

Housing Wealth Effects: The Long View

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

March 29, 2018

Abstract

We provide new, time-varying estimates of the the housing wealth effect back to the 1980s. We exploit systematic differential city-level exposure to regional house price cycles as an instrument for house price variation. Our main findings are that: 1) Large housing wealth effects are not new: we estimate substantial effects back to the mid 1980s; 2) There is no evidence that the elasticity to changes in house prices was particularly large in the 2000s; if anything, the elasticity was larger prior to 2000; and 3) We find no evidence of a boom-bust asymmetry. We compare these findings to the implications of a standard life-cycle model with borrowing constraints, uninsurable income risk, illiquid housing, and long-term mortgages. We show that the model can explain both the magnitude of these effects and their insensitivity to changes in the aggregate loan-to-value ratio, such as the dramatic rise in LTVs that occurred in the Great Recession. The wealth effect is insensitive to changes in LTVs both because low-LTV agents play an important role and because in a bust, the increase in constrained and sensitive agents is offset by an increase in underwater agents whose consumption is insensitive to house prices.

*Boston University, guren@bu.edu

†Boston University, Federal Reserve Bank of Minneapolis, and NBER, amckay@bu.edu

‡Columbia University and NBER, enakamura@columbia.edu

§Columbia University and NBER jsteinsson@columbia.edu

¶We would like to thank Massimiliano Cologgi, Hope Kerr, Jimmy Kuo, Jesse Silbert, Sergio Villar, and Xuyi Song for excellent research assistance. We would like to thank Aditya Aladangady, Adrien Auclert, Peter Ganong, Dan Greenwald, Erik Hurst, Virgiliu Midrigan, Pascal Noel, Chris Palmer, Jonathan Parker, Monika Piazzesi, Esteban Rossi-Hansberg, Raven Saks, 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), and 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

House wealth effects played an important role in both the boom of the early 2000s and the recession that followed (Mian and Sufi, 2011; Mian, Rao and Sufi, 2013; Mian and Sufi, 2014). We ask to what extent large housing wealth effects were a special artifact of the 2000s boom-bust cycle. It is often hypothesized that more households used their “houses as ATMs” in the 2000s than before due to automated underwriting, expanded credit, and increased access to home equity lines of credit (HELOCs). Moreover, household consumption may have been particularly responsive to house price changes in the bust because the decline in house prices pushed an unusually large number of households to high loan-to-value (LTV) ratios at which borrowing constraints bind. While there is substantial existing evidence on housing wealth effects, particularly for the boom-bust cycle of the 2000s, there is essentially no work that estimates whether housing wealth effects have changed over time using a consistent empirical methodology.¹

In this paper, we provide new, time-varying estimates on the housing wealth effect for the United States using a consistent empirical methodology going back to 1985. We then use a standard model to evaluate our findings, and in doing so elucidate the mechanisms underlying the housing wealth effect.

While national house price cycles were much smaller early in our sample than in the 2000s, there were substantial regional house price cycles. We exploit systematic differential exposure to these regional house price cycles across cities (formally, CBSAs) to identify our estimates. Our baseline measure of the housing wealth effect is the elasticity of retail employment with respect to house prices. We estimate this using a 10-year rolling window panel specification on quarterly data with annual changes.

We highlight three main empirical findings. First, large housing wealth effects are not new. We estimate large effects back to the 1980s. Second, there is no evidence that housing wealth effects were particularly large in the 2000s; if anything they were larger before 2000. Third, we find no evidence of a boom-bust asymmetry that might arise from households hitting borrowing

¹To our knowledge, two papers have looked at changes 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 (2016), and Liebersohn (2017).

constraints during housing busts. Our pooled estimate of the housing wealth effect for the sample period 1990-2015 is an elasticity of 0.060, which is roughly equivalent to a marginal propensity to consume out of housing wealth of 2.8 cents on the dollar.

To arrive at these estimates, we must confront several empirical challenges. House prices and economic activity are jointly determined and causation can run in both directions. Moreover, house prices are subject to substantial measurement error. The former concern is likely to impart an upward bias on OLS estimates of the effect of house prices on retail employment, while the latter concern will lead to a downward bias.

We overcome these empirical challenges by developing what we refer to as a “sensitivity instrument” for changes in house prices. The basic idea is to interact aggregate shocks with estimates of local exposure to these shocks, as in the case of the well-known “Bartik instrument” in labor economics. In our case, we exploit the fact that house prices in some cities are systematically more sensitive to regional house-price cycles than 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. Our instrument infers a housing wealth effect by evaluating whether consumption responds more in Providence than Rochester when house prices fluctuate in the Northeast. A natural interpretation of these different sensitivities is that they arise from current or perceived future constraints on housing supply. Our instrument is, therefore, related to the Saiz (2010) instrument used by Mian and Sufi (2014) and others. However, as we show, our instrument is much more powerful than the Saiz instrument, particularly prior to the 2000s.

Our approach requires that we estimate the sensitivity of local house prices to regional house price cycles. In doing so, it is crucial to account for “reverse causality,” i.e., the notion that local house prices may be more volatile because the local economic conditions are more volatile (as opposed to the reverse). We show that it is possible to account for reverse causality by directly including measures of local economic conditions in the estimating equation for the local house price sensitivity parameters. Hence, our instrument will not be affected by, say, differences in the industrial structure of cities that yields differences in the cyclicalities of the local economy and, in turn, induces differential sensitivity in house prices. Despite controlling for local economic conditions, our instrument turns out to be a very strong predictor of local house prices: Local and regional employment growth account for only 22% of the variation in house price growth over our sample, but including our instrument raises the R-squared to 65%. We are agnostic as to what

drives changes in regional house prices, a topic on which recent work has made substantial progress.² The power of our instrument is enhanced by the fact that house prices were less correlated with aggregate economic conditions in previous business cycles than in the Great Recession. Our panel data approach also allows us to control for city and region-time fixed effects as well as a number of time-varying controls.

Our finding that the housing wealth effect was not unusually large in the 2000s is most precisely estimated using our methodology, but it is not unique to it. OLS estimates of the housing wealth elasticity generate a similar time series pattern of estimates of the housing wealth effect as our IV strategy. The main difference between our IV estimates and OLS is that OLS estimates are somewhat larger due to an upward bias from reverse causality. One would also obtain a similar time series pattern of estimates from the Saiz instrument: the point estimates are not unusually high in the 2000s although it becomes imprecise prior to the 2000s.

We use retail employment as our main dependent variable. It is the best available proxy for consumption that is both geographically disaggregated and available for a long sample period. The BEA uses retail employment to impute local consumption in the regional NIPA accounts. Private sector datasets do the same (e.g., the “Survey of Buying Power”). Retail employment comoves strongly with the BEA’s PCE measure of consumption at the aggregate level. We analyze the relationship between local retail employment and local consumption from the Consumer Expenditure Survey, which is available for 17 cities. These turn out to comove roughly one-for-one once we account for measurement error. Retail employment is also an important component of non-tradable employment, which has been studied as a measure of local economic activity (e.g., Mian and Sufi, 2014) and is thus of interest in its own right.³

Theoretically-minded readers may find it hard to interpret a “housing wealth effect.” House prices are equilibrium variables that are affected by a myriad of shocks, many of which affect consumption through other channels. So, what do our empirical estimates capture? In Section 5, we show that in a simple general equilibrium model in which all markets are regional except for housing markets, our empirical approach yields an estimate of the partial equilibrium effect of

²The 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, and our empirical analysis is consistent with these sources of aggregate house price fluctuations (see, e.g. Landvoigt, Piazzesi, and Schneider, 2015; Favilukis, Ludvigson, and Van Nieuwerburgh, 2017; Kaplan, Mitman, and Violante, 2017).

³The existing literature has looked at the effects of house prices on various economic outcomes, including both consumption and employment. Some studies focus on particular consumption categories such as consumer packaged goods or cars (e.g., Mian and Sufi, 2011; Kaplan, Mitman, and Violante, 2016), although a few studies have used more holistic measures of consumption (e.g., Mian, Rao, and Sufi, 2013; Aladangady, 2017).

house prices on consumption. In this case, the direct effects of the shocks that drive aggregate variation in house prices and all general equilibrium effects are soaked up by the region-time fixed effects we include in our regressions. We also show that in a more realistic general equilibrium model with segmented markets across cities, our empirical approach yields an estimate of the partial equilibrium effect of house prices on consumption multiplied by a local general equilibrium multiplier. We furthermore show that this local general equilibrium multiplier can be approximated by estimates of the local fiscal multiplier (e.g., Nakamura and Steinsson, 2014).⁴ Since the recent empirical literature estimates the local general equilibrium multiplier to be somewhat larger than one, our estimate of the housing wealth effect is likely somewhat larger than the partial equilibrium effect of house prices on consumption.

Recent research has greatly advanced our understanding of the housing wealth effects in models with uninsurable income shocks, borrowing constraints, illiquid housing, and long-term mortgages (see, e.g., Agarwal et al., 2017; Berger et al., 2017; Chen, Michaux, and Roussanov, 2013; Davis and Van Nieuwerburgh, 2015; Gorea and Midrigan, 2017; Guren, Krishnamurthy, and McQuade, 2018; Kaplan, Mitman, and Violante, 2017; Li and Yao, 2007).⁵ In Sections 6 and 7, we lay out such a model — which we refer to as the “new canonical model” of housing wealth effects — and confront it with our empirical findings. We find that the model can explain both the level of the housing wealth effect and its insensitivity to the large changes in household LTV ratios observed over our sample period and in particular during the Great Recession.

