Review of Economic Studies (2021) **88**, 669–707 doi:10.1093/restud/rdaa018

Advance access publication 30 April 2020

Housing Wealth Effects: The Long View

ADAM M. GUREN
Boston University and NBER
ALISDAIR MCKAY
Federal Reserve Bank of Minneapolis
EMI NAKAMURA AND JÓN STEINSSON
University of California, Berkeley and NBER

First version received November 2018*; Editorial decision March* 2020*; Accepted April* 2020 (*Eds.*)

We provide new time-varying estimates of the housing wealth effect back to the 1980s. We use three identification strategies: ordinary least squares with a rich set of controls, the Saiz housing supply elasticity instrument, and a new instrument that exploits systematic differences in city-level exposure to regional house price cycles. All three identification strategies indicate that housing wealth elasticities were if anything slightly smaller in the 2000s than in earlier time periods. This implies that the important role housing played in the boom and bust of the 2000s was due to larger price movements rather than an increase in the sensitivity of consumption to house prices. Full-sample estimates based on our new instrument are smaller than recent estimates, though they remain economically important. We find no significant evidence of a boom–bust asymmetry in the housing wealth elasticity. We show that these empirical results are consistent with the behaviour of the housing wealth elasticity in a standard life-cycle model with borrowing constraints, uninsurable income risk, illiquid housing, and long-term mortgages. In our model, the housing wealth elasticity is relatively insensitive to changes in the distribution of loan-to-value (LTV) for two reasons: first, low-leverage homeowners account for a substantial and stable part of the aggregate housing wealth elasticity; second, a rightward shift in the LTV distribution increases not only the number of highly sensitive constrained agents but also the number of underwater agents whose consumption is insensitive to house prices.

Key words: House prices, Consumption, MPC, Housing wealth effect, Great recession

JEL Codes: E21, E32, R21

1. INTRODUCTION

Housing wealth effects are widely believed to have played an important role in the boom of the early 2000s and the recession that followed. Recent estimates indicate that the sensitivity of economic activity to house prices—which we refer to as the housing wealth elasticity—was quite large during this period (Mian and Sufi, 2011; Mian *et al.*, 2013; Mian and Sufi, 2014). The question we seek to answer in this article is whether this evidence from the 2000s boom–bust housing cycle is representative of the magnitude of housing wealth effects more generally or whether this episode was "special" in some way.

The editor in charge of this paper was Veronica Guerrieri.

The 2000s saw a large run up and subsequent decline in aggregate house prices, which led housing to play an unusually large role in driving the business cycle over this period. However, the 2000s also saw a variety of changes in housing markets that may have amplified the sensitivity of the economy to house prices. Lax credit standards during the boom and large increases in the number of constrained households as loan-to-value (LTV) ratios rose during the bust may have amplified the magnitude of the housing wealth elasticity over this period. Whether these changes had important implications for the aggregate housing wealth elasticity is unclear because prior work provides little guidance on how the housing wealth elasticity has varied over time.[1]

To shed light on this issue, we provide new time-varying estimates of the housing wealth elasticity back to the 1980s. These estimates indicate that the housing wealth elasticity was not larger during the 2000s boom–bust housing cycle than in other parts of our sample. If anything, it was smaller. These results provide no support for the notion that economic activity was more sensitive to house prices in the 2000s than before. The large role played by housing in the business cycle of the 2000s seems to have been exclusively a consequence of the large changes in house prices over this period. We also investigate whether the housing wealth elasticity is larger when house prices are falling than when house prices are rising, perhaps because more households hit a borrowing constraint during housing busts. We find no statistically significant evidence of such a boom–bust asymmetry. We show that these empirical results are consistent with the behaviour of the housing wealth elasticity in a standard life-cycle model with borrowing constraints, uninsurable income risk, illiquid housing, and long-term mortgages.

Estimating the housing wealth elasticity is challenging because house prices and economic activity are jointly determined and causation can run in both directions, potentially leading to a substantial upward bias of ordinary least squares (OLS) estimates. Measurement error in local house prices is a second potentially important source of bias that may offset the first. Recent work has addressed these challenges by using Saiz (2010) city-level estimates of housing supply elasticities as an instrument for the change in house prices in different cities during the 2000s boom or bust (e.g., Mian *et al.*, 2013; Mian and Sufi, 2014). This work has typically used an IV regression on a single cross-section to evaluate the housing wealth elasticity.

This empirical strategy has two potentially important shortcomings that we seek to address. First, the Saiz instrument has been shown to be correlated with other city characteristics (Davidoff, 2016). This raises the concern that cities with lower housing supply elasticities as measured by the Saiz instrument might be generally more cyclical due to differences in other characteristics. For instance, they may have different industrial composition, differential exposure to risk premia, or differential exposure to secular trends, such as an increase in housing demand in coastal areas with inelastic supply according to the Saiz instrument. We overcome this important challenge by employing a panel specification, which allows us to control for city specific trends, differential sensitivity to regional business cycles, and other controls including industry shares with time-specific coefficients.

A second weakness of the Saiz instrument is that it loses power before 2000, making it difficult to judge whether the housing wealth elasticity has changed over time. We address this challenge

1. To our knowledge, only two papers have looked at changes in the housing wealth elasticity over time. First, Case *et al.* (2013) find that the wealth effect was larger after 1986 than before using an OLS approach. Second, Aladangady (2017) finds that housing wealth effects pre-2002 are not significantly different from post-2002, although his estimates are imprecise. Finally, by comparing Case *et al.* (2005), which uses data for 1982–99, and Case *et al.* (2013), which covers 1978–2009 and has a higher estimate, one can attempt to back out the effect of adding the 2000s (along with 1978–82) to the sample. However, the two estimates are not in fact directly comparable, since both the econometrics and data are different between the two papers. Other empirical estimates for the recent period include Hurst and Stafford (2004), Campbell and Cocco (2007), Carroll *et al.* (2011), Attanasio *et al.* (2009), 2011 Attanasio *et al.* (2011), Calomiris *et al.* (2013), Cooper (2013), DeFusco (2018), Kaplan *et al.* (2020), and Liebersohn (2017).

by developing a new instrument for city-level house price changes. The Saiz instrument is based only on variation in land unavailability and regulation and is therefore a relatively weak predictor of house price movements. Our new instrument is based on a new proxy for housing supply elasticities that builds on earlier work of Palmer (2015) by exploiting the fact that house prices in some cities are systematically more sensitive to regional house-price cycles than are house prices in other cities. For example, when a house price boom occurs in the Northeast region, Providence systematically experiences larger increases in house prices than Rochester.

We construct our instrument by first estimating the systematic historical sensitivity of local house prices to regional housing cycles and then interacting these historical sensitivity estimates—which we interpret as proxies of housing supply elasticities—with today's shock to regional house prices. We refer to this instrument as a sensitivity instrument. The basic shift-share structure of our sensitivity instrument is the same as that of the Saiz instrument (and similar to the well-known Bartik instrument) but with a different proxy for the housing supply elasticity. This approach infers the housing wealth elasticity from the differential response of economic activity in cities like Providence relative to cities like Rochester when the Northeast region experiences a housing boom or bust.

We refine this approach to account for the fact that Providence and Rochester may exhibit systematic differences in sensitivity to aggregate shocks for non-housing reasons by estimating the sensitivity parameter using only the *residual* variation in house prices after controlling for local economic conditions. Importantly, our approach does not rely on regional house price variation being exogenous. In fact, regional house price variation can be driven by the same shocks that drive regional economic activity.[2] The main identifying assumption is that conditional on the many controls we discuss above, there is no unobserved factor that is both correlated with house prices in the time series and that differentially affects the same cities that are more historically sensitive to regional housing cycles in the cross section.

We use retail employment as our main dependent variable and proxy for consumer expenditures. This is a relatively standard choice in the measurement literature. For example, this is the approach taken by the BEA's regional income and product accounts and private sector organizations such as Moody's and the Survey of Buying Power. Retail employment comoves strongly with the BEA's PCE measure of consumption at the aggregate level. Indeed, the comovement is considerably stronger than between PCE and an aggregate of the Consumer Expenditure Survey. Changes in retail technology have had little impact on this relationship, as we show in Section 2—the role of retail employment as an intermediate input into purchases appears relatively stable over our sample period. Retail employment data are particularly well suited to our application because they provide long-term geographically disaggregated series, which are unavailable for other consumer expenditure proxies. Retail employment is also of interest in its own right as a measure of local non-tradeable economic activity (*e.g.* Mian and Sufi, 2014).[3]

Our main empirical finding about the evolution of the housing wealth elasticity over time holds for three different identification strategies: simple OLS with a rich set of controls, a panel version of the Saiz instrument, and our new sensitivity instrument. This result can therefore not be attributed to special features of any one identification strategy. For OLS and our sensitivity

2. The recent literature on general equilibrium models of house prices has emphasized shocks to current and expected future productivity, credit constraints, and risk premia as plausible sources of variation in house prices (see *e.g.* Landvoigt *et al.*, 2015; Favilukis *et al.*, 2017; Kaplan *et al.*, 2019).

3. The existing literature on housing wealth elasticities uses a variety of dependent variables, ranging from particular consumption categories such as consumer packaged goods or cars (*e.g.* Mian and Sufi, 2011; Kaplan *et al.* 2020), to credit card spending (*e.g.* Mian *et al.*, 2013) to broader measures based on the Current Expenditure Survey (Aladangady, 2017).

instrument, the housing wealth elasticity is statistically significantly *smaller* over the boom and bust (2000–12) period than for the rest of our sample.

While OLS, the Saiz instrument, and our sensitivity instrument yield similar results regarding changes over time in the housing wealth elasticity, they differ when it comes to the overall level of the housing wealth elasticity and the precision of these estimates. Estimates based on our sensitivity instrument are substantially smaller and more precisely estimated than those based on the Saiz instrument. Our sensitivity instrument yields a pooled elasticity estimate for retail employment over the sample period 1990–2017 of 0.072, while the Saiz instrument yields an estimate of 0.146 over this same sample period. These estimates are roughly equivalent to marginal propensities to consume out of housing wealth (MPCH) of 3.3 cents on the dollar and 6.5 cents on the dollar, respectively.[4]

Our sensitivity instrument is a more powerful predictor of local house prices than the Saiz instrument. As a consequence, it generates more precise estimates. The statistical power of our sensitivity instrument is a result of the fact that regional housing cycles explain roughly 40% of the variation in local house prices even after controlling for local economic conditions. Our sensitivity instrument implicitly captures many determinants of housing supply elasticity other than land-unavailability as measured by Saiz (2010). Since most potential confounders bias estimates of the housing wealth elasticity upward, it is comforting that our sensitivity instrument yields a lower estimate than both the Saiz instrument and OLS.

Theoretically-minded readers may find it hard to interpret the causal effect of house prices on consumption. House prices are equilibrium variables that are affected by many shocks which may affect consumption through other channels. So, what do our empirical estimates capture? In Section 5, we discuss how in a simple general equilibrium model in which all markets are regional except for housing markets, which are local, our empirical approach yields an estimate of the partial equilibrium effect of house prices on consumption. In this case, both the direct effects of the shocks that drive aggregate variation in house prices and all general equilibrium effects are soaked up by the region-time fixed effects in our regressions. We also discuss a more realistic general equilibrium model with segmented markets across cities (presented in more detail in Guren *et al.* (2019)) in which our empirical approach yields an estimate of the partial equilibrium effect of house prices on consumption multiplied by a local general equilibrium multiplier that can be obtained from the literature on fiscal stimulus (*e.g.* Nakamura and Steinsson, 2014).

We next develop a partial equilibrium model of housing wealth effects to help understand our empirical results. This model builds heavily on a recent literature that has incorporated illiquid housing and long-term mortgages into models with uninsurable income shocks and borrowing constraints.[5] In contrast to earlier models, this class of models can generate large housing wealth elasticities. Our calibrated model generates an average housing wealth elasticity of 0.09, roughly in line with what we estimate in the data.

We show that this model implies that the aggregate housing wealth elasticity is insensitive to the observed changes in household LTV ratios over our sample period and to large variation in credit constraints. Two features of the model are important to understand these theoretical results. First, the housing wealth elasticity is substantial and stable for households with relatively low LTVs. Moreover, the level of the housing wealth elasticity is relatively insensitive to LTV for low LTVs (below 0.6). The distribution of LTVs can, therefore, shift substantially within

4. For comparison, Mian *et al.* (2013) estimate an MPCH of 7.2 cents on the dollar during the bust of the 2000s housing cycle and Mian and Sufi's (2014) estimate an MPCH of between 4.1 and 7.3 cents on the dollar during this same bust.

5. See, *e.g.*, Agarwal *et al.* (2020), Berger *et al.* (2018), Roussanov *et al.* (2020), Davis and Van Nieuwerburgh (2015), Gorea and Midrigan (2018), Guren *et al.* (2019), Kaplan *et al.* (2019), Li and Yao (2007), Favilukis *et al.* (2017).

this low-LTV region without having a quantitatively significant effect on the aggregate housing wealth elasticity. The significant number of homeowners with low LTVs thus not only increases the aggregate housing wealth elasticity but also stabilizes it. As described by Berger *et al.* (2018), large housing wealth elasticities arise even at low levels of leverage in models with incomplete markets because households respond more strongly to the appreciation of their home than they do to the increase in implicit future rents.[6]

A second key point in understanding our theoretical results is that, because mortgages are long-term contracts, households are not forced to de-lever to satisfy an ongoing LTV constraint in a housing bust. Since negative equity households cannot access changes in their housing equity and are not forced to delever, their consumption is largely unresponsive to changes in home prices, as Ganong and Noel (2020) and Berger *et al.* (2018) have shown. We apply this idea to the large rightward shift in the LTV distribution that resulted from the fall in prices during the 2007–10 housing bust. This shift had two offsetting effects on the housing wealth elasticity. On the one hand, more households were pushed closer to their LTV constraint and consequently became more sensitive to changes in house prices. On the other hand, more households were pushed underwater on their mortgage to the point that they became insensitive to changes in house prices. Quantitatively, these two effects roughly offset to deliver a relatively stable aggregate elasticity in the housing bust. By contrast, in a model with short-term debt, the housing wealth elasticity rises sharply in the bust as households are forced to de-lever, which is at odds with our empirical results.

Some may find it surprising to learn that households were spending out of their home equity over a quarter century ago. However, the main tools used to extract housing equity—such as cash-out refinancing and HELOCs—have been available for several decades, and the HELOC share of mortgage debt only rose from 7% to 9% in the 2000s boom according to the Flow of Funds. Mortgage securitization was invented in the late 1960s and has been done on a large scale since the late 1970s. Others have argued that the major changes in mortgage debt availability occurred in the 1970s (see *e.g.* Foote *et al.*, 2012; Kuhn *et al.*, 2017). While certain mortgage products may have become available in the 2000s to segments of the population that did not have access to them before, our model shows that this is not likely to have materially affected the overall housing wealth effect. The following quote from Townsend–Greenspan's August 1982 client report written by Alan Greenspan illustrates well how much access households had to housing equity even before the start of our sample period.

> The combination of very rapidly rising prices for existing homes and a sharp increase in sales ... of these homes has created a huge increase in capital gains and purchasing power during the past two years ... by far the greater part has been drawn out of home equities and spent on other goods and services or put into savings. In fact, of the more than $60 billion ... increase in the market value of existing homes ... virtually the entire amount was monetized as mortgage debt extensions, creating nearly a 5% increase in consumer purchasing power.

A modern reader might be excused for thinking that this paragraph was written by Greenspan circa 2005.[7]

6. This contrasts with the well-known benchmark of Sinai and Souleles (2005) in which these two effects cancel exactly.

7. See Mallaby (2016) for further discussion of this point. We thank Sebastian Mallaby for helping us obtain the original copy of this report. Mallaby writes that Greenspan's calculations were based on direct estimates of home equity

The article proceeds as follows. Section 2 describes our main data sources. Section 3 describes our empirical methodology. Section 4 describes our empirical results. Section 5 makes explicit the link between our empirical analysis and the theoretical analysis that follows. Section 6 presents our partial equilibrium model. Section 7 analyses how changes in household balance sheets affect the housing wealth elasticity in the model. Section 8 concludes.

2. DATA

Our main measure of local economic activity is retail employment per capita. Retail employment is an interesting outcome variable in its own right. In addition, retail employment has long been viewed by measurement agencies as one of the best available proxies for consumer expenditures. For example, the BEA's Regional PCE measures and the private sector "Survey of Buying Power" both use retail employment data to impute consumer expenditures between economic census years. Private sector measures of consumer expenditures also use retail employment as a proxy. For example, Case *et al.* (2005, 2013) use data from Regional Financial Associates (now Moody's Economy.com) that is imputed in part from retail employment data.[8]

Figure 1 shows the relationship between the annual change in aggregate retail employment and the annual change in personal consumption expenditures from the BEA's NIPA. The latter is typically viewed as the gold-standard measure of aggregate consumption at the national level. The two series track each other closely. Intuitively, retail services are an intermediate input into household consumption, since consumers must purchase things to be able to consume them. At an aggregate level, retail employment actually does a better job capturing time-series variation in non-durable PCE than the CEX, which has displayed implausible negative growth rates in recent years (see, e.g., Heathcote *et al.*, 2010).

One might worry that the increasing prevalence of big box and online retailers might weaken the relationship between retail employment and the PCE. There is a very small downward trend in retail employment relative to real PCE in Figure 1 (hardly visible to the naked eye) that may reflect these forces. However, slow-moving trends will not affect our estimates, since our specification is formulated in growth rates and includes time fixed effects. Consistent with the figure, unreported rolling-window regressions suggest the time series relationship between retail employment and PCE is relatively stable over the time period we study.

