# Housing Wealth Effects: The Long View

Adam M. Guren[*], Alisdair McKay[†], Emi Nakamura[‡], and Jón Steinsson[§¶]

February 20, 2020

**Abstract**

We provide new time-varying estimates of the housing wealth effect back to the 1980s. We use three identification strategies: OLS 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 behavior 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 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.

[*]Boston University and NBER, guren@bu.edu

[†]Federal Reserve Bank of Minneapolis, and NBER, alisdair.mckay@mpls.frb.org

[‡]UC Berkeley and NBER, enakamura@berkeley.edu

[§]UC Berkeley and NBER jsteinsson@berkeley.edu

[¶]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.

# 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, Rao, and Sufi, 2013; Mian and Sufi, 2014). The question we seek to answer in this paper 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 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 behavior of the housing

[1]To our knowledge, only two papers have looked at changes in the housing wealth elasticity over time. First, Case, Shiller, and Quigley (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, Shiller, and Quigley (2005), which uses data for 1982-1999, and Case, Shiller, and Quigley (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, Otsuka, and Salacalek (2011), Attanasio et al. (2009, 2011), Calomiris, Longhofer, and Miles (2012), Cooper (2013); DeFusco (2016); Kaplan, Mitman, and Violante (2017), and Liebersohn (2017).

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's (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, Rao, and Sufi, 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 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]

[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, Piazzesi, and Schneider, 2015; Favilukis, Ludvigson, and Van Nieuwerburgh, 2017; Kaplan, Mitman, and Violante, 2017).

[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, Mitman, and Violante, 2016), to credit card spending (e.g., Mian, Rao, and Sufi, 2013) to broader measures based on the Current Expenditure Survey (Aladangady, 2017).

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 instrument, the housing wealth elasticity is statistically significantly *smaller* over the boom and bust (2000-2012) 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 percent 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

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

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 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 (2019) 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-

[5]See, e.g., Agarwal et al., 2017; Berger et al., 2018; Chen, Michaux, and Roussanov, 2018; Davis and Van Nieuwerburgh, 2015; Gorea and Midrigan, 2018; Guren, Krishnamurthy, and McQuade, 2019; Kaplan, Mitman, and Violante, 2017; Li and Yao, 2007; Favilukis, Ludvigson, and Van Nieuwerburgh, 2017).

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

2010 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 percent to 9 percent 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, Gerardi, and Willen, 2012; Kuhn, Schularick, and Steins, 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]

The paper proceeds as follows. Section 2 describes our main data sources. Section 3 describes

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

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 analyzes 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

[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 Asdrubali 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.

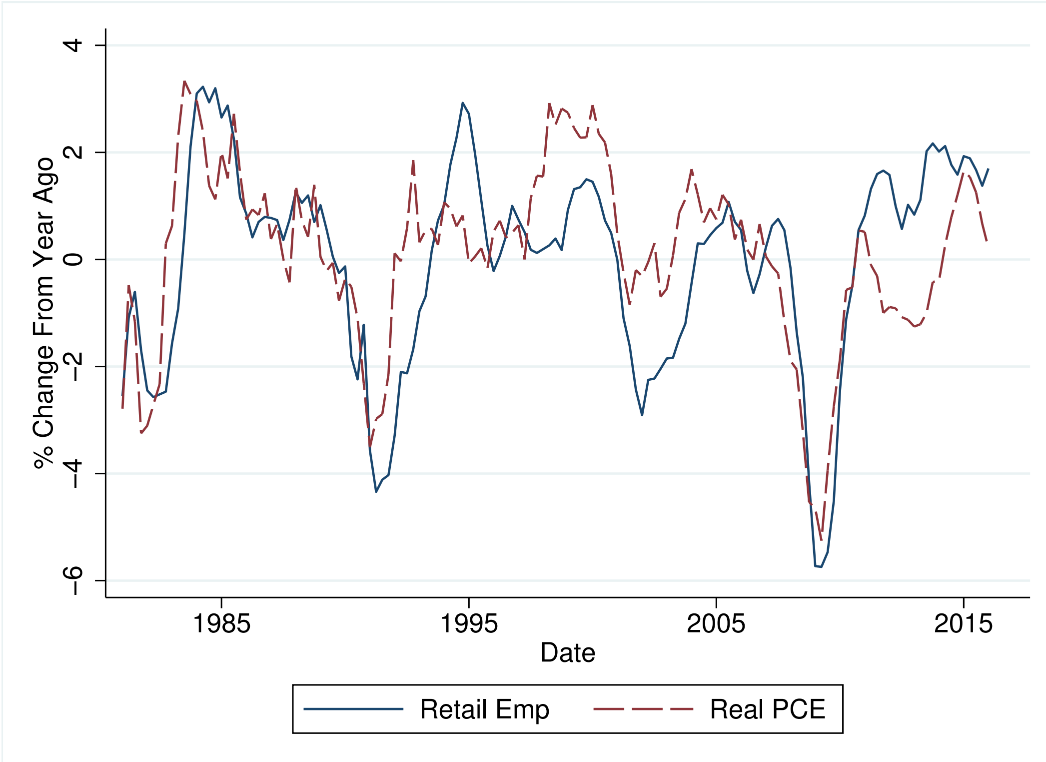


Figure 1: Growth of Retail Employment vs. Growth in Personal Consumption Expenditures

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 .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.

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 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 in raw form (Garett et al, 2005).[9] Some researchers have used data compiled by private sector data sources,

[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.

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 Appendix A.3 we analyze 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 to 2010 and inter-Censal population estimates for 2010 to 2017. These population estimates are available annually, and we interpolate to 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. 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 Appendix D.1.5, we redo our analysis using the CoreLogic house price index. The results are

[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.

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

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 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, $\psi_i$ 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 unmodeled influences on retail employment.

The coefficient of interest in equation (1) is $\beta$, which measures the housing wealth elasticity. Several challenges arise in estimating $\beta$. 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 $\beta$ upward since a strong economy will cause house prices to rise. On the other hand, house prices are measured with error, potentially biasing $\beta$ 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, Rao, and Sufi, 2013; Mian and Sufi, 2014). This work typically uses an 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 favoring 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 $\Delta Y$ is the log change in regional retail employment. In this equation, $\alpha_i$ 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 labor 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 Gilchirst and Zakrajesk'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

[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 centered in the mid-1990s and highly imprecise results with lower point estimates in the early 1990s.

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 kilometer radius of the center 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 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.

[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.

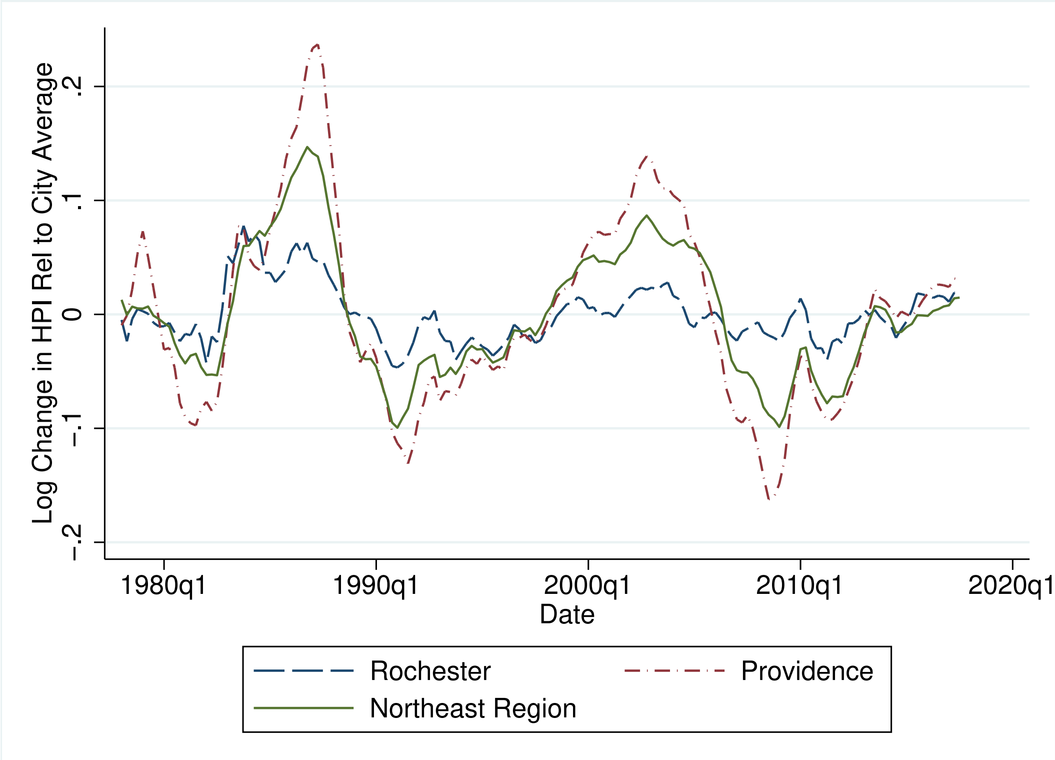


Figure 2: House Prices in Providence, Rochester, and the Northeast Region

Note: 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.

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 $\gamma_i$ 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 $\gamma_i$ 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, Rao, and Sufi (2013) and

[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 paper.

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 $\gamma_i$ 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 $\gamma_i$ 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 $\gamma_i$s 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 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 $\alpha_i$ 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 $\alpha_i$ is correlated with $\hat{\gamma}_i$ in the cross section. Since we are estimating $\beta$ 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 labor 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 flavor. 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

[15] One potential concern with this procedure is the role of measurement error in $\Delta y_{i,r,t}$ biasing the $\delta_i$ terms and thereby creating bias in the $\gamma_i$s. 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 $\delta_i$ is the same across CSBAs, but the $\delta$ we obtain is a causal elasticity. We can use this IV regression to estimate $\gamma_i$. This approach yields values for the $\gamma_i$ that are highly correlated with our baseline approach, and using these alternate $\gamma_i$s does not significantly alter our results.

[16] 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.

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 $\gamma_i$ 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 $\gamma_i$'s for a particular time period are estimated using data from all years except a seven 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 $\gamma_i$ 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. Panel A shows our $\hat{\gamma}_i$s, while panel B 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 kilometer radius of the center of a city that is not suitable for construction due to steep slopes or water, and

[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.

A. $\gamma_i$ at CBSA Level (Darker is Higher $\gamma_i$)

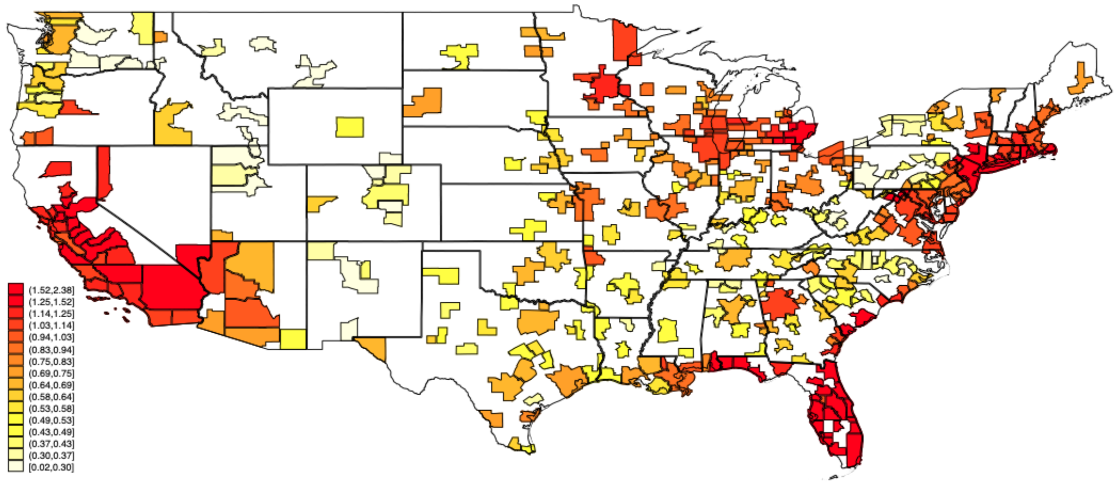


B. Saiz Estimated Housing Elasticity at City Level (Darker is More Inelastic)

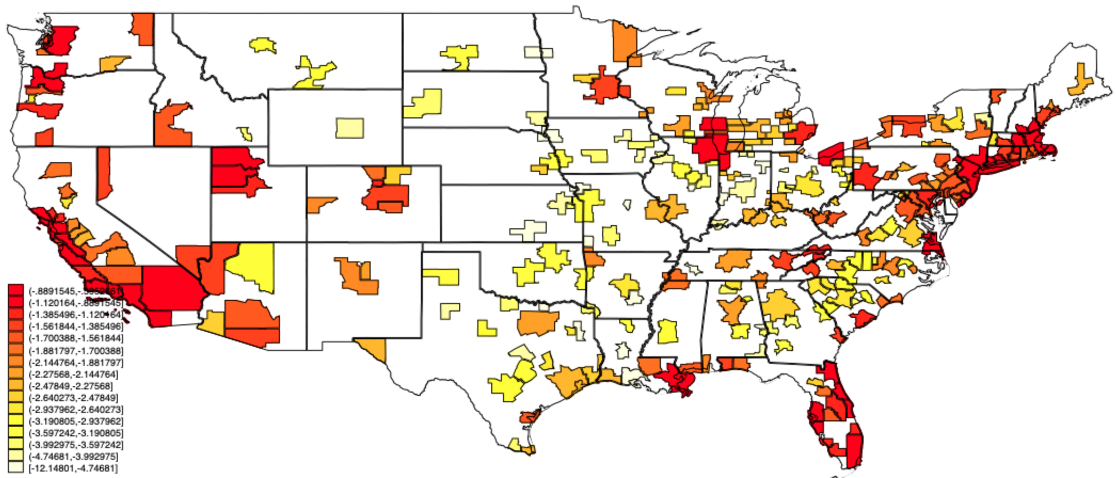


Figure 3: $\gamma_i$ and Saiz Elasticity by CBSA for Continental U.S.

Notes: These Figures provide heat maps for $\gamma_i$ and the Saiz elasticity. $\gamma_i$ 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 colors 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.

presumably this leaves out a number of important factors that determine land supply elasticities (for example, the 50 kilometer 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 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 Gyouko (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.

# 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 2006-2009. 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 $\beta$ 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 in Appendix A.3.[20] The OLS estimate implies an MPCH of 3.82 cents on the dollar, while the Saiz

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

[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

Table 1: Pooled Elasticity of Retail Employment Per Capita to House Prices

| Time Period | (1) 1978-2017 | (2) 1990-2017 | (3) 2000-2017 |
|---|---|---|---|
| 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) |

Note: Each column estimates equation (1) for the indicated time period. "OLS" uses no instrument. "Sensitivity Instrument" uses our sensitivity instrument with the $\gamma_i$s 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 Gilchirst-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 5% level, ** at the 1% level, and *** at the 0.1% level.

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 1978-1990 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.

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.

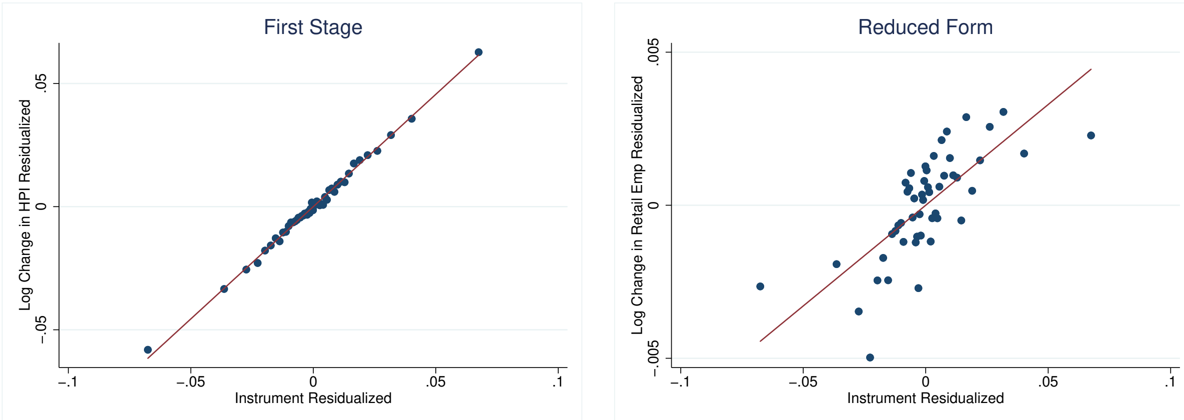


Figure 4: Sensitivity Instrument Pooled First Stage and Reduced Form Binned Scatter Plots

Note: 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 $\gamma_i$'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).

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, Rao, and Sufi's (2013) estimate using the Saiz instrument implies an elasticity of total consumer expenditures of between 0.13 and 0.26. They also estimate 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, Rao, and Sufi, (2013) and Mian and Sufi (2014)

[21] Mian, Rao, and Sufi report estimates of the elasticity of total consumer expenditures to housing net worth in the range 0.5-0.8. To convert Mian, Rao, and Sufi'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, Mitman, and Violante (2017) 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 percent. 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 UK to find an elasticity of 0.2 to 0.3.

Table 2: Comparison of Estimation Approaches for 2006-2009

| Specification | 2006-2009 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) |

Note: 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 $\gamma_i$'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 2006-2009 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 5% level, ** at the 1% level, and *** at the 0.1% level.

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, Rao, and Sufi, (2013) and Mian and Sufi (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 to account for long-run differences in growth rates across CBSAs in calculating the housing wealth elasticity. Davidoff (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 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 $\beta$ in equation (1) using the empirical strategies described in section 3. Panel A presents IV estimates with our sensitivity instrument along with OLS estimates, while panel B 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 centered in the mid-to-late 1990s. Appendix D presents results based on a number of alternate specifications, data sets, and methodologies and shows that the time series pattern in Figure 5 is highly robust. The 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

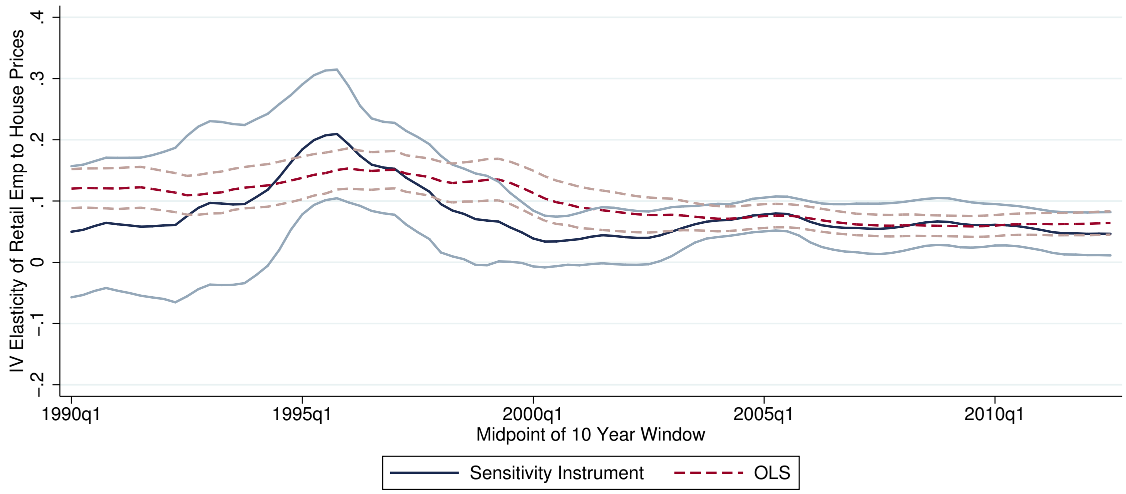


(a) Sensitivity Instrument

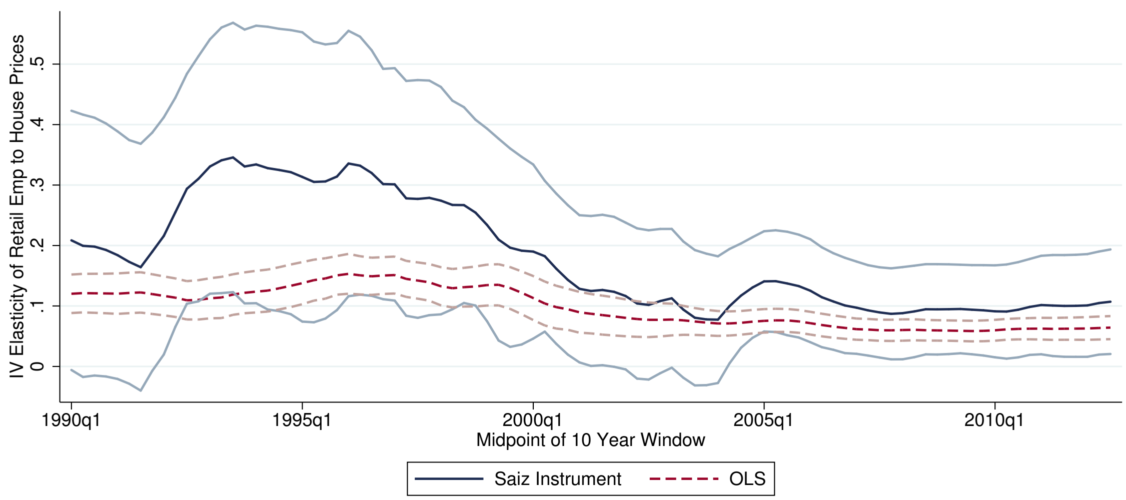


(b) Saiz Instrument

Figure 5: The Elasticity of Retail Employment Per Capita to House Prices Over 10 Year Windows

Note: 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 midpoint 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.

instrument, as well as several other specifications. 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 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's (2014) finding that house prices mainly affect 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 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

[22] Mian and Sufi (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.