Two features of the model are important to understand these theoretical results. First, incomplete markets models, such as the one we analyze, feature households that are impatient relative to the interest rate. As a result, households have substantial marginal propensities to consume out of extra wealth even when they are not near the LTV constraint. This, together with the large number of households with low LTVs, implies that a large fraction of the housing wealth effect in our model is due to households that have relatively low LTV ratios.

Second, in our model, households with negative equity are insensitive to changes in house prices. In the presence of long-term debt, underwater households are not forced to de-lever to meet an LTV constraint and are, furthermore, unable to sell their house or access any future home equity without an equity injection. They can only access their home equity after they pay their mortgage down but

⁴This formalizes intuitive arguments made by Mian and Sufi (2015).

⁵Earlier theoretical research suggested that the housing wealth effect might be zero because increased wealth from higher house prices was offset by higher implicit costs of living (Sinai and Souleles, 2005). This stark conclusion results from several simplifying assumptions including an assumption of complete markets and that household will live in the same house forever.

are typically liquidity constrained and thus discount states in which they can access their future home equity highly and are unresponsive to changes in house prices. As a consequence, the large rightward shift in the LTV distribution that resulted from the fall in prices during the 2007-2010 housing bust had two offsetting effects on the housing wealth effect. On one hand, more households were pushed closer to their LTV constraint and consequently became highly sensitive to changes in house prices. On the other hand, more households became underwater on their mortgage to the point that they became insensitive to price changes, as Ganong and Noel (2017) have emphasized. In our model, these two effects roughly offset to deliver a relatively stable elasticity despite a large rightward shift in the LTV distribution.

It may be surprising to learn that households were using their home equity to smooth consumption decades 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., Kuhn, Schularick, and Steins, 2017; Foote, Gerardi, and Willen, 2012). 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 2005.⁶

⁶See 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

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

2 Data

To estimate the housing wealth effect, we need a measure of local economic activity and a measure of house prices. For house prices, we use the Freddie Mac House Price Indices, which are a balanced panel of indices based on repeat sales for 381 CBSAs going back to 1975 (1976 is thus the first year for annual differences). Crucially, these house price indices do not impute any data from neighboring cities. Other house price indices that have comparable scope impute data for some cities for earlier time periods using data from other areas within the same region. Since our empirical approach relies on the differential response of one local area versus another in the same region, imputation would bias our estimates toward zero. A downside of the Freddie Mac data is that they make use of a combination of transaction and appraisal prices. Appraisal prices tend to be smoother than transaction prices. In the Appendix, we construct analogous results for the post-1992 period using the Federal Housing Finance Association transaction-only house price indices, and we obtain similar results.

Our main measure of local economic activity is retail employment per capita. Retail employment comes from the Quarterly Census of Employment and Wages (QCEW) which is available back to 1975 at the county level. We aggregate these data to the CBSA level create retail employment for 380 CBSAs.⁷ The QCEW infrequently has missing data for a county-quarter to protect the confidentiality of a dominant employer in a given industry. This almost exclusively occurs for very small counties within a CBSA. In our baseline specification we only use counties within each CBSA with no missing data. We also remove large and discontinuous jumps and changes due to county realignments from the data set. In the Appendix, we show that our results are robust to our cleaning of the data. We similarly create a measure of manufacturing employment per capita.

The QCEW data are available for SIC 1975-2000 and NAICS 1990-2016. The definition 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.

⁷We drop Dover, DE from our analysis because retail employment data is missing for the entire CBSA for a majority of years.

retail employment is different in NAICS than in SIC, primarily because NAICS split off wholesale employment into a separate sector. To merge the two series, we splice together log changes in 1993. In the Appendix, we show that the exact choice of splice date makes little difference because for the overlapping period, the two series are very similar in log changes even though they differ in levels. We supplement the QCEW data with annual data from the County Business Patterns (CBP). While the QCEW has very few counties with missing or censored data for large industries such as manufacturing and retail, there is enough missing and censored data in other industries that it is difficult to construct industry shares for use as a control. Consequently, we use CBP data for additional controls.⁸ Because there is not an overlap between the NAICS and SIC for the CBP, we harmonize the data to 2-digit SIC codes following Acemoglu et al. (2016) and create an annual series of industry shares at the city level. We then merge this into the quarterly QCEW data by interpolating. In the Appendix, we use the CBP data and an annual specification rather than the QCEW data and the results are little changed.

The population data we use to construct per-capital retail employment come from the Census Bureau’s post-Censal population estimates for 1970 to 2010 and inter-Censal population estimates for 2010 to 2015. These estimates are available annually, and we interpolate to a quarterly frequency. We use the same counties to calculate population that we use to calculate sectoral employment.

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 retail employment in the two cities. We do this using the following empirical specification:

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

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. We allow for CBSA fixed effects, ψ_i , which control for long-term trends in each CBSA, region-time fixed effects, $\xi_{r,t}$, which imply that our effects are identified only off of differential movements across

⁸The CBP also censors data but instead of simply omitting data, it provides a range. It is thus more suitable to construct a consistent series of 2-digit industry shares.

CBSAs a set of additional controls, $X_{i,r,t}$, and idiosyncratic shocks to retail employment, $\varepsilon_{i,r,t}$.

The coefficient of interest in equation (1) is β , which measures the housing wealth effect as an elasticity. Several challenges arise in estimating β . Causation runs both ways between local employment and house prices — implying that the error term in equation (1) may be correlated with the changes in house prices. This is likely to bias OLS estimates of β upward since a strong economy will cause house prices to rise. On the other hand, house prices are measured with error, potentially biasing β towards zero. To address these two sources of bias, we propose an approach based on variation in the sensitivity of house prices to aggregate shocks across CBSAs.

3.1 Simple Intuition for Identification

Before developing our identification strategy in detail, it is useful to consider an example. Figure 1 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 our identification strategy. 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 effect before 2000 when strategies using nation-wide variation 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 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. A likely reason for this differential sensitivity is variation in current and future expected housing supply constraints. We discuss this in more detail below.

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

$$\Delta p_{i,r,t} = \varphi_i + \zeta_{r,t} + \gamma_i \Delta P_{r,t} + \nu_{i,r,t}, \quad (2)$$




Figure 1: 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 1976 to 2015.

where $\Delta P_{r,t}$ denotes the log annual change in regional house prices and γ_i is a city-specific coefficient.⁹ Then use $z_{i,r,t} = \hat{\gamma}_i \Delta P_{r,t}$ as an instrument for $\Delta p_{i,r,t}$ in equation (1), where $\hat{\gamma}_i$ denotes the estimate of γ_i from equation (2). The intuition behind this instrument is essentially the same as for a difference-in-difference design: When house prices rise in the Northeast, they systematically rise 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 (2). The question, then, is whether this differential treatment translates into differential growth in retail employment. This empirical strategy is similar to 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.

⁹To keep our notation simple, we denote $\sum_{i \in I} \gamma_i \Delta P_{r,t}$, that is separate city-specific coefficients for each city i in the set of cities I , by $\gamma_i \Delta P_{r,t}$. We follow this simplified notation throughout the paper.

3.2 Refined Identification Strategy

The simple procedure described above runs into problems if retail employment responds differentially to regional shocks through other channels than local house prices. 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.¹⁰ This, in turn, would lead to biased estimates of β .

To address this problem, we refine the procedure described above by controlling for the local and regional changes in retail employment (allowing the coefficients on these variables to vary across CBSAs) when estimating γ_i :

$$\Delta p_{i,r,t} = \varphi_i + \delta_i \Delta y_{i,r,t} + \mu_i \Delta Y_{r,t} + \gamma_i \Delta P_{r,t} + \Psi X_{i,r,t} + \nu_{i,r,t}. \quad (3)$$

As before, $z_{i,r,t} = \hat{\gamma}_i \Delta P_{r,t}$ is the instrument we propose to use for $\Delta p_{i,r,t}$ in equation (1). This instrument captures the portion of local house price variation that is explained by differential sensitivity to the regional house price index holding $\Delta y_{i,r,t}$ and $\Delta Y_{r,t}$ fixed.

Our final specification allows for additional controls $X_{i,r,t}$ when estimating γ_i in equation (3). We exclude the CBSA in question from the construction of the regional house price index when running this regression, so as to avoid bias in γ_i due to the same price being on both the left and right hand side.¹¹ Finally, we estimate equation (3) using time periods other than the time period for which we are estimating equation (1). We do this to avoid $\hat{\gamma}_i$ reflecting endogenous variation in local house prices over the period we are estimating equation (1). For the rolling window estimates in section 4, we leave out all time periods outside a given 10-year window in estimating our rolling window coefficients. In the pooled estimates across time periods, we construct the instrument for each year using data excluding a 3 year buffer around the time period in question. In practice, these different leave-out procedures yield similar results for most time periods.

For this approach to yield a powerful instrument, there must be substantial variation in house

¹⁰Suppose, for example, that Providence has an industrial structure tilted towards highly cyclical durable goods relative to Rochester. In this case, a positive aggregate demand shock would lead retail employment to increase more in Providence than Rochester. If local economic booms raise house prices, this would induce a larger change in house prices in Providence than Rochester and, thus, imply that we would estimate a higher γ_i for Providence using equation (2) purely due to reverse causality.