There are relatively few alternative measures of consumer expenditures available at a sufficiently high frequency and with a sufficiently long panel to study housing wealth elasticities. Retail sales data are available at a geographically disaggregated level only every 5 years. Some recent work has used expenditure series for particular categories, such as AC Nielsen data or data on car purchases. These series are not available over the long time horizon required for our study. Moreover, the aggregate time series suggests that retail employment provides at least as good a measure of consumer expenditures (based on the production-based PCE measure) as these more specialized categories. Another possible source of data to consider might be retail sales tax data. However, retail sales tax data are only available for a subset of states and are incredibly noisy

extraction from mortgage data and the assumption that households spent the entire amount of money extracted from housing wealth in this way.

8. Unfortunately, the specific details of how the "consumption" series published by these private sector sources are constructed is not disclosed. However, it is clear that both the Regional Financial Associates data and the Survey of Buying Power data used by Asdrubaldi *et al.* (1996) rely substantially on retail employment in their data construction series. This is documented in Zhou (2010), and we have also verified this in private correspondence with the Survey of Buying Power.

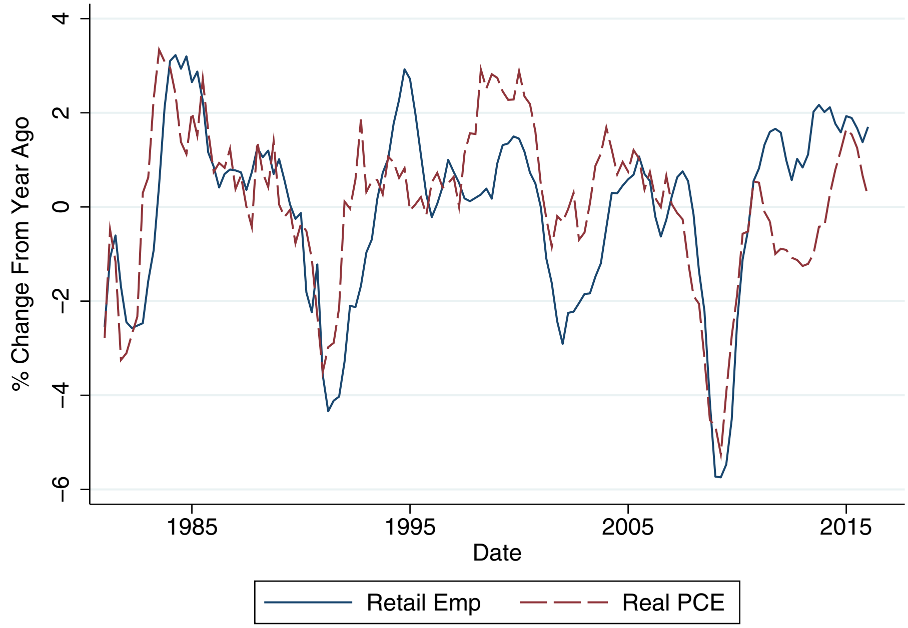

FIGURE 1

Growth of retail employment versus growth in Personal Consumption Expenditures

Note: The figure plots the 4-quarter change in aggregate retail employment (FRED series CEU4200000001) and the 4-quarter aggregate change in real personal consumption expenditures (FRED series PCECC96). We take out a linear time trend from both series to account for differential trend growth. The retail employment series has had a slightly larger secular decline than the real PCE series, falling 0.08% per year as opposed to 0.03% per year. Because our regressions include time fixed effects, we take out this differential trend growth from our analysis.

in raw form (Garrett *et al.*, 2005).[9] Some researchers have used data compiled by private sector data sources, such as the Regional Financial Associates or Survey of Buying Power data, but these sources do not introduce any additional microdata and are imputed from a combination of sources, including retail employment, as we describe above.

In Supplementary Appendix A.3, we analyse the relationship between city-level consumption and retail employment using data for 17 cities for which the BLS publishes city-level consumption using data from the Consumer Expenditure Survey. Both the CEX and retail employment have substantial sampling error. We use an instrumental variables approach to account for measurement error in retail employment per capita. Our instrumental variables estimates imply that consumer expenditures respond nearly one-for-one with changes in retail employment per capita, consistent with the aggregate time series in Figure 1, and we assume this elasticity is one when we interpret our empirical results as a consumption response.

Our data for retail employment come from the Quarterly Census of Employment and Wages (QCEW) which we use starting in 1978 at the county level.[10] The population data come from the Census Bureau's post-Censal population estimates for 1970–2010 and inter-Censal population estimates for 2010–17. These population estimates are available annually, and we interpolate to

9. Rodgers and Temple (1996) estimate that at a national level, the correlation between the growth rates of national retail sales and personal consumption is only 0.35.

10. The QCEW started in 1975, but the sample expanded in the early years to include more industries and the expansion was staggered across states (see Chodorow-Reich and Wieland, 2018). To limit the effect of the coverage expansion, we start our analysis with 1978 log differences.

a quarterly frequency. We aggregate the combined data set to the CBSA level and create retail employment per capita for 380 CBSAs.[11] Two issues that arise are how to handle missing data at the CBSA level and how to handle the change in industrial classifications from SIC to NAICS. Supplementary Appendix D.1.3 provide further detail on how we handle these issues, and show that alternative plausible approaches yield very similar results.

For house prices, our primary data source is the Freddie Mac House Price Indices, which are a balanced panel of indices based on repeat sales for 381 CBSAs from 1975 to 2017 (1976 is thus the first year for annual differences). We convert to a real house price index using the GDP deflator. The Freddie Mac House Price Indices have the advantage that they do not impute any data from neighboring cities. Imputation has the potential to bias our empirical estimates of the differential sensitivity of house price indexes to aggregate shocks across cities. A downside of the Freddie Mac data is that they are limited to conforming loans and makes use of a combination of transaction and appraisal prices. Appraisal prices tend to be smoother than transaction prices.

In Supplementary Appendix D.1.5, we redo our analysis using the CoreLogic house price index. The results are very similar to our baseline results. Unlike the Freddie Mac Price indices, the CoreLogic indices are all-transaction price indices and include homes purchased with non-conforming loans. The disadvantage is that CoreLogic has far more limited time coverage after dropping city-level indices that are imputed form state and regional indices. The similarity of the results shows that our results are not driven by including appraisals or dropping non-conforming loans.

We also use a variety of other data for controls, which we describe in Supplementary Appendix A.1.

3. EMPIRICAL APPROACH

The goal of our empirical analysis is to estimate the effect of a change in house prices in one city relative to another on relative per-capita retail employment in the two cities. We do this using the following empirical specification:

$$\Delta y_{i,r,t} = \psi_i + \xi_{r,t} + \beta \Delta p_{i,r,t} + \Gamma X_{i,r,t} + \varepsilon_{i,r,t}. \tag{1}$$

The subscript i denotes core-based statistical areas (CBSAs)—roughly speaking cities—r denotes Census regions, and t denotes time (measured in quarters). $\Delta y_{i,r,t}$ denotes the log annual change in retail employment per capita, while $\Delta p_{i,r,t}$ denotes the log annual change in house prices, ψ_i denotes a set of CBSA fixed effects, $\xi_{r,t}$ denotes a set of region-time fixed effects, $X_{i,r,t}$ denotes a set of additional controls, and $\varepsilon_{i,r,t}$ denotes other unmodelled influences on retail employment.

The coefficient of interest in equation (1) is β, which measures the housing wealth elasticity. Several challenges arise in estimating β. Causation runs both ways between local employment and house prices, implying that the error term in equation (1) will be correlated with the change in house prices. This is likely to bias OLS estimates of β upward since a strong economy will cause house prices to rise. On the other hand, house prices are measured with error, potentially biasing β towards zero.

Recent work has addressed these challenges by using Saiz's (2010) estimates of CBSA-level housing supply elasticities as an instrument for the change in house prices in different cities during the 2000s boom or bust (e.g., Mian *et al.*, 2013; Mian and Sufi, 2014). This work typically uses an

11. We drop Dover, DE from our analysis because retail employment data are missing for the entire CBSA for a majority of years.

IV regression on a single cross-section to evaluate the housing wealth elasticity. Davidoff (2016) has critiqued this approach, pointing out that the Saiz elasticity is correlated with measures of long-run demand growth: there has been a secular trend over several decades favouring coastal cities that have relatively high land-unavailability. Furthermore, the boom–bust house price cycle of the 2000s coincided closely with the overall business cycle making it difficult when using a single cross-section regression to distinguish between a city being generally more cyclical and the causal effect of house prices. In particular, it may be that coastal cities are simply more cyclically sensitive than inland cities with lower levels of land unavailability.

Our approach to addressing these weaknesses of earlier estimates is to employ a panel specification. This allows us to include a rich set of controls. Our inclusion of CBSA fixed effects mitigates Davidoff's (2016) concern that long-run demand factors are correlated with land unavailability. Since our regression is in log changes, the CBSA fixed effects will capture any differential long-run trends across CBSAs. The variation that we use to identify the housing wealth elasticity is therefore orthogonal to these trends. We also include a control for variation in CBSA cyclical sensitivities. We construct this control by estimating the following OLS regression:

$$\Delta y_{i,r,t} = \psi_i + \alpha_i \Delta Y_{r,t} + \varepsilon_{i,r,t}, \tag{2}$$

where ΔY is the log change in regional retail employment. In this equation, α_i reflects the differential sensitivity of retail employment in a given CBSA to regional retail employment. We then use $\hat{\alpha}_i \Delta Y_{r,\tau}$ as a control variable. The inclusion of this control implies that the variation that we use to identify the housing wealth elasticity is orthogonal to differential cyclical sensitivity across CBSAs.

Finally, many potential endogeneity concerns in our setting boil down to industrial structure being correlated with housing supply elasticities. To mitigate such concerns we control for local industry shares with separate coefficients for each time period. This accounts for all differential factors that are correlated in the cross-section with industry structure. For example, this control captures unobservable variables relating to some cities having more risky industries than others and therefore being differentially affected by shocks to labour demand or risk premia associated with industrial structure. We also include separate controls for differential city-level exposure to real 30-year mortgage rates and Gilchrist and Zakrajsek's (2012) measure of bond risk premia. These controls are constructed using analogous regressions to equation (2).[12]

Our interest in assessing the magnitude of the housing wealth elasticity over time raises a second challenge: Saiz's estimates of housing supply elasticities are relatively crude. These housing supply elasticity estimates are largely based on land-unavailability—the share of land within a 50 km radius of the centre of a city that is not suitable for construction due to steep slopes or water.[13] But housing supply elasticities are likely affected by a host of other factors. The crudeness of Saiz's estimates implies that housing wealth elasticity estimates based on the Saiz instrument are quite imprecise, especially in other time periods than the 2000s boom and bust. We overcome this challenge by developing a new instrument based on a new proxy for housing

12. We estimate the sensitivity of retail employment on regional retail employment in equation (2) on the "leave-out sample" to avoid overfitting concerns. However, we have also tried the more direct approach of including $\alpha_i \Delta Y_{r,t}$ as controls in equation (1) and the equivalent for the 30-year mortgage rate and the Gilchrist–Zakrajsek excess bond premium. Doing this for the latter two controls yields essentially the same results with slightly larger standard errors. Doing so for retail employment yields similar results starting with 10-year windows centred in the mid-1990s and highly imprecise results with lower point estimates in the early 1990s.

13. The elasticity is formally the predicted values from regression 6 in Table III of Saiz (2010). The Wharton Land Use Regulation Index and land unavailability in levels and interacted with log population are the only factors that are used to predict the elasticity. In practice, land unavailability is the dominant force.

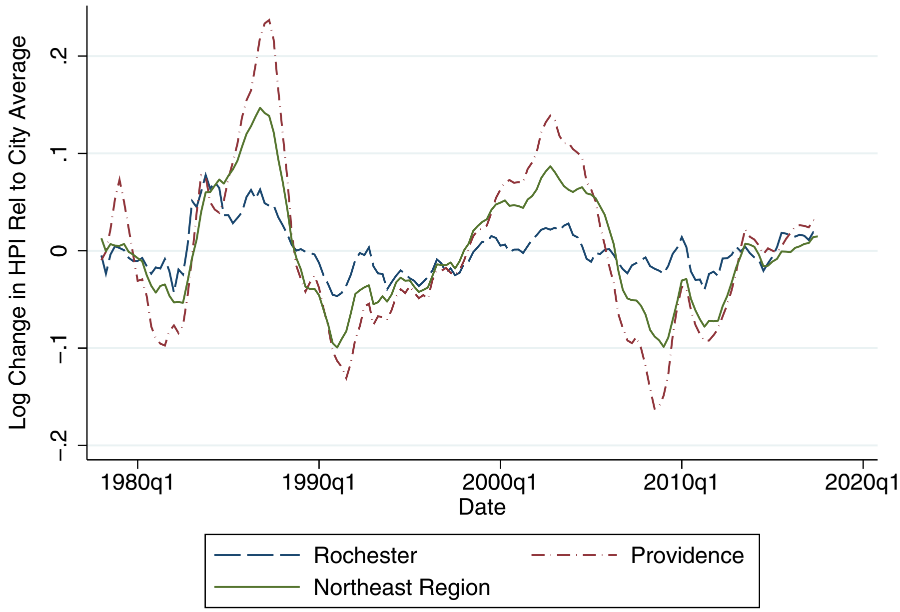

FIGURE 2

House prices in Providence, Rochester, and the Northeast Region

Notes: The figure shows house prices in the Providence CBSA, Rochester CBSA, and the Northeast Region. All data series are demeaned relative to the CBSA or region average from 1978 to 2015.

supply elasticities. One strategy for developing a better proxy would be to add more variables to an empirical model of housing supply elasticity such as the one that Saiz uses. We adopt a different strategy, which is to infer differences in housing supply elasticities across cities from systematic differences in the sensitivity of local house prices to regional house price variation. We refer to our new instrument as a sensitivity instrument.

3.1. *Simple intuition for the sensitivity instrument*

Before developing our sensitivity instrument in detail, it is useful to consider an example. Figure 2 plots the time series of house prices in Providence and Rochester as well as the Northeast region as a whole. Two features of this example are important for the construction of our sensitivity instrument. First, house prices in the Northeast have experienced large regional boom–bust cycles throughout our sample period. In particular, there was a large house-price cycle in the Northeast in the 1980s in addition to the house-price cycle of the 2000s. Regional house price cycles like the 1980s cycle in the Northeast occurred in several regions of the U.S. in the 1980s and 1990s. The timing of these regional cycles has varied, and they largely averaged out for the nation as a whole except for the nationwide boom–bust cycle of the 2000s. The existence of these regional cycles helps us estimate the housing wealth elasticity before 2000 when identification strategies using nation-wide variation in house prices lose power.

Second, the sensitivity of house prices in different CBSAs in the Northeast to the regional house price cycle varies systematically. When house prices boom in the Northeast, house prices in Providence respond much more than house prices in Rochester. This pattern of differential sensitivity is quite stable over the entire sample period, as noted by Sinai (2013). Furthermore, this pattern is a pervasive feature of house price data across different CBSAs and regions.

These two features of house price dynamics suggest the following simple identification strategy, which we will subsequently refine. First, estimate the sensitivity of house prices in different CBSAs to regional house price movements by running the regression:

$$\Delta p_{i,r,t} = \varphi_i + \gamma_i \Delta P_{r,t} + \nu_{i,r,t}, \tag{3}$$

where $\Delta P_{r,t}$ denotes the log annual change in regional house prices and γ_i is a city-specific coefficient.[14] Then use $z_{i,r,t} = \hat{\gamma}_i \Delta P_{r,t}$ as an instrument for $\Delta p_{i,r,t}$ in equation (1), where $\hat{\gamma}_i$ denotes the estimate of γ_i from equation (3). In this identification strategy, $\hat{\gamma}_i$ is our proxy for (the inverse of) the housing supply elasticity in city i. Equation (3) is not the first-stage regression. Rather it is the empirical model we use to generate a proxy for the housing supply elasticity in each city $\hat{\gamma}_i$. Our $\hat{\gamma}_i$ estimates, therefore, play the same role in our empirical strategy as Saiz's 2010 estimated housing supply elasticities play in the empirical strategy of, *e.g.*, Mian *et al.* (2013) and Mian and Sufi (2014).

Another way to describe our sensitivity instrument is that it is similar to a difference-in-difference design: when there is a housing boom in the Northeast, house prices systematically increase more in Providence than in Rochester, *i.e.*, Providence is differentially treated. Since we have panel data, we are able to estimate the systematic extent of differential treatment across CBSAs using equation (3). The question, then, is whether this differential treatment translates into differential growth in retail employment. This empirical strategy builds on work by Palmer (2015), who instruments for house prices in the Great Recession using the historical variance of a city's house prices interacted with the national change in house prices.

3.2. *Refined sensitivity instrument*

The simple procedure described above runs into problems if local house prices respond differentially to regional shocks through other channels than differences in housing supply elasticities. Suppose, for example, that there are differences in industrial structure across CBSAs that induce differences in the cyclical sensitivity of employment to the aggregate business cycle (for reasons other than housing). In this case, the heterogeneity in $\hat{\gamma}_i$ may arise from reverse causality. A hypothetical example is instructive: Suppose that Providence has an industrial structure tilted towards highly cyclical durable goods relative to Rochester. In this case, a positive aggregate demand shock would lead retail employment to increase more in Providence than Rochester. If local economic booms raise house prices, this would induce a larger change in house prices in Providence than Rochester and, thus, imply that we would estimate a higher γ_i for Providence using equation (3) purely due to reverse causality. In this case, variation in $\hat{\gamma}_i$ would reflect factors other than differences in housing supply elasticities across cities, potentially invalidating our sensitivity instrument.