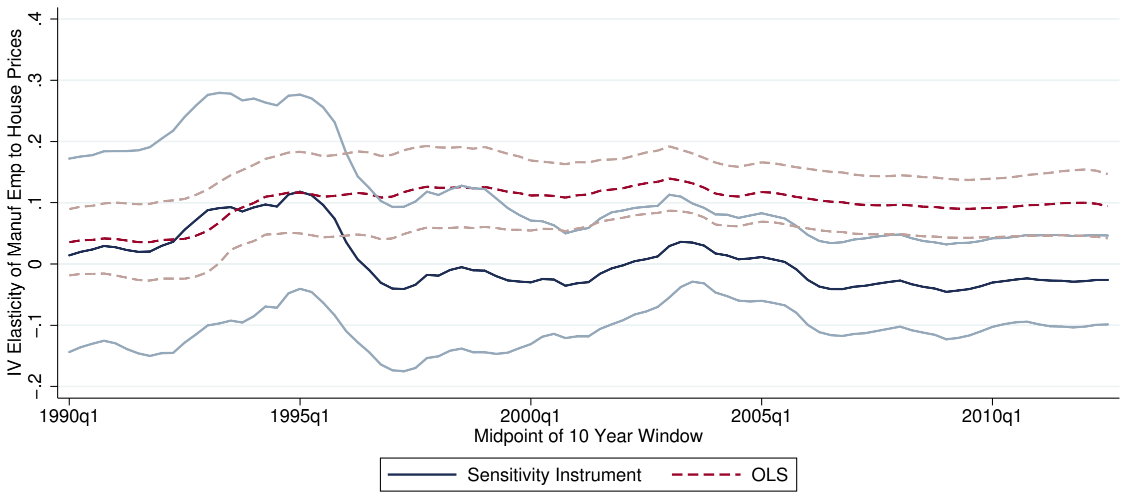


(a) Sensitivity Instrument

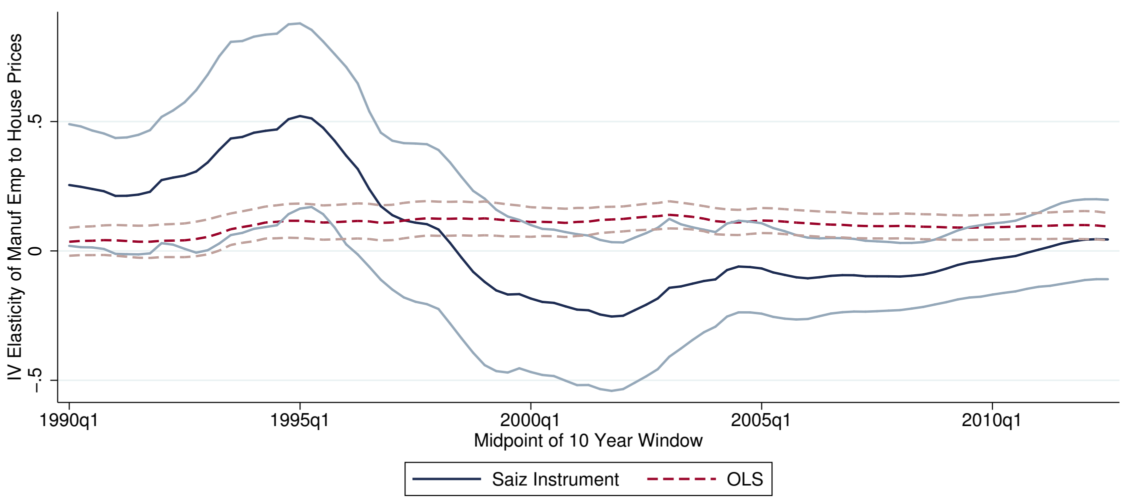


(b) Saiz Instrument

Figure 6: The Elasticity of Manufacturing Employment Per Capita to House Prices Over 10 Year Windows

Note: 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 midpoint 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.

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) | (0.030) | (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) |

Note: 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) (even columns) or an indicator for boom or bust (even columns). "OLS" uses no instrument. "Sensitivity Instrument" uses our sensitivity instrument with the $\gamma_i$s 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.

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 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

[23]The CoreLogic national home price index peaked in 2006q2 and troughed in 2012q1 to 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) |

Note: Each column estimates a version of equation (1) for 1990-20017. The odd-numbered columns replace $\Delta p$ with $\Delta p \times 1[\Delta p \geq 0]$ and $\Delta p \times 1[\Delta p < 0]$ 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 Instrument" uses our sensitivity instrument with the $\gamma_i$s 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 Gilchirst-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. 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.

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

[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. (2017) 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.

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 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 $\phi_x$ 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 labeled 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 $\beta$ 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 $\beta$ 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 $\beta$.

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 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 $\beta$ that results from our empirical specification can be expressed as:

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

where $\beta_{LFM}$ denotes the local fiscal multiplier and $\beta_{PE}$ 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,

[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 $\beta$ 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 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.

[26]We make certain simplifying assumptions to derive this result. One of these is to assume GHH preferences to avoid wealth effects on labor supply. We abstract from the collateral channel emphasized by Chaney, Sraer, and Thesmar (2012) and Adelino, Schoar, and Severino (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.

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 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:

$$\begin{aligned}u(c,h) &= \frac{1}{1-\gamma}\left(c^{(\varepsilon-1)/\varepsilon} + \omega h^{(\varepsilon-1)/\varepsilon}\right)^{(1-\gamma)\varepsilon/(\varepsilon-1)}\\ B(w) &= \frac{B_0}{1-\gamma}\left(w+B_1\right)^{(1-\gamma)}.\end{aligned}$$

Here $\gamma$ captures the curvature of the utility function, $\varepsilon$ is the elasticity of substitution between housing and non-durable consumption, $\omega$ is the taste for housing relative to non-housing 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 $\delta p$. This implies that the rent-price ratio is fixed and given by the parameter $\delta$. We consider alternative assumptions about the behavior of rents in 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 behavior associated with

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

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 $\psi^{\text{Sell}}$ of the value of the house in a transaction cost and to buy a house the individual must pay a fraction $\psi^{\text{Buy}}$.

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 $\theta$ 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 $\ell$ is a deterministic life-cycle component, $z$ is a persistent shock that follows an AR(1) process, and $\xi$ 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, $\gamma$, to 2. We set the elasticity of substitution between housing and non-durable consumption to 1.25 based on the estimates of Piazzesi, Schneider, and Tuzel (2007). We set the LTV limit, $\theta$, 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 percent 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 percent based on the difference between the long-run averages

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

of the 1-year Treasury rate and inflation. We set the cost of buying a house to 2 percent. 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 labor income is modeled 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 appendix provides further explanation of the income process and the parameter values. At retirement, a household faces no further labor 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 Smith (2014).

We set the remaining parameters by matching moments from the 2001 SCF. These parameters are the discount factor, $\beta$; the strength of the preference for housing, $\omega$; the strength of the bequest motive, $B_0$; the degree to which a bequest is a luxury, $B_1$; the rent-price ratio, $\delta$; the mortgage origination cost, $\psi^m$; and the transaction cost for selling a house, $\psi^{\text{Sell}}$. 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 percent 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 percent annual moving rate for owner occupiers based on March

Table 5: Parameter Values Set by Matching Moments from the SCF

| $\beta$ | $\omega$ | $B_0$ | $B_1$ | $\delta$ | $\psi^m$ | $\psi^{\text{Sell}}$ |
|---|---|---|---|---|---|---|
| 0.939 | 0.0795 | 85.0 | 1.75 | 0.0435 | 0.0203 | 0.110 |

2001 CPS data. Overall, 6.3 percent 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 percent 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. 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 $\Phi$ 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 $\Phi$. Over our sample period, household balance sheets changed substantially as households became much more leveraged.

[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. 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 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.

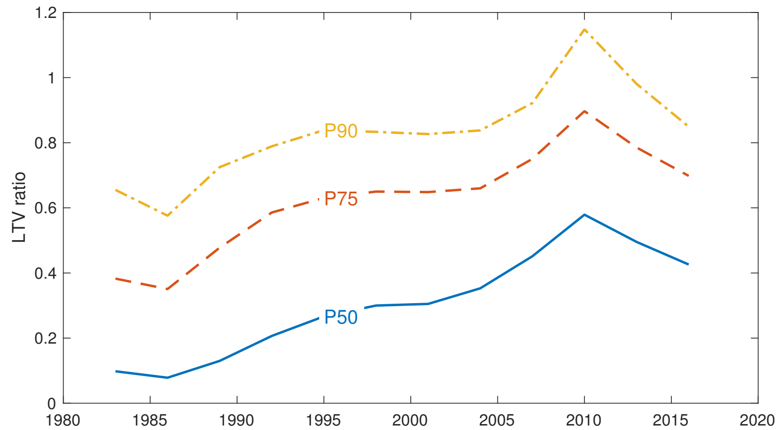


Figure 7: Percentiles of the LTV Distribution

Note: Data are from the SCF for 1983-2004 and CoreLogic for 2007-2016. The percentiles refer to all homeowners including those without a mortgage. For 2007-2016, 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.

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, panel A of Figure 8 reports the model-implied housing wealth elasticity using the distribution $\Phi$ 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

[31] Appendix B.3 describes how we estimate $\Phi$ 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.

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 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, panel B of Figure 8 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 $F_t^{-1}(F_{2001}(LTV_{2001}))$ 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. Panel B of Figure 8 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 Panel B of Figure 8 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 Panel B of Figure 8 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 sheet, a given percentage change in the value of housing becomes more important to consumption decisions, and the model implied housing wealth elasticity rises.

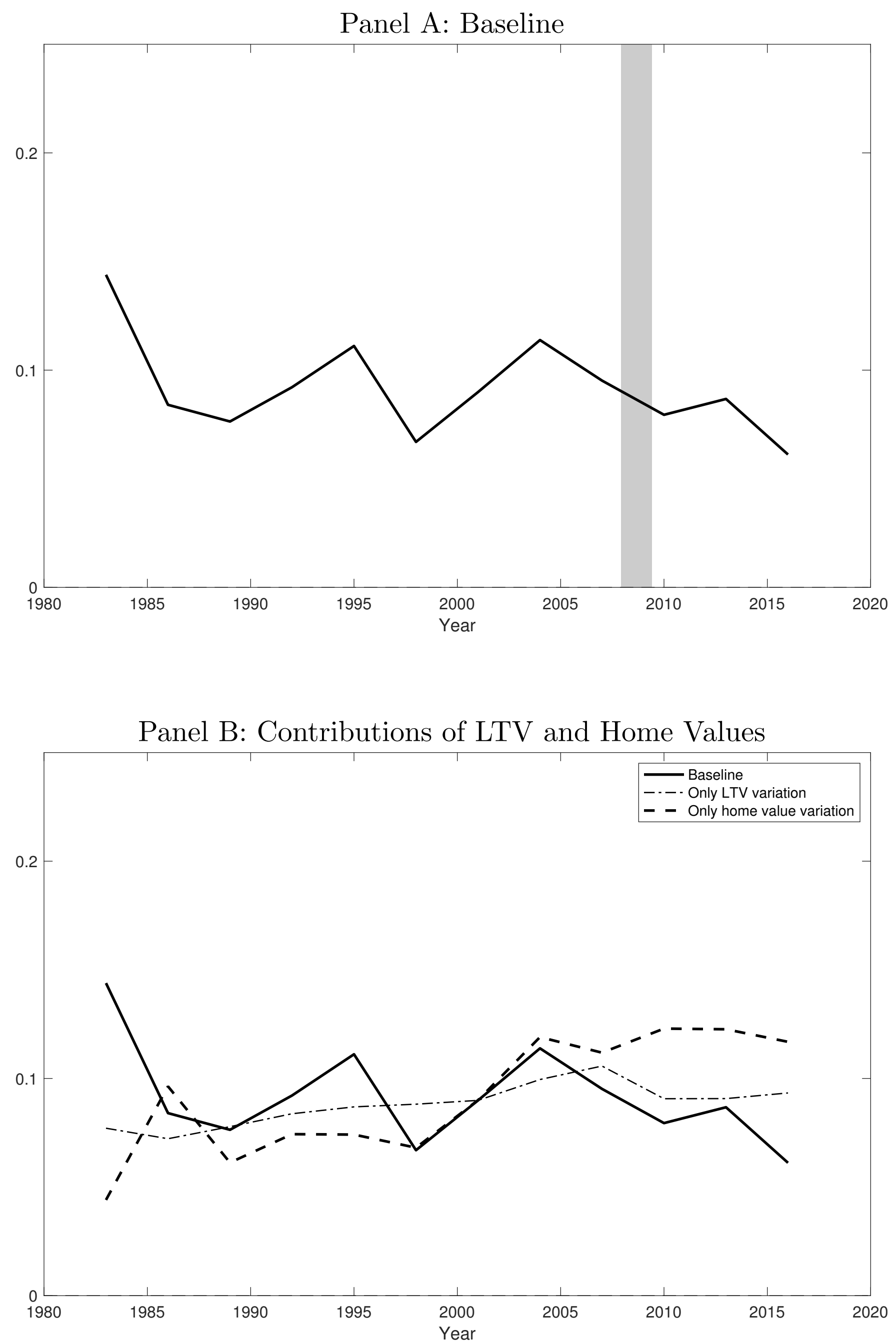


Figure 8: Model Housing Wealth Elasticity

Note: Panel (a) shows $(p/C)(dC/dp)$, 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 $\Phi_t$ is constructed from the SCF data for year $t$ adjusted to match the CoreLogic Homeowner Equity Reports as explained in Appendix B.3. We use a finite difference derivative that averages the values of plus and minus 10% price changes. Panel (b) repeats the same calculation with counterfactual $\Phi_t$'s constructed as described in the text.

### 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 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

[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.

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

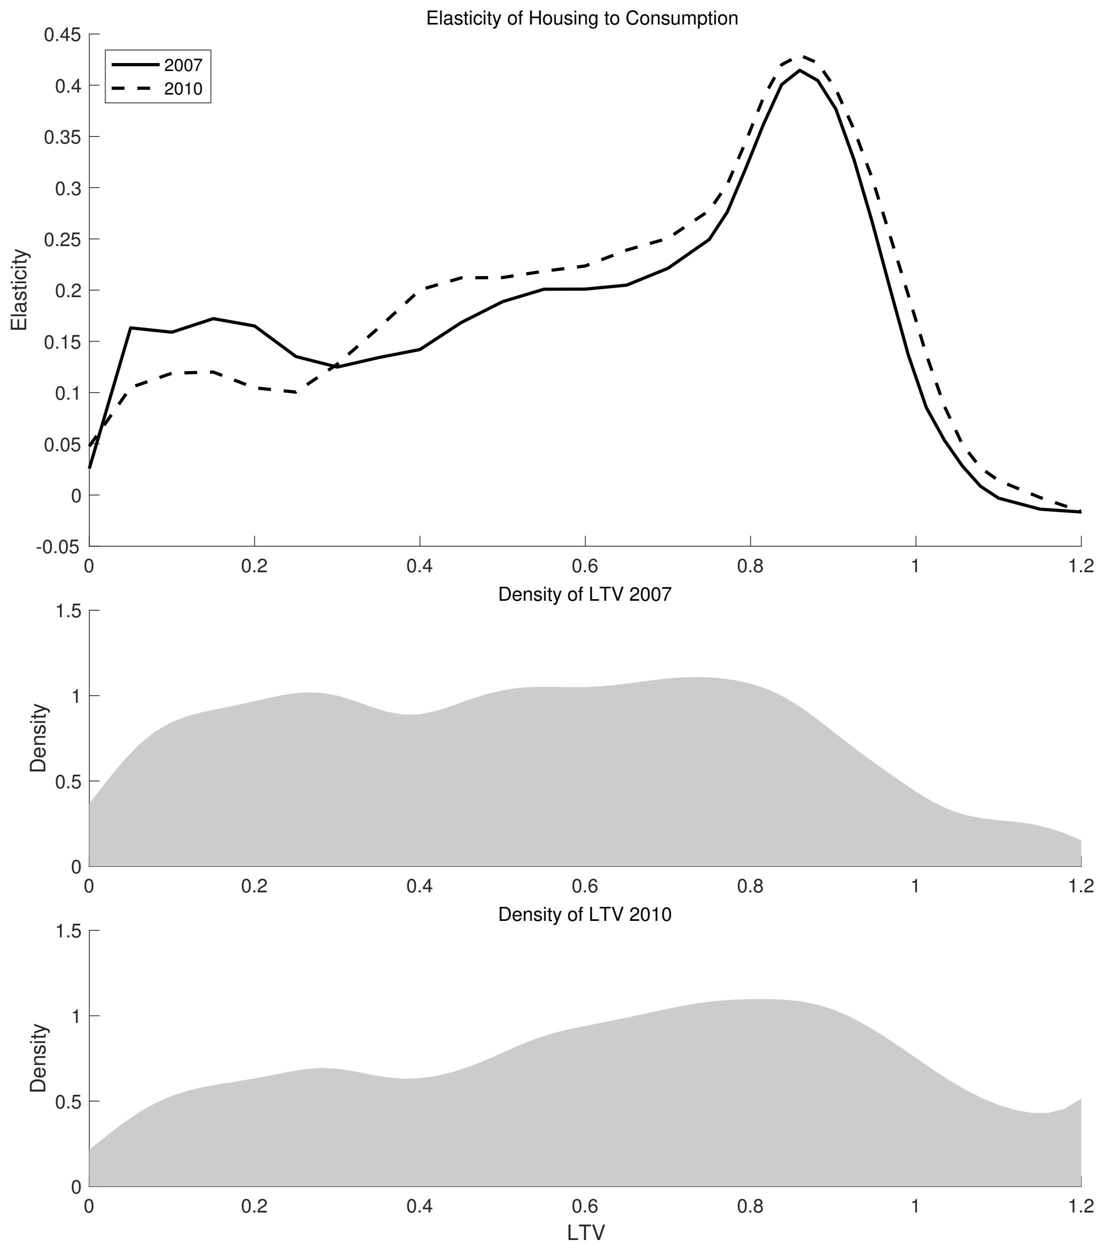


Figure 9: Housing Wealth Elasticity by LTV and Marginal Distributions of LTV in 2007 and 2010.

Note: 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.

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 $\leq 0.6$ | 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 $\geq 1.0$ | 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 |

Note: 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.