¹¹There is an arithmetic reason not to include region-time fixed effects in equation 3 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.

prices that is orthogonal to movements in local and regional retail employment. This is the case in our data: when we run regression (3) without the differential sensitivity term $\gamma_i \Delta P_{r,t}$, the R-squared is 0.22, but when $\gamma_i \Delta P_{r,t}$ is added, the R-squared rises to 0.65. In other words, our sensitivity instrument explains a large fraction of the total variation in local house prices, even conditioning on local and regional employment.¹²

We include a variety of controls $X_{i,r,t}$ in equations (1) and (3) in our baseline analysis beyond region-time and city-level fixed effects. We control for local industry shares with separate coefficients for each time period. This accounts for differential labor demand effects, as is captured by the original “Bartik” instrument. It also accounts for differential city-level exposure to unobservable risk premia associated with industrial structure. For example, if some cities have more risky industries than others and are therefore differentially affected by shocks to risk aversion, this control would capture this factor. Second, we include separate controls for the differential city-level exposure to regional retail employment, real 30-year mortgage rates, and Gilchrist and Zakrajek’s (2012) measure of bond risk premia. For each of these, we construct a control in an analogous fashion to our instrument by estimating an OLS regression:

$$\Delta y_{i,r,t} = \psi_i + \xi_{r,t} + \alpha_i \Delta X_{r,t} + \varepsilon_{i,r,t}, \quad (4)$$

where ΔX is either the log change in regional retail employment, the change in the 30-year fixed mortgage rate, or the change of the Gilchrist-Zakrajek excess bond premium. We then include $\hat{\alpha}_i \Delta X_{r,\tau}$ as a control.¹³ The ability to control for differential sensitivity of local retail employment to observable aggregates is a key advantage of our panel-data methodology. Finally, in equation (3) only, we control for changes in average wages as reported in the QCEW with CBSA-specific coefficients.

Our sensitivity instrument is a close cousin of the Bartik instrument, which instruments for city

¹²One potential concern with this procedure is the role of measurement error in $\Delta y_{i,r,t}$ biasing the δ_i terms and thereby creating bias in the γ_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 δ_i is the same across CSBAs, but the δ we obtain is a causal elasticity. We obtain an estimate for δ of 2.9. This estimate for δ can be used to subtract $\delta \Delta y_{i,r,t}$ from $\Delta p_{i,r,t}$, and then we can use this adjusted $\Delta p_{i,r,t}$ to estimate γ_i . This approach yields values for the γ_i that are highly correlated with our baseline approach, and using these alternate γ_i s does not significantly alter our results.

¹³We use this approach of estimating the sensitivity of retail employment to the controls on the “leave-out sample” to avoid overfitting concerns. However, we have also tried the more direct approach of including $\alpha_i \Delta X_{r,t}$ as controls in equation (1). Doing so for the 30-year mortgage rate or the Gilchrist-Zakrajek excess bond premium yields essentially the same results with slightly larger standard errors. Doing so for retail employment yields similar results starting in the early 1990s and highly imprecise results with lower point estimates in the 1980s.

labor demand by summing across industries the share of an industry in each city multiplied by the national change in employment in that industry, which is, in turn, a close cousin of difference-in-difference designs. The crucial idea in all of these strategies is that certain locations are differentially treated by an aggregate shock. In thinking about the validity of these strategies, it is important to understand that treatment intensity (in our case the estimated γ_i and in the case of the Bartik instrument the industry shares) need not be randomly assigned. This is in fact rarely the case. The key identifying assumption 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 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 correlated with regional house prices in the time series and α_i is correlated with $\hat{\gamma}_i$ in the cross section. Appendix B presents a more formal discussion of our identifying assumptions in the context of a two-equation simultaneous equations system from which we explicitly derive our estimating equations.

The analogy to the Bartik instrument is instructive in elucidating the intuition. Consider a Bartik instrument in which the key source of variation is differential exposure to oil shocks in Texas versus Florida. Texas and Florida obviously differ in other ways than just their exposure to oil shocks. This does not in and of itself invalidate the Bartik instrument. The key identifying assumption is that there is not some other factor that happens to differentially affect Texas at the same time as oil price go up.

What drives the heterogeneity in γ_i ? The simplest interpretation of the heterogeneity in γ_i is cross-sectional variation in the slope of current or perceived future housing supply curves across CBSAs. This is the source of variation emphasized by Mian and Sufi (2014) in their use of the Saiz (2010) instrument. Intuitively, suppose a region is hit by a shock that affects the demand for housing. This shock will have different effects on house prices in different CBSAs if the local housing supply curves have different slopes across CBSAs in the short run or in the longer run.¹⁴ One advantage of our instrument versus the Saiz instrument is that there are likely to be many

¹⁴In the Appendix, we discuss a model in which cities differ in their in medium-run elasticities, as opposed to short-run elasticities which we assume are zero for all cities. In this case, greater sensitivity of prices to a common regional shock arises from differences in medium-run housing supply elasticities. The reason is that house prices will be expected to revert back to normal faster in cities with higher medium-run housing supply elasticities. The expected capital loss on housing in these cities will temper the initial response of housing to the shock. These effects are hard to measure directly, not the least because house price fluctuations sometimes mean revert, implying that we never really see the “long-run.” Haughwout et al. (2013) provide some evidence that deviations in housing supply from a population trend during the early 2000s were indeed correlated with Saiz’s measure of the housing supply elasticity.

sources of variation in housing supply elasticities beyond those based on physical geography. These include land use regulation (Saiz, 2010) and future housing supply constraints (Nathanson and Zwick, 2017). Some regions may also be more “bubbly,” perhaps due to social connections to inelastic cities (Bailey et al., 2017) or credit (Favara and Imbs, 2015). Our instrument will capture all these sources of variation. This implies that it is substantially more powerful than the Saiz instrument. Also, our instrument can be calculated for any geographical area, while the Saiz instrument is available only for 269 Metropolitan Statistical Areas.

4 Empirical Estimates of Housing Wealth Elasticity

Figure 2 presents 10-year rolling window estimates of our measure of the housing wealth effect, the elasticity β in equation (1). Each point on the figure indicates the elasticity for a 10-year sample period starting in the quarter stated on the horizontal axis (e.g., the point for quarter 2005q1 is the estimate for the sample period 2005q1-2015q1). We start the figure in 1985 because the standard errors for our estimate become very large prior to that point, but we use data back to 1976 in creating our instrument.

Our estimates are calculated using the methodology described in section 3. We calculate a CBSA fixed effect once for the entire sample period and apply it to all 10-year windows rather than calculating a different CBSA fixed effect for each 10-year window. This avoids time variation in these fixed effects driving time variation in our coefficient of interest. Our baseline standard errors are constructed using two-way clustering by CBSA and region-time to allow for arbitrary time series correlations for a given CBSA and for correlations across CBSAs within a region at a particular time. The standard errors do not account for sampling error associated with the generated instrument, and in the Appendix, we consider alternative bootstrap standard errors that account for sampling variation in the instrument.

The 2000s do not exhibit particularly large housing wealth effects: if anything, the effects have declined over time since the 1990s. This time pattern of effects is insensitive to whether or not we include controls in the regression, though the level is slightly lower including controls. Given the large standard errors for the pre-1990 estimates (which often cannot be statistically distinguished from zero), we focus on the post-1990 period in constructing a pooled estimate of housing wealth effects. For the period 1990-2015, the pooled elasticity is 0.060, with a standard error of 0.014, as indicated in specification (1) in Table 1. In other words, a 10% decline in house prices in CBSA i




Figure 2: 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 house prices at the CBSA level for rolling 10-year sample periods. Each point indicates the elasticity for a 10-year sample period that begins in the quarter stated 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 region-time.

relative to other CBSA's leads to with roughly a 0.6% decline in retail employment. This pooled estimate is equivalent to a marginal propensity to consume out of housing wealth of 2.76 cents on the dollar.¹⁵

Figure 3 presents binned scatter plots for the “first stage” and “reduced form” of our instrumental variables specification for the 1990-2015 pooled sample. These plots show a tight first stage relationship and clear reduced form relationship in both sample periods.

To evaluate whether housing wealth effects are particularly potent in housing busts — perhaps due to debt-deleveraging — we consider non-linear regression specifications in Table 1. Specification (2) in Table 1 includes separate coefficients for positive and negative house price changes, while

¹⁵To convert our elasticity to a marginal propensity to consume out of housing wealth requires dividing the elasticity of consumption to house prices by the ratio of housing wealth to consumption. The average ratio of H/C over 1985 to 2016 where H is measured as the market value of owner-occupied real estate from the Flow of Funds and C is measured as total personal consumption expenditures less PCE on housing services and utilities, is 2.17. Hence, we obtain the marginal propensity to consume out of housing wealth as $0.060/2.17 = 2.76$ cents for each additional dollar of housing wealth.




Figure 3: 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 house prices at the CBSA level for the pooled 1990-2015 sample. These correspond to the same regressions as in Table 1 for different sample periods. For these estimates, we first construct our instrument for each quarter by estimating the γ_i 's in equation (3) for each quarter, leaving out a three-year buffer around the quarter in question. We then estimate equation (1) pooling across all years in the sample in question.

specification (3) includes a quadratic term in house price changes. We find no evidence of a boom-bust asymmetry in house price elasticities. In specification (2), the coefficient on positive house price changes is somewhat larger, but we cannot reject equality of these coefficients. In specification (3) the quadratic term is both statistically insignificant and quantitatively small.

How do these results compare to OLS and the Saiz instrument? Figure 4 presents 10-year rolling window estimates from OLS. The lower elasticity in the 2000s and the general time-series pattern in the elasticity found with our instrument are also clearly evident in OLS. OLS does, however, yield slightly higher elasticity estimates than our IV approach in most periods. This suggests that our instrument corrects for reverse causality of house price changes that might otherwise bias upward estimates of the housing wealth elasticity and that this source of bias outweighs the countervailing influence of measurement error.