To address this problem, we refine the procedure described above for estimating γ_i by controlling for local and regional changes in retail employment with city-specific coefficients as well as other controls $X_{i,r,t}$:

$$\Delta p_{i,r,t} = \varphi_i + \delta_i \Delta y_{i,r,t} + \mu_i \Delta Y_{r,t} + \gamma_i \Delta P_{r,t} + \Psi X_{i,r,t} + \nu_{i,r,t}. \tag{4}$$

In this case, we estimate the γ_is using only the variation in local house prices that is orthogonal to $\Delta y_{i,r,t}$, $\Delta Y_{r,t}$, and $X_{i,r,t}$. This implies that our $\hat{\gamma}_i$ estimates are not driven by the type of

14. To keep our notation simple, we denote $\Sigma_{i \in I} \gamma_i \Delta P_{r,t} I_i$ where I_i is an indicator for city i (that is separate city-specific coefficients for each city i) by $\gamma_i \Delta P_{r,t}$. We use this simplified notation throughout the article.

reverse causation described above. We use all the same controls when estimating equation (4) as we do when estimating equation (1). This implies that $X_{i,r,t}$ includes (among other variables) two-digit industry shares multiplied by time dummies. We therefore non-parametrically control for all variation that is correlated with industry structure in the cross section. In equation (4), we additionally control for changes in average wages as reported in the QCEW with CBSA-specific coefficients.[15]

The key identifying assumption for our sensitivity instrument is that, conditional on controls, there are no other aggregate factors that are both correlated with regional house prices in the time series and that differentially impact retail employment per capita in the same CBSAs that are sensitive to house prices as captured by $\hat{\gamma}_i$. In other words, to bias our results, there must exist a confounding factor with the structure $\alpha_i \mathcal{E}_{r,t}$, where $\mathcal{E}_{r,t}$ is an aggregate shock and α_i reflects the differential sensitivity of retail employment in a given CBSA to this aggregate shock, such that $\mathcal{E}_{r,t}$ is correlated with regional house prices in the time series and α_i is correlated with $\hat{\gamma}_i$ in the cross section. Since we are estimating β using panel data, in which we observe many time periods, with many aggregate shocks, we are able to directly control for differential sensitivity of local retail employment to a variety of observable aggregate variables. This has the important advantage that it allows us to rule out many potential confounding factors with a $\alpha_i \mathcal{E}_{r,t}$ structure.[16]

Our sensitivity instrument is a close cousin of the Bartik instrument, which instruments for city labour demand with city industry shares interacted with national changes in employment in each industry. For example, consider a Bartik instrument in which the key source of variation is differential exposure to oil shocks in Texas versus Florida. The identifying assumption is that there is not some other factor that happens to differentially affect Texas at the same time as oil prices go up. Our identifying assumption that there is no aggregate factor that is correlated with regional house prices in the time series and that differentially impacts retail employment in a way correlated with $\hat{\gamma}_i$ has a similar flavour. It is important to understand that for these strategies to be valid, treatment intensity (in our case $\hat{\gamma}_i$ and in the case of the Bartik instrument the industry shares) need not be randomly assigned. This is in fact rarely the case. In the Bartik example, Texas and Florida obviously differ in other ways than just their exposure to oil shocks, but as long as we can attribute any differential effects that occur at the time of oil price shocks to differences in oil exposure, this does not invalidate the instrument. Another important point is that measurement error in our generated instrument will show up in our standard errors; unlike generated regressors, generated instruments do not present inference issues.

Our panel data approach allows us eliminate sources of mechanical correlation. In particular, we exclude the CBSA in question from the construction of the regional house price index when running regression (4), so as to avoid bias in γ_i due to the same price being on both the left- and right-hand side.[17] In our rolling-window analysis, we also estimate equation (4) using time periods other than the time period for which we are estimating equation (1), while in the full-sample analysis the γ_i's for a particular time period are estimated using data from all years except a seven

15. One potential concern with this procedure is the role of measurement error in $\Delta y_{i,r,t}$ biasing the δ_i terms and thereby creating bias in the γ_is. To assess the severity of this concern, we have also considered a specification in which we instrument for $\Delta y_{i,r,t}$ using a 2-digit Bartik instrument for local economic conditions. For power reasons, we must assume that δ_i is the same across CSBAs, but the δ we obtain is a causal elasticity. We can use this IV regression to estimate γ_i. This approach yields values for the γ_i that are highly correlated with our baseline approach, and using these alternate γ_is does not significantly alter our results.

16. Supplementary Appendix C presents a more formal discussion of these identifying assumptions in the context of a two-equation simultaneous equations system from which we explicitly derive our estimating equations.

17. There is an arithmetic reason not to include region-time fixed effects in equation 4 that arises as a consequence of this leave-out procedure. Since a leave-out mean appears in this regression, arithmetically, it is possible to perfectly predict local house prices if region-time fixed effects are included.

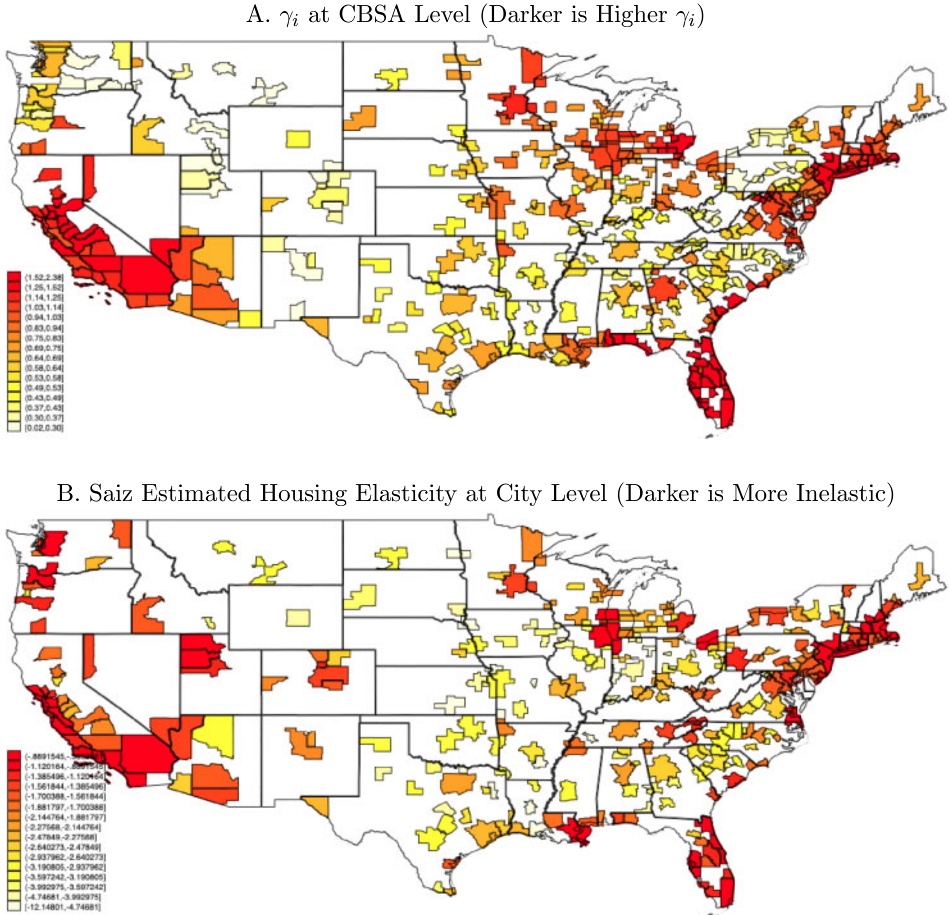

FIGURE 3

γ_i and Saiz elasticity by CBSA for continental U.S.

Notes: These figures provide heat maps for γ_i and the Saiz elasticity. γ_i is estimated in a single pooled regression that does not leave out any years from 1978 to 2017. The Saiz instrument is adjusted so that darker colours represent inelasticity rather than elasticity so that darker regions in both figures are where prices tend to move by more in response to a shock.

year window around the point in question. We do this to avoid $\hat{\gamma}_i$ reflecting contemporaneous or nearly contemporaneous variation in local house prices to the variation used to estimate equation (1). In practice, these different leave-out procedures yield similar results.

3.3. *Inspecting the variation in $\hat{\gamma}_i$*

The goal of estimating our sensitivity measure γ_i in equation (4) is to generate a new proxy for housing supply elasticities of different cities that captures a more comprehensive set of the determinants of housing supply than the estimates of Saiz (2010) and can be used to construct a more powerful instrument for variation in house prices. It is, therefore, instructive to compare our $\hat{\gamma}_i$'s with Saiz's (2010) estimates of housing supply elasticities. Figure 3 does this using two heatmaps. Figure 3A shows our $\hat{\gamma}_i$s, while Figure 3B shows the inverse Saiz elasticity.

At a broad-brush level, Figure 3 shows significant similarity between our $\hat{\gamma}_i$'s and Saiz's elasticity estimates. Both measures indicate that many CBSAs on the California coastline, in

Florida, and along the Northeast seaboard have inelastic housing supply, while many cities in the interior of the US, especially in Texas and on the Great Plains, have elastic housing supply. However, a closer look at Figure 3 reveals substantial differences across the two measures. For example, Saiz's estimates suggest much lower housing supply elasticities in the Pacific Northwest, the Rocky Mountains, and near Lake Erie and Lake Ontario than our $\hat{\gamma}_i$s. In fact, the R-squared of a regression of our $\hat{\gamma}_i$s on Saiz's elasticity is only 0.13.

Why might these differences arise? First, Saiz's estimates of housing supply elasticities are relatively crude, as we discuss above. They are based on the share of land within a 50 km radius of the centre of a city that is not suitable for construction due to steep slopes or water, and presumably this leaves out a number of important factors that determine land supply elasticities (*e.g.* the 50 km radius isn't appropriate for all cities). In addition, it is important to recognize that the amplitude of house price cycles is determined not only by current housing supply elasticities but also by expectations about future housing supply elasticities. Many cities with an intermediate degree of land unavailability are not currently constrained but may become constrained in the future. Whether these cities become constrained in the future depends on their expected long-run growth rate. Indeed, Nathanson and Zwick (2018) emphasize that the amplitude of housing cycles in such cities can depend heavily on both expectations about future long-run growth and the degree of disagreement about future long-run growth prospects. The existence of a group of people that are very optimistic about the long-run prospects of a city with an intermediate degree of land constraints can create particularly large housing cycles in Nathanson and Zwick's model.

These types of differences can potentially contribute to explaining the discrepancies between our $\hat{\gamma}_i$s and Saiz's estimated elasticities. Consider, for example, Las Vegas and Pittsburgh. Both have an intermediate degree of land unavailability, but our $\hat{\gamma}_i$ for Las Vegas is very large, while our $\hat{\gamma}_i$ for Pittsburgh is among the smallest among all large cities (see Table A.4 in the Supplementary Appendix for a list of cities with large and small $\hat{\gamma}_i$s in each region). One way to make sense of this large difference in $\hat{\gamma}_i$s is that Las Vegas is a high-growth city with an industrial structure that may be particularly conducive to high degrees of disagreement about future long-run growth (in particular wild optimism), while Pittsburgh's growth is much slower and few people are wildly optimistic about its long-run prospects. Similar arguments can be made for the discrepancies between our $\hat{\gamma}_i$s and Saiz's elasticity estimates for many other cities such as Orlando, Phoenix, and the California Central Valley, on the one hand, and Cleveland, Rochester, Buffalo, New Orleans, and Salt Lake City, on the other hand.[18]

Detroit is another interesting example. Both our $\hat{\gamma}_i$ and Saiz's elasticity estimate indicate that housing supply is relatively inelastic in Detroit. However, our $\hat{\gamma}_i$ for Detroit is large relative to Saiz's elasticity estimate for the city. A distinctive feature of Detroit is that it has been in steep decline throughout much of our sample period. Glaeser and Gyourko (2005) argue that cities in decline have particularly inelastic housing supply because houses are very durable. Essentially, the growth rate of the housing stock in Detroit is stuck at the rate of depreciation, making housing supply particularly unresponsive to economic conditions. The high value of $\hat{\gamma}_i$ we estimate for Detroit seems to capture this better than Saiz's estimate. Other factors that may play a role are that some regions are more "bubbly" due to social connections to inelastic cities (Bailey *et al.*, 2018) or credit (Favara and Imbs, 2015).

18. It is important to note that given our empirical strategy, our empirical estimates will not pick up increases in consumption that arise directly from, say, Las Vegas having higher long-run trend growth than Pittsburgh. These trend differences will be picked up in the city fixed effect. Also, differences in the growth loading on other shocks will be captured by the fact that we control for differential exposure to regional employment growth and industrial structure with time specific coefficients. For our procedure to estimate a high $\hat{\gamma}_i$ for Las Vegas, it must be that residual house price growth conditional on all of there controls is high when regional house prices boom.

TABLE 1
Pooled elasticity of retail employment per capita to house prices

Time period	(1) 1978–2017	(2) 1990–2017	(3) 2000–17
OLS	0.083***	0.081***	0.068***
	(0.007)	(0.008)	(0.008)
Sensitivity instrument	0.058***	0.072***	0.055***
	(0.017)	(0.015)	(0.014)
Saiz instrument	0.084	0.141***	0.134***
	(0.047)	(0.038)	(0.035)

Notes: Each column estimates equation (1) for the indicated time period. "OLS" uses no instrument. "Sensitivity Instrument" uses our sensitivity instrument with the γ_is estimated using equation (4) for each quarter, using a sample period that leaves out a three-year buffer around the quarter in question. Saiz uses an instrument that interact's Saiz's elasticity with the national change in house prices. All three approaches use the same control variables: two-digit industry shares with date-specific coefficients, the cyclical sensitivity control estimated using equation (2), and the analogously constructed controls for differential city exposure to interest rates and the Gilchrist–Zakrajsek excess bond premium along with CBSA and division-time fixed effects. Standard errors are two-way clustered at the time and CBSA level. ** indicates statistical significance at the 1% level, and *** indicates statistical significance at the 0.1% level.

4. EMPIRICAL ESTIMATES OF HOUSING WEALTH ELASTICITY

We present three sets of results on the housing wealth elasticity in this section. In Section 4.1, we present pooled estimates and relate these to earlier work using estimates based on single cross-section over the housing bust of 2006–9. In Section 4.2, we present time-varying estimates based on 10-year rolling window regressions. In Section 4.3, we explore whether the housing wealth elasticity changed in the boom and/or the bust of the early 2000s and whether the housing wealth elasticity more generally displays an asymmetry between periods of price increases and price decreases.

4.1. *Full-sample and single cross-section estimates of housing wealth elasticity*

Table 1 presents estimates of the elasticity β in equation (1) for our full sample period as well as for several sub-periods (across columns). For each sample period, we present OLS estimates as well as IV estimates with our sensitivity instrument and the Saiz instrument (across rows). We estimate CBSA fixed effects once for the entire sample period and apply them to all sample periods rather than estimating a different set of CBSA fixed effects for each sample period.[19] This avoids time variation in these fixed effects driving time variation in our coefficient of interest. We report standard errors that are constructed using two-way clustering by CBSA and time to allow for arbitrary time series correlations for a given CBSA and for correlations across CBSAs at a particular time.

The first column of Table 1 reports our estimates for the full sample period 1978–2017. OLS yields an estimate of 0.083 with a standard error of 0.007, while IV with our sensitivity instrument yields an estimate of 0.058 with a standard error of 0.017, and IV with the Saiz instrument yields an estimate of 0.086 with a standard error of 0.047. To put the magnitudes in context, the estimate based on our sensitivity instrument implies that a 10% decline in house prices in a CBSA relative to other CBSA's leads to a 0.58% decline in retail employment. This is equivalent to a marginal propensity to consume out of housing wealth (MPCH) of 2.67 cents on the dollar assuming a one-to-one relationship between retail employment and consumption as suggested by regressions

19. We regress all variables on CBSA fixed effects for the full sample and use the residuals from these regressions in our main analysis.

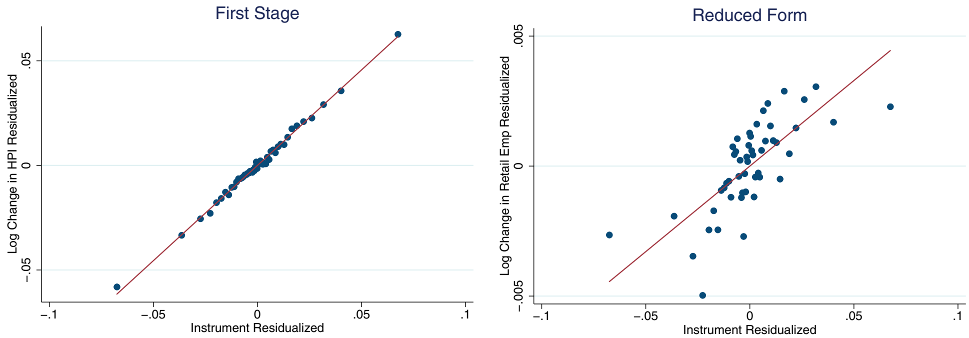

FIGURE 4

Sensitivity instrument pooled first stage and reduced form binned scatter plots

Notes: The figure shows binned scatter plots of the first stage and reduced form of the IV elasticity of retail employment per capita to real house prices at the CBSA level for the pooled 1990–2017 sample. These correspond to specification (1) in Table 1. For these estimates, we first construct our instrument for each quarter by estimating the γ_i's in equation (4) for each quarter, leaving out a three-year buffer around the quarter in question. We then estimate equation (1) pooling over the sample period 1990–2017. Both the x and y variables are residualized against all fixed effects and controls to create a two-way relationship that can easily be plotted (the Frisch-Waugh theorem).

in Supplementary Appendix A.3.[20] The OLS estimate implies an MPCH of 3.82 cents on the dollar, while the Saiz instrument implies an MPCH of 3.96 cents on the dollar.