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 (2019) 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 ($\leq 0.6$) stayers category fell by 23 percent. These people were mostly pushed into the high-LTV $(0.8, 1.0]$ and underwater ($\geq 1.0$) categories. The high-LTV group grew by 9 percent of the population, while the underwater group grew by 11 percent of the population. The comparison between 2007 and 2010 is even more stark. The number of low-LTV households decreased by 7 percent and the number of underwater households increased by 7 percent, 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,

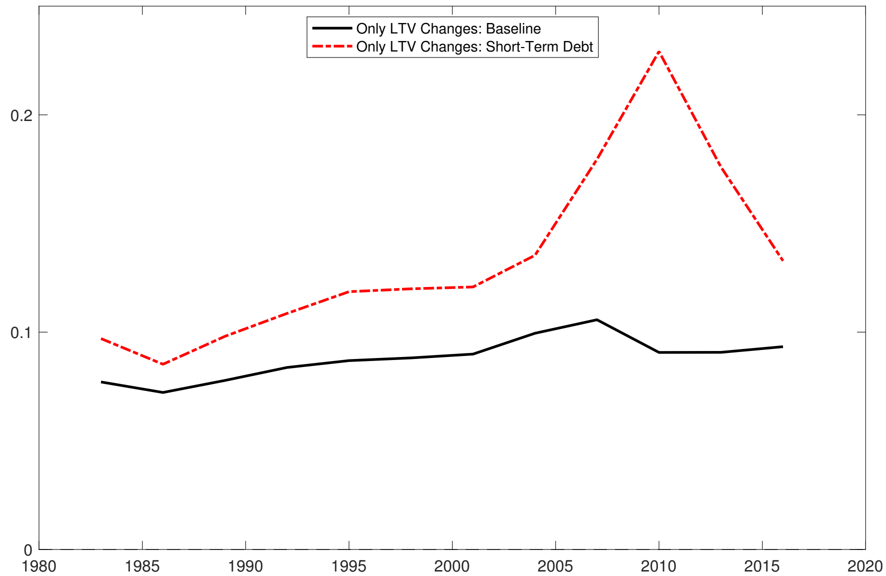


Figure 10: Short-Term Debt Counterfactual

Note: The figure shows the changes in the housing wealth elasticity that result from changes in the LTV distribution as in Panel (b) of 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.

in contrast to both our empirical analysis and our model with long-term debt. 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 each year. This implies that the difference between the housing wealth elasticity in the baseline and the boom

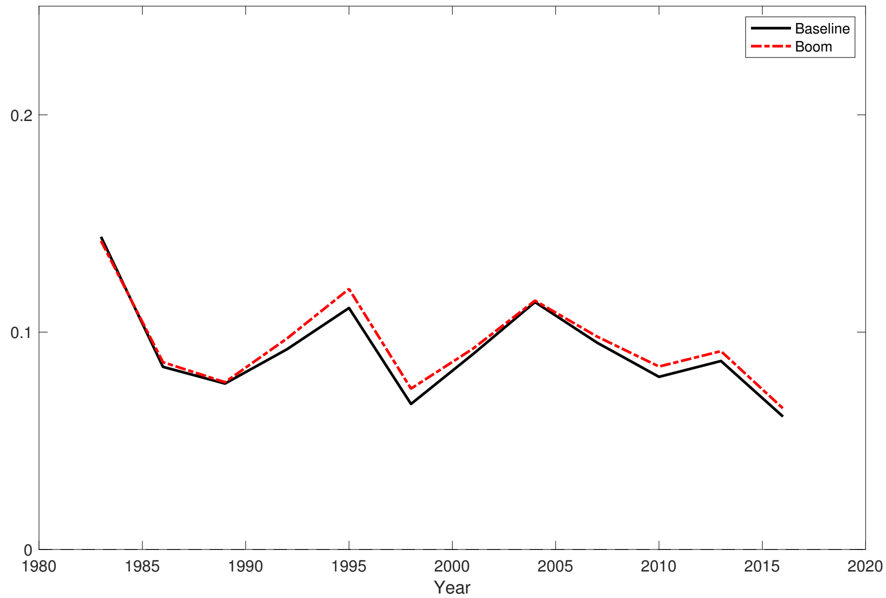


Figure 11: The Housing Wealth Elasticity Under Boom Parameters

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

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, Piazzesi, and Schneider (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

[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.

in the 2000s boom-bust cycle. It is worth noting, however, that our model does not provide an 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 off – 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 Appendix compares the model-implied housing wealth elasticity for this alternate experiment to our baseline model-implied housing wealth elasticity. The demand for

[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, Mitman, and Violante (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.

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. 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 modest anticipated capital gains during the boom years is able to match the behavior 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 percent.

We explore how our results are affected by several other modeling choices in the 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 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 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 Figure A.28).

## 8 Conclusion

In this paper, 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.

# References


Acemoglu, D., D. Autor, D. Dorn, G. H. Hanson, and B. Price (2016). Import Competition and the Great U.S. Employment Sag of the 2000s. *Journal of Labor Economics 34*(S1), S141–S198.

Adelino, M., A. Schoar, and F. Severino (2015). House Prices, Collateral, and Self-Employment. *Journal of Financial Economics 117*(2), 288–306.

Agarwal, S., G. Amromin, S. Chomsisengphet, T. Landvoigt, T. Piskorski, A. Seru, and V. Yao (2017). Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinancing Program.

Aladangady, A. (2017). Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata. *American Economic Review 107*(11), 3415–3446.

Asdrubaldi, P., B. E. Sorensen, and O. Yosha (1996). Channels of Interstate Risk Sharing: United States 1963-1990. *The Quarterly Journal of Economics 111*(4), 1081–1110.

Attanasio, O., L. Blow, R. Hamilton, and A. Leicester (2009). Booms and Busts: Consumption, House Prices, and Expectations. *Economica 76*(301), 20–50.

Attanasio, O., A. Leicester, and M. Wakefield (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*(3), 399–435.

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2018). The Economic Effects of Social Networks: Evidence From the Housing Market. *Journal of Political Economy 126*(6), 2224–2276.

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2018). House Prices and Consumer Spending. *The Review of Economic Studies 85*(3), 1502–1542.

Bhutta, N. and B. J. Keys (2016). Interest Rates and Equity Extraction During the Housing Boom. *American Economic Review 106*(7), 1742–1774.

Calomiris, C. W., S. D. Longhofer, and W. Miles (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 J. F. Cocco (2003). Household Risk Management and Optimal Mortgage Choice. *Quarterly Journal of Economics 118*(4), 1449–1494.

Campbell, J. Y. and J. F. Cocco (2007). How Do House Prices Affect Consumption? *Journal of Monetary Economics 54*(3), 591–621.

Carroll, C. D., M. Otsuka, and J. Salacalek (2011). How Large Are Housing and Financial Wealth Effects? A New Approach. *Journal of Money, Credit and Banking 43*(1), 55–79.

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

Case, K. E., R. J. Shiller, and J. M. Quigley (2013). Wealth Effects Revisited, 1975-2012. *Critical Finance Review 2*(1), 101–128.

Chaney, T., D. Sraer, and D. Thesmar (2012). The Collateral Channel: How Real Estate Shocks Affect Corporate Investment. *American Economic Review 102*(6), 2381–2409.

Chodorow-Reich, G. and J. Wieland (2018). Secular Labor Reallocation and Business Cycles.

Cloyne, J., K. Huber, E. Ilzetzki, and H. Kleven (2019). The Effect of House Prices on Household Borrowing: A New Approach. *American Economic Review 109*(6), 2104–2136.

Cooper, D. (2013). House Price Fluctuations: The Role of Housing Wealth as Borrowing Collateral. *The Review of Economics and Statistics 95*(4), 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*(2), 177–206.

Davis, M. and S. Van Nieuwerburgh (2015). Housing, Finance, and the Macroeconomy. In G. Duranton, J. V. Henderson, and W. C. Strange (Eds.), *Handbook of Regional and Urban Economics, Volume 5*, pp. 753–811. Elsevier.

DeFusco, A. A. (2018). Homeowner Borrowing and Housing Colalteral: New Evidence from Expiring Price Controls. *Journal of Finance 73*(2), 523–573.

Deng, Y., J. M. Quigley, and R. Van Order (2000). Mortgage Terminations, Heterogeneity, and the Exercise of Mortgage Options. *Econometrica 68*(2), 275–307.

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

Favilukis, J., S. C. Ludvigson, and S. Van Nieuwerburgh (2017). The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium. *Journal of Political Economy 125*(1), 140–223.

Flavin, M. and S. Nakagawa (2008). A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence. *American Economic Review 98*(1), 474–495.

Foote, C. L., K. S. Gerardi, and P. S. Willen (2012). Why Did So Many People Make So Many Ex Post Bad Decisions? The Causes of the Foreclosure Crisis. In A. S. Blinder, A. W. Lo, and R. M. Solow (Eds.), *Rethinking the Financial Crisis*. Russel Sage Foundation.

Ganong, P. and P. Noel (2019). The Effect of Debt on Default and Consumption: Evidence from Housing Policy in the Great Recession.

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

Gilchrist, S. and E. Zakrajsek (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review 102*(4), 1692–1720.

Glaeser, E. L. and J. Gyourko (2005). Urban Decline and Durable Housing. *Journal of Political Economy 113*(2), 345–375.

Gorea, D. and V. Midrigan (2018). Liquidity Constraints in the U.S. Housing Market.

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

Guerrieri, V. and G. Lorenzoni (2017). Credit Crises, Precautionary Savings, and the Liquidity Trap. *The Quarterly Journal of Economics 132*(3), 1427–1467.

Guner, N., R. Kaygusuz, and G. Ventura (2014). Income Taxation of U.S. Households: Facts and Parametric Estimates. *Review of Economic Dynamics 17*(4), 559–581.

Guren, A. M., A. Krishnamurthy, and T. McQuade (2019). Mortgage Design in an Equilibrium Model of the Housing Market.

Guren, A. M., A. McKay, E. Nakamura, and J. Steinsson (2019). What Can We Learn From Cross-Sectional Empirical Estimates in Macroeconomics?

Guvenen, F. and J. Anthony A. Smith (2014). Inferring Labor Income Risk and Partial Insurance From Economic Choices. *Econometrica 82*(6), 2085–2129.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2019). What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?

Heathcote, J., F. Perri, and G. L. Violante (2010). Unequal We Stand: An Empirical Analysis of EconomicInequality in the United States, 1967-2006. *Review of Economic Dynamics 13*(1), 15–51.

Hurst, E. and F. P. Stafford (2004). Home Is Where the Equity Is: Mortgage Refinancing and Household Consumption. *Journal of Money, Credit, and Banking 36*(6), 985–1014.

Kaplan, G., K. Mitman, and G. L. Violante (2017). Non-Durable Consumption and Housing Net Worth in the Great Recession: Evidence From Easily Accessible Data.

Kaplan, G., K. Mitman, and G. L. Violante (2019). The Housing Boom and Bust: Model Meets Evidence.

Kaplan, G., G. L. Violante, and J. Weidner (2014). The Wealthy Hand-To-Mouth. *Brookings Papers on Economic Activity* (Spring), 77–138.

Kuhn, M., M. Schularick, and U. I. Steins (2017). The Great American Debt Boom, 1948-2013.

Landvoigt, T., M. Piazzesi, and M. Schneider (2015). The Housing Market(s) of San Diego. *American Economic Review 105*(4), 1371–1407.

Li, W. and R. Yao (2007). The Life-Cycle Effects of House Price Changes. *Journal of Money, Credit and Banking 39*(6), 1375–1409.

Liebersohn, J. (2017). Housing Demand, Regional House Prices, and Consumption.

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

Mian, A., K. Rao, and A. Sufi (2013). Household Balance Sheets, Consumption, and the Ecoomic Slump. *Quarterly Journal of Economics 128*, 1687–1726.

Mian, A. and A. Sufi (2011). House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis. *American Economic Review 101*(5), 2132–2156.

Mian, A. and A. Sufi (2014). What Explains the 2007-2009 Drop in Employment? *Econometrica 82*(6), 2197–2223.

Nakamura, B. E. and J. Steinsson (2014). Fiscal Stimulus in a Monetary Union: Evidence from US Regions. *American Economic Review 104*(3), 753–792.

Nathanson, C. G. and E. Zwick (2018). Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market. *Journal of Finance 73*(6), 2587–2633.

Palmer, C. (2015). Why Did So Many Subprime Borrowers Default During the Crisis : Loose Credit or Plummeting Prices ?

Piazzesi, M., M. Schneider, and S. Tuzel (2007). Housing, Consumption, and Asset Pricing. *Journal of Financial Economics 83*, 531–596.

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

Roussanov, N., M. Michaux, and H. Chen (2018). Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty.

Saiz, A. (2010). The Geographic Determinants of Housing Supply Elasticity. *Quarterly Journal of Economics 125*(3), 1253–1296.

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

Wong, A. (2019). Refinancing and the Transmission of Monetary Policy to Consumption.

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

# Appendices

# A Data

## A.1 Data Construction Details

Our main data sources are the Quarterly Census of Employment and Wages (QCEW) from 1975 to 2017, the County Business Patterns (CBP) from 1975 to 2016, the 1970-2010 post-Censal county population estimates, the 2010-2017 inter-Censal county population estimates, and the Freddie Mac House Price Indices. We also merge in data from a number of other data sources for controls. Throughout, we use consistent CBSA definitions and assign each CBSA to a census region based on the majority of its population in the 2000 Census.

The QCEW data provides county-level data for each SIC industry from 1975-2000 and for each NAICS industry from 1990-2017 and is publicly available on the BLS website. The QCEW is sometimes missing data due to BLS disclosure requirements, as described by the BLS at: https://www.bls.gov/osmr/pdf/st040100.pdf. This missing data problem primarily affects smaller and more narrowly-defined industries in smaller counties where there are few enough employers within an industry that the BLS's disclosure criteria are not met. In our baseline analysis, we only use data for counties within a CBSA that have no missing data for the industry in question (either retail or manufacturing) over the whole sample. In practice, this means that we drop a few small outlying counties for a few CBSAs along with the entire Dover, Delaware CBSA. We show in Appendix D.1.3 that dropping these counties for the entire sample does not affect the results because they are so small.

The QCEW reports monthly employment and aggregate wages by industry. We take quarterly averages and use employment for retail or manufacturing employment. We also use the total wage bill for all employees. We merge in county-level population as estimated annually by the Census, linearly interpolating to quarters. We then drop counties with missing data at any point in the sample, aggregate both employment in each industry and population to the CBSA level, and calculate log changes in employment per capita. As mentioned in the text, the QCEW begins in 1975 but there are data issues the first two years. The log changes we use thus begin in 1978 with the difference between 1978 and 1977 log employment per capita. We create average wages by dividing total wages by total employment. We then clean the data by removing observations where we observe an unusually-large jump in employment or employment per capita over a single quarter. We do so because this likely reflects changes such as county realignments or a large employer being recategorized across industries rather than an actual changes in employment. Appendix D.1.3 shows our results are not sensitive to this data cleaning.

The shift from SIC to NAICS changed the definitions of retail, manufacturing, construction, and real estate employment. For instance, wholesale employment was included as part of retail under SIC but separated into its own sector under NAICS. This causes discrete jumps in sectoral employment. However, for the 1990-2000 period where the BLS provides both SIC and NAICS data, the two series are almost the same in log changes. We thus splice together the SIC and NAICS data in Q1 1993. We choose Q1 1993 because this is the first date for which we can splice together one year and three year log changes. In Appendix D.1.3, we show our results are robust to splice date

We also use a number of additional variables in constructing controls. We merge in 2-digit industry shares from the CBP. We use the CBP rather than the QCEW for this because whereas the QCEW simply omits data when the BLS cannot disclose data, the Census provides employment ranges for industries that do not satisfy this, which include some 2-digit industries. Because the CBP data do not provide an overlapping period for SIC and NAICS, we deal with the SIC to NAICS transition by harmonizing all of the data to consistent 2-digit industry classifications using an algorithm developed by Acemoglu et al. (2016). We then aggregate to the CBSA level and create 2-digit industry shares. Because the CBP data are only available through 2016, we assume that industry shares are constant from 2016 to 2017. Our results are not dependent on this one year. Since the CBP data are annual, we linearly interpolate to quarters to get a 2-digit industry share series by quarter. In a robustness check, we also use the annual CBP data rather than the quarterly QCEW data for the analysis.

We then merge our data set with the Freddie Mac House Price Index. We convert it to a real index using the GDP deflator downloaded from FRED. In a robustness check, we use a proprietary house price index from CoreLogic which uses only transactions but has less time coverage.

We then merge in the Gilchrist and Zakrajsek (2012) excess bond premium. GZ give two different time series on their website. We use the "ebp_oa" measure. This is the excess bond premium which subtracts a fitted value for "distance to default" from options. We also merge in

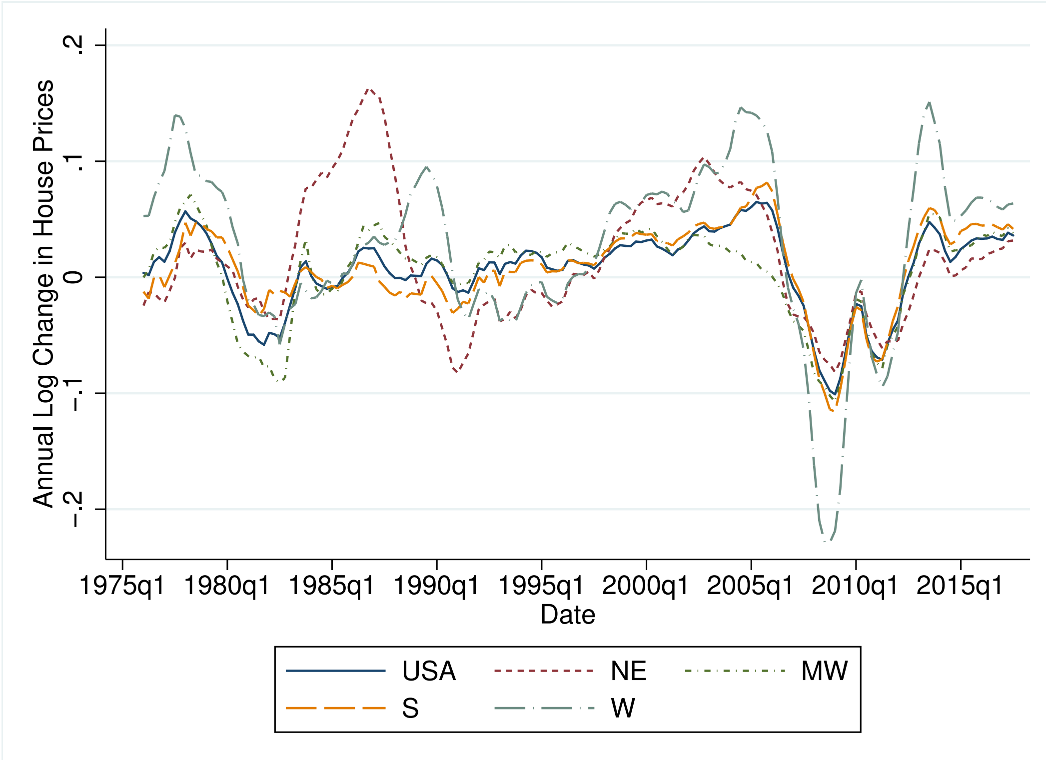


Figure A.1: Regional House Price Indexes

Note: This figure presents regional house price indexes for each of the Census regions we use in our analysis.

the real 30-year mortgage rate, which we create by taking the average 30-year fixed mortgage rate from FRED and adjusting for inflation using 1-year inflation measured by the BEA's GDP deflator.

We create the regional log changes in employment per capita and house prices, and subtract off the log change in the GDP deflator to get real house prices. We do so by using the average log change in employment per capita or house prices weighting by 2000 population and leaving out each individual CBSA from the calculation of the aggregate.

Finally, we merge in the Saiz housing supply elasticities. We do so by matching the central city with Saiz, who uses older MSA definitions. In some cases, two CBSAs are assigned to the same Saiz MSA. All results are robust to dropping the second match.

## A.2 Regional Home Price Indexes

Figure A.1 shows time series plots of the annual log change in housing prices for the United States as a whole and each of the four Census regions we use as the regional aggregates for our empirical approach. One can see that prior to the 2000s boom and bust, the national house price index exhibited relative small fluctuations. However, there were regional house price cycles. In particular, there was a small bust in the Midwest in the early 1980s, a boom and busts in the Northeast from the mid-1980s to the mid-1990s, and a boom and bust in the West in the early 1990s. We use this variation to help us identify the housing wealth effect prior to 2000.

## A.3 Cross-City Evidence on Retail Employment vs. Consumption

In this section, we provide additional evidence on the relationship between retail employment and consumer expenditures beyond the aggregate evidence presented in the main text by studying the relationship between city-level consumption and retail employment using data for 17 cities for which the BLS publishes city-level average consumption per person using data from the Consumer Expenditure Survey (CEX). Since the aggregated city-level CEX data are available in two-year averages back to 1986 (e.g., 1986-1987, 1987-1988, and so on), so we construct both the left-hand side and right-hand side variables as 2-year log differences of the 2-year averaged data.

We estimate:

$$\Delta \log \bar{C}_{i,t} = \xi_i + \zeta_t + \beta \Delta \log \bar{Y}_{i,t} + \varepsilon_{i,t}, \tag{7}$$

where $\Delta \log \bar{C}_{i,t}$ is the 2-year log change of 2-year averaged consumption, $\Delta \log \bar{Y}_{i,t}$ is the 2-year log change of 2-year averaged retail employment, and $\xi_i$ and $\zeta_t$ are city and time fixed effects respectively. We estimate the equation over the sample period 1986-2014. To construct real consumption, we deflate consumer expenditures from the CEX by the city-level CPI, which is available from the BLS for the cities we consider at an annual frequency.

Table A.1 presents the results of this analysis. We estimate this equation both with total consumption as the dependent variable (columns 1-2) and total consumption excluding imputed rent (columns 3-4). We first present estimates using OLS. The OLS estimates (column 1 and 3) show that a 1% increase in retail employment is associated with 0.45-0.55% in total consumption both including and excluding imputed rent.