Figure 5 compares the IV estimates based on our sensitivity instrument to estimates using Saiz's housing supply elasticity as an instrument.¹⁶ The Saiz instrument yields a broadly similar pattern of declining elasticities over time with slightly higher elasticity estimates for most of the sample period. Given that we expect our instrument to reflect, to a substantial extent, variation in current or future expected housing supply elasticities, it is perhaps not surprising that our estimates of γ_i over the full sample are correlated with Saiz's measures of the housing supply elasticity and land unavailability, as well as the Wharton Land Use Regulation Index. We find that the R-squared of

¹⁶This instrument in Figure 5 interacts Saiz's estimated housing supply elasticity with the national change in house prices. In the Appendix, we experiment with a regional version of this instrument.

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

	(1)	(2)	(3)
$\Delta \log (P)$	0.060*** (0.014)		
$\Delta \log (P) -$		0.052*** (0.019)	
$\Delta \log (P) +$		0.077*** (0.019)	
P Test for Equality		0.340	
$\Delta \log (P)$			0.066*** (0.015)
$\Delta \log (P)^2$			0.035 (0.040)

Note: For these estimates, we first construct our instrument for each quarter by estimating the γ_i 's in equation (3) 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 region-time and CBSA level.

univariate regressions is 16 % with the elasticity only, 14% with unavailability only, 15% with land use regulations only and 19% with all three together. However, the Saiz instrument is also much less powerful than our instrument, leading to larger standard errors. Prior to the mid-1990s, the Saiz instrument has sufficiently large standard errors that it is essentially uninformative.

In contrast to the substantial effects we estimate of house prices on retail employment, we find no effect on manufacturing employment. Figure 6 shows the results. Our point estimates are close to zero for most of the sample period, although the estimates are fairly imprecise. This is consistent with our interpretation that the effects on retail employment we observe are driven by the effects of a housing wealth effect, which one would expect to affect local spending, but not the 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 through their effect on local demand.¹⁷

How do our estimates of the housing wealth elasticity compare to others in the literature? For the period 1990-2016, our estimate of the pooled elasticity is 0.060, with a standard error of

¹⁷Mian and Sufi (2014) use "non-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 very difficult to create a consistent time series of non-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.




Figure 4: Housing Wealth Effect: Sensitivity Instrument vs. OLS

Note: The red dashed line plots the point estimates of the housing wealth effect based on 10-year rolling windows using our instrument (same as in Figure 2). The light red dashed lines plot the upper and lower bounds of 95% confidence intervals for these estimates. The dark blue line plots the point estimates of the housing wealth effect estimated using OLS with the same controls as our baseline IV specification. The lighter blue lines plot the upper and lower bounds of 95% confidence intervals. Standard errors are two-way clustered by region-time and CBSA.

0.014. This is equivalent to a marginal propensity to consume out of housing wealth of 2.76 cents on the dollar, as we describe above. This is at the low end of the range that has been reported in the literature. Aladangady (2017) estimates a dollar-for-dollar MPCH of 4.7 cents for owners and zero for renters using an instrument constructed by interacting the Saiz instrument with the real interest rate. Multiplying by a long-run homeownership rate of 65 percent implies an overall MPCH of roughly 3.1 cents. Mian, Rao, and Sufi (2013) estimate an MPCH out of housing wealth of 5.4 cents for total consumption, but 2.3 cents of that comes from automobiles and only 1.6 cents for non-durable goods. 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 implies a range for the elasticity of retail employment to house prices of between 0.09 and 0.16.¹⁸ We provide a more extensive discussion of the broader literature in the Appendix.

¹⁸To convert Mian and Sufi's elasticity with respect to total net worth to a housing wealth elasticity, we must divide by the ratio of housing net worth to total net worth, which is between three and four (Berger et al., 2017).




Figure 5: Housing Wealth Effect: Sensitivity Instrument vs. Saiz Instrument

Note: The red dashed line plots the point estimates of the housing wealth effect based on 10-year rolling windows using our instrument (same as in Figure 2). The light red dashed lines plot the upper and lower bounds of 95% confidence intervals for these estimates. The dark blue line plots the point estimates of the housing wealth effect based on 10-year rolling windows estimated using an instrument based on the estimated housing supply elasticity of Saiz (2010) interacted with the national annual log change in house prices. The lighter blue lines plot the upper and lower bounds of 95% confidence intervals for the estimates based on the Saiz instrument. The lighter blue lines plot the upper and lower bounds of 95% confidence intervals. Standard errors are two-way clustered by region-time and CBSA.

A recent literature on housing wealth effects has estimated the elasticity using a pure cross-section specification for 3-year growth rates from 2006 to 2009 (e.g., Mian, Rao, and Sufi, 2013). To provide an apples-to-apples comparison between our specification and those that have been analyzed in this literature, Table 2 presents results for several variants of this “long difference” specification. All of these versions include the full set of controls and region fixed effects that we include in our baseline specification.

The first row of Table 2 shows that the long-difference approach for 2006-2009 yields an estimate of the elasticity of retail employment per capita to house prices to our baseline specification of 0.055, which is just below our pooled post-1990 estimate of 0.060. In this specification, we demean the left and right-hand side variables as well as all controls by CBSA over the entire 1976-2015 sample before running the regression. This is similar to our baseline specification with CBSA fixed effects



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 house prices at the CBSA level for rolling 10-year sample periods. Each point indicates the elasticity for a 10-year sample period that begins in the quarter stated 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 region-time.

and removes differential long-run average growth rates across CBSAs for the 1976-2015 period. The existing literature has not used this specification since it has typically focused on data for the 2006-2009 period without a long panel. The second row (and subsequent rows) shows that if we do not remove CBSA-level averages, the estimated coefficient is 0.092. This suggests that it is important to account for long-run differences in growth rates across CBSAs in calculating the elasticity, highlighting a virtue of our panel data approach. The existing literature also typically uses employment rather than employment per capita. If we do not adjust for population or remove CBSA-level averages, the elasticity rises from 0.092 to 0.120. Finally, the existing literature for this period typically uses the Saiz instrument. To be comparable with the sample of CBSAs for which one can use the Saiz instrument, the fourth row of Table 2 limits the sample to the cities for which this instrument is available. This raises the elasticity slightly, to 0.134. The Saiz instrument yields a slightly higher elasticity of 0.167, although the difference is not statistically significant.

Table 2: Comparison of Estimation Approaches for 2006-2009

Specification	2006-2009 Elasticity	
Baseline Instrument (Per Capita), CBSA FE	0.055**	(0.020)
Baseline Instrument (Per Capita)	0.092***	(0.020)
Baseline Instrument	0.120***	(0.021)
Baseline Instrument, Saiz Sample	0.134***	(0.026)
Saiz Elasticity Instrument	0.167**	(0.061)
OLS	0.119***	(0.013)

Note: This table compares our regional sensitivity instrument to the Saiz Instrument and OLS for the 2006 to 2009 long difference. For the sensitivity instrument, we first construct our instrument for the three-year window estimating the γ_i 's in equation (6), leaving out a three-year buffer around the quarter in question. We then estimate $\Delta y_{i,r,t} = \xi_r + \beta \Delta p_{i,r,t} + \Gamma X_{i,r,t} + \varepsilon_{i,r,t}$, where $X_{i,r,t}$ 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 1976-2015 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 with the cyclical sensitivity control instrumenting with the elasticity rather than our sensitivity instrument. OLS runs the second-stage regression by OLS with the same controls but without taking out CBSA fixed effects or using per-capita variables. Standard errors are two-way clustered at the region-time and CBSA level for all specifications.

As described above, the standard error is much larger. As with our main results, our instrument typically generates slightly lower estimates of housing wealth elasticities than OLS, while the Saiz instrument generates slightly higher estimates than OLS.

4.1 Retail Employment vs. Consumption

In principle, there are a variety of measures of local economic activity that one could consider, and the existing literature has analyzed the effects of house price changes on both consumption and employment for the case of the boom-bust cycle in the 2000s. Moreover, production-based estimates of economic activity, such as retail employment, are typically viewed as higher quality than those based on household surveys. For example, aggregate consumption based on the consumer expenditure survey has displayed implausible negative growth rates in recent years in contrast to production based estimates such as employment, GDP, or PCE.

Employment data such as the QCEW provide the only high quality measures of local economic activity at a high frequency available at the CBSA level that we are aware of going back over a substantial time period.¹⁹ Moreover, retail employment has long been viewed by measurement

¹⁹Some of the recent literature has made use of specialized consumption series that are available at a regionally disaggregated level (e.g., from Nielsen or car purchases). However, it is evident from aggregate data that e.g., car purchases have quite different dynamics from non-durable consumption. Broader based data on consumer purchases is available only for the very recent period, and even these data typically focus only on goods as opposed to services.

agencies as one of the best local indicators of local consumption. Intuitively, retail employees are a crucial input for households to be able to consume. As a consequence, retail employment is often used to impute consumption in government and private sector attempts to construct regional consumption measures. For example, BEA’s Regional PCE measures and the private sector “Survey of Buying Power” both use QCEW retail employment data to impute consumption in between economic census years.²⁰

We present an analysis of the relationship between BEA non-durable consumption and retail employment in Appendix D. At an aggregate level, retail employment comoves strongly with real personal consumption expenditures in the time series. Since our analysis allows for time trends, it will not be affected by long-term trends such as shifts in retailing toward online retailers or big-box stores.

We then study 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. We use an instrumental variables approach to correct for measurement error in retail employment per capita. Once we address measurement error, consumer expenditures respond nearly one-for-one with to retail employment per capita.²¹

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. In contrast, at the aggregate level or city level, house prices are an endogenous variable. House prices are affected by a myriad of shocks and these shocks may affect consumption not only through house prices, but also directly or through other channels. So what does it mean to estimate the causal effect of house prices on consumption at the city level?