IV with the Saiz instrument yields a very noisy estimate of the housing wealth elasticity over the 1978–90 sample. The second column of Table 1, which limits the sample to 1990–2017, shows that the statistical imprecision of the full sample IV estimates with the Saiz instrument are due to large amounts of noise in the early part of our sample. Limiting the sample to 1990–2017 causes the precision of the IV estimates with the Saiz instrument to improve and the point estimate to rise. For this sample period, IV with the Saiz instrument yields an estimate of the housing wealth elasticity of 0.142 with a standard error of 0.037, which is equivalent to an MPCH of 6.54 cents on the dollar. The precision of the IV estimate with our sensitivity instrument also improves and the point estimate increases when we limit to 1990–2017, but by much less. In particular, we obtain a point estimate of 0.072 with a standard error of 0.015, equivalent to an MPCH of 3.32 cents on the dollar. OLS, by contrast, is virtually unchanged. In what follows, we focus on the 1990–2017 sample.

To elucidate the results based on the sensitivity instrument, Figure 4 presents binned scatter plots for the first stage and reduced form for the 1990–2017 pooled sample. These plots show that neither the first-stage nor the reduced-form relationships are driven by outliers. The first stage is strong, reflecting the statistical power of our approach.

The pooled estimate obtained with our sensitivity instrument is somewhat smaller than estimates in the recent literature. For example, Mian and Sufi's (2014) results imply an elasticity of retail employment to house prices between 0.09 and 0.16, which corresponds to an MPCH between 4.1 and 7.3 cents on the dollar. Mian *et al.* (2013) estimate using the Saiz instrument implies an elasticity of total consumer expenditures of between 0.13 and 0.26. They also estimate

20. To convert our elasticity to a marginal propensity to consume out of housing wealth requires dividing the elasticity of consumption to house prices by the ratio of housing wealth to consumption. The average ratio of H/C over 1985 to 2016 where H is measured as the market value of owner-occupied real estate from the Flow of Funds and C is measured as total personal consumption expenditures less PCE on housing services and utilities, is 2.17. Hence, we obtain a marginal propensity to consume out of housing wealth of 0.058/2.17 = 2.67 cents for each additional dollar of housing wealth.

TABLE 2
Comparison of estimation approaches for 2006–9

Specification	2006–9 Elasticity	
Sensitivity instrument (per capita), CBSA FE	0.060**	(0.019)
Sensitivity instrument (per capita)	0.096***	(0.018)
Sensitivity instrument (not per capita)	0.116***	(0.020)
Sensitivity instrument, Saiz sample (not per capita)	0.126***	(0.024)
Saiz elasticity instrument (not per capita)	0.165	(0.093)
OLS (not per capita)	0.118***	(0.013)

Notes: This table compares our sensitivity instrument to the Saiz instrument and OLS for a single cross section long-difference from 2006 to 2009. For the sensitivity instrument, we construct our instrument for the three-year window estimating the γ_i's in equation (6), on the full sample but leaving out a three-year buffer around the quarter in question. We then estimate the single cross section $\Delta y_{i,r} = \xi_r + \beta \Delta p_{i,r} + \Gamma X_{i,r} + \varepsilon_{i,r}$, where $X_{i,r}$ includes the control for city-level exposure to regional retail employment and 2-digit industry share controls, and region fixed effects. For the CBSA fixed effects specification, we first take out CBSA fixed effects (or equivalently demean) for the entire 1978–2017 period for all variables, but we do not do so for other specifications. The full sample includes 379 CBSAs (excluding Dover, DE and The Villages, FL, which has a suspicious jump in employment for the 2006-9 window). The Saiz sample is limited to the 270 CBSAs for which we have land unavailability from Saiz (2010) instead of the full 379 CBSA sample. For the Saiz elasticity instrument, we run the same regression but instrument with the Saiz (2010) elasticity rather than our sensitivity instrument. OLS runs the second-stage regression by OLS. Robust standard errors are in parenthesis. ** indicates statistical significance at the 1% level, and indicates statistical significance *** at the 0.1% level.

the MPCH directly as 7.2 cents on the dollar.[21] Recall that our estimate based on the sensitivity instrument implies an MPCH of 3.3 cents on the dollar.

One theme that emerges from of Table 1 is that estimates based on the sensitivity instrument tend to be somewhat smaller than OLS, while estimates based on the Saiz instrument tend to be somewhat larger than OLS. Earlier work by Mian *et al.* (2013) and Mian and Sufi's (2014) has also found that housing wealth elasticities are larger using the Saiz instrument than OLS. To understand what drives this, it is useful to consider elasticity estimates based on a single cross section of 3-year growth rates from 2006 to 2009, which is the type of specification that Mian *et al.* (2013) and Mian and Sufi's (2014) use. Table 2 presents results for several variants of this type of specification. All of these specifications include region fixed effects and the full set of controls that we include in our baseline specification.

The specification in the first row is analogous to our baseline panel specification and yields an estimate of 0.060, which is slightly larger than our full-sample estimate of 0.058, but smaller than our post-1990 pooled estimate of 0.072. The second row presents a specification without CBSA fixed effects, *i.e.*, without demeaning all variables using means over the entire 1978–2017 sample period. This raises the estimated elasticity to 0.096, which suggests that it is important

21. Mian *et al.* report estimates of the elasticity of total consumer expenditures to housing net worth in the range 0.5–0.8. To convert Mian *et al.*'s elasticities with respect to total net worth to housing wealth elasticities, one must multiply by the mean housing wealth to total wealth ratio in their data, which is between 0.25–0.33 (Berger *et al.*, 2018). This yields a range for the elasticity of retail employment to house prices of between 0.13 and 0.26. Mian and Sufi (2014) estimate an elasticity of restaurant and retail employment to total net worth of between 0.37 and 0.49 for 2006–9, which must be adjusted using a similar procedure. This yields a range for the elasticity of retail employment to house prices of between 0.09 and 0.16. They do not adjust for population flows, which they find are unimportant in their sample. Similarly, Kaplan *et al.* (2020) estimate the elasticity of non-durable consumption with respect to net worth using the Saiz instrument and find estimates between 0.34 and 0.38 which implies an elasticity with respect to house prices of between 0.085 and 0.13. Aladangady (2017) who estimates an MPCH of 4.7 cents for homeowners and zero for renters, which corresponds to an MPCH of roughly 3.1 cents overall given a homeownership rate of 65%. Other studies estimate a marginal propensity to borrow out of housing wealth. For instance, Cloyne *et al.* (2019) use quasi-experimental variation in refinancing timing due to expiring prepayment penalties in the U.K. to find an elasticity of 0.2 to 0.3.

to account for long-run differences in growth rates across CBSAs in calculating the housing wealth elasticity. Davidoff's (2016) has pointed out that housing supply constraints are correlated with long-run demand growth and argued that this poses a problem for cross-sectional analysis of the housing wealth elasticity based on the Saiz instrument. The fact that we can control for such long-run differences in growth rates using CBSA fixed effects is an important virtue of our panel data approach relative to the single cross section specification prevalent in the recent literature.

The third row of Table 2 presents results for a specification in which we follow the common practice of not adjusting for population (*e.g.* Mian and Sufi, 2014). This raises the elasticity from 0.096 to 0.116, indicating that some of the non-per-capita response is due to population flowing towards regions with increasing house prices. The fourth row of Table 2 limits the sample to the cities for which the Saiz instrument is available. This raises the elasticity slightly to 0.126. The fifth row of Table 2 presents results based on the Saiz instrument. This yields an elasticity of 0.165. Moving from our sensitivity instrument to the Saiz instrument also increases the size of the standard errors by more than a factor of three. The final row of Table 2 presents results based on OLS, which yields an elasticity of 0.118. Our sensitivity instrument gives an estimate of housing wealth elasticity that is close to or slightly lower than OLS, while the Saiz instrument gives higher estimates than OLS.

4.2. *Time-varying estimates of the housing wealth elasticity*

Figure 5 presents 10-year rolling window estimates of the elasticity β in equation (1) using the empirical strategies described in Section 3. Figure 5A presents IV estimates with our sensitivity instrument along with OLS estimates, while Figure 5B presents IV estimates with the Saiz instrument along with OLS. Each point in the figure gives the elasticity for a 10-year sample period with its midpoint in the quarter stated on the horizontal axis (*e.g.* the point for quarter 2010q1 is the estimate for the sample period 2005q1–2015q1). We start the figure with the 10-year window from 1985q1 to 1995q1 because the standard errors for our estimates are very large prior to that point, but we use data back to 1978 in creating our instrument. As with Table 1, we take out a single CBSA fixed effect for the whole sample and two-way cluster by CBSA and time.

Figure 5 indicates that the housing wealth elasticity was not particularly large in the 2000s relative to earlier years. If anything, the elasticity has declined since the 1990s. This is true for all three estimation methods. This suggests that the time-series pattern for the housing wealth elasticity that we estimate is not an idiosyncratic feature of a particular identification strategy. For the sensitivity instrument, there is a noticeable increase in the estimated elasticity for 10-year periods centred in the mid-to-late 1990s. Supplementary Appendix D presents results based on a number of alternate specifications, datasets, and methodologies and shows that the time series pattern in Figure 5 is highly robust. The Supplementary Appendix focuses on the sensitivity instrument, since it provides the most precise estimate and is new. We present results without controls, based on 5-year rolling windows, weighting by population, excluding the "sand states," using 3-year differences rather than annual differences, using housing data from CoreLogic, using a fixed set of $\hat{\gamma}_i$s for the sensitivity instrument, as well as several other specifications. Supplementary Appendix D.1.2. also presents 10-year rolling window estimates of the first stage and reduced form for the sensitivity instrument. The main time series patterns are clearly evident in the reduced form, and although the first stage is stronger after 2000, it still has a high F-statistic (above 100) prior to 2000.

It is also instructive to consider whether changes in house prices affect manufacturing employment. Figure 6 plots results analogous to those presented in Figure 5 except that the

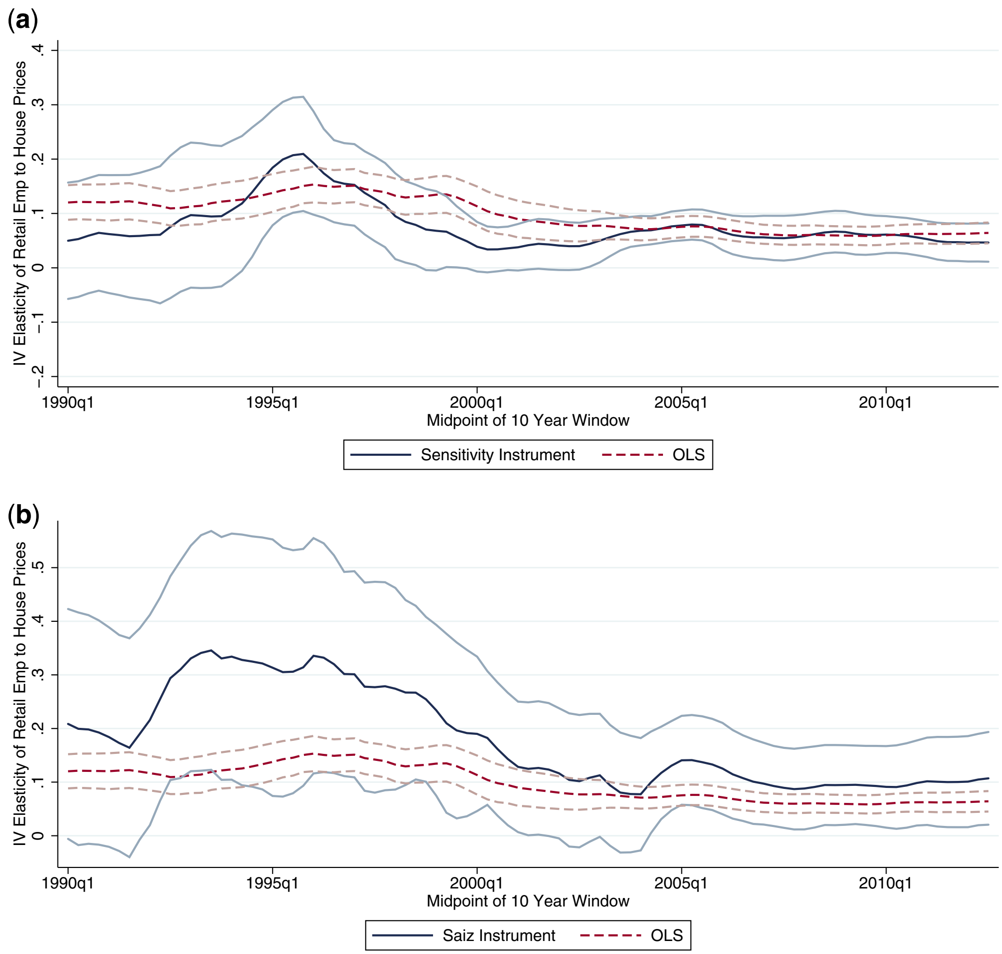

FIGURE 5

The elasticity of retail employment per capita to house prices over 10 year windows

Notes: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods for three different methods. Each point indicates the elasticity for a 10-year sample period with it's mid-point in the quarter stated on the horizontal axis. Panel A uses the sensitivity instrumental variable estimator that is described in Section 3 with ordinary least squares overlaid in red dashed lines. Panel B uses an instrument that interacts the estimated housing supply elasticity from Saiz (2010) with the national annual log change in house prices with ordinary least squares overlaid in red dashed lines. All three specifications use the same controls and CBSA fixed effects as described in the main text. Sensitivity and OLS also include region-time fixed effects, while Saiz uses only time fixed effects. The figure reports 95% confidence intervals in addition to point estimates for the elasticity. The standard errors are constructed using two-way clustering by CBSA and time for OLS and sensitivity and CBSA and time for Saiz.

dependent variable in the analysis is manufacturing employment. In contrast to the results for retail employment, the IV estimates with our sensitivity instrument yield point estimates for manufacturing employment that are close to zero for most of the sample period, although the estimates are fairly imprecise. The absence of an effect on manufacturing employment is consistent with our interpretation that the effects on retail employment we observe are driven by a housing wealth effect. One would expect a housing wealth effect to affect local spending, but not demand for manufacturing goods which are presumably largely consumed in other cities. This result is similar to Mian and Sufi (2014) finding that house prices mainly affect

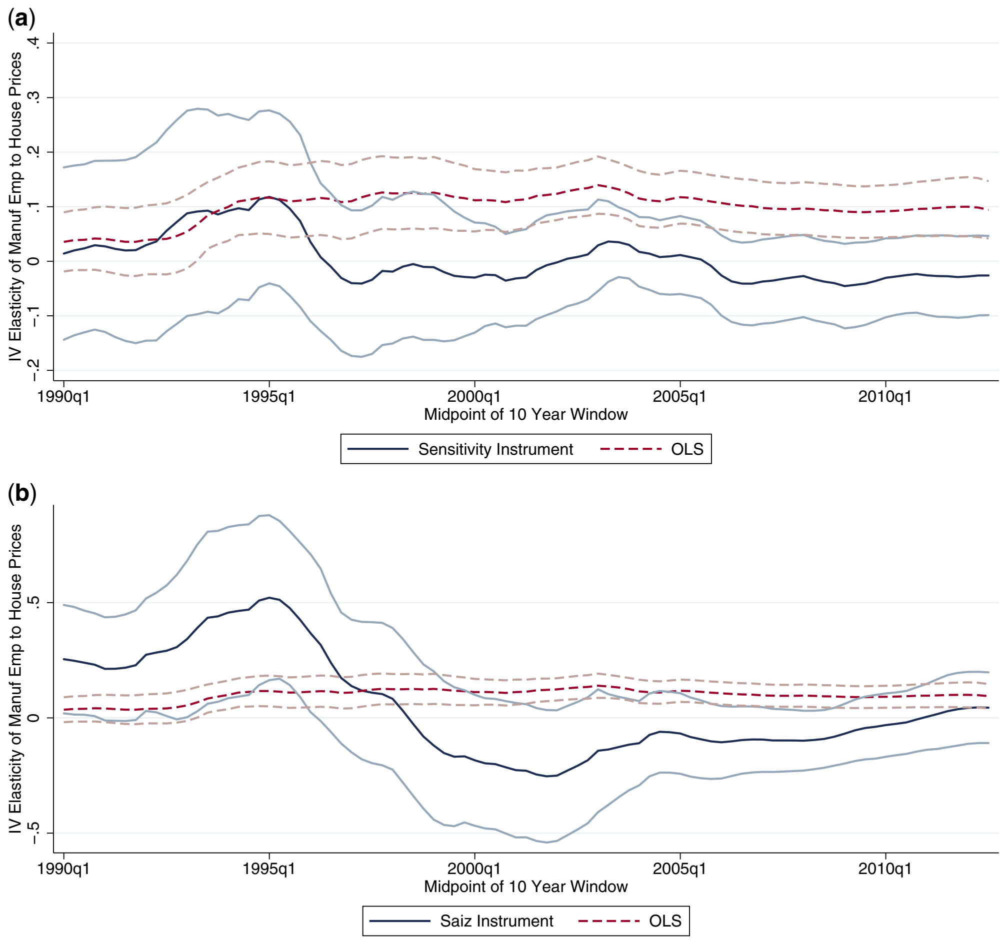

FIGURE 6

The elasticity of manufacturing employment per capita to house prices over 10 year windows

Notes: The figure plots the elasticity of manufacturing employment per capita to real house prices at the CBSA level for rolling 10-year sample periods for three different methods. Each point indicates the elasticity for a 10-year sample period with it's mid-point in the quarter stated on the horizontal axis. Panel A uses the sensitivity instrumental variable estimator that is described in Section 3 with ordinary least squares overlaid in red dashed lines. Panel B uses an instrument that interacts the estimated housing supply elasticity from Saiz (2010) with the national annual log change in house prices with ordinary least squares overlaid in red dashed lines. All three specifications use the same controls and CBSA fixed effects as described in the main text. Sensitivity and OLS also include region-time fixed effects, while Saiz uses only time fixed effects. The figure reports 95% confidence intervals in addition to point estimates for the elasticity. The standard errors are constructed using two-way clustering by CBSA and time for OLS and sensitivity and CBSA and time for Saiz.

non-tradeable production—presumably through an effect on local demand—but does not affect tradeable employment.[22]

IV estimates with the Saiz instrument for manufacturing employment are considerably more volatile than those with our sensitivity instrument. Post-2000 the point estimates from this

22. Mian and Sufi's (2014) use "tradeable employment" which is dominated by manufacturing. We use manufacturing instead because we are faced with the SIC to NAICS transition in 2000, which makes it difficult to create a consistent time series of tradeables using Mian and Sufi's approach for identifying such industries at the 4-digit level. By contrast, for manufacturing we can handle the transition by splicing together log changes for the manufacturing series under SIC and NAICS as we do for retail employment.