However, an important issue in estimating this equation is that the right-hand side variable (retail employment growth) is measured with substantial error. Hence, the OLS estimate is likely to be biased downward due to attenuation bias. To account for this, we also present results based on an IV estimation strategy where we instrument for retail employment using city-level house price growth. Since the consumption and employment growth rates are for 2-year log changes in 2-year averages (i.e., the growth rate between say 1997/1996 and 1994/1995) we make use of 3-year log changes in house prices as the instrument (i.e., the house price growth from 1994 to 1997 in this

Table A.1: City-Level Consumption vs. Retail Employment Regressions

| | 1 | 2 | 3 | 4 |
|---|---|---|---|---|
| | OLS<br>Total Cons | IV<br>Total Cons | OLS<br>Ex Imputed Rent | IV<br>Ex Imputed Rent |
| Retail Emp.<br>Per Capita Growth | 0.454**<br>(0.164) | 0.947**<br>(0.297) | 0.512*<br>(0.210) | 0.967**<br>(0.378) |
| CBSA FE | X | X | X | X |
| Time FE | X | X | X | X |
| N | 423 | 408 | 423 | 408 |

Note: This table shows regressions of the elasticity of growth total consumption and total consumption excluding imputed rent to retail employment growth estimated from equation (7). The analysis uses 17 CBSAS for which the Consumer Expenditure Survey CBSA-level data is available. Because the aggregated city-level CEX data are available in two-year averages back to 1986 (e.g., 1986-1987, 1987-1988 and so on), so we construct both the left-hand side and right-hand side variables as 2-year log differences of the 2-year averaged data. Consumption is deflated by the city-level CPI. The IV specification instruments with 3-year log changes in house price growth as indicated in the text.

example). It is important to recognize that we are not making any assumption here about whether variation in the house price index is endogenous or exogenous to business cycle shocks — only that the measurement error in house price indexes is likely to be orthogonal to the measurement error in employment growth (because the two statistics are calculated from entirely separate samples).

Our IV estimates of the relationship between retail employment and consumer expenditures are presented in in columns 2 and 4. A 1% increase in city-level retail employment is associated with roughly a 1% increase in city-level consumption both including and excluding imputed rent. Both the CEX and retail employment presumably provide a noisy measure of true consumer expenditures. Hence, it is not surprising that the IV estimates are higher than their OLS counterparts (as we would expect in the presence of attenuation bias). However, the IV analysis suggests that once we account for measurement error, retail employment per capita varies roughly one-for-one with real consumption. This is consistent with the aggregate evidence we present in the text on the time series behavior of aggregate consumption measures and retail employment.

The IV coefficient presented in the table above gives the ratio of the covariances of retail employment with house prices, and consumer expenditures with house prices, and hence provides insight into the relative responsiveness of consumer expenditures relative to retail employment associated with house price shocks. The comovement between retail employment and consumer expenditures may differ in response to other sources of variation, aside from house price shocks. However, since our analysis focuses on how retail employment responds to changes in house prices, we believe that the IV estimates we report above are the "right" kind of variation to focus on for

Table A.2: Data Targets and Model Moments for Income Risk

| Moment | Data | Model |
|---|---|---|
| St. dev. of 1-year log earnings growth | 0.51 | 0.63 |
| St. dev. of 5-year log earnings growth | 0.78 | 0.76 |
| Growth of cross-sectional variance of log earnings over the life-cycle | 0.66 | 0.66 |
| Fraction of 1-year log earnings growth in [-1.0,1.0] | 0.94 | 0.92 |
| Fraction of 1-year log earnings growth in [-0.1,0.1] | 0.53 | 0.52 |
| Fraction of 5-year log earnings growth in [-1.0,1.0] | 0.90 | 0.89 |
| Fraction of 5-year log earnings growth in [-0.1,0.1] | 0.27 | 0.27 |

the purpose of our analysis.

# B Calibration and Numerical Methods

## B.1 Model Income Process

We use an income process that captures salient features of the earnings dynamics reported in Guvenen et al. (2016), hereafter GKOS. Specifically we model log annual income as $\log y = \ell + z + \xi$, where $\ell$ is a deterministic life-cycle component, $z$ is a persistent shock that follows an AR(1) process, and $\xi$ is a transitory shock. The deterministic life-cycle is from Figure 3 of the July 2015 version of GKOS. GKOS model the transitory income shocks as a "non-employment shock" and we mimic this specification. The transitory component $\xi$ is zero with some probability and is equal to $\log(1 - x)$ with complementary probability, where $x$ is drawn from an exponential distribution on the interval $(0, 1)$. We use data from the 2002 March CPS on hours worked in the prior year to estimate the parameter of the exponential distribution to be 2.25 and the probability that $\xi$ is zero to be 0.75 using maximum likelihood. We fix the persistence of the AR(1) component $z$ to 0.97 because a near unit-root persistence is needed to match the near linear growth of the cross-sectional earnings variance over the life-cycle. The innovations to $z$ are drawn from a mixture of two normals, which allows us to capture the leptokurtic nature of earnings growth rates reported by GKOS. We fix the mean of the mixture components to zero and estimate the mixture probability and the standard deviations using a simulated method of moments procedure. Table A.2 lists the target moments and the model-implied values. All empirical moments are taken from GKOS. The resulting parameters of the innovations to $z$ are a first component with a mixture weight of 0.984 and a standard deviation of 0.071 and a second mixture component with a weight of 0.016 and a standard deviation of 1.60.

The data used by GKOS is on earnings before taxes. We use the "log" tax function estimated by Guner et al. (2014) for all households to approximate the US tax system including state and local taxes, which states that the average tax rate is $0.135 + 0.062(Y/\bar{Y})$.

### B.2 Model Calibration

We next describe the procedure we use to calibrate the discount factor, $\beta$; the strength of the preference for housing, $\omega$; the strength of the bequest motive, $B_0$; the degree to which a bequest is a luxury, $B_1$; the rent-price ratio, $\delta$; the mortgage origination cost, $\psi^M$; and the transaction cost for selling a house, $\psi^{\text{Sell}}$. The method we use is to minimize a quadratic objective function. We begin by describing our empirical targets, then discuss the objective function, and conclude with an assessment of the model's fit.

The broad overview of our empirical targets is provided in the main text and here we provide some additional information. Starting with the 2001 SCF, we first compute home-value-to-income for households with heads aged 25 to 60. We compute this ratio as the value of real estate held to household income. Next we compute LTV as the sum of all housing debt relative to the value of all real estate. Liquid assets are defined as the sum of liquid accounts ("liq" in the SCF extracts sums checking, savings, and money market accounts), directly held mutual funds, stocks, and bonds less revolving debt. Following Kaplan et al. (2014), liquid account holdings are scaled by 1.05 to reflect cash holdings. We normalize the model and the data such that median income among 40 year-olds is 1.0. To compute life-cycle profiles, we use rolling 5-year windows by age (i.e., moments at age 30 include heads 28 to 32 years old).

We now turn to our refinancing frequency target. Deng et al. (2000) estimate a statistical model of refinancing behavior that controls for the difference between existing mortgage interest rates and the market rate. Their model allows for unobserved heterogeneity with three household types who differ in the propensity to refinance. Their Figure 2 reports the time-varying refinancing rate for each of the groups. Using this information and the relative sizes of the groups estimated by Deng et al., we simulate a population of households and compute the fraction of mortgages refinanced each year, which yields 9.3 percent.

Lastly, we target a 3.2 percent moving rate for owner occupiers based on March 2001 CPS data. Overall, 6.3 percent of owner occupiers reported living in a different house one year earlier. 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. Table A.3 lists the reasons-for-moving responses, whether we included or excluded these moves, and the frequency of these responses.

Table A.3: Reasons for Moving in 2001 March CPS

| Reason for moving | Frequency | Included |
|---|---|---|
| Did not move | 93.63 | |
| Change in marital status | 0.38 | no |
| To establish own household | 0.41 | no |
| Other family reason | 0.57 | no |
| New job or job transfer | 0.63 | no |
| To look for work or lost job | 0.03 | no |
| To be closer to work/easier | 0.14 | no |
| Retired | 0.11 | no |
| Other job-related reason | 0.06 | no |
| Wanted to own home, not rent | 1.9 | yes |
| Wanted new or better house | 1.12 | yes |
| Wanted better neighborhood | 0.23 | no |
| Cheaper housing | 0.11 | yes |
| Other housing reason | 0.44 | no |
| Attend/leave college | 0.04 | no |
| Change of climate | 0.08 | no |
| Health reasons | 0.06 | no |

Our quadratic objective function is constructed as follows. For life-cycle profile targets we average the squared difference between the model and the data over the life-cycle. We normalize the weight of the LTV target to 1.0. The weights for liquid assets, home-to-income, homeownership rate, the aggregate refinancing rate, and the aggregate moving rate are 0.25, 1, 25, 500, and 250, respectively.

Figures A.2 to A.5 show the model's fit to the marginal distributions of liquid assets, LTV, housing values and homeownership across ages. The calibrated model predicts a refinancing rate of 8.6% and a moving rate of 2.9%.

## B.3 Constructing the Distribution of Idiosyncratic States from SCF and CoreLogic Data

Our model has five state variables: cash on hand, mortgage debt, housing position, persistent income, and age. We create analogous variables using each wave of the SCF from 1983 to 2016. We equate cash on hand to liquid assets plus annual wage income, where liquid assets are defined as in Section B.2. Wage income is set to X5702 + X5704 + X5716 + X5718 + X5720 + X5722 in the SCF, which is the sum of income from wages/salaries, sole proprietorships, unemployment insurance and workers' compensation, child support and alimony, welfare assistance, Social Security, or other

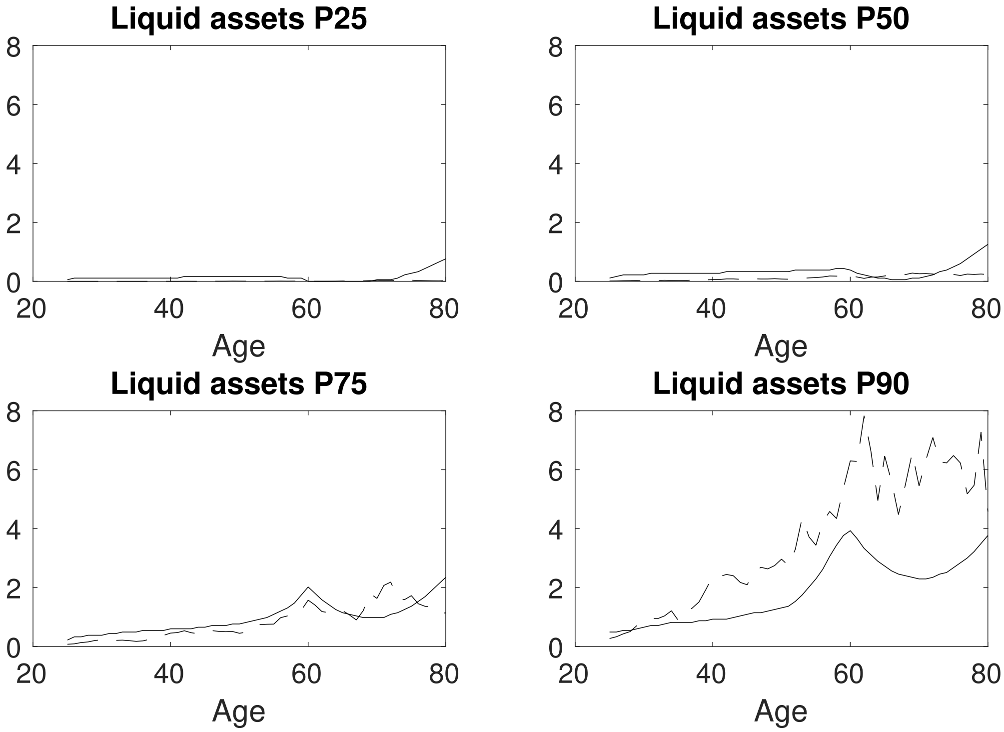


Figure A.2: Liquid Asset Holdings By Age for Model (Solid) and Data (Dashed)

Note: Data refer to SCF 2001.

pensions. We remove taxes from income in the same manner as described in Section B.1. Mortgage debt is set to the sum of all loans backed by housing, which includes home equity lines of credit when this information is available (1989 and onwards). Our model is written in terms of physical units of housing, $h_{i,t}$, that trade at price $p_t$ per unit. We define a unit of housing as a dollar of housing in 2001. We deflate the value of housing based on the evolution of the FHFA national price index relative to the trend of disposable income per capita. We use disposable income from the BEA's Personal Income and Outlays release and divide by the civilian non-institutional population reported by the BLS. We smooth the log of the quarterly series with the HP filter with coefficient 1600 and time aggregate to annual observations. Normalizing by disposable income per capita is a simple way of adjusting for changes in nominal income and rendering the price index roughly stationary. We set the persistent income state based on wage income. Finally, age is simply the age of the household head. We create a product grid on the state space and allocate the mass of the SCF observations to the grid points in a manner that preserves the means of variables.

We are concerned that the SCF may understate the decline in home values during the Great Recession. The Flow of Funds reports a 24% drop in the nominal value of owner occupied real estate between 2007 and 2010. Similarly, the FHFA expanded data house price index reports a 21%

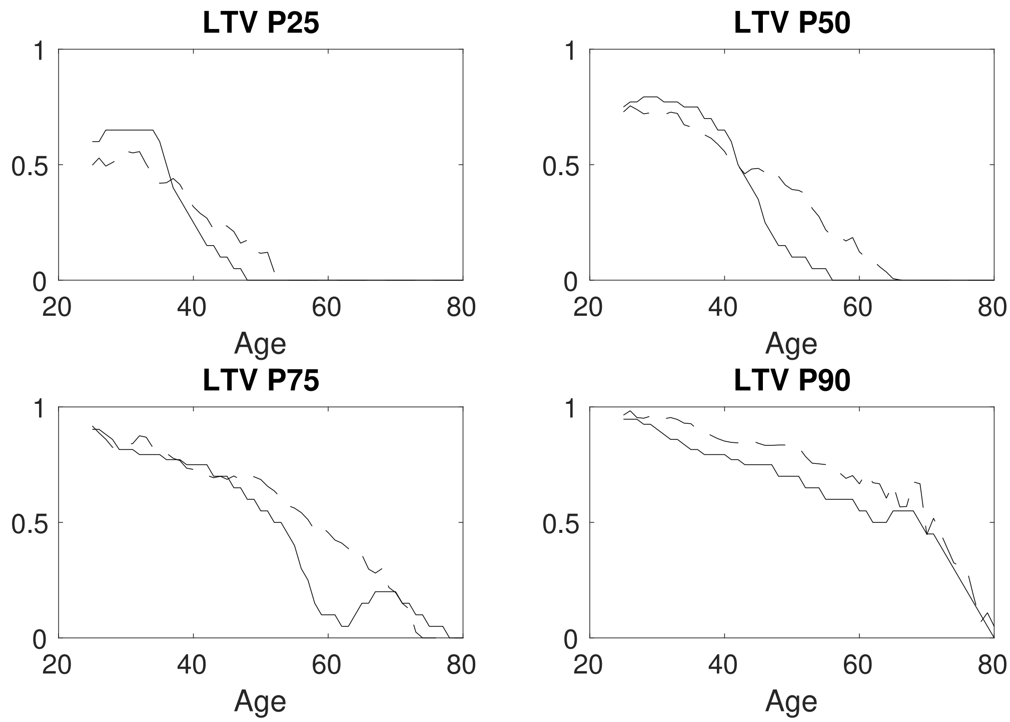


Figure A.3: LTV By Age for Model (Solid) and Data (Dashed)

Note: Data refer to SCF 2001.

drop in house prices, which is similar to other repeat sales house price indices. By contrast, the drop in the SCF is only 14%. It may be that households are slow to recognize (or admit) that the value of their homes has fallen, leading to systematic misreporting in the SCF during the housing bust in the Great Recession.

To address this concern, we use data from CoreLogic's Homeowner Equity Reports to adjust the SCF home values to match CoreLogic's estimated distribution of equity. Since 2007, CoreLogic has reported the CDF of the nationwide LTV distribution at a given set of percentiles. For example, the fraction of households with a mortgage with LTVs less than 50%, 50% to 55%, 55% to 60%, and so on. Starting from this information we construct a marginal distribution of LTVs. To do so we use the SCF to determine the share of all households with a mortgage and the conditional distribution of LTVs within the 0 to 50% bin. We then linearly interpolate the CoreLogic CDF within these 5 percentage point intervals. In making this adjustment, we maintain the order of households in the LTV distribution, but change the values of the LTVs to match the marginal from CoreLogic.

For our calculations, we need to know the mortgage balance and home value separately. We assume that the SCF values for mortgage balances are correct and adjust the home values to match

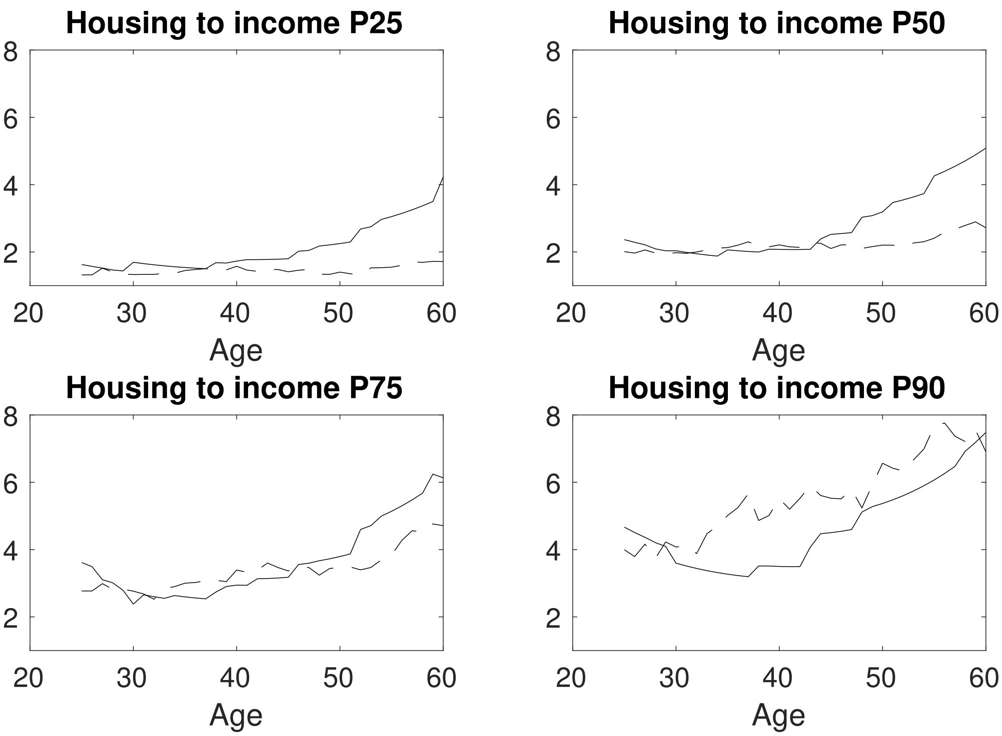


Figure A.4: Housing-to-Income By Age for Model (Solid) and Data (Dashed)

Note: Data refer to SCF 2001. Sample restricted to working-age homeowners.

the LTV calculated in the previous step. This reflects the fact that most households can easily look up their mortgage balance during their SCF interview but cannot easily establish their house's market value. Similarly, we assume that the share of homeowners with a mortgage in the SCF is correct in making our adjustment.

As described in footnote 31, CoreLogic changed its methodology in 2010 to better account for loan amortization and HELOC draw-down. This led to a reduction in their estimated share of homeowners who are underwater on their loan. We use the old methodology for 2007 when the new methodology is not available and the new methodology for 2013 when the old methodology is not available. For 2010, when both are available, we report results with both, although our baseline results use the new methodology.