Consider a simple model of an economy consisting of several regions with many cities in each region. Suppose housing markets are local to each city and the cities differ in their housing supply elasticities. All other markets are fully integrated across cities within a region (and may in some cases be integrated across regions). The cities are initially in identical steady states before being hit

²⁰Another commonly used data series is Moody’s Economy.com measure of retail sales, which is constructed from state sales tax data and national retail-sales data from the Census Bureau and benchmarked every 5 years using data from the Census of Retail Trade. Although this does not use retail employment, it is only available at the state level, and sales tax data is notoriously noisy.

²¹We have also verified, in unreported work, that changes in CBSA-level retail employment are highly correlated with changes in retail sales over the 5-year intervals at which retail sales are available in the Economic Census.

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, and 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.

Take a log-linear approximation to the aggregate consumption function around the initial steady state and then take an annual difference. This yields:

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

where $c_{i,r,t}$ denotes the logarithm of consumption and ϕ_x denotes the elasticity of $c(\cdot \cdot \cdot)$ 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 β captures the response of consumption to a house price change holding these other variables constant. This shows that if we are able to identify variation in local house prices that is orthogonal to the error terms $\xi_{r,t}$ and $\varepsilon_{i,r,t}$ and the assumptions stated above about market structure hold the coefficient β will estimate the partial equilibrium effect of house prices on consumption.²²

²²If non-linearities are important, the fixed effects in equation (5) will not fully absorb the general equilibrium

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. In other words, the differential house price movements will result in local general equilibrium effects. Since these local general equilibrium effects will differ across cities within a region, they will not be absorbed by the region-time fixed effects in our empirical specification and will affect our estimate of β .

Local general equilibrium effects result from changes in local demand affecting local wages, prices, and incomes. This suggests that evidence from other local demand shocks might be useful in pinning down the effect of local general equilibrium on our empirical estimates. In an Online Appendix, 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 as a measure of local general equilibrium effects. More specifically, we show that the housing wealth effect estimate β that results from our empirical specification can be expressed as:

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

where β_{LFM} denotes the local fiscal multiplier and β_{PE} denotes the partial equilibrium effect of house prices on consumption.²³ Intuitively, a dollar of spending triggers the same local general equilibrium response regardless of whether it arises from the housing wealth effect or government spending. Nakamura and Steinsson (2014) estimate that the local government spending multiplier is roughly 1.5.

An additional consideration of local general equilibrium effects is that an increase in house prices might stimulate an increase in income in the construction sector. Part of the consumption response to home price changes might therefore reflect the consumption response to this change in income. In the Online Appendix, we discuss how an estimate of the elasticity of construction employment to home prices can be used to assess the strength of this channel. After removing the local income

price effects. For example, if consumption growth responds importantly to $\Delta p_{i,r,t} \times \Delta \Omega_{r,t}$ or to $\Delta p_{i,r,t} \times \Delta R_{r,t}$, then our estimated β will reflect these interactions in addition to the housing wealth effect. In the next section we present a fully non-linear model of the housing wealth effect and we show in Appendix F that the model implies these interaction effects are small. In particular, the housing wealth effect is close to linear in the magnitude of the price change and symmetric with respect to positive and negative price changes.

²³We make certain simplifying assumptions to derive this result. One of these is to assume GHH preferences to avoid wealth effects on labor supply. We also abstract from the collateral channel emphasized by Chaney, Sraer, and Thesmar (2012) and Adelino, Schoar, and Severino (2015). Finally, we assume that the government and households both buy the same consumption good.

multiplier and the construction channel, the partial equilibrium response of consumption to the change in house prices is about half of our estimated housing wealth effect.

6 Theory: Local Consumption Response to House Prices

We now present our version of the new canonical 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 the Appendix.

6.1 Assumptions

Households live for T periods and have preferences for 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 the wealth of offspring.

We parameterize household preferences as:

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

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

Here γ captures the curvature of the utility function, ε is the elasticity of substitution between housing and non-durable consumption, B_0 captures the strength of the warm-glow bequest motive, and B_1 captures non-homotheticity in bequest motives.²⁴

An individual can consume housing either by owning or renting. A unit of housing can be purchased at price p and rented for one period at cost δp , with a fixed rent-price ratio given by the parameter δ . In our baseline model, people expect home prices will remain at their current level. Renting h units of housing delivers the same utility as buying that amount of housing, but the rent is more expensive than the user cost of owner occupied housing, which makes owning attractive despite its associated transaction costs. We consider alternative assumptions about the behavior

²⁴In the presence of illiquid durable goods such as housing, the parameter γ is related to, but not equivalent to, the coefficient of risk aversion (see Flavin and Nakagawa, 2008).

of rents in the Appendix.²⁵ To sell a house the individual must pay ψ^{Sell} of the value of the house in a transaction cost and to buy a house the individual must pay ψ^{Buy} .

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

$$m' \leq \theta p h', \quad (6)$$

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

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

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

6.2 Calibration

A household is born at age 25, works for 36 years retiring at 61, and dies deterministically after age 80. We set most of the parameters through external calibration, which we describe first, and then we set a small set of parameters through internal calibration. We set the curvature of the utility function, γ , to 2. We set the elasticity of substitution between housing and non-durable consumption to 1.25 based on the estimates of Piazzesi, Schneider, and Tuzel (2007).

We set the LTV limit, θ , to 0.80 based on GSE guidelines for conforming mortgages without private mortgage insurance. We set the after-tax, real interest rate on mortgage debt to 3 percent per year based on the long-run averages of nominal mortgage rates and inflation.²⁶ We set the

²⁵In an environment without anticipated capital gains or losses, and if depreciation, taxes, and insurance premia are proportional to the home value, then the user cost will be a constant fraction of the home value. We calibrate the rental cost to exceed the user cost of owner-occupied housing. One interpretation is that the depreciation rate of rental property is higher due to moral hazard.

²⁶Between 1971 and 2017 the average CPI inflation rate was 4.1 percent, the average 30-year fixed rate mortgage

real return on liquid assets to 1 percent based on the difference between the long-run averages of the 1-year Treasury rate and inflation. We set the cost of buying a house to 2 percent, which reflects some of the closing costs associated with a purchase. Furthermore, a home-buyer taking a mortgage also pays the fixed cost of obtaining a new mortgage, which could capture time and other non-pecuniary costs of the transaction.

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. (2016). We conceive of the transitory income shocks as non-employment shocks motivated by the income process in Guvenen et al. (2016). 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 and the fraction of the year 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 last 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., 2016). 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 through internal calibration. The parameters are the discount factor, β ; the strength of the preference for housing, ω ; the strength of the bequest motive, B_0 ; the degree to which a bequest is a luxury, B_1 ; the rent-price ratio, δ ; the mortgage origination cost, ψ^m ; and the transaction cost for selling a house, ψ^{Sell} . Our target moments are the life-cycle profiles of home value to income for working-age homeowners (we target the 25th, 50th, 75th, and 90th percentiles of the distribution at each age), mortgage LTV (P25, P50, P75, P90 by age), liquid assets (P25, P50, P75, P90 by age), and the homeownership rate by age. These empirical moments are calculated from the 2001 SCF. 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 2018; Bhutta and Keys 2016) but some refinancing activity results from changes in interest rates, which are not

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.

Table 3: Internally Calibrated Parameter Values

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

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 no difference in interest rates and compute the fraction of mortgages that are refinanced each year. Finally, 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 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, these are 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. The Appendix explains our empirical moments and objective function in more detail, and the resulting parameter values appear in Table 3.

7 Model Simulations

7.1 Defining Our Elasticity “Experiment”

Let us begin by defining the object in the model that corresponds to the IV estimates of the housing wealth elasticity we presented earlier in the paper. In the model, we can write aggregate consumption in city i as:

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

where $c(\dots)$ is the consumption function of an individual, p_i is the price of a unit of housing in city i , x is a vector of idiosyncratic state variables, and Φ is the distribution of households over idiosyncratic states. The idiosyncratic states in the model are units of housing, mortgage debt, liquid assets, income, and age. Our object of interest is the elasticity of C_i with respect to p_i , which will depend on Φ .

We take Φ directly from the SCF data for years 1983, 1986, ..., 2016. This allows us to assess how the empirically observed changes in household balance sheets, formally represented by Φ , have

changed the housing wealth effect. We obtain empirical estimates of Φ from the SCF, as we describe in the Appendix. 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. 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 feed in the values of the state variables described above for each wave of the SCF, which occurs every three years. Over the last thirty years, household balance sheets have changed substantially. The 75th percentile of the LTV distribution rose from about 0.4 in 1983 to nearly 0.8 in 2010, and spiked during the Great Recession due to the decline in home values. Moreover, the median home value relative to median income rose from a ratio of about two times income in 1983 to nearly four times income in 2007. We calculate the local elasticity of consumption with respect to house prices described above, given these changes in state variables. Panel (a) of Figure 7 shows our baseline results.²⁷ The model generates a smooth local consumption elasticity, despite these large changes in the distribution of LTV's and home values.

7.2 Why So Stable?

Why don't these large changes in the distribution of individual states lead to larger variation in the housing wealth elasticity? To unpack our result, we alter one variable at a time. We begin by altering LTV values to reflect the marginal distribution in year t while keeping all of the other state variables at their 2001 level. To do so, we start with the 2001 SCF data and for each wave of the SCF 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 and alter the LTV according to the marginal distribution of LTV in year t , holding all other state variables fixed. We construct the house value counterfactual in the same manner using the marginal distribution of house values in year t .