TABLE 3
Evaluation of housing wealth elasticity over the 2000s boom–bust cycle

	OLS		Sensitivity instrument		Saiz instrument	
	(1)	(2)	(3)	(4)	(5)	(6)
Elasticity	0.108***	0.107***	0.158***	0.158***	0.261*	0.273*
	(0.013)	(0.013)	(0.034)	(0.034)	(0.114)	(0.111)
Elasticity × boom	−0.030		−0.118**		−0.109	
(2000q2–2006q2)	(0.021)		(0.042)		(0.130)	
Elasticity × bust	−0.047**		−0.116**		−0.163	
(2006q3–2012q2)	(0.016)		(0.040)		(0.112)	
Elasticity × boom or bust		−0.041∗		−0.117**		−0.157
		(0.015)		(0.038)		(0.111)

Notes: Each column estimates equation (1) over the sample period 1990–2017 and includes a term that interacts the main regressor with the Boom (2000q2-2006q2) and the Bust (2006q3–2012q2) (odd columns) or an indicator for boom or bust (even columns). "OLS" uses no instrument. "Sensitivity Instrument" uses our sensitivity instrument with the γ_is estimated using equation (4) for each quarter, using a sample period that leaves out a three-year buffer around the quarter in question. Saiz uses an instrument that interact's Saiz's elasticities with the national change in house prices. All three approaches use the same control variables: two-digit industry shares with date-specific coefficients, the cyclical sensitivity control estimated using equation (2), and the analogously constructed controls for differential city exposure to interest rates and the Gilchirst–Zakrajsek excess bond premium along with CBSA and division-time fixed effects. For the Saiz and sensitivity specifications, instruments are interacted with time dummies for the indicated period. Standard errors are two-way clustered at the time and CBSA level. * indicates statistical significance at the 5% level, ** at the 1% level, and *** at the 0.1% level.

specification tend to be negative but are not significantly different from zero. Prior to 1995, the point estimates from this specification are large and positive, but rather imprecisely estimated. OLS yields relatively stable positive estimates on manufacturing employment, perhaps reflecting endogeneity bias.

4.3. *Testing for changes in the housing wealth elasticity*

The idea that housing wealth elasticities may have been particularly large in the Great Recession is related to the idea that housing wealth effects are particularly potent in housing busts—perhaps due to powerful debt-deleveraging during downturns. Tables 3 and 4 assess this possibility.

In Table 3, we directly test for a change in the housing wealth elasticity during the boom and the bust of the large house price cycle in the 2000s. We do this by adding to our baseline regression specification an interaction of our main regressor of interest with a dummy for the boom period (2000q2–2006q2) and separately for the bust period (2006q3–2012q2) in the odd-numbered columns.[23] The even-numbered columns include a single interaction of our main regressor of interest with a dummy for the boom and bust periods (2000q2–2012q2). We find negative coefficients for the interaction of the elasticity and the boom and bust indicators. These coefficients are statistically significant for OLS and the sensitivity instrument for the bust and the boom and bust. Furthermore, we find no strong evidence that the housing wealth elasticity is particularly large during the bust relative to the boom. These results provide statistical tests that validate our results from Section 4.2.

In Table 4, we consider non-linear regression specifications to assess more generally whether the housing wealth elasticity is different during periods when house prices are increasing versus decreasing. The odd numbered columns in this table report results for specifications that include

23. The CoreLogic national home price index peaked in 2006q2 and troughed in 2012q1–q2. We thus choose a six-year window on both sides of the 2006q2 peak for the boom and bust. Our results are not sensitive to the exact begin date of the boom window or end date of the bust window.

TABLE 4
Evaluation of nonlinearity in the housing wealth elasticity

	OLS		Sensitivity IV		Saiz IV	
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta\log(P)-$	0.085***		0.087***		0.135***	
	(0.012)		(0.021)		(0.039)	
$\Delta\log(P)+$	0.077***		0.052**		0.148**	
	(0.011)		(0.018)		(0.046)	
P test for equality	0.581		0.189		0.746	
$\Delta\log(P)$		0.082***		0.071***		0.138***
		(0.008)		(0.015)		(0.037)
$\Delta\log(P)^2$		0.017		−0.012		−0.073
		(0.032)		(0.040)		(0.174)

Notes: Each column estimates a version of equation (1) for 1990-20017. The odd-numbered columns replace Δp with $\Delta p \times 1\left[\Delta p \geq 0\right]$ and $\Delta p \times 1\left[\Delta p < 0\right]$ as regressors. The even-numbered columns add a quadratic term in the log change in house prices to equation (1). "OLS" uses no instrument. "Sensitivity IV" uses our sensitivity instrument with the γ_is estimated using equation (4) for each quarter, using a sample period that leaves out a three-year buffer around the quarter in question. "Saiz IV" uses an instrument that interact's Saiz's elasticity with the national change in house prices. All three approaches use the same control variables: two-digit industry shares with date-specific coefficients, the cyclical sensitivity control estimated using equation (2), and the analogously constructed controls for differential city exposure to interest rates and the Gilchrist–Zakrajsek excess bond premium along with CBSA and division-time fixed effects. In estimating the Saiz and Sensitivity instruments, we instrument with $Z \times 1[Z \geq 0]$ and $Z \times 1[Z < 0]$ for the odd-numbered columns and we instrument with Z and Z^2 for the even-numbered columns. ** indicates statistical significance at the 1% level, and indicates statistical significance *** at the 0.1% level.

separate coefficients for positive and negative house price changes, while the even numbered columns reports results for specifications that include a quadratic term in house price changes. We find no statistically significant evidence of a boom–bust asymmetry in house price elasticities. The coefficients on negative house price changes are slightly larger as debt-deleveraging in busts would suggest, but the *p*-value of a test for equality is large in all cases. Likewise, the quadratic terms are statistically insignificant in all three cases.[24]

5. DATA TO THEORY

In the decision problem of a household, house prices are exogenous. The "causal effect" of house prices on household consumption in such a partial equilibrium setting is therefore straightforward to interpret. By contrast, at the aggregate level or city level, house prices are an endogenous variable. House prices are affected by a myriad of shocks and these shocks may affect consumption not only through house prices but also directly or through other channels. So what does it mean to estimate the causal effect of house prices on consumption at the city level?

Consider a simple model of an economy consisting of several regions with many cities in each region. Suppose housing markets are local to each city and the cities differ in their housing supply elasticities. All other markets are fully integrated across cities within a region (and may in some cases be integrated across regions). The cities are initially in identical steady states before being hit by a one-time, unexpected, and permanent aggregate shock that alters the demand for housing. This shock leads house prices to respond differently across cities due to the difference

24. Previous evidence is mixed on whether there is an asymmetry in responses to house price increases versus decreases. Case *et al.* (2005) find an asymmetry, but Case *et al.* (2013) reject this initial finding with additional years of data. Cloyne *et al.* (2019) find a large elasticity if the collateral constraint is relaxed but nothing if it is tightened. Guerrieri and Iacoviello (2017) find an asymmetry in CBSA-level data for services employment using CBSA data. Finally, Liebersohn (2017) finds a large asymmetry for durables but not for consumption overall.

in housing supply elasticities, but all other prices respond symmetrically within region because all other markets are integrated within region. It is not important for our argument exactly what the nature of the aggregate shock is. It could be an aggregate productivity shock, an aggregate demand shock (*e.g.* monetary, fiscal, or news shock), or an aggregate housing specific shock such as a shock to the preference for housing or to construction costs.

Consumption in city i, in region r, and at time t can be written as $c_{i,r,t} = c(p_{i,r,t}, \omega_{i,r,t}, \Omega_{r,t}, R_{r,t})$, where $\omega_{i,r,t}$ is a vector of idiosyncratic shocks, $\Omega_{r,t}$ is a vector of regional or national shocks, $R_{r,t}$ is a vector of prices such as interest rates and wages. One can interpret $R_{r,t}$ as including not only current prices, but also prices for future-dated goods. Since all markets other than the housing market are integrated across cities within region, $R_{r,t}$ does not have an i subscript. All cities have the same aggregate consumption function. Consumption only differs across cities to the extent that they experience different home prices and different shocks. In a companion paper (Guren *et al.*, 2019), we provide an example of a fully specified general-equilibrium model of the type described above that yields a consumption function of this form.

Taking a log-linear approximation to the aggregate consumption function around the initial steady state and then taking an annual difference yields:

$$\Delta c_{i,r,t} = \underbrace{\phi_p}_{\beta} \Delta p_{i,r,t} + \underbrace{\phi_\Omega \Delta\Omega_{r,t} + \phi_R \Delta R_{r,t}}_{\xi_{r,t}} + \underbrace{\phi_\omega \Delta\omega_{i,r,t}}_{\varepsilon_{i,r,t}}, \tag{5}$$

where $c_{i,r,t}$ denotes the logarithm of consumption and ϕ_x denotes the elasticity of $c(\cdots)$ with respect to the variable x evaluated at the steady state. These elasticities should be understood as vectors of elasticities where appropriate. Equation (5) is labelled to show how it relates to equation (1) in our empirical analysis.

Suppose we ran the empirical specification described in Section 3 on data from this model. Equation (5) shows that the general equilibrium impact of changes in prices other than house prices as well as the direct effect of aggregate and regional shocks will be absorbed by the region-time fixed effects $\xi_{r,t}$. Our coefficient of interest β captures the response of consumption to a house price change holding these other variables constant. This shows that if we are able to identify variation in local house prices that is orthogonal to the error term $\varepsilon_{i,r,t}$ and the assumptions stated above about market structure hold, the coefficient β will estimate the *partial equilibrium* effect of house prices on consumption.[25]

The simple general equilibrium model discussed above makes the strong assumption that all markets except the housing market are fully integrated across cities within a region. If we relax this assumption, the differential response of house prices across cities will result in differential responses in other markets as well. For example, the local spending response to house prices could put upward pressure on wages resulting in a further increase in consumer spending. In other words, the differential house price movements will result in local general equilibrium effects. Since these local general equilibrium effects will differ across cities within a region, they will not be absorbed by the region-time fixed effects in our empirical specification and will affect our estimate of β.

Local general equilibrium effects result from changes in local demand affecting local wages, prices, and incomes. This suggests that evidence from other local demand shocks might be useful

25. If non-linearities are important, the fixed effects in equation (5) will not fully absorb the general equilibrium price effects. For example, if consumption growth responds importantly to $\Delta p_{i,r,t} \times \Delta\Omega_{r,t}$ or to $\Delta p_{i,r,t} \times \Delta R_{r,t}$, then our estimated β will reflect these interactions in addition to the housing wealth effect. In the next section we present a fully non-linear model of the housing wealth effect and we show in Supplementary Appendix E.1 that the model implies these interaction effects are small. In particular, the housing wealth elasticity is close to linear in the magnitude of the price change and symmetric with respect to positive and negative price changes.

in pinning down the effect of local general equilibrium on our empirical estimates. In Guren *et al.* (2019), we present a general-equilibrium regional business cycle model with heterogeneous housing supply elasticities that allows for local general equilibrium effects. In this model, we show that the local government spending multiplier can be used to quantify local general equilibrium effects. More specifically, we show that the housing wealth elasticity estimate β that results from our empirical specification can be expressed as:

$$\beta \simeq \beta_{\text{LFM}} \beta_{\text{PE}},$$

where β_{LFM} denotes the local fiscal multiplier and β_{PE} denotes the partial equilibrium elasticity of house prices on consumption.[26] Intuitively, a dollar of spending triggers the same local general equilibrium response regardless of whether it arises from a housing wealth effect or government spending. Nakamura and Steinsson (2014) estimate that the local government spending multiplier is roughly 1.5 at the state level but 1.8 at the region level. Since our analysis is at the CBSA level, the relevant local government spending multiplier for our analysis is likely somewhat smaller than 1.5.

6. A MODEL OF THE LOCAL CONSUMPTION RESPONSE TO HOUSE PRICES

We now present a partial equilibrium model of housing and consumption. The key features of the model are a life cycle, uninsured idiosyncratic income risk, borrowing constraints, illiquid housing, and long-term mortgage debt subject to an LTV constraint. We keep our model purposefully simple and evaluate its robustness to some of our starker assumptions in Supplementary Appendix E.

6.1. *Assumptions*

Households live for T periods and have preferences over non-durable consumption and housing services given by:

$$\mathbb{E}_0 \left[\sum_{t=1}^{T} \beta^t u(c_t, h_{t+1}) + \beta^{T+1} B(w_{T+1}) \right],$$

where c is consumption, h is housing, $B(\cdot)$ is a bequest motive, and w_{T+1} is wealth left to offsprings. We parameterize household preferences as:

$$u(c,h) = \frac{1}{1-\gamma} \left(c^{(\varepsilon-1)/\varepsilon} + \omega h^{(\varepsilon-1)/\varepsilon} \right)^{(1-\gamma)\varepsilon/(\varepsilon-1)}$$

$$B(w) = \frac{B_0}{1-\gamma} (w + B_1)^{(1-\gamma)}.$$

Here, γ captures the curvature of the utility function, ε is the elasticity of substitution between housing and non-durable consumption, ω is the taste for housing relative to non-housing

26. We make certain simplifying assumptions to derive this result. One of these is to assume GHH preferences to avoid wealth effects on labour supply. We abstract from the collateral channel emphasized by Chaney *et al.* (2012) and Adelino *et al.* (2015). We assume that the government and households both buy the same consumption good. Finally, we assume that construction employment does not respond to house prices. In Guren *et al.* (2019), we assess how relaxing this last assumption affects our results.

consumption, B_0 captures the strength of the warm-glow bequest motive, and B_1 captures non-homotheticity in bequest motives.[27]

An individual can consume housing either by owning or renting. A unit of housing can be purchased at price p or rented for one period at cost δp. This implies that the rent-price ratio is fixed and given by the parameter δ. We consider alternative assumptions about the behaviour of rents in Supplementary Appendix E. In our baseline model, people expect home prices will remain constant at their current level. In extensions, we consider cases with expected capital gains on housing. Throughout, we abstract from home price risk and the precautionary behaviour associated with those risks. Renting h units of housing delivers the same utility as buying that amount of housing, but the rent is higher than the user cost of owner occupied housing, which makes owning attractive despite its associated transaction costs. To sell a house the individual must pay ψ^{Sell} of the value of the house in a transaction cost and to buy a house the individual must pay a fraction ψ^{Buy}.

Households can take out mortgages. We denote the mortgage principal that a household brings into the period by m. At origination, mortgage debt must satisfy,

$$m' \leq \theta p h', \tag{6}$$

where θ is the maximum LTV and primes denote next period values. The mortgage interest rate is R_m and a household must pay a transaction cost of $\psi^m m'$ to originate a mortgage. We model mortgages as long-term debt that households can refinance at any time. To refinance, a household must pay the same transaction cost as when a mortgage is initiated ($\psi^m m'$ where m' is the new mortgage balance). The repayment schedule requires a payment such that $m' = G(a) R_m m$, where a is the age of the household. Following Campbell and Cocco (2003), $G(a)$ is defined so that the loan amortizes over the rest of the homeowner's lifetime. The amortization schedule is given by:

$$G(a) \equiv 1 - \frac{1 - R_m^{-1}}{1 - R_m^{-(T-a+1)}}.$$

The household can save, but not borrow, in liquid assets with return $R_a < R_m$. Finally, we model log annual income as $\log y = \ell + z + \xi$, where ℓ is a deterministic life-cycle component, z is a persistent shock that follows an AR(1) process, and ξ is a transitory shock.

6.2. *Calibration*

A household is born at age 25, works for 36 years, retiring at 61, and dies deterministically after age 80. We set the curvature of the utility function, γ, to 2. We set the elasticity of substitution between housing and non-durable consumption to 1.25 based on the estimates of Piazzesi *et al.* (2007). We set the LTV limit, θ, to 0.80 based on GSE guidelines for conforming mortgages without private mortgage insurance. We set the after-tax, real interest rate on mortgage debt to 3% per year based on the long-run averages of nominal mortgage rates and inflation.[28] We set the real return on liquid assets to 1% based on the difference between the long-run averages of

27. In the presence of illiquid durable goods such as housing, the parameter γ is related to, but not equivalent to, the coefficient of risk aversion (see Flavin and Nakagawa, 2008).