### B.4 Value Functions and Model Solution

The household's problem can be written as follows. If a household buys a home it solves:

$$V^H(w,z,a,p) = \max_{c,m',h',A'} \left\{ u(c,h') + \beta \mathbb{E}\left[ V(\ell',h',m',z',a+1,p) \right] \right\},$$

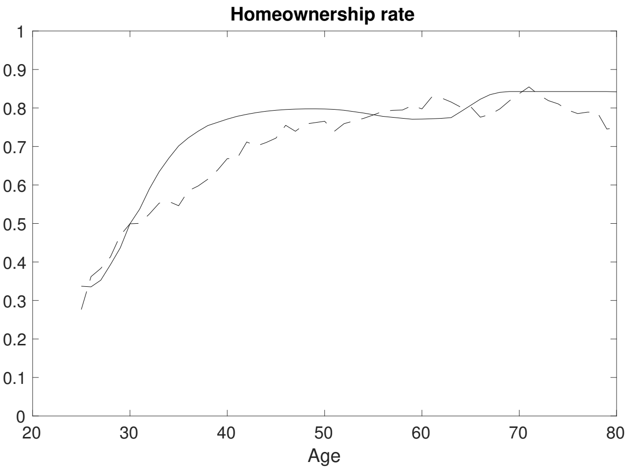


Figure A.5: Homeownership By Age for Model (Solid) and Data (Dashed)

Note: Data refer to SCF 2001.

where $w$ is wealth defined below, $z$ is the persistent income shock, $a$ is age, $c$ is consumption, $h$ is housing, $m$ is mortgage debt, $A$ is liquid savings, $\ell$ is liquid cash on hand. This maximization problem is subject to the LTV constraint:

$$c + A' - \left(1 - \psi^{\text{m}}\right) m' + (1 + \psi^{\text{Buy}})ph' = w,$$

where $w$ is defined as:

$$w = (1 - \psi^{\text{Sell}})ph - R_m m + \ell.$$

In this Appendix we will use the notation $R_m m$ to denote the balance due on the outstanding mortgage including interest, but this is shorthand for a more complicated expression that reflects the fact that the interest rate on the mortgage depends on the LTV at origination. Liquid cash on hand is determined according to:

$$\ell = R_A A + y.$$

Households that previously rented have $h = m = 0$. If a household refinances its mortgage it solves:

$$V^m(f, h, z, a, p) = \max_{c, m', A'} \left\{ u(c, h) + \beta \mathbb{E} \left[ V(\ell', h, m', z', a + 1, p) \right] \right\},$$

subject to the LTV constraint and:

$$c + A' - (1 - \psi^m) m' = f,$$

where $f$ is financial wealth defined as:

$$f = \ell - R_m m.$$

If an owner-occupier household neither refinances nor sells its house it solves:

$$V^0(\ell, h, m, z, a, p) = \max_{c, A'} \left\{ u(c, h) + \beta \mathbb{E} \left[ V(\ell', h, G(a)m, z', a+1, p) \right] \right\},$$

subject to:

$$c + A' - m' = \ell - R_m m.$$

A renter solves:

$$V^R(w, z, a, p) = \max_{c, h', A'} \left\{ u(c, h') + \beta \mathbb{E} \left[ V(\ell', 0, 0, z', a+1, p) \right] \right\}.$$

Entering the next period, the household has a discrete choice over the adjustment costs:

$$\begin{aligned} V(\ell, h, m, z, a) = \max \Big\{ & V^H \left( ((1 - \psi_H^{\text{Sell}})ph - R_m m + \ell, z, a, p \right), \\ & V^m \left( \ell - R_m m, h, z, a, p \right), \\ & V^0(\ell, h, m, z, a, p), \\ & V^R \left( (1 - \psi_H^{\text{sell}} - \delta)ph + \ell - R_m m, z, a, p \right), \\ & V^D \left( z, a, p \right) \Big\}, \end{aligned}$$

with $V^m$ and $V^0$ unavailable to households that previously rented. $V^D(z, a, p) = V^R(\epsilon, z, a, p) - \phi$ is an option to default on the mortgage, which leaves the household as a renter with a small liquid asset position $\epsilon$ and incurs a utility cost $\phi$. Defaults play very little role in our analysis (we set $\phi = 4$ and homeowners are loathe to default), but it is useful to allow this option for homeowners without alternatives.

We solve the household's problem using value function iteration. In solving the model we place a grid on LTV as opposed to mortgage debt. We also specify grids for wealth, financial wealth, liquid wealth, and income. We allow the household to make continuous choices of consumption, liquid savings, and mortgage debt, but we restrict housing to discrete values. The output of each iteration of our Bellman equation is the value on the grid points for $(\ell, h, m, z, a)$. The most obvious way of solving this problem is to solve for the optimal actions for each of the discrete adjustment options for each combination $(\ell, h, m, z, a)$. A more efficient approach makes use of the fact that, for example, all households with a certain level of wealth who buy a house will make the same choice so we can solve the problem on the more compact space of $(W, z, a)$ and then interpolate the value onto $(\ell, h, m, z, a)$. This works well for the value functions but there is a small complication for the decision rules because the housing quantity choice is discrete and so we cannot easily interpolate the decision rules onto $(\ell, h, m, z, a)$. To find the decision rules, we cannot avoid solving the problem for the specific combinations of $(\ell, h, m, z, a)$, but we only need to do this for the households who choose to buy a new house or rent a house and are therefore making a choice over $h'$. For households who refinance, $h'$ is fixed so there is no problem interpolating the decision rules.

# C Empirical Approach

## C.1 Empirical Approach in a Structural Simultaneous Equations Framework

This Appendix explains our identification strategy using a simple simultaneous equations econometric framework. It derives the equation we use to create our instrument structurally and formalizes our identification assumption.

Consider the following empirical model for the determination of retail employment and house prices:

$$\Delta y_{i,r,t} = \psi_i + \xi_{r,t} + \beta \Delta p_{i,r,t} + \alpha_i \mathcal{E}_{r,t} + \varepsilon_{i,r,t}, \tag{8}$$

$$\Delta p_{i,r,t} = \varphi_i + \zeta_{r,t} + \delta \Delta y_{i,r,t} + \gamma_i \mathcal{V}_{r,t} + \nu_{i,r,t}. \tag{9}$$

We allow for CBSA fixed effects ($\psi_i$ and $\varphi_i$) and region-time fixed effects ($\xi_{r,t}$ and $\zeta_{r,t}$). $\mathcal{V}_{r,t}$ and $\nu_{i,r,t}$ denote regional and idiosyncratic shocks that affect house prices, respectively. $\mathcal{E}_{r,t}$ and $\varepsilon_{i,r,t}$ denote regional and idiosyncratic shocks that affect retail employment, respectively. These shocks should be viewed as vectors of more primitive shocks and may be correlated with each other (e.g., $\mathcal{V}_{r,t}$ and $\mathcal{E}_{r,t}$ may be correlated). Measurement error in retail employment and house prices would show up in this model as a correlation between $\nu_{i,r,t}$ and $\varepsilon_{i,r,t}$. The model allows for heterogeneity in sensitivity to regional shocks across CBSAs within region (the $i$ subscripts on $\alpha_i$ and $\gamma_i$). This feature will play an important role. Equation (8) is the analog of equation (1) in the main text and the coefficient of interest is $\beta$, the causal effect of house prices on retail employment measured as an elasticity, which we call the housing wealth elasticity.[38]

Equations (8) and (9) form a system of simultaneous equations. Changes in local house prices affect local retail employment through the $\beta\Delta p_{i,r,t}$ term in equation (8). However, changes in local employment also affect local house prices through the $\delta\Delta y_{i,r,t}$ term in equation (9). Since causation between local employment and house prices runs both ways, estimating equation (8) by OLS will yield a biased estimate of $\beta$. The classic approach to solving this problem is to look for a variable that shows up in equation (9) but not in equation (8) and to use this variable as an instrument for $\Delta p_{i,r,t}$ when estimating equation (8). In a panel data context, there is another related possibility for identification: differential sensitivity to aggregate shocks.

As we discuss in the main text, a simple implementation of this idea would be to estimate equation (3) and use $z_{i,r,t} = \hat{\gamma}_i \Delta P_{r,t}$ as an instrument for $\Delta p_{i,r,t}$ in equation (8). The simple procedure runs into problems if retail employment responds differentially to regional shocks through other channels than local house prices. Suppose for simplicity that there is no actual variation in the $\gamma_i$s in equation (9), but that local retail employment does respond differentially to regional shocks through heterogeneity in $\alpha_i$s in equation (8). In this case, the differential response of local retail employment to regional shocks induces differential responses of local house prices to these same shocks through the $\delta\Delta y_{i,r,t}$ term in equation (9). Were we to estimate equation (3) in this case, we would estimate heterogeneous $\gamma_i$s. The source of these estimated $\gamma_i$s would, however, be the $\alpha_i \mathcal{E}_{r,t}$ term in equation (8). In this case, therefore, $z_{i,r,t} = \hat{\gamma}_i \Delta P_{r,t}$ would clearly be correlated with $\alpha_i \mathcal{E}_{r,t}$. Intuitively, the differential response of local house prices to regional house prices in this example arises from reverse causation and cannot be used to identify $\beta$.

To address this problem, consider the following more sophisticated identification strategy. First aggregate equation (9) to the regional level. Since the cross-sectional average of $\nu_{i,r,t}$ is zero, this yields:

$$\Delta P_{r,t} = \zeta_{r,t} + \delta\Delta Y_{r,t} + \gamma_r \mathcal{V}_{r,t},$$

where $\Delta Y_{r,t}$ denotes the log annual change in regional retail employment, and $\gamma_r$ denotes the regional average of $\gamma_i$. Next, use this equation to rewrite equation (9) as:

$$\Delta p_{i,r,t} = \varphi_i + \tilde{\zeta}_{r,t} + \delta\Delta y_{i,r,t} + \frac{\gamma_i}{\gamma_r}\Delta P_{r,t} - \frac{\gamma_i}{\gamma_r}\delta\Delta Y_{r,t} + \nu_{i,r,t}. \tag{10}$$

[38]The empirical model also allows for heterogeneity in the sensitivity to idiosyncratic shocks. This feature is captured through heterogeneity in the variances of $\varepsilon_{i,r,t}$ and $\nu_{i,r,t}$ in the cross-section. Because this is less important for our empirical approach, the notation in equation (8) and (9) is not as explicit about this heterogeneity in sensitivity.

where $\tilde{\zeta}_{r,t} = \zeta_{r,t}(1-1/\gamma_r)$. Estimating this equation yields estimates of each city's relative sensitivity to regional house prices $\gamma_i/\gamma_r$, which we can again denote $\hat{\gamma}_i$. Finally, use $z_{i,r,t} = \hat{\gamma}_i \Delta P_{r,t}$ as an instrument in equation for $\Delta p_{i,r,t}$ in equation (8). The logic for this procedure is similar to the simpler procedure described above, but it has the advantage that it eliminates the reverse causality problem by directly controlling for $\Delta y_{i,r,t}$ and $\Delta Y_{r,t}$ in equation (10). A slightly more general version of equations (8) and (9) allows for heterogeneity in the response of house prices to local employment, which replaces $\delta$ in equation (9) with $\delta_i$. In this case, the coefficient on $\Delta y_{i,r,t}$ in equation (10) is $\delta_i$. Our empirical specification in the main text allows for this generalization as well as additional controls.

The identifying assumption implicit in the procedure described above is that $z_{i,r,t}$ is uncorrelated with $\alpha_i \mathcal{E}_{r,t} + \varepsilon_{i,r,t}$, the error term in equation (8). Because we control for $\Delta y_{i,r,t}$ and $\Delta Y_{i,r,t}$ in defining the instrument in equation (10), such a correlation cannot result from reverse causation. Furthermore, purely idiosyncratic variation ($\varepsilon_{i,r,t}$) will not be correlated with $z_{i,r,t}$ either in the time-series or cross-section. The remaining concern is that there is some component of $\alpha_i \mathcal{E}_{r,t}$ — call it $\alpha_i^j \mathcal{E}_{r,t}^j$ — that is correlated with $z_{r,t}$. To be a threat to identification, $\alpha_i^j \mathcal{E}_{r,t}^j$ must have two features. First, $\mathcal{E}_{r,t}^j$ must be correlated with regional house price cycles. Second, $\alpha_i^j$ must be correlated with $\hat{\gamma}_i$ in the cross-section. An assumption that is sufficient to rule out endogeneity of our instrument is therefore that $\alpha_i^j \perp \hat{\gamma}_i$, i.e., that the same CBSAs whose house price indexes are relatively more sensitive to regional house price cycles do not also have local employment that is differentially more sensitive to $\mathcal{E}_{r,t}^j$. With additional controls, these identifying assumptions must only hold conditional on the controls.

### C.2 Additional Details on $\gamma_i$ Variation Across Cities

Table A.4 shows the top 5 and bottom 5 $\gamma_i$ cities of over 500,000 population in 2000 for each Census region. As discussed in the main text, many cities that have similar values of the Saiz elasticity but have significantly different values of $\gamma_i$, such as Providence and Rochester.

## D Empirical Robustness

### D.1 Robustness of Rolling Windows Analysis

This section analyzes the robustness of the results, focusing on the rolling windows analysis in Figure 5. We focus on the robustness of the sensitivity instrument because it is most novel and

Table A.4: Highest and Lowest $\gamma_i$ CBSAs by Census Region (Pop > 500,000)

| Region | Northeast | Midwest | South | West |
|---|---|---|---|---|
| Lowest | Pittsburgh, PA | Wichita, KS | Greensboro, NC | Albuquerque, NM |
| | Rochester, NY | Omaha, NE | Greenville, SC | Colorado Springs, CO |
| | Harrisburg, PA | Indianapolis, IN | Winston-Salem,NC | Salt Lake City, UT |
| | Buffalo, NY | Columbus, OH | Raleigh, NC | Denver, CO |
| | Scranton, PA | Youngstown, OH | Jackson, MS | Portland, OR |
| ⋮ | | | | |
| | Worcester, MA | Toledo, OH | Jacksonville, FL | Fresno, CA |
| | Bridgeport, CT | Milwaukee, WI | Tampa, FL | Sacramento, CA |
| | New Haven, CT | Chicago, IL | Orlando, FL | Las Vegas, CA |
| | New York-Newark, NY | Minneapolis, MN | Sarasota, FL | Riverside, CA |
| Highest | Providence, RI | Detroit, MI | Miami, FL | Stockton, CA |

Notes: The table shows the top five and bottom five CBSAs with a population over 500,000 in each census region sorted by $\gamma_i$. $\gamma_i$ is estimated in a single pooled regression that does not leave out any years from 1975 to 2017.

provides for the tightest confidence intervals. The section is organized as follows:

D.1.1 Controls vs. No Controls

D.1.2 First Stage and Reduced Form

D.1.3 Alternate Specifications: First Stage and Reduce Form, 3-Year Differences, Population Weighting, Single Gamma, No Sand States, Pre-Period Gamma, Dropping Nearby Cities for Regional HPIs, Controlling For Distance to Region Centroid, 5-Year Windows, Alternative Controls and Variable Construction

D.1.4 Single Cross Section Results

D.1.5 Alternate House Price and Employment Data

### D.1.1 The Role of Controls in Our Baseline Specification.

In our baseline specification, we include a number of controls in the estimation of equations (1) and (4). These include controls for differential city-level exposure to regional retail employment, real 30-year mortgage rates, and Gilchirst and Zakrajek's (2012) measure of bond risk premia, controls for 2-digit industry shares with time-specific coefficients, and, in equation (4), the log change in average wages with city-specific coefficients.

To evaluate the role of these controls in our results, Figure A.6 shows the point estimates for the case with controls and the case with no controls (the standard errors do not change significantly).

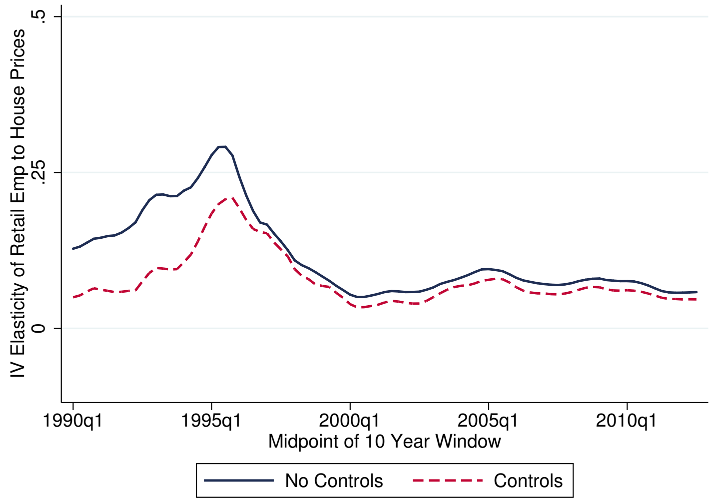


Figure A.6: Elasticity of Retail Employment Per Capita to House Prices: No Controls vs. Controls
Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods as in Figure 5a in red (dashed) and a version without the controls used in our baseline specification in blue. Each point indicates the elasticity for a 10-year sample period with its midpoint in the quarter indicated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. Confidence intervals are similar to Figure 5a and are not shown so that the comparison between the two specifications is clearer.

One can see that the elasticity is about 25% bigger without controls in the 10-year windows with their midpoints after 1996. For the 10-year windows with their midpoints prior to 1996, the controls do much more to reduce the estimated elasticity.

Which controls matter the most? In unreported results, we find that the industry shares have the largest effect, although the control for city-level exposure to regional retail employment does also explain a significant portion of the gap between the controls and no controls specifications in the 2000s.

### D.1.2 First Stage and Reduced Form

Figure A.7 shows the the first stage and the reduced form of the main results in Figure 5a. The instrument is somewhat stronger in the later period, but consistently has an F statistic above 100. The main time series pattern we observe in our IV regression is clearly evident in the reduced form.

### D.1.3 Alternate Specifications

In this section, we evaluate several alternate specifications. All of the specifications yield a pattern of declining housing wealth elasticities in rolling windows since 1995. The most important form of variation across specifications is that several specifications yield lower estimates of the elasticities


Figure A.7: First Stage and Reduced Form

Note: Panel A plots the first stage and panel B plots the reduce form of Figure 5a. Each point indicates the coefficient for a 10-year sample period with its midpoint in the quarter indicated on the horizontal axis. The instrumental variables estimator is described in Section 3. Confidence intervals are similar to Figure 5a and are not shown so that the comparison between the two specifications is clearer. The figure reports 95% confidence intervals constructed using two-way clustering by CBSA and time.

in 10-year windows with their midpoint in the early 1990s than in our baseline analysis (though with large standard errors). This is why we do not put too much emphasis on the 1980s and early 1990s results.

Figure A.8 shows three-year differences rather than one-year differences. The time pattern we find in our main figure remains the same, and the central elasticity is also similar. The main difference is that with three-year differences, the point estimate is slightly lower at the very end of the sample.

Figure A.9 shows results weighting by CBSA population in 2000 rather than unweighted. The time series pattern is similar to the pattern in our baseline analysis. However, the weighted elasticity is higher in the 10-year windows centered in the early 1990s and somewhat lower for the 10-year windows starting in the late 2000s and 2010s.

Figure A.10 shows results that estimate $\gamma_i$, the sensitivity of each city to regional house prices, once for all time periods rather than separately for each 10-year window leaving out the periods in that 10-year window. This specification does not therefore incorporate time-variation in $\gamma_i$ across windows (which may be partly real and partly due to sampling error). Since the 10-year window with its midpoint in 1996, this specification yields a declining elasticity, in line with our baseline analysis. The elasticity is lower both in the early 1990s and in the late 2000s and 2010s than in our baseline analysis.

Figure A.11 shows results dropping the “sand states” of California, Arizona, Nevada, and

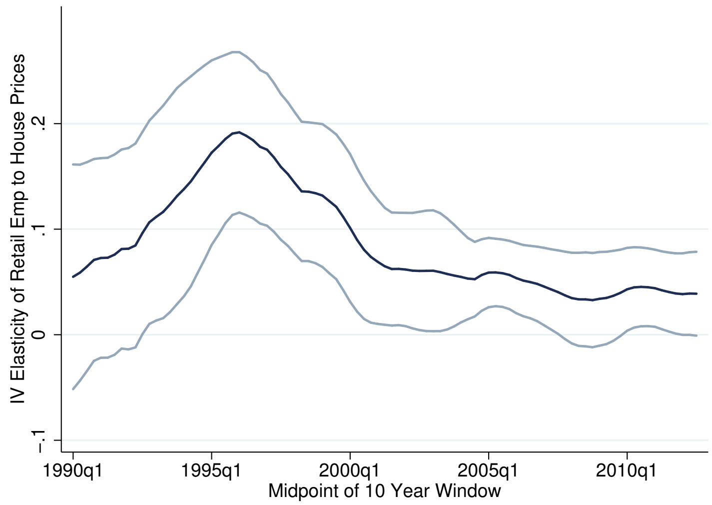


Figure A.8: Elasticity of Retail Employment Per Capita to House Prices: 3-Year Differences

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods as in Figure 5a using 3-year instead of 1-year log differences for all variables. Each point indicates the elasticity for a 10-year sample period with its midpoint in the quarter indicated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. 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.

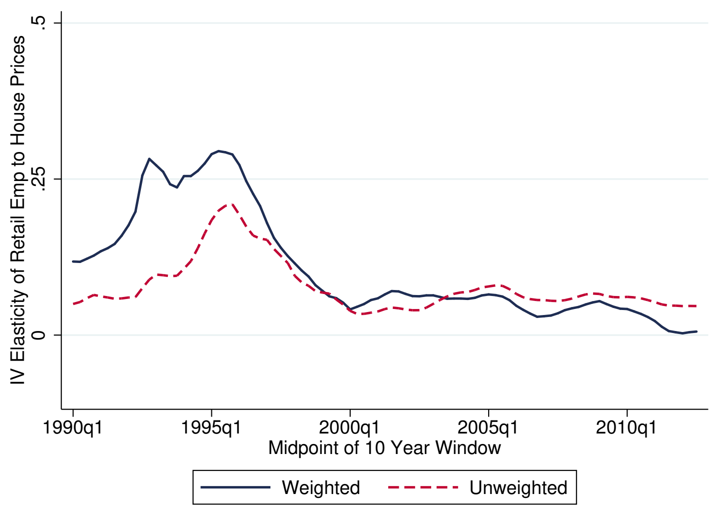


Figure A.9: Elasticity of Retail Employment Per Capita to House Prices: Population Weighting

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods as in Figure 5a in red (dashed) and a version with all regressions weighted by 2000 population in blue. Each point indicates the elasticity for a 10-year sample period with its midpoint in the quarter indicated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. Confidence intervals are similar to Figure 5a and are not shown so that the comparison between the two specifications is clearer.