The decomposition of these effects presented in panel (b) of Figure 7. The dash-dot line shows the effect of changing LTV's alone, and the dashed line shows the contribution of changing home values alone. Evidently, changes in the LTV distribution have very limited effects on the housing wealth elasticity through the lens of the new canonical model. The LTV counterfactual increases

²⁷We compute the elasticity by averaging together the responses to a 10% positive and 10% negative shock. We consider larger shocks in Appendix Figure 28 and find 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 that there is no such asymmetry.




Figure 7: Model Housing Wealth Effect

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 Φ_t is constructed from the SCF data for year t as explained in Appendix F.4. 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 Φ_t 's constructed as described in the text.

Table 4: Decomposition of Housing Wealth Effect

	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	-0.75	0.05
Downsizers	0.30	0.04	0.49	0.07	0.48	0.04
Stayers						
LTV ≤ 0.6	0.17	0.54	0.14	0.43	0.17	0.37
LTV $\in (0.6, 0.8]$	0.26	0.05	0.21	0.09	0.27	0.08
LTV $\in (0.8, 1.0]$	0.26	0.01	0.33	0.07	0.39	0.09
LTV ≥ 1.0	0.02	0.00	0.15	0.01	0.11	0.05
Total	0.085	1.00	0.107	1.00	0.112	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.

only modestly between 1983 and 2012 during the “great leveraging,” and it does not show any apparent spike in the 2000s despite the substantial boom and bust. This differs dramatically from the common narrative that the housing wealth effect rose in the Great Recession as LTVs rose sharply preventing households from accessing home equity through cash-out refinancing. It is crucial to emphasize that the results are for the *elasticity* of consumption with respect to housing wealth; not the level of consumption, which is likely more sensitive to changes in the LTV distribution.

What variation in the housing wealth elasticity the model does generate comes largely from changes in home values. Intuitively, as houses become a bigger part of the household balance sheet during the early 2000s, changes in the value of housing became more important to consumption decisions and the housing wealth elasticity rises, in line with the sufficient statistic put forward by Berger et al. (2017).

To understand what drives the stability of the housing wealth effect in the new canonical model, Table 4 decomposes the housing wealth effect for selected years. Specifically, the table shows the average elasticity within different groups of the population together with the relative size of each group. We report $(p/C)(dc/dp)$, where C is average consumption in the population and c is group consumption, so that the total elasticity can be calculated as the dot product of the group elasticities and the group sizes. Note that this elasticity is very related to the marginal propensity to consume out of housing wealth, but it is written as an elasticity rather than a marginal propensity to be consistent with our empirical results.

The table shows that the a substantial fraction of the overall housing wealth effect is driven




Figure 8: Housing Wealth Effect with Different Discount Rates

Note: The figure shows the same calculation of $(p/C)(dC/dp)$ as in Panel (a) of Figure 7 using an alternate β parameters for the household's decision problem.

by low-LTV owners who remain in their homes despite the change in home prices. This group has outsized importance because it accounts for 43% of the population or equivalently 72% of homeowners in 2007. This group has an average elasticity between 0.1 and 0.2. Individuals in these states are not near the LTV constraint, but they still have a fairly high elasticity for several reasons. First, a fundamental feature of models of incomplete markets with liquidity constraints is that households are impatient in the sense that β is substantially below R^{-1} . As a consequence, unlike in a model in which the permanent income hypothesis holds, households have a relatively high marginal propensity to consume out of extra wealth even if they are not constrained. Second, increases in house prices lead to increases in wealth and an offsetting increase in the implicit cost of rent emphasized by Sinai and Souleles (2005). The former dominates the latter because the gain in assets is immediate while the increase in liabilities occurs over many years and is discounted by homeowners.

A key parameter in determining the housing wealth effect is consequently the discount rate β . Figure 8 shows that raising β by 0.01 reduces the level of the housing wealth effect by approximately the same amount without changing the time series pattern. When agents are more patient, fewer agents are borrowing constrained even taking as given their state variables, and moreover

unconstrained households spend less of any increase in wealth.²⁸

The other factor that plays an important role in the stability of the housing wealth elasticity implied by the new canonical model is that while households near the LTV constraint have a high elasticity, households that are underwater have a near-zero elasticity. To elucidate this point, the top panel of Figure 9 shows the elasticity as a function of LTV. The bottom panels of the figure show kernel density estimates of the LTV distribution of homeowners in 2007 and 2010. As with Table 4, we report the elasticity as $(p/C)(dc/dp)$ so the aggregate elasticity for homeowners with a mortgage can be found by integrating the elasticity against the LTV density in the lower panels.

The key feature of Figure 9 is that the elasticity of the LTV ratio does not rise monotonically and instead features a “hump” near the constraint. Households that have high LTV ratios 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 elasticity jumps substantially and remains high until households reach an LTV near 0.95. Intuitively, the households in this region tend to be highly financially constrained. When house prices rise, they respond by refinancing their mortgage, downsizing their house, or selling to rent. Once the LTV ratio rises above roughly 0.95, however, the elasticity drops rapidly. Households that are underwater on their mortgage are highly liquidity constrained and not able to access changes in housing wealth for many years, so house prices do not affect their consumption decision today. They cannot sell or refinance unless they have other liquid wealth to help pay off their mortgage. Due to the presence of long-term debt, these households are not forced to de-lever unless they move or refinance; they can simply pay their mortgage down over time without the LTV constraint binding. This intuition for the insensitivity of high LTV households’ consumption has been documented and analyzed by Ganong and Noel (2017).

Table 4 shows the fraction of households in low (≤ 0.6), medium (0.6, 0.8], and high (0.8, 1.0] elasticity bins for 1986, 2007 and 2010. If we compare 1986 and 2010 we see that 17 percent of the population moved out of the low-LTV stayer category mostly and were pushed into the higher-LTV categories. It is useful to break this change into two parts: 1986-2007 and 2007-2010. For the first period, we see that 11 percent of the population moved out of the low-LTV category and into the medium and high categories in approximately equal shares. The elasticity is higher for the medium and high LTV regions, but the net increase relative to the low-LTV elasticity is about 0.2. So we have an increase of $0.11 \times 0.2 = 0.022$. As the differences in the elasticity across LTVs and the

²⁸In our model, the housing wealth effect is stronger for young homeowners than older homeowners. This is likely due to them having higher LTVs and lower liquid asset holdings.

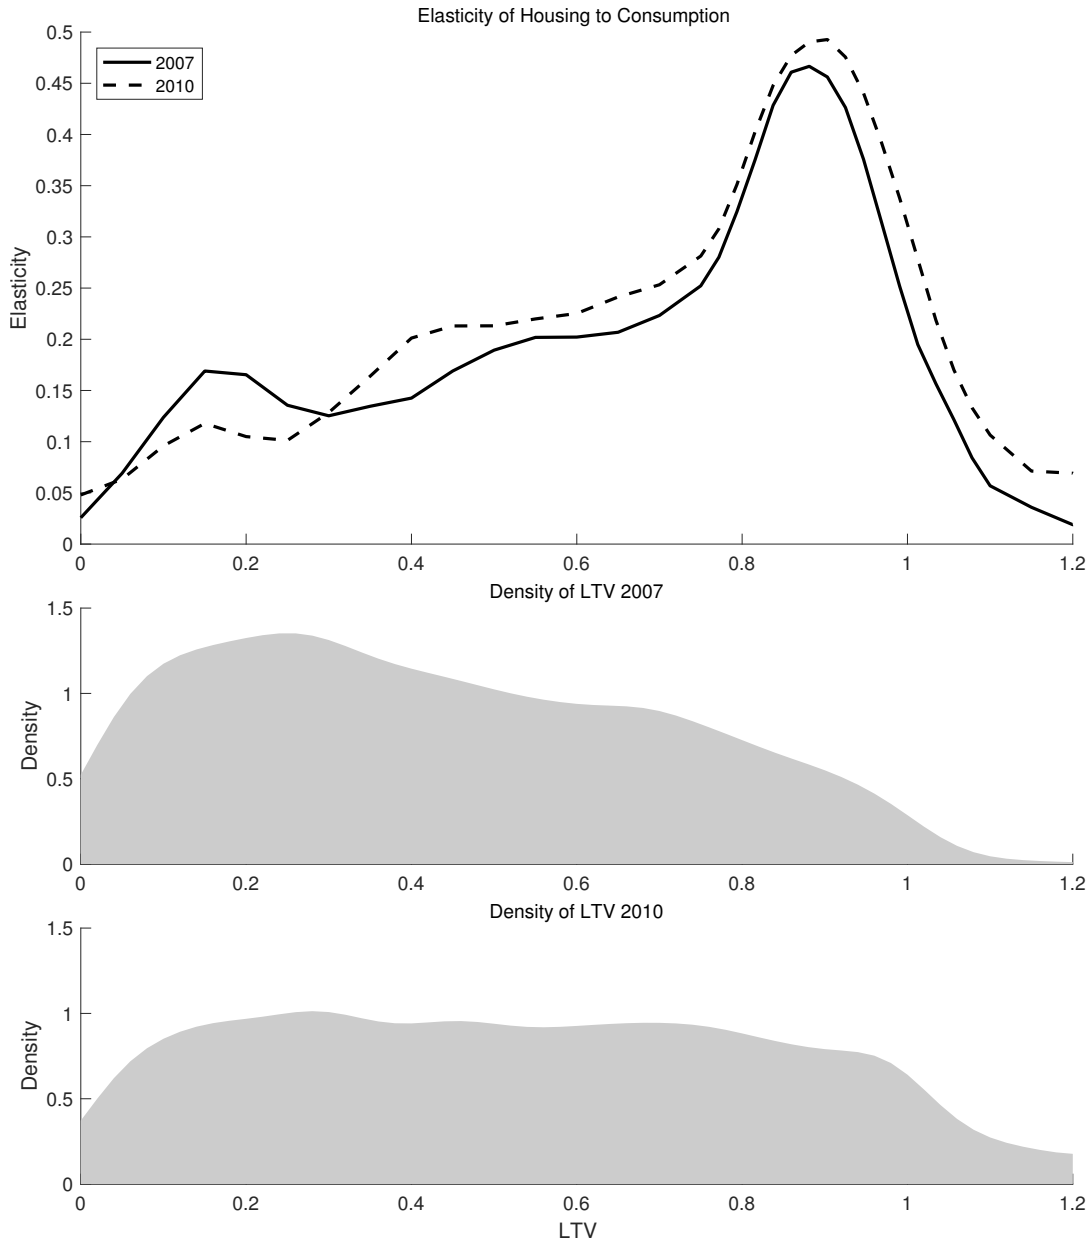


Figure 9: Housing Wealth Effect by LTV and Marginal Distributions of LTV in 2007 and 2010.