28. Between 1971 and 2017 the average CPI inflation rate was 4.1%, the average 30-year fixed rate mortgage rate was 8.2%, and the average 1-year treasury rate was 5.3%. Our choice of a 3% real interest rate on mortgage debt is meant to capture the tax-deductibility of mortgage interest.

TABLE 5
Parameter values set by matching moments from the SCF

β	ω	B_0	B_1	δ	ψ^m	ψ^{Sell}
0.939	0.0795	85.0	1.75	0.0435	0.0203	0.110

the 1-year Treasury rate and inflation. We set the cost of buying a house to 2%. This is meant to reflect closing costs associated with a home purchase.

During the household's working years, we model log annual income as the sum of a life-cycle component, a transitory component, and a persistent component. The life-cycle component is taken from Guvenen *et al.* (2019). We conceive of the transitory income shocks as non-employment shocks motivated by the income process in Guvenen *et al.* (2019). With some probability the household is employed for the full year and the (log) transitory income shock is zero. With the remaining probability, the household spends part of the year out of work. The fraction of the year the household spends non-employed is drawn from an exponential distribution truncated to the interval $(0,1)$. The probability of a non-zero non-employment shock and the parameter of the exponential distribution are estimated by maximum likelihood using the distribution of weeks worked in the prior year reported in the 2002 March CPS. The persistent component of labour income is modelled as an AR(1) with an AR coefficient of 0.97 and innovations drawn from a mixture of two normals, which allows us to capture the leptokurtic nature of income growth rates (see Guvenen *et al.*, 2019). The Supplementary Appendix provides further explanation of the income process and the parameter values. At retirement, a household faces no further labour income risk and is paid a social security benefit based on their final working-life income, which is calculated in the manner proposed by Guvenen and Anthony Smith (2014).

We set the remaining parameters by matching moments from the 2001 SCF. These parameters are the discount factor, β; the strength of the preference for housing, ω; the strength of the bequest motive, B_0; the degree to which a bequest is a luxury, B_1; the rent-price ratio, δ; the mortgage origination cost, ψ^m; and the transaction cost for selling a house, ψ^{Sell}. Our target moments are the life-cycle profiles of home value to income for working-age homeowners (we target the 25th, 50th, 75th, and 90th percentiles of the distribution at each age), mortgage LTV (P25, P50, P75, P90 by age), liquid assets (P25, P50, P75, P90 by age), and the homeownership rate by age. In addition, we target a 9.3% refinancing rate per year. Empirically, the refinancing rate is higher than this target (see *e.g.* Wong, 2019; Bhutta and Keys, 2016) but some refinancing activity results from changes in interest rates, which are not part of our analysis. Our target is based on Deng *et al.* (2000) who estimate a model of refinancing probability as a function of mortgage age and the difference between the mortgage interest rate and the market rate. We simulate their model for an environment with constant interest rates and compute the fraction of mortgages that are refinanced each year. Finally, we target a 3.2% annual moving rate for owner occupiers based on March 2001 CPS data. Overall, 6.3% of owner occupiers reported living in a different house one year ago. The CPS asks for the reason for the move and many of the movers report moving for reasons that are outside of the scope of our model; for example, due to a change in marital status. We exclude these moves resulting in the 3.2% moving rate. In total, we have 650 moments for seven parameters, so our model is highly over-identified, and we seek to minimize a weighted sum of the squared difference between the model-implied and empirical moments. Supplementary Appendix B explains our empirical moments and objective function in more detail. The resulting parameter values appear in Table 5.

7. MODEL SIMULATIONS

We begin by defining the object in the model that corresponds to the IV estimate of the housing wealth elasticity we present earlier in the paper.[29] In the model, we can write aggregate consumption in city i as:

$$C_i = \int c(x, p_i) d\Phi(x),$$

where $c(\cdots)$ is the consumption function of an individual, p_i is the price of a unit of housing in city i, x is a vector of idiosyncratic state variables, and Φ is the distribution of households over idiosyncratic states. The idiosyncratic states in the model are liquid assets, units of housing owned, mortgage debt, income, and age. The object of interest is the elasticity of C_i with respect to p_i, which we refer to as the model-implied housing wealth elasticity. We compute this elasticity by averaging together the responses to a 10% positive and 10% negative change in p_i.[30]

The model-implied housing wealth elasticity depends on the distribution Φ. Over our sample period, household balance sheets changed substantially as households became much more leveraged. Figure 7 shows the evolution of household leverage using data from the SCF and CoreLogic. We use the CoreLogic equity estimates from 2007 onwards out of a concern that SCF respondents may have overstated the value of their homes during the housing bust of the late 2000s. The data show two big increases in leverage, one from the mid 1980s through the late 1990s and one in the Great Recession, and a period of deleveraging after the Great Recession. In the boom, there was an increase in debt accompanied by an increase in prices, so LTVs did not rise much until the tail end of the boom. There has also been a great deal of variation in median home values relative to median income over our sample period. This ratio was about two in 1983, but rose to nearly four in 2007. We are interested in assessing the degree to which these large changes in household balance sheets affect the model-implied housing wealth elasticity.

To this end, Figure 8A reports the model-implied housing wealth elasticity using the distribution Φ that we observe for each wave of the Survey of Consumer Finances (SCF) from 1983 to 2016 (with a post-2007 adjustment based on LTV estimates from CoreLogic).[31] The variation observed across years in the figure therefore represents the extent to which our model implies that the observed changes in household balance sheets and demographics over the period 1983 to 2016 have lead to variation in the size of the housing wealth elasticity. The most striking feature of the figure, in our minds, is how small the variation in the model-implied housing elasticity is across years. The model generates a relatively smooth housing wealth elasticity across years despite very substantial changes in household balance sheets. For example, the large increases

29. Our empirical strategy uses retail employment as a proxy for consumption. In our analysis of the model we do not model the retail sector explicitly and assume that retail employment responds one for one with consumption. Supplementary Appendix A.3 provides evidence using city-level CEX data in support of the one-to-one relationship between consumption and retail employment.

30. We consider larger shocks in Supplementary Appendix Figure A.21 and find that the elasticity is stable with respect to the size of the price change and similar for negative and positive price changes, which is consistent with our empirical finding regarding boom–bust asymmetry.

31. Supplementary Appendix B.3 describes how we estimate Φ from the SCF and CoreLogic estimates. We adjust home values from 2007 onwards to match the distribution of LTVs for homeowners with a mortgage in CoreLogic's Homeowner Equity Reports. We do so by assuming that households correctly report their mortgage balance in the SCF but misreport their house value so that the error is rank-preserving in LTV. We make this adjustment because we are concerned that the SCF data understates the fall in home values during the Great Recession and therefore understates the fraction of households with high LTVs. CoreLogic did not produce equity estimates for earlier years. However, it turns out that our results are little affected by this adjustment.

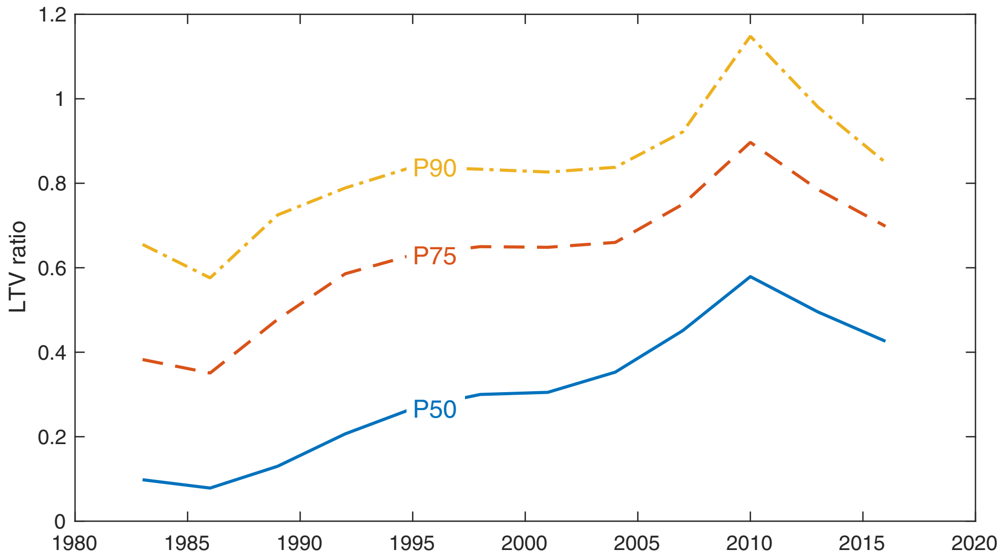

FIGURE 7

Percentiles of the LTV distribution

Notes: Data are from the SCF for 1983–2004 and CoreLogic for 2007–16. The percentiles refer to all homeowners including those without a mortgage. For 2007–16, the raw CoreLogic data refer to the distribution of LTV among homeowners with a mortgage and we use the SCF to calculate the fraction of homeowners with a mortgage and adjust the CoreLogic data.

in leverage that occurred between 2007 and 2010 and between 1983 and 1992 do not result in significant changes in the housing wealth elasticity implied by our model.

To unpack this result, Figure 8B presents results where we vary the distribution of one state variable at a time holding the distribution of the other state variables constant. First, we vary the LTV distribution to reflect the marginal distribution in year t while keeping all of the other state variables at their 2001 level. We start with the 2001 SCF data and for each year t we replace the LTV values with $F_t^{-1}(F_{2001}(\mathrm{LTV}_{2001}))$ where $F_t(\cdot)$ is the CDF of the marginal distribution of LTV for year t. Intuitively, we preserve each household's rank in the 2001 LTV distribution, but alter the LTV distribution according to the marginal distribution of LTV in year t, holding all other state variables fixed. Figure 8B also presents analogous results where we vary the marginal distribution of house values holding the distribution of all the other state variables fixed.

The dash-dot line in Figure 8B shows the effect of changing only the distribution of LTVs on the housing wealth elasticity implied by the model. This line is quite flat. Evidently, even the substantial changes in the LTV distribution illustrated in Figure 7 have a limited effect on the housing wealth elasticity implied by the model. The housing wealth elasticity in this LTV counterfactual increases modestly between 1983 and 2010 and, in particular, it does not spike during the Great Recession despite a sharp increase in LTVs. These results do not support the view that the housing wealth elasticity was particularly high during the Great Recession due to sharp deleveraging by households in the face of large increases in LTVs. It is crucial to emphasize that these results are for the *elasticity* of consumption with respect to housing wealth; not the level of consumption.

The dashed line in Figure 8B shows the effect of changing only the distribution of house values on the model-implied housing wealth elasticity. In contrast to changes in the LTV distribution, changes in house values do generate noticeable changes in the housing wealth elasticity. As Berger *et al.* (2018) have emphasized, when houses become a bigger part of the household balance

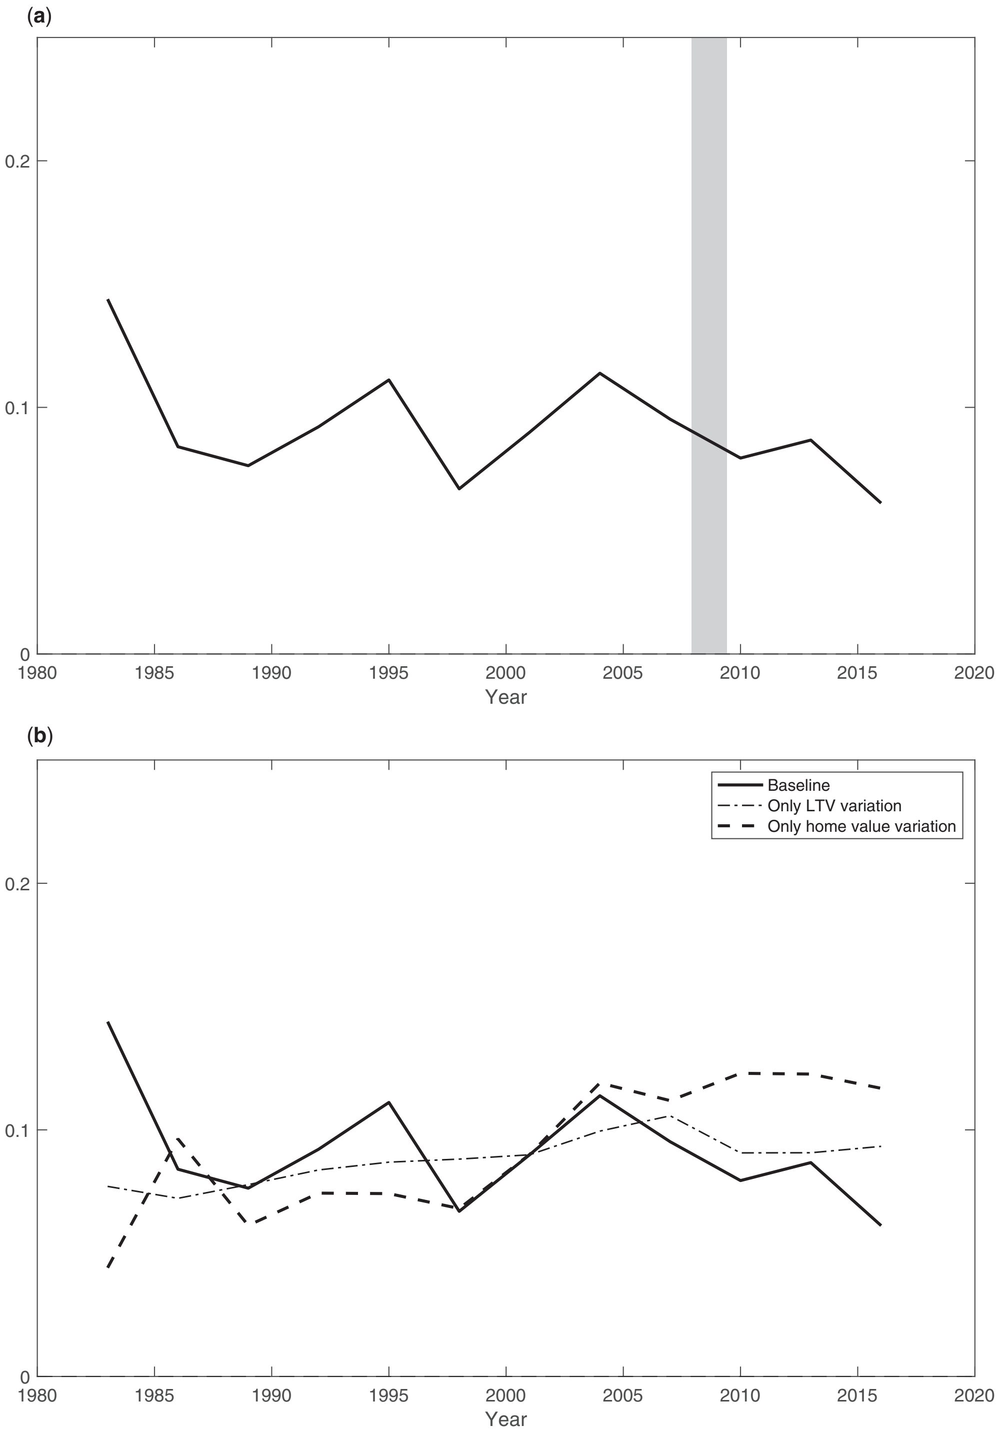

FIGURE 8

Model housing wealth elasticity

Notes: (A) (p/C)(dC/dp), where C is aggregate consumption in the population. C is calculated from $\int c(x,p)d\Phi_t(x)$, where the consumption function is the solution to the household's decision problem for a given relative price of housing and Φ_t is constructed from the SCF data for year t adjusted to match the CoreLogic Homeowner Equity Reports as explained in Supplementary Appendix B.3. We use a finite difference derivative that averages the values of plus and minus 10% price changes. (B) The same calculation with counterfactual Φ_t's constructed as described in the text.

TABLE 6
Decomposition of the housing wealth elasticity

	1986		2007		2010	
	Elas.	Group size	Elas.	Group size	Elas.	Group size
Renters (not moving)	0.01	0.30	0.00	0.30	0.00	0.31
Upsizers	−0.64	0.06	−0.89	0.04	-0.78	0.05
Downsizers	0.29	0.04	0.46	0.06	0.44	0.04
Stayers						
LTV ≤ 0.6	0.17	0.54	0.15	0.38	0.17	0.31
LTV $\in (0.6, 0.8]$	0.26	0.05	0.18	0.10	0.22	0.08
LTV $\in (0.8, 1.0]$	0.23	0.01	0.28	0.09	0.29	0.10
LTV ≥ 1.0	0.01	0.00	0.01	0.04	0.03	0.11
Total	0.084	1.00	0.095	1.00	0.079	1.00

Notes: We classify people according to their housing tenure, moving decisions, and LTV and then compute aggregate consumption within each cell before and after a price change to compute the dc for the group. We report $(p/C)(dc/dp)$ where C is average consumption in the population. The aggregate elasticity is the dot-product of the group elasticities and group sizes. Group sizes refer to shares of the population in the SCF adjusted to match equity estimates for homeowners with a mortgage in the CoreLogic Homeowner Equity Reports.

sheet, a given percentage change in the value of housing becomes more important to consumption decisions, and the model implied housing wealth elasticity rises.

7.1. *Why so stable?*

Why do large changes in the LTV distribution not lead to larger variation in the housing wealth elasticity in our model? To understand this, we decompose the aggregate model-implied housing wealth elasticity into a weighted average of the elasticities for groups with different LTVs. This is done in Table 6 and Figure 9. Table 6 decomposes the model-implied housing wealth elasticity for selected years by showing the average elasticity for renters, those moving into a larger house (upsizers), those moving into a smaller house (downsizers), and several LTV bins of those not moving (stayers) together with the relative size of each group.[32] The top panel of Figure 9 reports the model-implied housing wealth elasticity as a function of LTV in 2007 and 2010. The bottom panel of Figure 9 plots the density of households across LTV ratios for these years. The overall housing wealth elasticity for homeowners with a mortgage can be found by integrating the function in the top panel of Figure 9 with respect to the density in the lower panel of this figure.