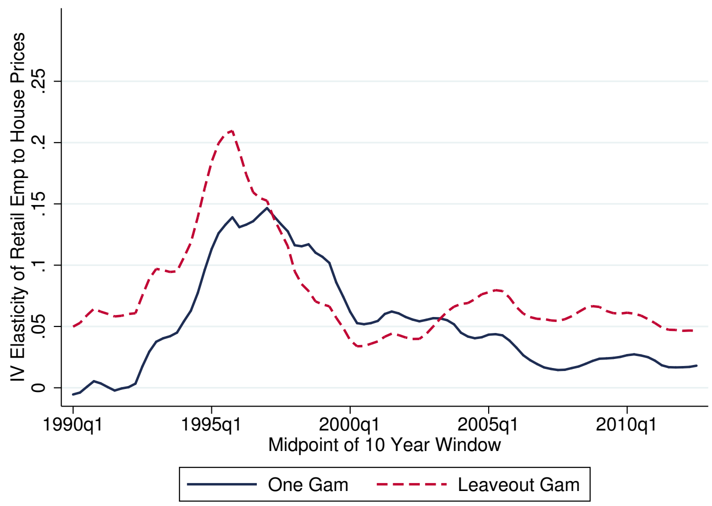


Figure A.10: Elasticity of Retail Employment Per Capita to House Prices: Single Gamma

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods as in Figure 5a in red (dashed). The specification in blue is the same except that $\gamma_i$ is estimated for each CBSA by equation (4) once for all periods (including those in the 10-year window) rather than separately for each 10-year window leaving out that 10-year window. Each point indicates the elasticity for a 10-year sample period with its midpoint in the quarter indicated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. Confidence intervals are similar to Figure 5a and are not shown so that the comparison between the two specifications is clearer.

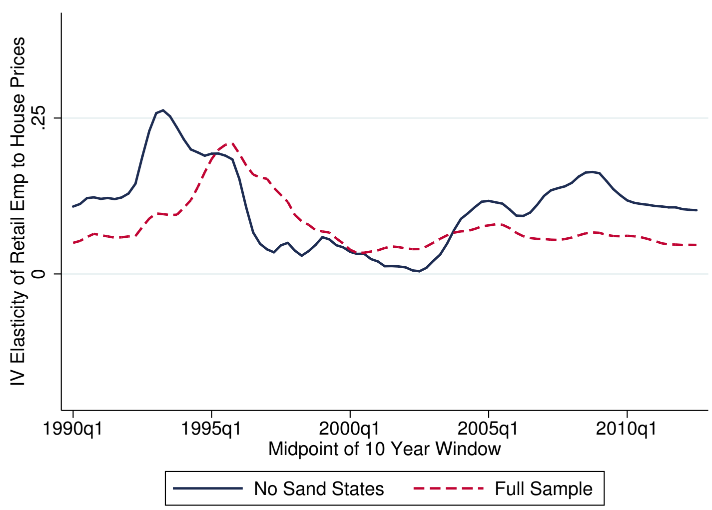


Figure A.11: Elasticity of Retail Employment Per Capita to House Prices: No Sand States

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods as in Figure 5a in red (dashed) and a version dropping the “sand states” of California, Nevada, Arizona, and Florida in blue. Each point indicates the elasticity for a 10-year sample period with its midpoint in the quarter indicated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. Confidence intervals are similar to Figure 5a and are not shown so that the comparison between the two specifications is clearer.

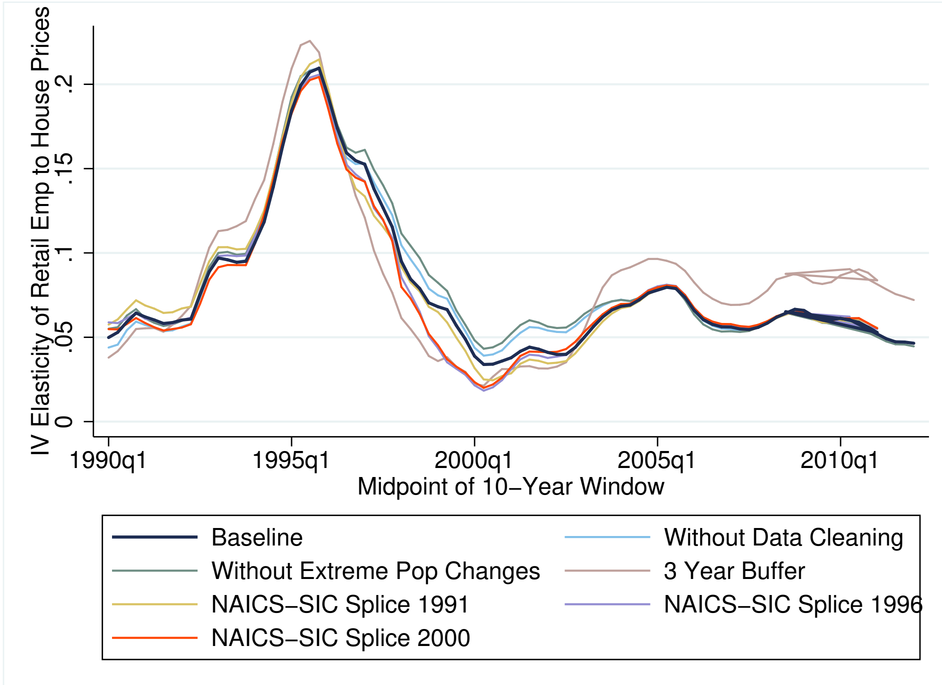


Figure A.12: Elasticity of Retail Employment Per Capita to House Prices: Misc. Robustness Tests

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods as in Figure 5a in dark blue and with seven other specifications that do not substantially affect the results. Each point indicates the elasticity for a 10-year sample period with its midpoint in the quarter indicated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. Confidence intervals are similar to Figure 5a and are not shown so that the comparison between the specifications is clearer. The "without data cleaning" specification does not drop counties that have bad observations in the QCEW and also does not remove periods in which a CBSA has an unusual jump in employment. The "without extreme pop changes" specification drops periods with extreme population changes. The "3 year buffer" specification drops a three year buffer around the 10-year window in question for regression (4) and for the regressions as in equation (2) used to create the controls for differential city-level exposure to regional retail employment, real 30-year mortgage rates, and Gilchirst and Zakrajek's measure of bond risk premia. The "NAICS-SIC Splice 1991," "NAICS-SIC Splice 1996," and "NAICS-SIC Splice 2000" uses these three alternate dates rather than Q1 1993 as the date we use to splice the NAICS and SIC retail employment series together.

Florida from the analysis. Critics of the Saiz instrument such as Davidoff (2017) often argue that much of the identification is driven by these states. This figure shows that the declining pattern of elasticities since the mid 1990s is not affected by these states. Indeed, the quantitative results since the mid 1990s are similar whether or not one includes these states.

Figure A.12 shows seven different robustness checks that do not change the results substantially (in light colors) together with the baseline specification (in dark blue). The first specification leaves the data "raw" rather than dropping counties within a CBSA with bad observations in the QCEW and cleaning the data to remove time periods with jumps as described in Appendix A.1. This has essentially no impact on the results. The second specification drops observations with particularly large population changes. Again, this has no impact on the results. The third specification uses

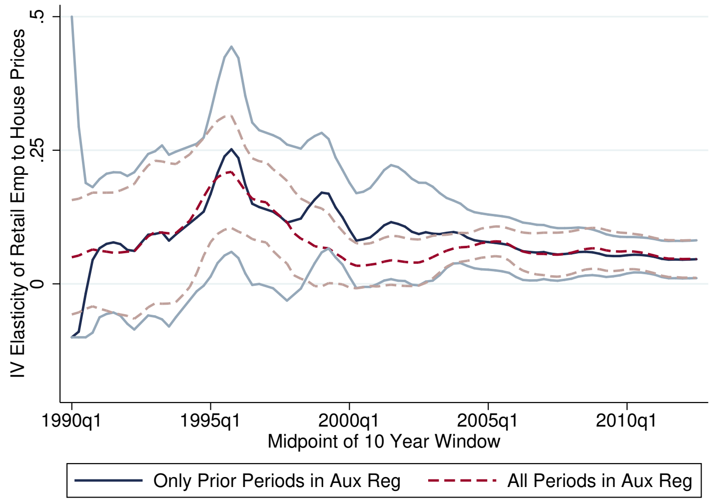


Figure A.13: Elasticity of Retail Employment Per Capita to House Prices: Only Prior Periods To Create Instrument

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods as in Figure 5a in red (dashed) and a version in which only periods prior to each 10-year window are used to estimate the instrument in blue. Each point indicates the elasticity for a 10-year sample period with its midpoint in the quarter indicated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. The standard errors are constructed using two-way clustering by CBSA and time.

a three-year buffer around the 10-year window in constructing the instrument using equation (4) and in constructing the controls for differential city-level exposure to regional retail employment, real 30-year mortgage rates, and the Gilchrist-Zakrajsek excess bond premium using equation (2). This has a slight effect on the results for a few 10-year windows in the mid-to-late 1990s, the mid 2000s, and the 2010s, but the difference is not significant economically or statistically and the main time series pattern remains. Finally, the last three specifications show results that change the date at which we splice together the SIC and NAICS retail employment data from 1993Q1 to 1991Q1, 1996Q1, and 2000Q1, respectively. How we splice together SIC and NAICS also has essentially no impact on the results.

Figure A.13 shows results for a version that uses only periods prior to the 10-year window for the instrument estimation rather than also using periods after the 10-year window. The results are similar, although the standard errors are generally larger earlier in the sample and the point estimates are slightly larger for the 10-year windows with midpoints in the early 1990s.

Figure A.14 addresses concerns about the regional house price index and employment being driven by nearby cities that share common shocks. Rather than creating the regional index and employment for each city using a leave out mean, this figure drops all cities within 250 miles when

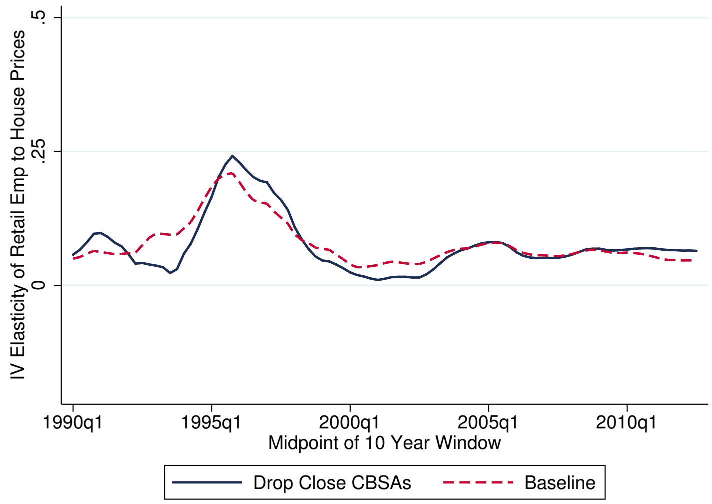


Figure A.14: Elasticity of Retail Employment Per Capita to House Prices: Dropping CBSAs Within 250 Miles For Regional

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods as in Figure 5a in red (dashed) and a version for which the regional indices used for each CBSA do not include any CBSAs within 250 miles rather than only dropping the CBSA in question in blue. We use an instrumental variables estimator that is described in Section 3. Each point indicates the elasticity for a 10-year sample period with its midpoint in the quarter indicated on the horizontal axis. Confidence intervals are similar to Figure 5a and are not shown so that the comparison between the two specifications is clearer.

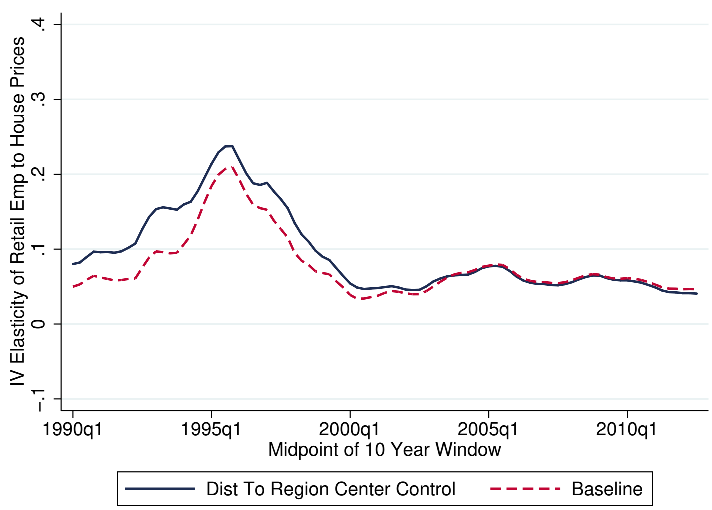


Figure A.15: Elasticity of Retail Employment Per Capita to House Prices: Controlling For Distance to Region's Population Centroid

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods as in Figure 5a in red (dashed) and a version that controls for the distance between the centroid of the CBSA's principal city and the region's population centroid computed from the 2000 Census interacted with time fixed effects. We use an instrumental variables estimator that is described in Section 3. Each point indicates the elasticity for a 10-year sample period with its midpoint in the quarter indicated on the horizontal axis. Confidence intervals are similar to Figure 5a and are not shown so that the comparison between the two specifications is clearer.


Figure A.16: Elasticity of Retail Employment Per Capita to House Prices: Rolling 5-Year Window
Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level as in Figure 5a except for rolling 5-year windows instead of 10-year windows. Each point indicates the elasticity for a 5-year sample period with its midpoint in the quarter indicated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. 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.

creating the regional average. The results are essentially the same.

Figure A.15 addresses concerns that the more sensitive cities may be closer to the economic "core" of a region by including controls for the distance to the centroid of the census region interacted with time fixed effects. The centroid is computed from the 2000 Census population. The results do not change significantly. We have found essentially unchanged results if we define the population centroid based on measures of economic activity rather than population or distance to the largest city in the region.

Finally, Figure A.16 shows results for 5-year windows This finer resolution does not yield evidence that the elasticity rose significantly during the 2000s. This assuages concerns that by using a 10-year window we are obscuring significant higher-frequency variation. One notable feature of this plot is that the elasticity falls below zero for several years around 2013 before swinging back to positive at the end of the sample, although the standard errors increase significantly in this period.

### D.1.4 Single Cross-Section Results

Figure A.17 shows the point estimates and standard errors for repeated 3-year cross-sections of the type that have been used in recent analyses of the Great Recession. Our specification here is the same as our baseline in equation (1) except that we use a single cross section and replace the region-time fixed effect with a region fixed effect. We also demean all variables including controls once for the whole period from 1978-2017.

The results are much more volatile, and the standard errors are sufficiently large in many periods to make the estimates essentially uninformative. These tend to be time periods where the 3-year difference in the regional house price index that is used to construct our instrument is near zero (e.g., a peak or a trough), so the instrument loses power. Aside from these periods, though, one can see a tendency for the housing wealth elasticity to be greater from 1990 to 2003 than it was in the 2000s.

#### D.1.5 Alternate Data: CoreLogic House Prices and County Business Patterns Employment

Our baseline analysis uses house price data from Freddie Mac. Figures A.18 and A.19 instead use the CoreLogic house price index. The CoreLogic index is an arithmetic repeat sale house price index that has two advantages. First, it includes a broader sample of homes bought with non-conforming loans. Second, it includes only transactions white the Freddie Mac index includes appraisals. Because appraisers tend to look backwards, this would create to a "smoothed" index that would may cause an upward bias in our estimates of the house price elasticity, because we would be observing the same retail employment change for a smaller smoothed change in house prices. The CoreLogic index does, however, have a downside: many CBSAs are imputed from a higher geography prior to 2000. This would create issues with our estimation strategy because it would create an artificial correlation between the house prices in a CBSA and nearby cities that are used to impute the CBSA's house prices. Consequently, we take two approaches to the CoreLogic data. In Figure A.18, we use only the CoreLogic data dropping any imputed observations in an unbalanced panel. In Figure A.19, by contrast, we use the Freddie Mac data in prior to 2000 and the CoreLogic data after 2000, when many fewer house prices are imputed, in a balanced panel. The results of both specifications are similar to our baseline analysis.

Figure A.20 uses County Business Pattern (CBP) employment data rather than QCEW data for retail employment, which has somewhat different sampling frames and industry definitions than our baseline QCEW data. The CBP is only available annually, and so we report annual results rather than quarterly. As a consequence, the standard errors are larger. Nonetheless, the same general time pattern we observe with the QCEW is evident with the CBP.

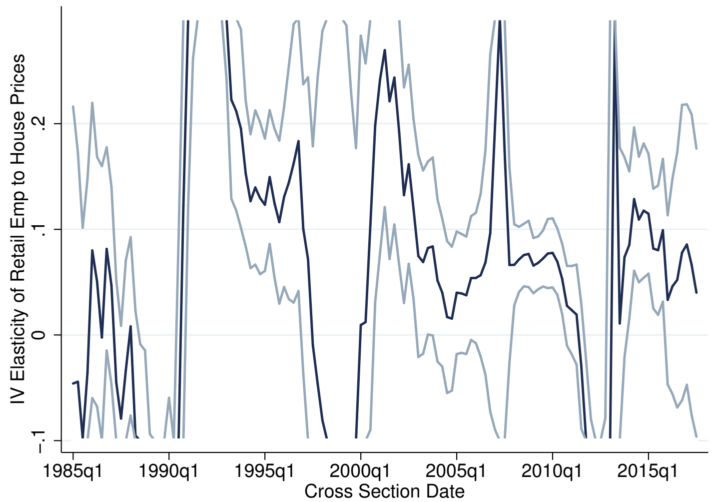


Figure A.17: Repeated Cross-Sections: 3 Year Differences

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for repeated 3-year difference cross sections. Each point indicates the elasticity for a single cross section between the indicated date and three years prior. We use an instrumental variables estimator that is described in Section 3 but with only region FE instead of region-time FE. We take out the CBSA fixed effect (or equivalently demean) once for the full 1978-2017 sample. The figure reports 95% confidence intervals in addition to point estimates for the elasticity. We report robust standard errors.

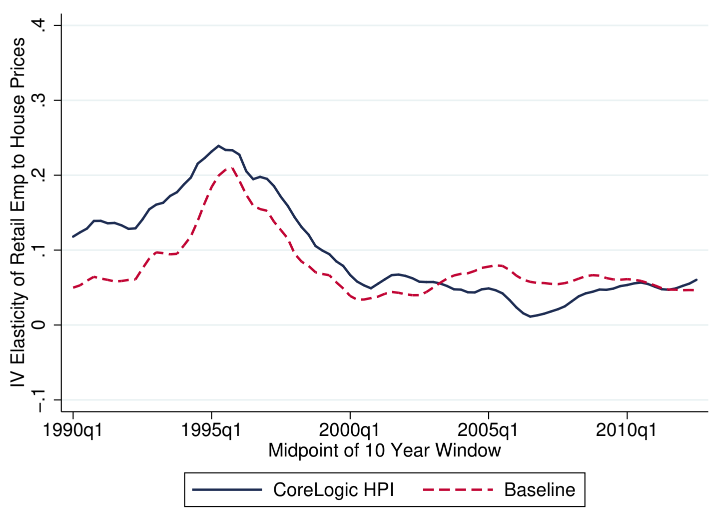


Figure A.18: Alternate Data: CoreLogic House Price Data (Unbalanced Panel)

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods as in Figure 5a except using the CoreLogic house price index instead of the Freddie Mac house price index. The CoreLogic index does not include appraisals and includes a broader sample of homes purchased with non-conforming mortgages. However, it is not available as far back for every geography. This figure shows results using an unbalanced panel that adds each CBSA as it becomes available. Each point indicates the elasticity for a 10-year sample period with its midpoint in the year indicated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. Confidence intervals are similar to Figure 5a and are not shown so that the comparison between the two specifications is clearer.

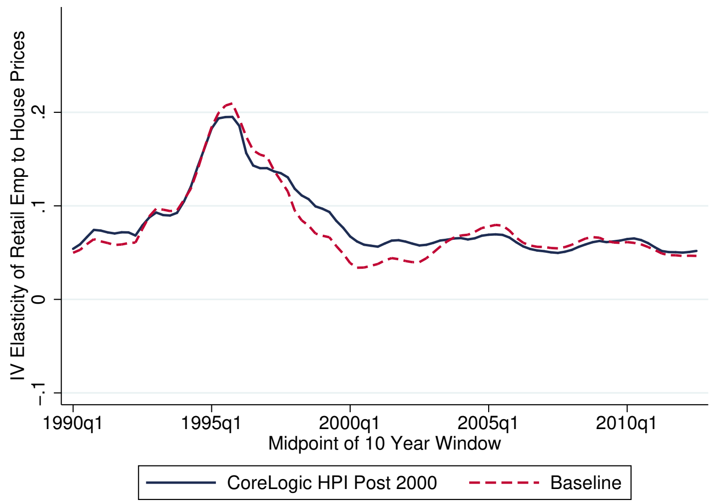


Figure A.19: Alternate Data: Freddie Mac Pre-2000, CoreLogic Post (Balanced Panel)

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods as in Figure 5a except using the CoreLogic house price index instead of the Freddie Mac house price index after 2000 but uses the Freddie Mac house price index before 2000. This allows us to create a balanced panel. This figure shows results using an unbalanced panel that adds each CBSA as it becomes available. Each point indicates the elasticity for a 10-year sample period with its midpoint in the year indicated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. Confidence intervals are similar to Figure 5a and are not shown so that the comparison between the two specifications is clearer.