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.

change in the LTV distribution are both modest, the combination of the two does not amount to a large change in the overall housing wealth effect despite the “great leveraging” from 1986 to 2007.

For the 2007-2010 episode, there was an abrupt change in the LTV distribution as home values fell, which reduced the denominator of the LTV ratio and caused a rightward shift in the LTV distribution, as documented in the bottom panel of Figure 9. The increase in LTVs pushed some mass into the hump of the top panel of Figure 9 which is a region with very high elasticities. This is the usual intuition about an increase in constrained agents in the recession. However, the increase in LTVs also pushed some mass that was in the hump past the hump into the area with low elasticities. On net, these two effects roughly offset. Between 2007 and 2010, six percent of the population moved out of the the low-LTV group in Table 4. However, the mass in the high LTV group where the hump is largest only grew by two percent of the population, while the underwater group grew by four percent of the population. One can see the offsetting effects visually in Figure 9 by integrating the elasticity as a function of LTV in the top panel by the distributions in the bottom panel to obtain the aggregate elasticity for homeowners.

This intuition helps reconcile the findings of papers that argue that in the aggregate most of the housing wealth effect is due to income effects (e.g., Kaplan, Mitman, and Violante, 2017) with papers that argue that constrained households had particularly high MPCs (e.g., Mian and Sufi, 2011, 2014; Mian, Rao, and Sufi, 2013). While it is true that there was an increase in the number of high-MPC constrained households in the bust, there was also an offsetting increase in the number of zero-MPC underwater households, and these roughly cancel out.

The role of underwater households we describe above arises only in the presence of long-term debt. With only short-term debt, the LTV constraint binds on underwater households that have to roll-over their mortgage each period, which forces these households to de-lever. This deleveraging when house prices fall creates a powerful housing wealth effect in the bust. Indeed, the Appendix shows a version of Figure 9 for the short-term debt model that shows that the elasticity remains elevated for underwater homeowners. To illustrate these effects quantitatively, Figure 10 compares our baseline model to one with only short-term debt for the counterfactual where we allow only the marginal distribution of LTV to change over time. The increase in LTVs in the Great Recession now leads to a sharp increase in the housing wealth elasticity during the bust, in contrast to both our empirical analysis and our model with long-term debt.




Figure 10: Short-Term Debt Counterfactual

Note: The figure shows the changes in the housing wealth effect that result from changes in the LTV distribution as in Panel (b) of Figure 7. 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.

7.3 Credit Conditions and the Housing Wealth Effect

Most of the housing wealth effect is driven by the consumption response of low-LTV owners who are not typically financially constrained. As a result, changes in credit conditions do not much change the housing wealth effect. To demonstrate this, consider a scenario in which the LTV limit rises from 80% to 95% and the mortgage origination cost is reduced to zero. For reasons that will become clear we call this the “boom” parameterization. Figure 11 shows that the housing wealth effect is very similar to the baseline results presented above. One way to read the figure is that, taking as given the distribution of individual states, the housing wealth effect would barely change when credit conditions become looser or tighter.²⁹ These findings may seem surprising in light of analyses such as Guerrieri and Lorenzoni (2017), which shows tightening credit conditions can lead to an economic contraction, or Landvoigt, Piazzesi, and Schneider (2015), which shows cheap credit for poor households was a driving force in the house price boom in San Diego. First, note that our object of interest is not the level of consumption but the *sensitivity* of consumption to home price

²⁹Conditional 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.




Figure 11: Housing Wealth Effect Under Boom Parameters

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

changes. Second, 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. Third and finally, with long-term debt, homeowners are not forced to de-lever when credit conditions tighten.

7.4 No Short-run Housing Adjustment

In our partial equilibrium model, changes in house prices lead to substantial changes in the demand for housing in large part due to a substitution effect between housing and other consumption goods when the relative price changes.³⁰ These changes in housing demand are unlikely to be consistent with general equilibrium because in the short run housing supply is quite inelastic both because construction of new houses involves time to build and also because each year’s construction of new houses represents a small fraction of the overall housing stock.

In an equilibrium model of house prices, the price of housing would need to evolve so as to clear the housing market. As a step in that direction, we suppose that in the “long run” home prices rise by, say, 10%, but in the short run house prices are endogenously determined so that housing demand is unchanged. In practice this means that house prices rise less than 10% initially leaving

³⁰Our baseline model implies a price elasticity of housing demand of 0.15.

an expected capital gain that motivates people to own housing despite the substitution motive that would cause them to reduce their demand when the relative price increases.³¹ We find that a small expected capital gain is sufficient to stabilize short-run housing demand. In this context, the housing wealth effect is very similar to the baseline calculation (see Figure 34 in the Appendix), but is slightly lower as the substitution effect has been muted.

7.5 Is the Model Consistent with the Evolution of Household Balance Sheets?

Our theoretical analysis uses the observed distribution of idiosyncratic states instead of the distribution that is generated by the model. Two considerations motivate this approach. First, it allows us to analyze the consequences of changes in household balance sheets while remaining agnostic as to the driving forces behind those changes. Second, this approach allows us to match the distributions of individual states as closely as possible. On the other hand, a reasonable question to ask is whether our model is in fact consistent with the observed evolution of household balance sheets.

To address this question, we simulate the model's predictions for the evolution of the LTV distribution in response to the observed changes in home prices. We begin the simulation at each wave of the SCF and ask how the model's predictions compare to the next wave of the SCF. For example, we start with the 1983 data and simulate the model using home prices for 1983-1986 and compare the model's predictions to the SCF data in 1986. Figures 35 and 36 in the Appendix show results for both our baseline model and more complex versions that incorporate the boom parameterization and news shocks discussed above, which Kaplan, Mitman, and Violante (2017) argue are central to fitting the evolution of house prices during the 2000s boom-bust episode. The baseline model largely succeeds in fitting the evolution of leverage from the late 1980s until the mid-1990s and during the Great Recession. The baseline model is less successful during the 1998-2007 period when the LTV distribution was stable despite a large run-up in house prices, but the model does quite well under the boom parameterization with news shocks about future price appreciation. Intuitively, both features make households more aggressive in borrowing against their houses, and without these modifications LTV's are predicted to fall during this period as home prices rise. However, modifying the model to incorporate the boom parameters and the news shocks does not materially affect our predictions about the housing wealth elasticity. As explained in Section 7.3, the housing wealth effect is virtually identical under the boom parameters. Adding expected

³¹This experiment was motivated in by the belief shocks of Kaplan, Mitman, and Violante (2017), 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.

capital gains for 2004-2007 slightly increases the housing wealth effect in 2007 from 0.107 to 0.115. Appendix F.11 provides additional details of these simulations.

7.6 Extensions and Robustness

We explore how our results are affected by several other modeling choices in the Appendix. We show that the housing wealth effect is modestly increasing in mortgage interest rates, which we set at 3% in our baseline specification, but the effect is quantitatively small although more pronounced in recent years (see Figure 31). 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. 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 prices because 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, perhaps because buyers expected capital gains, implying a low user cost as in Kaplan, Mitman, and Violante (2017). Making the polar opposite (and also unrealistic) assumption that rents remain constant when home prices change yields somewhat higher housing wealth effects because it leads renters to defer buying a house and spend more on non-housing consumption. This alternate rent assumption does not change the time series pattern of the wealth effect (see Figure 32).

In our model, households can default subject to a utility cost. We set the utility cost at a level such that the default option is rarely exercised, but our results are essentially unchanged if we make default more attractive. This finding is related to the insensitivity of our results to credit constraints: the housing wealth effect is heavily influenced by low-LTV households.

8 Conclusion

We present new evidence on housing wealth elasticities going back to the 1980s. To do so, we develop a new identification approach that exploits the differential sensitivity of house prices in cities to regional house price cycles. This allows us to obtain estimates of the housing wealth elasticity using a consistent methodology over a long time period. There is no evidence that the elasticity to changes in house prices was particularly large in the 2000s; if anything, the elasticity was larger prior to 2000.

