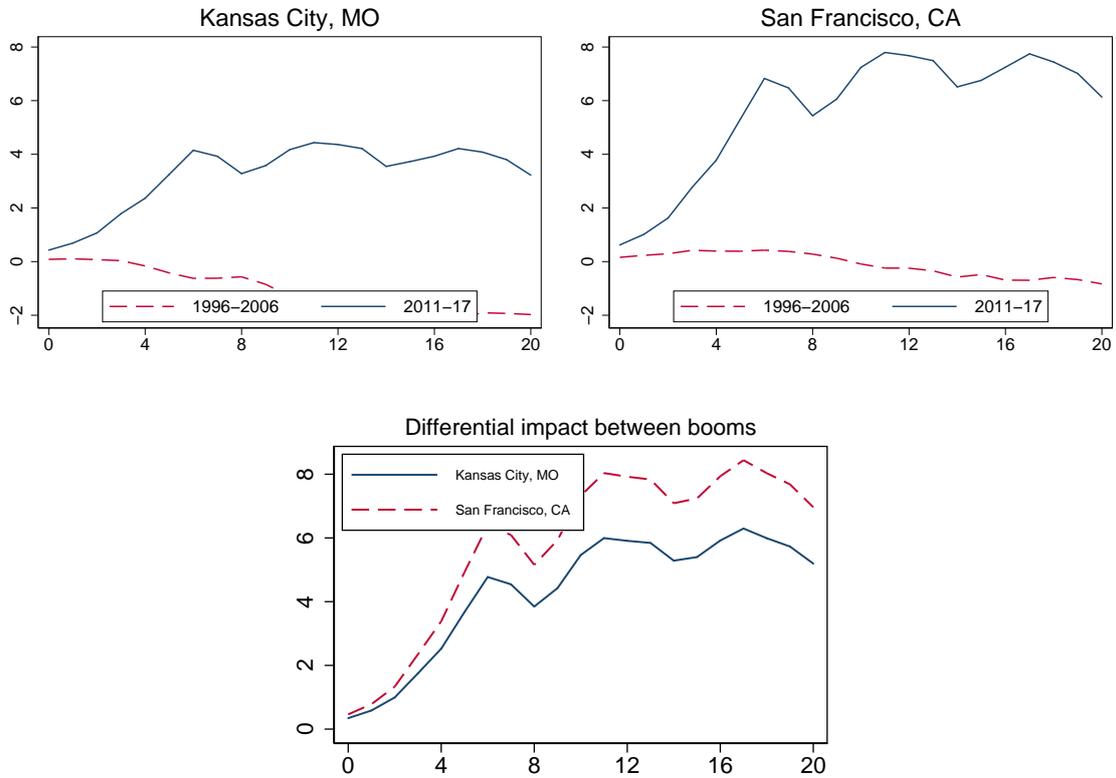
Notes: Cumulative impulse responses of real house prices to a 1-standard deviation decline in the branching deregulation index constructed by Rice and Strahan (2010), together with the associated 90th/10th confidence bands, assessed at the sample mean elasticity.

Figure B.4: House prices response to a credit supply shock across booms



Notes: Cumulative impulse responses of real house prices to a 1-standard deviation decline in the branching deregulation index constructed by Rice and Strahan (2010), assessed at the sample mean elasticity for each housing boom period. The right-hand chart depicts the difference in the impact on house prices of a credit supply shock between the 2012-17 and the 1996-2006 booms.

Figure B.5: House prices response to a credit supply shock for selected MSAs



Notes: Cumulative impulse responses of real house prices to a 1-standard deviation decline in the branching deregulation index constructed by Rice and Strahan (2010) for selected MSAs and for each housing boom. Kansas City represents a high-supply elasticity MSA, while New York a low-supply elasticity MSA.