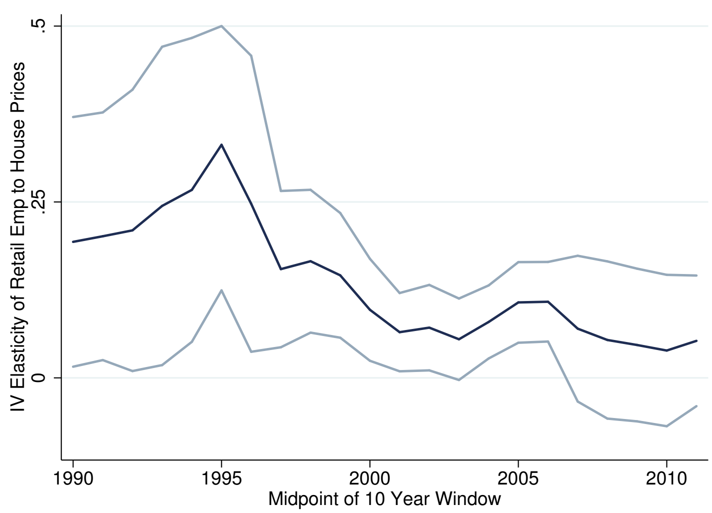


Figure A.20: Alternate Data: County Business Patterns Employment Data

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods as in Figure 5a except using County Business Patterns data for retail employment rather than the QCEW. The CBP is available annually, and so the figure is annual. Each point indicates the elasticity for a 10-year sample period with its midpoint in the year indicated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. 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.

Table A.5: 1978-2017 Elasticity of Retail Employment Per Capita to House Prices: Weighted

| | (1) | (2) | (3) |
|---|---|---|---|
| $\Delta \log (P)$ | 0.067** | | |
| | (0.022) | | |
| $\Delta \log (P) -$ | | 0.089** | |
| | | (0.031) | |
| $\Delta \log (P) +$ | | 0.042 | |
| | | (0.025) | |
| P Test for Equality | | 0.188 | |
| $\Delta \log (P)$ | | | 0.063** |
| | | | (0.023) |
| $\Delta \log (P)^2$ | | | -0.040 |
| | | | (0.053) |

Note: For these estimates, we first construct our instrument for each quarter by estimating the $\gamma_i$'s in equation (4) for each quarter, using a sample period that leaves out a three-year buffer around the quarter in question. We then estimate equation (1) pooling across all years. Specification 1 does so for all price changes, specification 2 does so by comparing positive and negative house price changes, and specification 3 uses a quadratic in the log change in house prices. For specification 2, we instrument with $Z \times 1\,[Z \geq 0]$ and $Z \times Z\,[< 0]$ and for specification 3 we instrument with $Z$ and $Z^2$. The estimating equation is the same as equation (1) except for $\Delta \log (H)$ being interacted with indicators for $\Delta \log (H) \geq 0$ and $\Delta \log H < 0$ in specification 2 and the addition of the quadratic term in specification 3. Standard errors are two-way clustered at the time and CBSA level. All regressions are weighted by 2000 population.

Table A.6: 1978-2017 Elasticity of Retail Employment Per Capita to House Prices: No Controls

| | (1) | (2) | (3) |
|---|---|---|---|
| $\Delta \log (P)$ | 0.068*** | | |
| | (0.015) | | |
| $\Delta \log (P) -$ | | 0.071*** | |
| | | (0.023) | |
| $\Delta \log (P) +$ | | 0.065*** | |
| | | (0.020) | |
| P Test for Equality | | 0.858 | |
| $\Delta \log (P)$ | | | 0.067*** |
| | | | (0.016) |
| $\Delta \log (P)^2$ | | | -0.008 |
| | | | (0.042) |

Note: For these estimates, we first construct our instrument for each quarter by estimating the $\gamma_i$'s in equation (4) for each quarter, using a sample period that leaves out a three-year buffer around the quarter in question. We then estimate equation (1) pooling across all years. Specification 1 does so for all price changes, specification 2 does so by comparing positive and negative house price changes, and specification 3 uses a quadratic in the log change in house prices. For specification 2, we instrument with $Z \times 1\,[Z \geq 0]$ and $Z \times Z\,[< 0]$ and for specification 3 we instrument with $Z$ and $Z^2$. The estimating equation is the same as equation (1) except for $\Delta \log (H)$ being interacted with indicators for $\Delta \log (H) \geq 0$ and $\Delta \log H < 0$ in specification 2 and the addition of the quadratic term in specification 3. Standard errors are two-way clustered at the time and CBSA level. All regressions do not include the standard controls in our baseline specification.

### D.2 Pooled Results: With vs. Without Controls, Weighted vs. Unweighted

Table A.6 presents the results of our pooled regression This section presents results on the robustness of the pooled specifications weighting by 2000 population. Table A.6 presents the same results without controls. The point estimates for the weighted regression are somewhat lower, and the point estimates for the version without controls are somewhat higher. The weighted results show stronger evidence of a boom-bust asymmetry, but the difference remains statistically insignificant.

## E Model Extensions and Robustness

This section analyzes the robustness of the model results to various extensions and relaxations of assumptions. It is organized as follows:

E.1 Linearity and Interaction Effects

E.2 Changes to $\beta$

E.3 Changes in Credit Constraints

E.4 Interest Rate Changes

E.5 Rental Cost of Housing

E.6 Short-Term Debt

E.7 Housing Transaction Costs

E.8 No Short-Run Housing Adjustment

E.9 Accounting for the Evolution of Household Balance Sheets

### E.1 Linearity and Interaction Effects

Figure A.21 shows model estimates of the housing wealth elasticity for both positive and negative changes larger than the 10% changes we use in our main analysis. The figure shows that the housing wealth elasticity does not change meaningfully as we change the magnitude of the home price change nor does it show any meaningful asymmetry between positive and negative price changes.

In Section 5, we explain how aggregate shocks are absorbed by the time fixed effects in our empirical specification. In light of this, one way to interpret our theoretical experiments is in terms

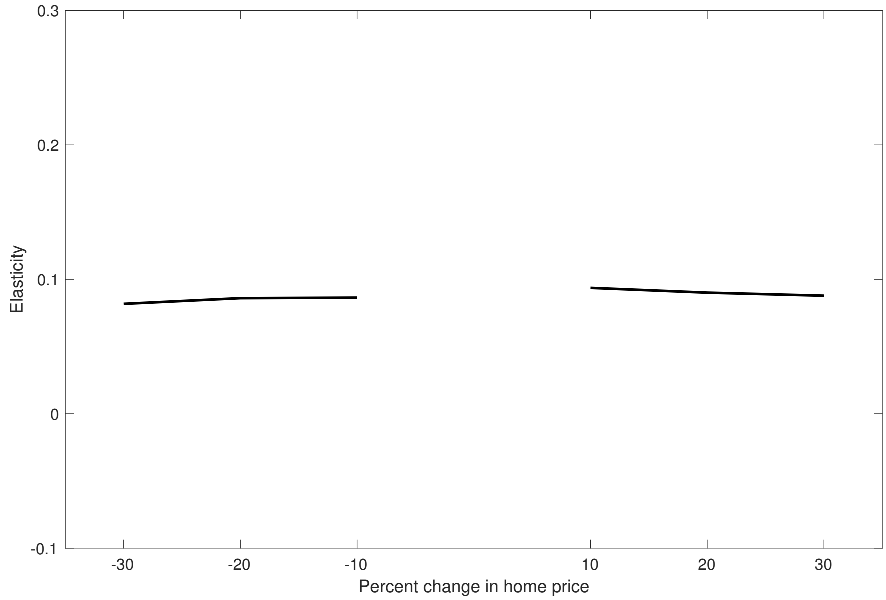


Figure A.21: Linearity of the Housing Wealth Elasticity

Note: The figure shows $(p/C)(dC/dp)$ where $dC/dp$ is computed from a +10% change, a +20% change and so on.

of two cities with different housing supply elasticities being subjected to an aggregate shock to housing demand. The demand shock itself is absorbed by the fixed effect so we do not model it explicitly and instead focus on the differential reaction of home prices in the two cities. This argument allows us to remain agnostic about the shocks driving home prices. However, the argument assumes that the second-order interaction of home prices and the demand shock does not have important consequences for consumption. Let us next evaluate the validity of this argument for a specific shock to housing demand: an increase in the preference for housing, $\omega$. Figure A.22 shows that the housing wealth elasticity does not change meaningfully as we change $\omega$. This implies that the cross derivative $d^2C/(dpd\omega)$ is small.

### E.2 Changes in the Discount Factor $\beta$

Figure A.23 shows that raising $\beta$ by 0.01 reduces the level of the housing wealth elasticity by approximately the same amount without changing the time series pattern. Increasing $\beta$ reduces the speed at which homeowners spend the resources freed up by changes in home prices.

### E.3 Changes in Credit Constraints

Our baseline analysis assumes that credit conditions remain constant as households change their balance sheets, yet an important part of the narrative of the housing boom and bust was an

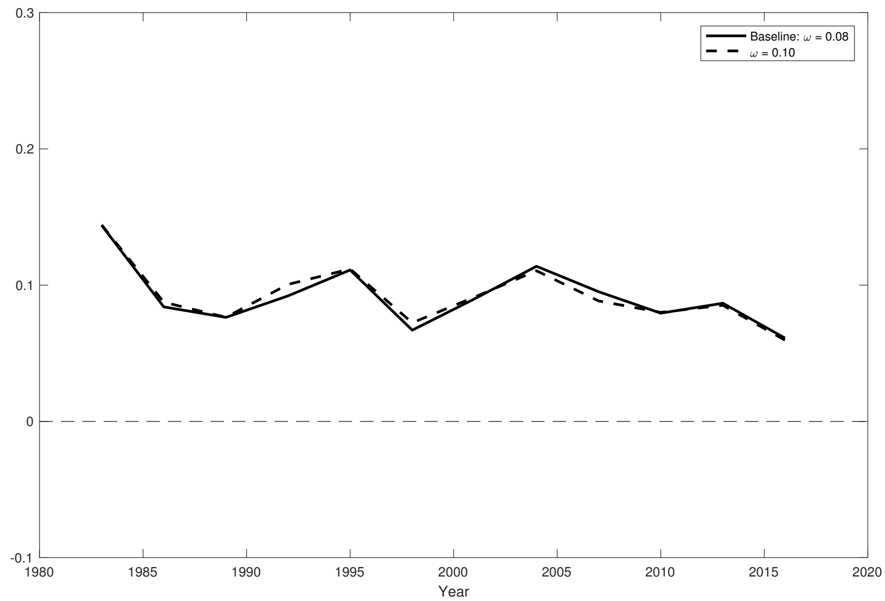


Figure A.22: The Housing Wealth Elasticity For Alternative $\omega$

Note: The figure shows the housing wealth elasticity elasticity when households have a stronger preference for housing. Our calculation takes current states as given by the SCF data so the figure shows the effect of $\omega$ on the consumption decision rule.

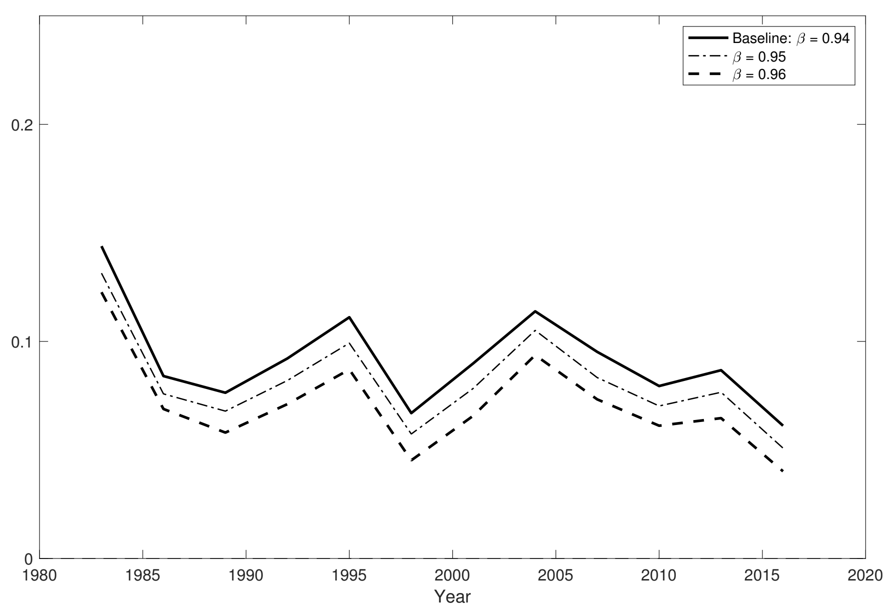


Figure A.23: The Housing Wealth Elasticity with Different Discount Rates

Note: The figure shows the same calculation of $(p/C)(dC/dp)$ as in Panel (a) of Figure 8 using an alternate $\beta$ parameters for the household's decision problem.

expansion and contraction in household credit (e.g., Favilukis et al., 2017). To analyze how looser credit conditions in the housing boom and tighter credit conditions in the Great Recession affected the housing wealth elasticity, we consider two alternative parameterizations of the LTV constraint, one with a maximum LTV of $\theta = 0.90$ and one with a maximum LTV of $\theta = 0.70$ (both assumed to remain constant in the future).[39]

Figure A.24 shows that the housing wealth elasticity is barely changed by shifts in the LTV constraint. The intuition in Table 6 is useful to understand these results: a large part of the housing wealth elasticity comes from households who are far from the LTV constraint and their behavior is little affected by the details of the constraint. We should also emphasize that our focus is on the impact of the credit constraint on the *elasticity* of consumption with respect to home prices as opposed to the *level* of consumption. This is an important difference from other analyses that focus on how the level of consumption reacts to changes in borrowing constraints. For example, Table 2 of Favilukis et al. (2017) shows that relaxing the collateral constraint from a 25 percent downpayment requirement to a 1 percent requirement raises the consumption share of GDP by two percentage points. This is entirely consistent with our finding that the housing wealth elasticity was relatively invariant to credit constraints.

### E.4 Changes in Interest Rate

Figure A.25 shows results with alternate values for the real mortgage interest rate. We find that the housing wealth effect is increasing in the real mortgage interest rate especially from 2004 onwards, which is a reflection of the increase in mortgage balances relative to income and non-housing assets in those years. As we describe in the main text, the interest rate affects the level of the wealth elasticity because at higher interest rates, households are more likely to downsize their homes and downsizers have large elasticities. The decline in real rates over time may have, to some degree, countered the upward movements in the housing wealth elasticity coming from higher home values and leverage.

[39] Some analyses of changing credit conditions (e.g., Guerrieri and Lorenzoni, 2015) take the initial distribution of individual states from a model simulation (e.g., a steady state). In that type of analysis, if there is a tightening of the credit constraint, households are forced to de-lever. Our analysis differs in that we are taking the distribution of idiosyncratic states from the data and conditional on these states the constraints that households faced in the past are irrelevant. To the extent that households were forced to delever, this should be reflected in the data we see.

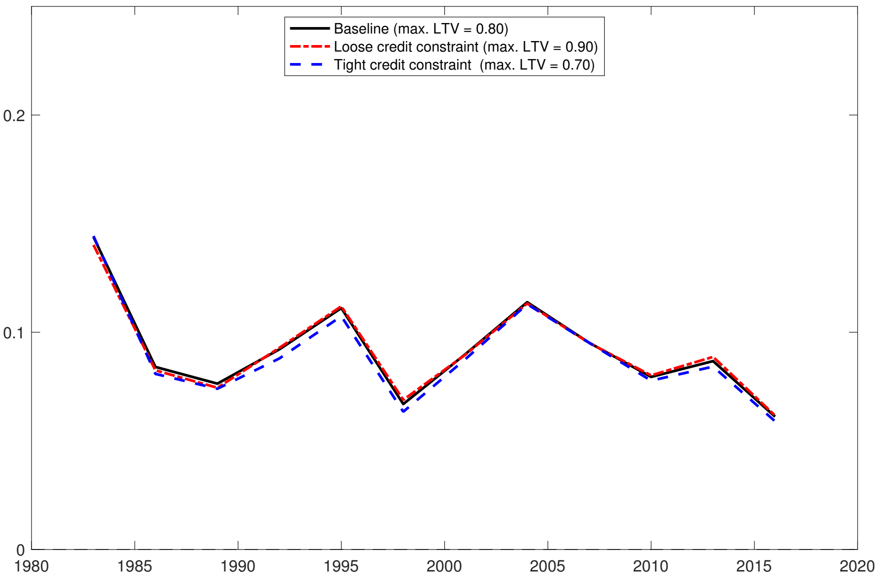


Figure A.24: Impact of Borrowing Constraints on the Housing Wealth Elasticity

Note: Housing wealth elasticity for alternative calibrations of the LTV constraint.

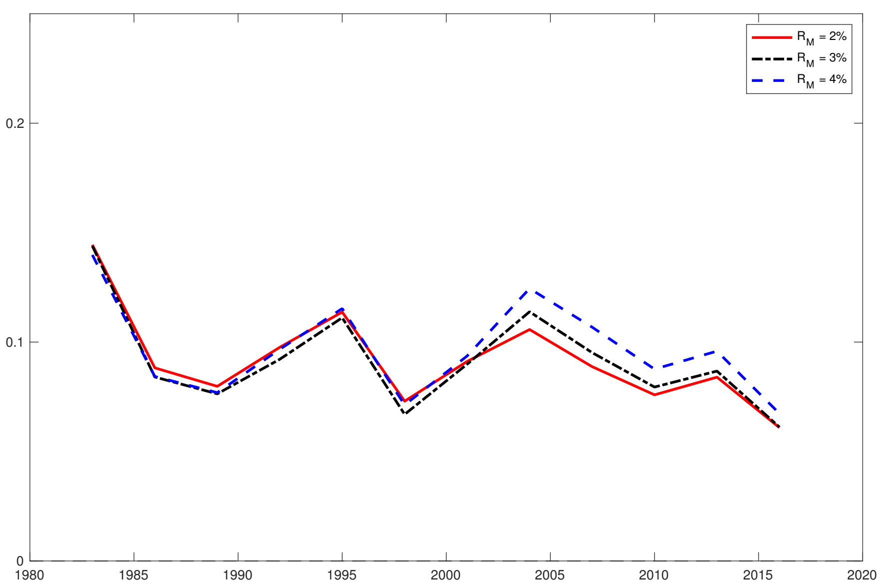


Figure A.25: Sensitivity to Interest Rates

Note: Housing wealth elasticity for alternative calibrations of the mortgage interest rate.

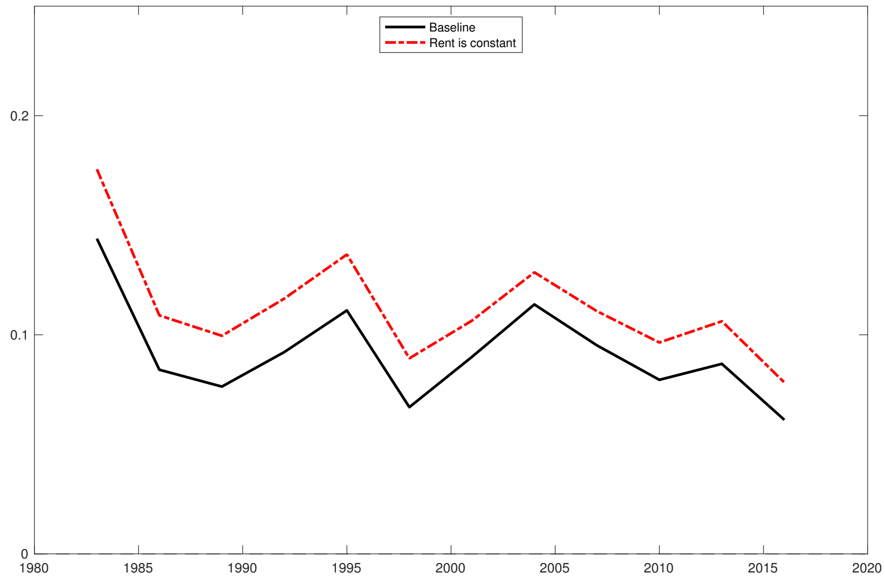


Figure A.26: Sensitivity to Rent-Price Ratio

Note: Housing wealth elasticity assuming that rents are constant as opposed to our baseline assumption that the rent-price ratio is constant.