There are two main forces that contribute to the stability of the model-implied housing wealth elasticity. First, a substantial fraction of the overall housing wealth elasticity is driven by stayers with relatively low LTVs (below 60%). The average housing wealth elasticity for this group is quite substantial—between 0.15 and 0.2—and this group is very large, accounting for 38% of the population in 2007 (62% of homeowners).[33] Looking at the top panel of Figure 9, it is also evident that the housing wealth elasticity is relatively constant for values of LTV below 70%. This implies that shifts in the LTV distribution over this range have little impact on the overall

32. We report $(p/C)(dc/dp)$, where C is average consumption in the population and c is group consumption. This means that the total elasticity can be calculated as the sum of the group elasticities multiplied by the group size.

33. At first blush, the figures in Table 6 look different from many estimates of the prevalence of negative equity in the bust. The reason for this is that Table 6 reports group sizes as a fraction of the population, while most equity figures are for homeowners with a mortgage. For instance, CoreLogic estimates that in 2010Q4 23.2% of homeowners with a mortgage had negative equity. We use the share of households with a mortgage from the SCF to obtain that 11.7% of the population was a homeowner with negative equity. Table 6 shows 11% because some of the underwater households are allocated to the "downsizers" bin if they downsize, switch to renting, or default.

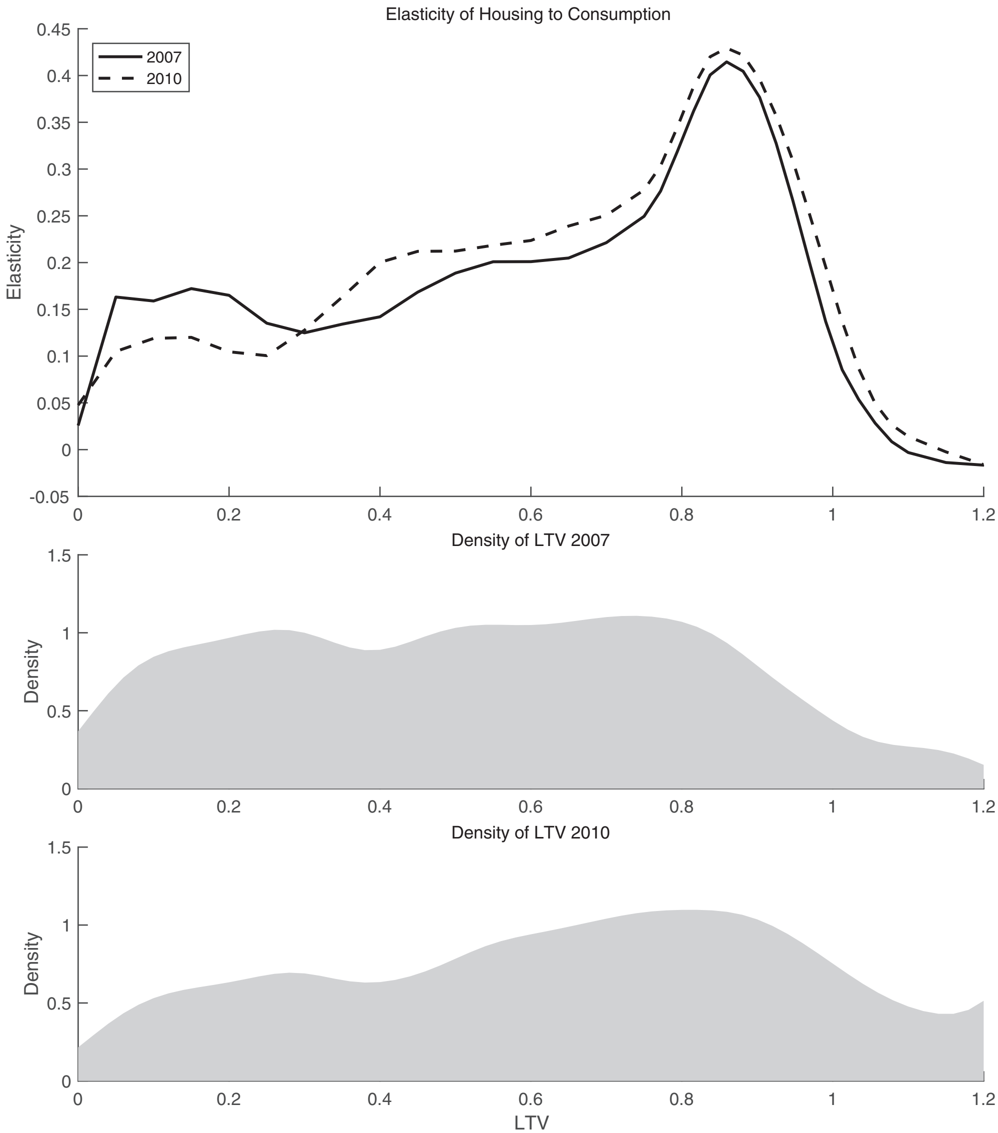

FIGURE 9

Housing wealth elasticity by LTV and marginal distributions of LTV in 2007 and 2010.

Notes: The top panel shows the elasticity implied by the model for a particular LTV. For a given LTV, we weight households by their distance from that LTV using Gaussian kernel with bandwidth 0.05 and report the weighted average elasticity. The results of this calculation depending on the (conditional) distribution of other state variables for a given LTV and this accounts for the difference between 2007 and 2010. The lower panels show kernel density estimates of the LTV distribution using the same kernel and bandwidth.

housing wealth elasticity. In other words, the overall housing wealth elasticity for this group is relatively stable in the face of shifts in the LTV distribution. Since this is a large group, its stability contributes to stabilizing the aggregate housing wealth elasticity.

It is useful to understand why the low-LTV households in our model increase their consumption substantially in response to house price appreciation. Berger *et al.* (2018) show using an incomplete markets model that a household's consumption response to changes in housing wealth at low levels of leverage are driven by a substitution effect, an endowment effect, and an income effect.[34] When house prices increase, households shift consumption away from housing and towards other goods because housing has become more expensive (substitution effect), they feel wealthier, which tends to lead them to increase non-housing consumption (endowment effect), but they perceive higher implicit rents from occupying their more valuable house, which tends to decrease non-housing consumption (income effect). Berger *et al.* show that in their incomplete markets model, unlike the permanent income hypothesis model, the endowment effect of a house price increase may be substantially larger in magnitude than the (negative) income effect. Households effectively discount future increases in implicit rental costs relative to immediate increase in the value of their home. This is true even for many households with low LTVs because they still face liquidity constraints as home equity is illiquid and many of them have few liquid assets. Furthermore, for low-LTV households the magnitude of the housing wealth elasticity is not much affected by changes in leverage because the substitution, endowment, and income effects do not depend strongly on leverage.

The other factor that plays an important role in the stability of the model-implied housing wealth elasticity—particularly in the Great Recession—is the fact that it does not rise monotonically with LTV but rather features a "hump" around the LTV constraint. This hump stands out in the top panel of Figure 9. Households that have LTV ratios close to the LTV constraint tend to have low liquid assets and have a high marginal propensity to consume for precautionary reasons. At an LTV of 0.8, the LTV constraint binds, and the model implied housing wealth elasticity jumps and remains high until households reach an LTV of about 0.95. Intuitively, the households in this region tend to be highly financially constrained, and changes in the house price tighten or loosen these constraints. When house prices rise, these households respond by refinancing their mortgage, downsizing their house, or selling to rent, all of which allow them to increase consumption. Once the LTV ratio rises above roughly 0.95, however, the model-implied housing wealth elasticity drops rapidly. As Ganong and Noel (2020) emphasize, households that are underwater on their mortgage are not able to access changes in housing wealth induced by changing house prices. Their LTV is too high for them to be able to refinance or sell their house unless they have other liquid wealth to help pay off their current mortgage. At the same time, these households are not forced to de-lever when their mortgage debt is long-term. They can simply pay their mortgage down over time. Their consumption is consequently highly insensitive to house price changes.

The hump in the model-implied housing wealth elasticity as a function of LTV means that as house prices fell during the Great Recession some households were pushed from lower LTVs into the hump, but at the same time other households were pushed out the other side of the hump to even higher LTVs. The lower panels of Figure 9 illustrate this by plotting kernel density estimates of the LTV distribution of homeowners in 2007 and 2010. Recall that one can obtain the aggregate wealth elasticity for homeowners with a mortgage by integrating the function in the top panel with respect to these distributions. On net, the two effects roughly cancel, causing the overall housing wealth elasticity in the model to be insensitive to changes in the LTV distribution. While the extent to which these two effects cancel depends on the exact distributions, these two countervailing forces will generally help stabilize the housing wealth elasticity in a housing bust.

Returning to Table 6, we see that between 1986 and 2010 the fraction of the population in the low-LTV (≤ 0.6) stayers category fell by 23%. These people were mostly pushed into the high-LTV $(0.8, 1.0]$ and underwater (≥ 1.0) categories. The high-LTV group grew by 9% of the population, while the underwater group grew by 11% of the population. The comparison between

34. More generally, there is also a collateral effect. But this effect is less important for households far from the LTV constraint.

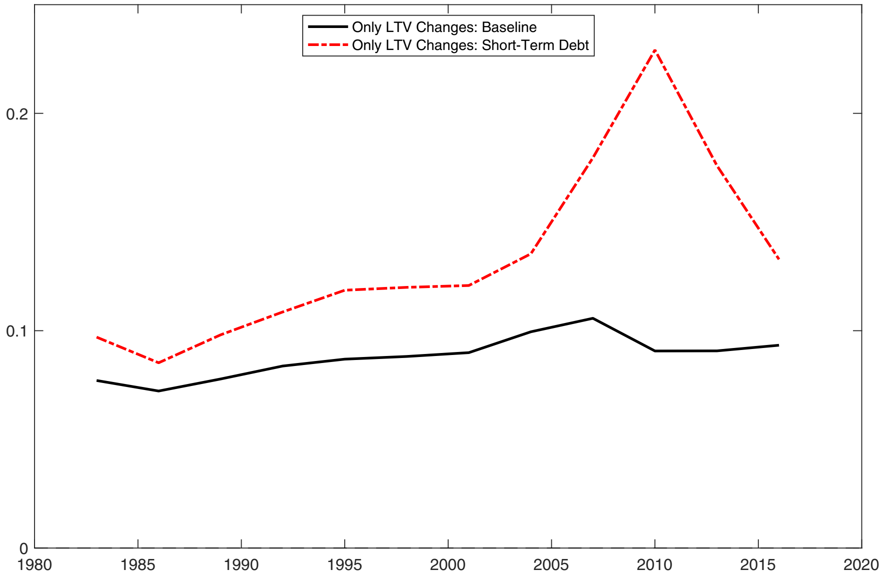

FIGURE 10

Short-term debt counterfactual

Notes: The figure shows the changes in the housing wealth elasticity that result from changes in the LTV distribution as in Figure 8. The figure repeats the results from our baseline model and adds results from an alternative model in which the LTV constraint must hold every period not just at the origination of a mortgage.

2007 and 2010 is even more stark. The number of low-LTV households decreased by 7% and the number of underwater households increased by 7%, resulting in a net decline in the housing wealth elasticity.

The presence of long-term debt in our model is important for these results. Figure 10 compares our baseline model to one with only short-term debt where all homeowners are subject to the LTV constraint every period. With short-term debt, the increase in LTVs in the Great Recession leads to a sharp increase in the model-implied housing wealth elasticity and a boom–bust asymmetry, in contrast to both our empirical analysis and our model with long-term debt. Supplementary Appendix E.6 shows a version of the top panel of Figure 9 for a model with short-term debt. In this case, the consumption of homeowners with high LTVs is much more sensitive to home prices because they are subject to the LTV constraint period by period. As a result, an increase in the number of high-LTV households has a much more pronounced effect on the aggregate housing wealth elasticity. Moreover, underwater households are forced to delever. Their consumption, therefore, remains sensitive to home prices even at very high LTVs. This deleveraging when house prices fall drives up the housing wealth elasticity in the bust.

7.2. *Credit conditions and the housing wealth elasticity*

Our analysis above abstracts from variation in credit constraints over time. Part of the reason why household leverage rose as much as it did in the early 2000s may have been because of increased credit availability. To assess how such changes might affect our results, consider a scenario in which the LTV limit rises from 80% to 95% and the mortgage origination cost is reduced to zero. We refer to this parameterization as the "boom" parameterization. Figure 11 compares the housing wealth elasticity for our baseline parametrization of our model and the boom parametrization. In both cases, we are using the distribution of state variables (LTV, etc) observed in the SCF for

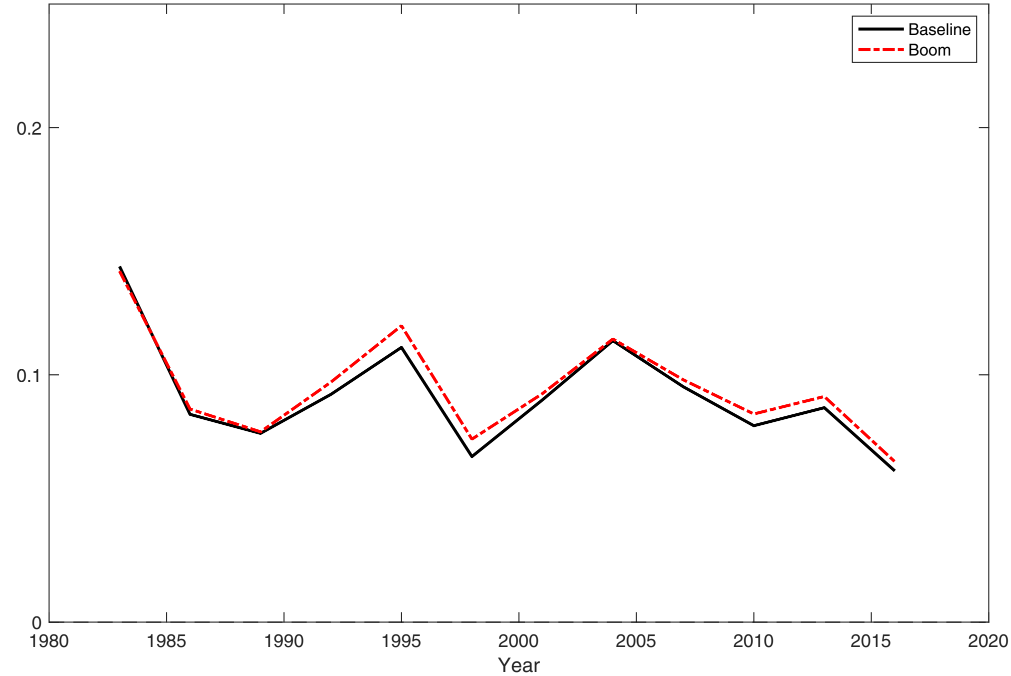

FIGURE 11

The housing wealth elasticity under boom parameters

Notes: The figure shows the housing wealth elasticity under the boom parameters: no mortgage origination cost, and an LTV limit of 95%.

each year. This implies that the difference between the housing wealth elasticity in the baseline and the boom case measures the effect of relaxing credit constraints substantially for a given set of households.[35]

The striking result in Figure 11 is that relaxing credit constraints has minimal effect on the housing wealth elasticity in our model. The key intuition is very similar to the intuition for the stability of the wealth elasticity when the LTV distribution changes: a substantial part of the wealth elasticity is driven by unconstrained households. Furthermore, shifting the location of the credit constraint shifts the location of the "hump," leading to offsetting effects as households move in and out of the hump.

This finding may seem surprising in light of analyses such as Guerrieri and Lorenzoni (2017), who show that tighter credit conditions can lead to an economic contraction, or Landvoigt *et al.* (2015), who show that cheap credit for poor households was a driving force in the house price boom in San Diego. Note, however, that our object of interest is not the *level* of consumption—which is sensitive to credit conditions in our model—but the *sensitivity* of consumption to home price changes, which is not. Also, our analysis takes the distribution of individual states as given from the data, so we are not changing these state variables as we change credit conditions.

We have emphasized that the insensitivity of the housing wealth elasticity in our model to changes in the LTV distribution and to credit conditions is consistent with our empirical results in the 2000s boom–bust cycle. It is worth noting, however, that our model does not provide an

35. Conditional on a household's current state variables, the constraints faced in the past are irrelevant to current decisions. Therefore the two lines in Figure 11 show the effect of a permanent change in credit conditions for the immediate consumption response to home prices.

explanation for the increase and decrease in the housing wealth elasticity we estimate in the 1990s when using the sensitivity instrument. The large increase in leverage that occurred 1980s and 1990s may seem like a plausible cause of these unusually large housing wealth elasticity estimates, butour model does not support this notion. Furthermore, the timing is of—the major expansion in credit occurred *before* the 10-year window from 1992 to 2002 that has a particularly high elasticity using the sensitivity instrument.

7.3. *No short-run housing adjustment*

A potential concern with our partial equilibrium analysis above is that changes in house prices lead to non-trivial changes in the demand for housing. Intuitively, households would like to substitute away from houses and towards other consumption goods when the relative price of houses increases.[36] In the short run, however, housing supply is quite inelastic both because construction of new houses takes time and also because each year's construction of new houses represents a small addition to the overall housing stock. This implies that the change in housing demand implied by our partial equilibrium model is unlikely to be consistent with market clearing in the housing market.