## E.5 Alternate Assumptions on the Cyclicality of the Rental Cost of Housing

We assume that rents are a constant fraction of home prices. The logic underlying that assumption is that in the absence of expected capital gains, the user cost of housing is roughly proportional to the home price as the main component of the user cost is the foregone interest. Nevertheless, during the housing boom of the 2000s, the rent-price ratio fell considerably. One interpretation is that in cities where home prices were rising sharply, rents remained low because the user cost was kept down by expected capital gains.

Figure A.26 shows that the level of the housing wealth elasticity rises but its time series remains unchanged if we make the polar opposite assumption that rents remain steady when home prices change. In this scenario, an increase in home prices makes renting relatively more attractive. Some renters delay purchasing a home and no longer need to accumulate savings for a downpayment, which allows them to increase their consumption. This force raises the housing wealth elasticity in the aggregate.

## E.6 Short-Term Debt

Figure A.27 shows the model implied elasticity of consumption to house prices by LTV as in Figure 9 for a short-term debt model in which households must satisfy the LTV constraint each period in order to roll over their debt. The elasticity is much higher for high-LTV homeowners and remains elevated even for underwater homeowners (note the difference in y-axis scales relative to Figure 9). This is the case because households are "margin called" when house prices fall in order to meet the LTV constraint. The model consequently generates an increase in the wealth elasticity in the Great Recession as well as a substantial boom-bust asymmetry.

### E.7 Housing Transaction Costs

Figure A.28 shows the housing wealth elasticity for an alternative calibration in which we double the cost of selling a house, $\psi^{Sell}$. In this alternative calibration, the housing wealth elasticity is lower but the time series pattern is unchanged. With larger transaction costs, homeowners are less likely to realize capital gains on their homes and consumption is more insulated from home price changes. The alternative calibration under-predicts the fraction of households buying a home each year at 2.2% as compared to our empirical target of 3.2%.

### E.8 No Short Run Housing Adjustment

To incorporate the inelastic nature of short-run housing supply, we consider an alternative experiment in which there is no change in the demand for housing in the short run. Specifically, we consider two cities with different long-run housing supply elasticities, but both cities have a zero short-run housing supply elasticity. In both cities, a development occurs that shifts housing demand out and in the long-run this will have a larger effect on prices in the less-elastic city. At higher prices, this city will demand less housing than the more-elastic city. We assume that the price of housing rises by 10 percent more in the less elastic city in the long run. In the short run, which we take to be the first year after the news arrives, the price must adjust so that neither city changes its housing demand. In practice, this means that the price differential is initially less than 10 percent because the less elastic city requires an expected capital gain in order to induce people to hold more housing in the short-run. To put it formally, we can write the demand for housing in city $i$ as:

$$H_{i,t} = \int h(x, p_{i,t}, p_{i,t+1}) d\Phi_t(x),$$

where we assume that the price is constant from $t+1$ onwards. In the more elastic city, we assume the price remains constant at $\bar{p}_t$. This should be interpreted as a normalization, since we focus on

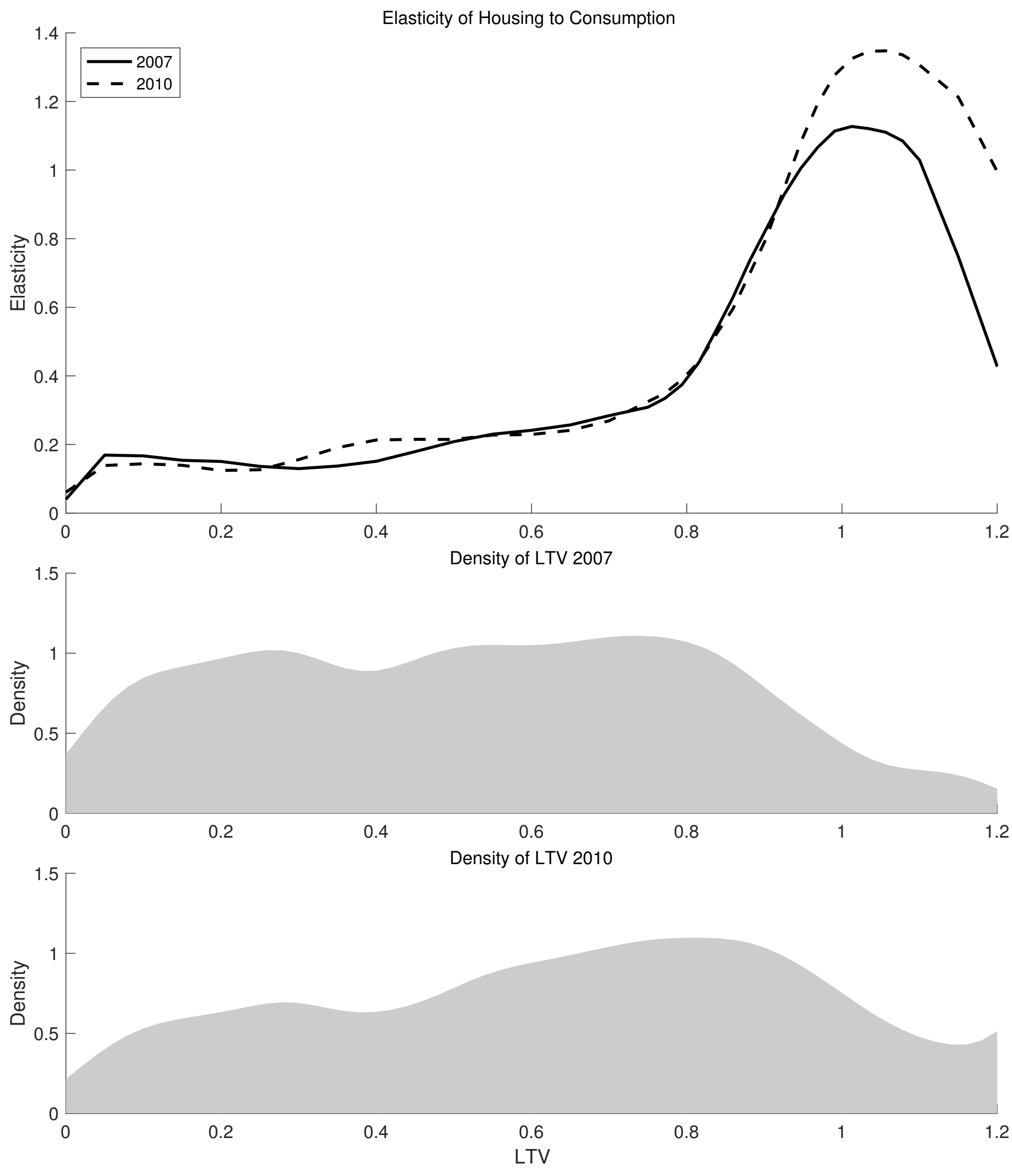


Figure A.27: Elasticity by LTV for Short-Term Debt Model

Note: Housing wealth elasticity across LTVs under the short-term debt model.

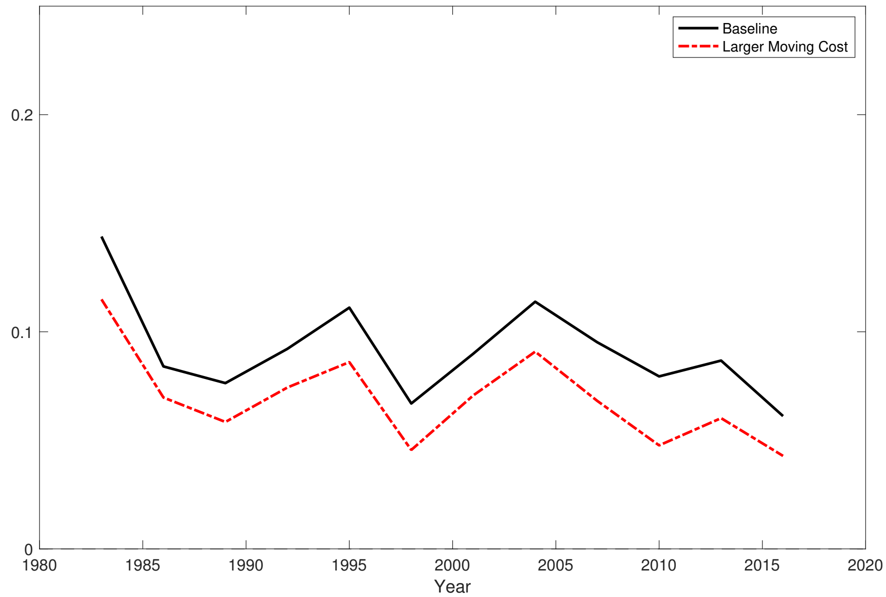


Figure A.28: Sensitivity to Housing Transaction Cost

Note: Housing wealth elasticity with $\psi^{Sell} = 0.22$, double its value in the baseline.

the differential behavior of the two cities. In the more-elastic city, housing demand is given by:

$$\bar{H}_t = \int h(x, \bar{p}_t, \bar{p}_t) d\Phi_t(x).$$

In the less-elastic city, we assume the price will rise by 10 percent relative to the more-elastic city in the long-run and in the short-run the price evolves so that the demand for housing in the two cities is equal. That is we solve for the $p_{i,t}$ that satisfies:

$$\bar{H}_t = \int h(x, p_{i,t}, 1.1 \times \bar{p}_t) d\Phi_t(x).$$

Finally, we compare consumption across cities, which we calculate from

$$C_{i,t} = c(x, p_{i,t}, p_{i,t+1}) d\Phi_t(x).$$

Figure A.29 shows the housing wealth elasticity in the short-run (i.e., it compares consumption in the two cities on the date the news arrives expressed as an elasticity with respect to the short-run difference in prices across the cities). The housing wealth elasticity is very similar to in our baseline case, though slightly lower, because the less-elastic city is no longer substituting out of housing towards consumption in the short-run. The difference from the baseline case is only minor

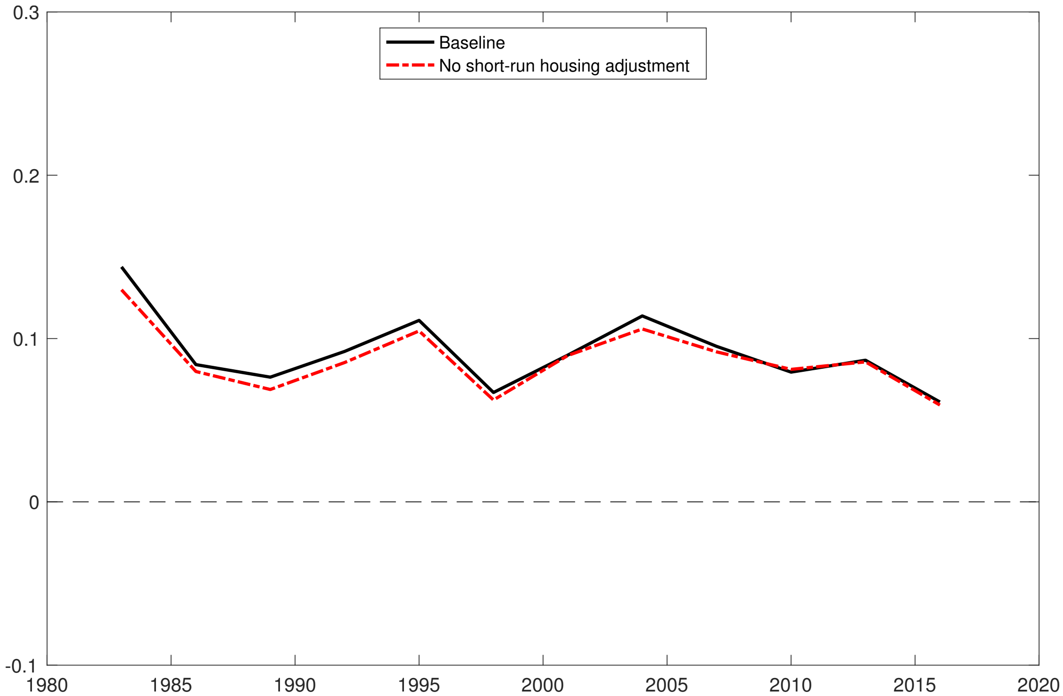


Figure A.29: No Short-Run Housing Adjustment

Note: The figure shows the housing wealth elasticity when expected capital gains on housing adjust so as to stabilize housing demand in the first period after the shock.

because the demand for housing is sensitive to expected capital gains. Small expected capital gains are sufficient to obtain no short-run housing adjustment (and presumably could also explain increases, as opposed to decreases, in housing supply in response to a house price increase).

## E.9 Accounting for the Evolution of Household Balance Sheets

Our baseline analysis takes the distribution of individual states from the SCF. In this Appendix, we show that the model does a fairly good job explaining the year-to-year changes in the distribution of LTVs except for the housing boom years of the early 2000s. We then show how the model can be extended to allow for a relaxation of credit constraints and news about future capital gains in the boom to better explain the evolution of the LTV distribution during those years without substantially affecting the housing wealth elasticity.

Given the observed distribution of individual states at the start of year $t$, and a sequence of home prices, $\{P_{t+k}\}_{k=0}^{K-1}$, the model implies an evolution of distributions of states for years $t+1$,..., $t+K$. We begin a simulation with each wave of the SCF (i.e. 1983, 1986,...) and simulate four years of data using the observed evolution of home prices. In this analysis we do not make the CoreLogic adjustment to the SCF data because doing so creates a sharp break in the LTV distribution in

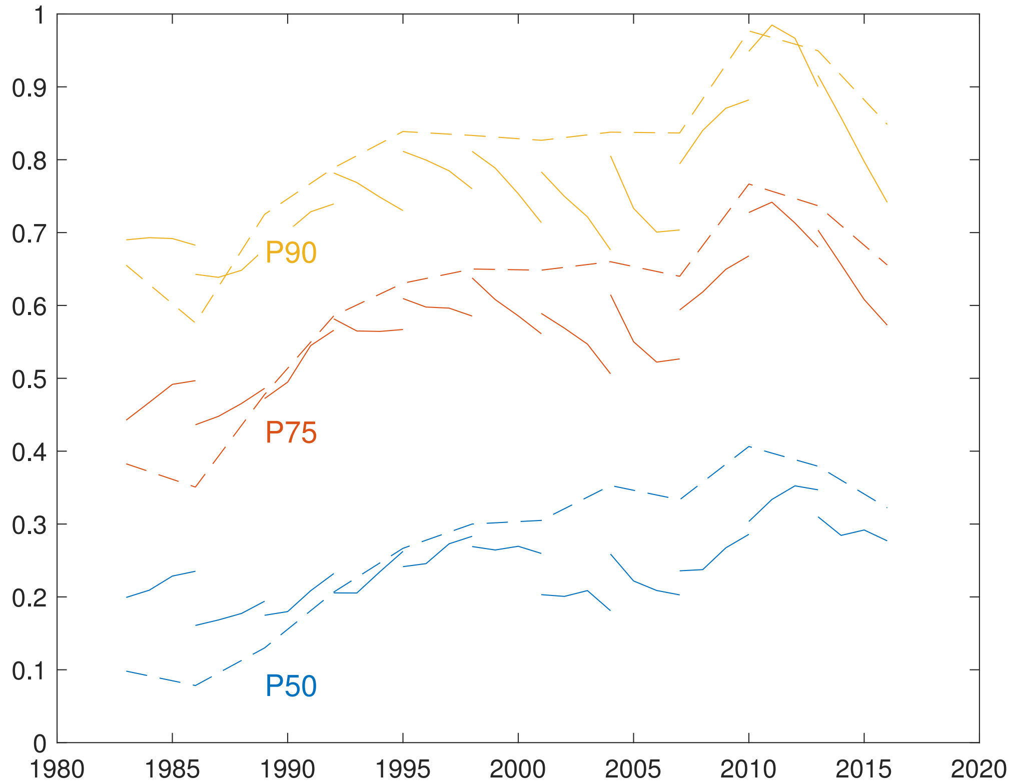


Figure A.30: Simulated LTV Distribution

Note: The dashed lines show the median, 75th percentile and 90th percentile of the LTV distribution among homeowners using data from the SCF. The short solid lines show simulated data, initializing the distribution of individual states from SCF and then feeding into the model the observed evolution of home prices.

2007 that comes from a methodological change, and we should not expect the model to reproduce that pattern.

Figure A.30 plots quantiles of the LTV distribution both in the data and implied by the model. The model succeeds on three dimensions. First, it captures the increase in leverage in the late 1980s and early 1990s. Second, it is consistent with the increase in leverage in the Great Recession. Finally, it is consistent with the deleveraging observed at the end of the sample. Where the model fails is during the housing boom. During those years, the increases in home values would push LTV down if mortgage debt remained constant, but in the data there is no evident fall in LTV as mortgage debt rose in line with home values leaving LTVs roughly flat over this period. The model does not predict this increase in mortgage debt, so LTVs fall during these years according to the model.

Next we introduce the "boom" parameterization described in the main text that differs from the baseline parameters in that the LTV limit increases from 80 percent to 95 percent and refinancing is free. We use this calibration to simulate the years 1998 to 2007. One interpretation of free

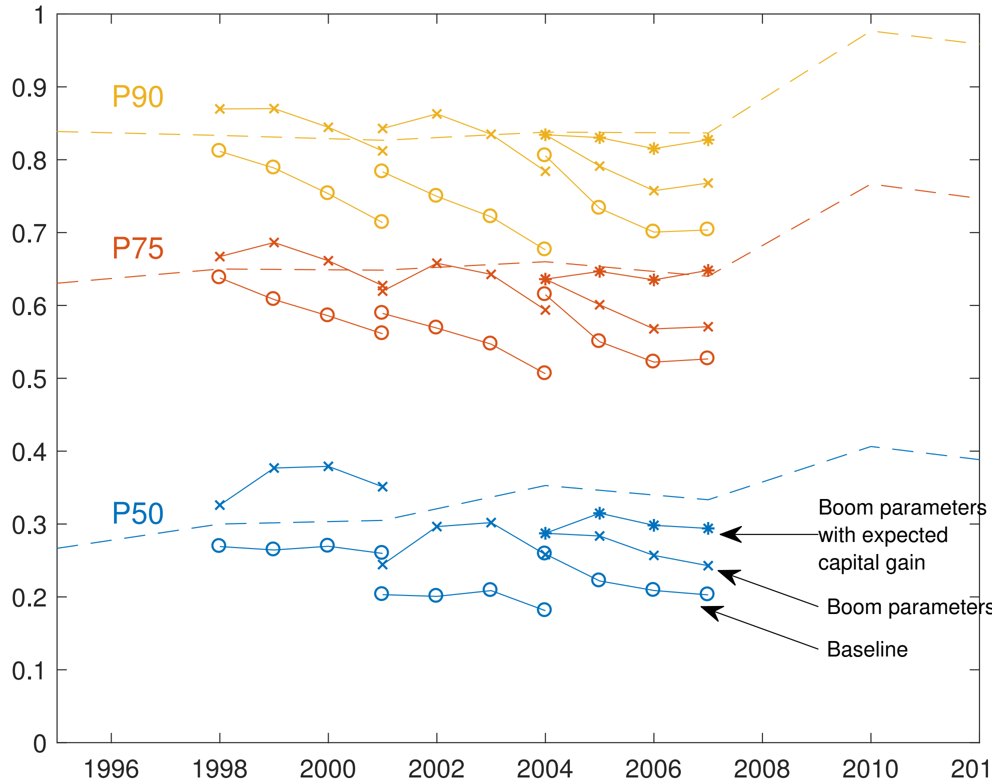


Figure A.31: Simulated LTV Distribution

Note: The figure repeats the simulation from Figure A.30 (baseline) and also shows results for the boom parameters and the boom parameters with expected capital gains for the period 2004-2007.

refinancing is that the decline in mortgage interest rates following the 2001 recession created strong incentives for refinancing that offset the transaction costs of doing so. Second, we allow for an increasing sequence of capital gains, with the one-year-ahead expected capital gain rising from 0 in 2004 to 2% in 2007 (i.e., 67, 133, and 200 basis points in 2005, 2006, and 2007). This is motivated by Kaplan, Mitman, and Violante (2019) who argue that expected capital gains are central to fitting the evolution of house prices and leverage during the 2000s boom-bust episode. Expected capital gains increase leverage for two reasons: first, homeowners feel richer and increase consumption due to a wealth elasticity and, second, the expected return on housing lowers the user cost of housing and prompts an increase in the demand for housing financed with mortgage debt.

Figure A.31 shows that these changes to the model give a fairly good account of the LTV distribution during the housing boom. As described in the main text, these changes do not significantly alter the housing wealth elasticity. Indeed, in 2007, the baseline parameters lead to a housing wealth elasticity of 0.095, the boom parameters lead to 0.098, and the boom parameters with the expected capital gain leads to 0.114.