To address this concern, we consider an alternate experiment in which there is no change in housing demand in the short run. We compare two cities that both have completely inelastic housing supply in the short run, but have different long-run housing supply elasticities. These cities are hit by an aggregate shock that leads house prices to rise by 10% more in the less-elastic city in the long run (*e.g.* an expectation that demand for housing will be higher in the long run). In the short run, house prices are endogenously determined so that housing demand is unchanged, i.e., the housing market clears with no change in the housing stock. The short-run equilibrium increase in house prices in the less-elastic city relative to the more-elastic city is less than 10%, leaving an expected relative capital gain that dissuades households in that city from reducing their relative housing demand despite its higher price.[37]

Figure A.29 in the Supplementary Appendix compares the model-implied housing wealth elasticity for this alternate experiment to our baseline model-implied housing wealth elasticity. The demand for housing is highly sensitive to expected capital gains. Even a small expected capital gain is sufficient to equilibrate the housing market without any change in quantity of housing. The housing wealth elasticity in this alternate experiment is virtually identical to our baseline. This result demonstrates that the aggregate housing wealth effect in our baseline analysis does not stem from a substitution out of housing in the aggregate, which would be difficult to reconcile with an inelastic supply of housing in the short run.

7.4. *Extensions and robustness*

Our theoretical analysis uses the observed distribution of idiosyncratic states instead of the distribution that is generated by the model. Supplementary Appendix E.9 explores how well the model is able to explain the evolution of household balance sheets from one wave of the SCF to the next. The model's predictions align with the data fairly well except for the period of the housing boom where the model predicts declining LTVs while the observed LTV distribution remained quite stable. An extended model with the "boom parameterization" discussed above and some

36. Our baseline model implies a price elasticity of housing demand of 0.15.

37. This experiment was motivated in part by the belief shocks of Kaplan *et al.* (2019), which raise expectations of future housing demand without affecting current preferences. In their general equilibrium model, the belief shock raises current prices through expectations of future capital gains.

modest anticipated capital gains during the boom years is able to match the behaviour of leverage in these years. These modifications have a very limited effect on the housing wealth elasticity. Credit constraints have little effect on the elasticity, as we discuss in Section 7.2. Allowing for a modest increase in expected capital gains also has a small effect, increasing the housing wealth elasticity in 2007 by only about 16%.

We explore how our results are affected by several other modelling choices in the Supplementary Appendix. We show that the housing wealth elasticity is modestly increasing in mortgage interest rates, which we set at 3% in our baseline specification (see Supplementary Figure A.25). At higher interest rates, households are more likely to downsize their homes and downsizers have large elasticities. This is particularly true of high-LTV households, who were more common in recent years.

Our baseline analysis assumes that rents are proportional to home prices. During the housing boom of the 2000s, the rent-price ratio fell considerably. Making the polar opposite (and also unrealistic) assumption that rents remain constant when home prices change yields a somewhat higher housing wealth elasticity because it leads renters to defer buying a house and spend more on non-housing consumption. However, this alternate rent assumption does not change the time series pattern of the housing wealth elasticity (see Supplementary Figure A.26).

A larger cost of selling a home reduces the housing wealth elasticity in our model. As the transaction cost increases, homeowners become less willing to realize the capital gain or loss on their houses and consumption becomes more insulated from price fluctuations. While the moving cost affects the level of the housing wealth elasticity, it does not change the time series pattern (see Supplementary Figure A.28).

8. CONCLUSION

In this article, we provide new evidence on the housing wealth elasticity going back to the 1980s. These estimates indicate that the housing wealth elasticity was if anything smaller post-2000 than earlier in our sample. Our results indicate that the outsized role of housing in the economy during this period arose exclusively from the large magnitude of house price movements during this period, as opposed to economic activity being more sensitive to house prices in the 2000s than before.

Our empirical findings are based on three methods: OLS with a rich set of observables, the Saiz housing supply elasticity instrument, and a new sensitivity instrument we develop that exploits the differential sensitivity of house prices in cities to regional house price cycles. Importantly, we use a panel approach for all three identification strategies, which allows us to include fixed effects and time-varying controls to account for differential city trends and cyclical sensitivities. All three approaches come to the same conclusions regarding the trajectory of the housing wealth elasticity over our sample period. We also find no evidence of a boom–bust asymmetry. Our new sensitivity instrument yields substantially smaller and more precise estimates than those based on the Saiz instrument, though they remain economically important.

We develop a theoretical framework to interpret our empirical estimates. We show that our empirical approach yields an estimate of the partial equilibrium effect of house prices on consumption multiplied by a local general equilibrium multiplier that can be approximated by the local fiscal multiplier. All other general equilibrium effects are soaked up by fixed effects in our regressions. Our empirical approach thus allows us to draw inferences about the effects of house price fluctuations while remaining agnostic about fundamental shocks that drive house prices. We then analyze the partial equilibrium housing wealth elasticity in the context of a life-cycle model with borrowing constraints, uninsurable income risk, illiquid housing, and long-term mortgages.

In this canonical model, there is no increase in the housing wealth elasticity associated with rightward shifts in the LTV distribution such as those that occurred during the Great Recession. This arises for two reasons in the model. First, much of the housing wealth elasticity arises from impatient low-LTV households who have a substantial wealth elasticity that is insensitive to LTV and create a stabilizing force for the aggregate housing wealth elasticity. Second, there are offsetting effects of the increase in the number of highly constrained households during the Great Recession on the one hand (which tended to raise the elasticity) and the increase in the number of underwater households on the other hand (which tended to lower the elasticity).

Our empirical and theoretical analyses together indicate that the substantial housing wealth elasticities observed in the 2000s were not driven by special features of the 2000s boom–bust cycle. Instead, a substantial wealth elasticity is instead a fundamental feature of the economy going back to at least the 1980s.

Acknowledgments. We would like to thank Massimiliano Cologgi, Hope Kerr, Jimmy Kuo, Joao Fonseca Rodrigues, Jesse Silbert, Xuiyi Song, Yeji Sung, and Sergio Villar for excellent research assistance. We would like to thank Aditya Aladangady, Adrien Auclert, James Cloyne, Masao Fukui, Peter Ganong, Dan Greenwald, Jonathon Hazell, Erik Hurst, Virgiliu Midrigan, Raven Molloy, Pascal Noel, Chris Palmer, Jonathan Parker, Monika Piazzesi, Esteban Rossi-Hansberg, Martin Schneider, Johannes Stroebel, Stijn Van Nieuwerburgh, Joseph Vavra, Gianluca Violante, Ivan Werning, and seminar participants at various institutions and conferences for useful comments. Guren thanks the National Science Foundation (grant SES-1623801) and the Boston University Center for Finance, Law, and Policy. Nakamura thanks the National Science Foundation (grant SES-1056107). Nakamura and Steinsson thank the Alfred P. Sloan Foundation for financial support. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Supplementary Data

Supplementary data are available at Review of Economic Studies online.

REFERENCES

ADELINO, M., SCHOAR, A. and SEVERINO, F. (2015), “House Prices, Collateral, and Self-Employment”, *Journal of Financial Economics*, **117**, 288–306.

AGARWAL, S., AMROMIN, G., CHOMSISENGPHET, S. *et al.* (2020), “Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinancing Program”. NBER Working Paper No. 21512.

ALADANGADY, A. (2017), “Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata”, *American Economic Review*, **107**, 3415–3446.

ASDRUBALDI, P., SORENSEN, B. E. and YOSHA, O. (1996), Channels of Interstate Risk Sharing: United States 1963-1990”, *The Quarterly Journal of Economics*, **111**, 1081–1110.

ATTANASIO, O., BLOW, L., HAMILTON, R. *et al.* (2009), Booms and Busts: Consumption, House Prices, and Expectations”, *Economica*, **76**, 20–50.

ATTANASIO, O., LEICESTER, A. and WAKEFIELD, M. (2011), Do House Prices Drive Consumption Growth? The Coincident Cycles of House Prices and Consumption in the UK”, *Journal of the European Economic Association*, **9**, 399–435.

BAILEY, M., CAO, R., KUCHLER, T. *et al.* (2018), The Economic Effects of Social Networks: Evidence From the Housing Market”, *Journal of Political Economy*, **126**, 2224–2276.

BERGER, D., GUERRIERI, V., LORENZONI, G. *et al.* (2018), House Prices and Consumer Spending”, *The Review of Economic Studies*, **85**, 1502–1542.

BHUTTA, N. and KEYS, B. J. (2016), Interest Rates and Equity Extraction During the Housing Boom”, *American Economic Review*, **106**, 1742–1774.

CALOMIRIS, C. W., LONGHOFER, S. D. and MILES, W. (2013), The Housing Wealth Effect: The Crucial Roles of Demographics Wealth Distribution, and Wealth Shares”, *Critical Finance Review 2*, 49–99.

CAMPBELL, J. Y. and COCCO, J. F. (2003), Household Risk Management and Optimal Mortgage Choice”, *Quarterly Journal of Economics*, **118**, 1449–1494.

CAMPBELL, J. Y. and COCCO, J. F. (2007), How Do House Prices Affect Consumption? *Journal of Monetary Economics*, **54**, 591–621.

CARROLL, C. D., OTSUKA, M. and SALACALEK, J. (2011), How Large Are Housing and Financial Wealth Effects? A New Approach”, *Journal of Money, Credit and Banking*, **43**, 55–79.

CASE, K. E., SHILLER, R. J. and QUIGLEY, J. M. (2005), Comparing Wealth Effects: The Stock Market vs. the Housing Market”, *Advances in Macroeconomics*, **5**, 1–34.

CASE, K. E., SHILLER, R. J. and QUIGLEY, J. M. (2013), Wealth Effects Revisited, 1975-2012", *Critical Finance Review*, **2**, 101–128.

CHANEY, T., SRAER, D. and THESMAR, D. (2012), The Collateral Channel: How Real Estate Shocks Affect Corporate Investment", *American Economic Review*, **102**, 2381–2409.

CHODOROW-REICH, G. and WIELAND, J. (2018), "Secular Labor Reallocation and Business Cycles" (Working Paper #21864, NBER).

CLOYNE, J., HUBER, K., ILZETZKI, E. *et al.* (2019), The Effect of House Prices on Household Borrowing: A New Approach", *American Economic Review*, **109**, 2104–2136.

COOPER, D. (2013), House Price Fluctuations: The Role of Housing Wealth as Borrowing Collateral", *The Review of Economics and Statistics*, **95**, 1183–1197.

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**, 177–206.

DAVIS, M. and VAN NIEUWERBURGH, S. (2015), "Housing, Finance, and the Macroeconomy", in Duranton, G., Henderson, J. V. and Strange, W. C. (eds) *Handbook of Regional and Urban Economics*, Vol. 5 (Amsterdam: Elsevier) 753–811.

DEFUSCO, A. A. (2018), Homeowner Borrowing and Housing Colalteral: New Evidence from Expiring Price Controls", *Journal of Finance*, **73**, 523–573.

DENG, Y., QUIGLEY, J. M. and VAN ORDER, R. (2000), Mortgage Terminations, Heterogeneity, and the Exercise of Mortgage Options", *Econometrica*, **68**, 275–307.

FAVARA, G. and IMBS, J. (2015), Credit Supply and the Price of Housing", *American Economic Review*, **105**, 958–992.

FAVILUKIS, J., LUDVIGSON, S. C. and VAN NIEUWERBURGH, S. (2017), The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium", *Journal of Political Economy*, **125**, 140–223.

FLAVIN, M. and NAKAGAWA, S. (2008), A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence", *American Economic Review*, **98**, 474–495.

FOOTE, C. L., GERARDI, K. S. and WILLEN, P. S. (2012), "Why Did So Many People Make So Many Ex Post Bad Decisions? The Causes of the Foreclosure Crisis", in Blinder, A. S. Lo, A. W., and Solow, R. M. (eds) *Rethinking the Financial Crisis* (New York, NY: Russel Sage Foundation).

GANONG, P. and NOEL, P. (2020), "The Effect of Debt on Default and Consumption: Evidence from Housing Policy in the Great Recession" (Working Paper #24964, NBER).

GARRETT, T. A., HERNANDEZ-MURILLO, R. and OWYANG, M. T. (2005), Does Consumer Sentiment Predict Regional Consumption? *Federal Reserve Bank of St. Louis Review* (March/April Part 1), 123–135.

GILCHRIST, S. and ZAKRAJSEK, E. (2012), Credit Spreads and Business Cycle Fluctuations", *American Economic Review*, **102**, 1692–1720.

GLAESER, E. L. and GYOURKO, J. (2005), Urban Decline and Durable Housing", *Journal of Political Economy*, **113**, 345–375.

GOREA, D. and MIDRIGAN, V. (2018), "Liquidity Constraints in the U.S. Housing Market" (Working Paper #23345, NBER).

GUERRIERI, L. and IACOVIELLO, M. (2017), Collateral Constraints and Macroeconomic Asymmetries", *Journal of Monetary Economics 90*, 28–49.

GUERRIERI, V. and LORENZONI, G. (2017), Credit Crises, Precautionary Savings, and the Liquidity Trap", *The Quarterly Journal of Economics*, **132**, 1427–1467.

GUREN, A. M., KRISHNAMURTHY, A. and MCQUADE, T. (2019), "Mortgage Design in an Equilibrium Model of the Housing Market" (Working Paper #24446, NBER).

GUREN, A. M., MCKAY, A., NAKAMURA, E. *et al.* (2019), "What Can We Learn From Cross-Regional Empirical Estimates in Macroeconomics?" (Working Paper #26881, NBER).

GUVENEN, F. and ANTHONY, A. SMITH, J. (2014), Inferring Labor Income Risk and Partial Insurance From Economic Choices", *Econometrica*, **82**, 2085–2129.

GUVENEN, F., KARAHAN, F., OZKAN, S. *et al.* (2019), " What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?" (Working Paper # 20913, NBER).

HEATHCOTE, J., PERRI, F. and VIOLANTE, G. L. (2010), Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006", *Review of Economic Dynamics*, **13**, 15–51.

HURST, E. and STAFFORD, F. P. (2004), Home Is Where the Equity Is: Mortgage Refinancing and Household Consumption", *Journal of Money, Credit, and Banking*, **36**, 985–1014.

KAPLAN, G., MITMAN, K. and VIOLANTE, G. L. (2020), "Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence From Easily Accessible Data" (Working Paper #22232, NBER).

KAPLAN, G., MITMAN, K. and VIOLANTE, G. L. (2019), "The Housing Boom and Bust: Model Meets Evidence" (Working Paper #23694, NBER).

KUHN, M., SCHULARICK, M. and STEINS, U. I. (2017), The Great American Debt Boom, 1948-2013 (Working Paper, University of Bonn).

LANDVOIGT, T., PIAZZESI, M. and SCHNEIDER, M. (2015), The Housing Market(s) of San Diego", *American Economic Review*, **105**, 1371–1407.

LI, W. and YAO, R. (2007), The Life-Cycle Effects of House Price Changes", *Journal of Money, Credit and Banking*, **39**, 1375–1409.

LIEBERSOHN, J. (2017), “Housing Demand, Regional House Prices, and Consumption” (Working Paper, The Ohio State University).

MALLABY, S. (2016), *The Man Who Knew: The Life and Times of Alan Greenspan* (New York: Penguin).

MIAN, A., RAO, K. and SUFI, A. (2013), Household Balance Sheets, Consumption, and the Economic Slump”, *Quarterly Journal of Economics 128*, 1687–1726.

MIAN, A. and SUFI, A. (2011), House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis”, *American Economic Review*, **101**, 2132–2156.

MIAN, A. and SUFI, A. (2014), What Explains the 2007-2009 Drop in Employment? *Econometrica*, **82**, 2197–2223.

NAKAMURA, B. E. and STEINSSON, J. (2014), Fiscal Stimulus in a Monetary Union: Evidence from US Regions”, *American Economic Review*, **104**, 753–792.

NATHANSON, C. G. and ZWICK, E. (2018), Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market”, *Journal of Finance*, **73**, 2587–2633.

PALMER, C. (2015), “Why Did So Many Subprime Borrowers Default During the Crisis: Loose Credit or Plummeting Prices?” (Working paper, Massachusetts Institute of Technology).

PIAZZESI, M., SCHNEIDER, M. and TUZEL, S. (2007), Housing, Consumption, and Asset Pricing”, *Journal of Financial Economics 83*, 531–596.

RODGERS, J. D. and TEMPLE, J. A. (1996), “Sales Taxes, Income Taxes, and Other Non-Property Tax Revenues”, in Aronson, J. and Schwartz, E. (eds) *Management Policies in Local Government Finance* (Washington D.C.: International City/County Management Association for the ICMA University) 229–259.

ROUSSANOV, N., MICHAUX, M. and CHEN, H. (2020), “Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty”, *Journal of Finance*, **75**, 323–375.

SAIZ, A. (2010), The Geographic Determinants of Housing Supply Elasticity”, *Quarterly Journal of Economics*, **125**, 1253–1296.

SINAI, T. (2013), “House Price Moments in Boom-Bust Cycles”, in Glaeser, E. L. and Sinai, T. (eds) *Housing and the Financial Crisis* (Chicago: University of Chicago Press) 19–68.

SINAI, T. and SOULELES, N. S. (2005), “Owner-Occupied Housing as a Hedge Against Rent Risk”, *The Quarterly Journal of Economics*, **120**, 763–789.

WONG, A. (2019), “Refinancing and the Transmission of Monetary Policy to Consumption” (Working paper, Princeton University).

ZHOU, X. (2010), *Essays on U.S. State-Level Financial Wealth Data and Consumption Data* (Ph. D. Thesis, Johns Hopkins University).