Theoretically-minded readers may find it hard to interpret what a “housing wealth effect” means given that house prices are an endogenous object. We develop a theoretical framework to 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 next turn to modeling household consumption behavior at a local level. We present a standard life-cycle model with borrowing constraints, uninsurable income risk, illiquid housing, and long-term mortgages. In the model, there is little response of the housing wealth effect to shifts in the LTV distribution — even the large run-up in household leverage during the Great Recession. This arises for two reasons. First, much of the housing wealth elasticity arises from features of consumer behavior in incomplete markets models that are insensitive to the LTV distribution; and second, in a model with long-term debt, the increase in the elasticity associated with highly constrained households during the Great Recession was offset by an increase in underwater households with close to zero responsiveness to movements in house prices.

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.
- 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 (2017). The Economic Effects of Social Networks: Evidence From the Housing Market.
- Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2017). House Prices and Consumer Spending.
- Bhutta, N. and B. J. Keys (2016). Interest Rates and Equity Extraction During the Housing Boom. *American Economic Review* 106(7), 1742–1774.
- Blanchard, O. and L. F. Katz (1992). Regional Evolutions. *Brookings Papers on Economic Activity* 1, 1–75.
- Bostic, R., S. Gabriel, and G. Painter (2009). Housing Wealth, Financial Wealth, and Consumption: New Evidence From Micro Data. *Regional Science and Urban Economics* 39(1), 79–89.
- 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.
- Cloyne, J., K. Huber, E. Ilzetzki, and H. Kleven (2017). The Effect of House Prices on Household Borrowing: A New Approach.
- 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. (2013). Supply Elasticity and the Housing Cycle of the 2000s. *Real Estate Economics* 41(4), 793–813.
- 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. (2016). Homeowner Borrowing and Housing Colateral: New Evidence from Expiring Price Controls.
- 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.
- Ganong, P. and P. Noel (2017). The Effect of Debt on Default and Consumption: Evidence from Housing Policy in the Great Recession.
- Gilchrist, S. and E. Zakrajsek (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review* 102(4), 1692–1720.
- Gorea, D. and V. Midrigan (2017). 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 (2018). Mortgage Design in an Equilibrium Model of the Housing Market.
- 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 (2016). What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?
- Haughwout, A., R. Peach, J. Sporn, and J. Tracy (2013). The Supply Side of the Housing Boom and Bust of the 2000s. In E. L. Glaeser and T. Sinai (Eds.), *Housing and the Financial Crisis*, pp. 69–104. Chicago: University of Chicago Press.
- 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 (2015). Consumption and House Prices in the Great Recession: Model Meets Evidence.

- Kaplan, G., K. Mitman, and G. L. Violante (2017). The Housing Boom and Bust: Model Meets Evidence.
- Kaplan, G., G. L. Violante, and G. Kaplan (2016). Non-Durable Consumption and Housing Net Worth in the Great Recession: Evidence From Easily Accessible Data.
- 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, 1949-2013.
- 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 Economic 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.
- Mian, A. and A. Sufi (2015). *House of Debt: How They (and You) Caused the Great Recession, and How We Can Prevent It from Happening Again*. Chicago: University of Chicago Press.
- 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 (2017). Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market.
- 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.
- Roussanov, N., M. Michaux, and H. Chen (2013). 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.
- Sinai, T. and N. S. Souleles (2005). Owner-Occupied Housing as a Hedge Against Rent Risk. *Quarterly Journal of Economics* 120(2), 763–790.
- Wong, A. (2018). Population Aging and the Transmission of Monetary Policy to Consumption.

A Data Appendix

Our main data sources are the Quarterly Census of Employment and Wages (QCEW), the County Business Patterns (CBP), the 1970-2010 post-Censal county population estimates, the 2010-2015 inter-Censal county population estimates, and the Freddie Mac House Price Indices. We also merge in data from a number of other data sources. 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 1990-2015 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 total employment, retail, manufacturing, real estate, or construction) over the whole sample. In practice, this means that we drop a few small outlying counties along with the entire Dover, Delaware CBSA. We show in Appendix C 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 employment, manufacturing, real estate, and construction (the later two for the Online Appendix). 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 merge in county-level population, 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. The log changes thus begin in 1976 with the difference between 1976 and 1975 log employment. We create average wages by dividing total wages by total employment. We then clean the data by removing observations where we observe a sudden 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 C 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 C, we show our results are robust to splice date.

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 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. Since the CBP data is 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. In a robustness check, we use FHFA price indices, which are available for an alternate level of CBSA aggregation. We aggregate everything to this alternative CBSA definition when using this price index.

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_0a” measure, which is the excess bond premium which is subtracting a fitted value for “distance to default” from options. We also merge in 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 using the BEA’s GDP deflator.

We create the regional log changes in employment and house prices. We do so by using the average log change in employment 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.

B Identification Explained with Simultaneous Equations

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

$$\Delta y_{i,r,t} = \psi_i + \xi_{r,t} + \beta \Delta p_{i,r,t} + \alpha_i \mathcal{E}_{r,t} + \varepsilon_{i,r,t}, \quad (7)$$

$$\Delta p_{i,r,t} = \varphi_i + \zeta_{r,t} + \delta \Delta y_{i,r,t} + \gamma_i \mathcal{V}_{r,t} + \nu_{i,r,t}. \quad (8)$$

We allow for CBSA fixed effects (ψ_i and φ_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 α_i and γ_i). This feature will play an important role. Equation (7) is the analog of equation (1) in the main text and the coefficient of interest is β , the causal effect of house prices on retail employment measured as an elasticity.³²

Equations (7) and (8) 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 (7). However, changes in local employment also affect local house prices through the $\delta \Delta y_{i,r,t}$ term in equation (8). Since causation between local employment and house prices runs both ways, estimating equation (7) by OLS will yield a biased estimate of β . The classic approach to solving this problem is to look for a variable that shows up in equation (8) but not in equation (7) and to use this variable as an instrument for $\Delta p_{i,r,t}$ when estimating equation (7). 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 (2) and use $z_{i,r,t} = \hat{\gamma}_i \Delta P_{r,t}$ as an instrument for $\Delta p_{i,r,t}$ in equation (7). 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 γ_i s in equation (8), but that local retail employment does respond differentially to regional shocks through heterogeneity in α_i s in equation (7). In this case, the differential response of local

³²The 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 (7) and (8) is not as explicit about this heterogeneity in sensitivity.

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 (8). Were we to estimate equation (2) in this case, we would estimate heterogeneous γ_i s. The source of these estimated γ_i s would, however, be the $\alpha_i\mathcal{E}_{r,t}$ term in equation (7). 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 β .

To address this problem, consider the following more sophisticated identification strategy. First aggregate equation (8) 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 γ_r denotes the regional average of γ_i . Next, use this equation to rewrite equation (8) as:

$$\Delta p_{i,r,t} = \varphi_i + \tilde{\zeta}_{r,t} + \delta\Delta y_{i,r,t} + \frac{\gamma_i}{\gamma_r}\Delta P_{r,t} - \frac{\gamma_i}{\gamma_r}\delta\Delta Y_{r,t} + \nu_{i,r,t}. \quad (9)$$

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 γ_i/γ_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 (7). 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 (9). A slightly more general version of equations (7) and (8) allows for heterogeneity in the response of house prices to local employment, which replaces δ in equation (8) with δ_i . In this case, the coefficient on $\Delta y_{i,r,t}$ in equation (9) is δ_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 (7). Because we control for $\Delta y_{i,r,t}$ and $\Delta Y_{r,t}$ in defining the instrument in equation (9), 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, α_i^j must be correlated with $\hat{\gamma}_i$ in the cross-section. An assumption that is sufficient to rule out endogeneity of




Figure 12: Regional House Price Indexes

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

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 Empirical Robustness Appendix

This appendix discusses the robustness of our empirical results. For clarity, it is broken into several sub-sections. First, we present a time series of the regional house prices we use for identification. Second, we we present a robustness analysis for the rolling windows results in Figure 2. Third, we discuss the robustness of the pooled results in Table 1.

C.1 Regional House Price Indexes

Figure 12 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, there was not a substantial national house price cycle. 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




Figure 13: 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 house prices at the CBSA level for rolling 10-year sample periods as in Figure 2 in blue and a version without the controls used in our baseline specification in red (dashed). Each point indicates the elasticity for a 10-year sample period that begins in the quarter stated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. Confidence intervals are similar to Figure 2 and are not shown so that the comparison between the two specifications is clearer.

us identify the housing wealth effect prior to 2000.

C.2 Robustness of Rolling Windows Analysis

This section analyzes the robustness of the rolling windows analysis in Figure 2.

C.2.1 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 (3). These include three different controls for the 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 (3) the log change in average wages with city-specific coefficients.

To evaluate the role of these controls in our results, Figure 13 shows the point estimates for the case with controls and the case with no controls (the standard errors do not change significantly). One can see that the effect is about 20% bigger without controls, although the difference is slightly stronger in the earlier period. Without controls, the central estimate of the elasticity across years is 0.082 rather than 0.060.

Which controls matter the most? In unreported results, we find that the industry shares have



Figure 14: Elasticity of Retail Employment Per Capita to House Prices: 3-Year Differences

Note: The figure plots the elasticity of retail employment per capita to house prices at the CBSA level for rolling 10-year sample periods as in Figure 2 using 3-year instead of 1-year log differences for all variables. Each point indicates the elasticity for a 10-year sample period that begins in the quarter stated 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 region-time.

the largest effect, although the control for city-level exposure to regional retail employment does explain a significant portion of the gap between the controls and no controls specifications in the 2000s.

C.2.2 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 1990. The most important form of variation across specifications is that several specifications yield lower estimates of the elasticities in the 1980s than in our baseline analysis (though with large standard errors). This is why we do not put too much emphasis on the 1980s results.

Figure 14 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 15 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




Figure 15: Elasticity of Retail Employment Per Capita to House Prices: Population Weighting

Note: The figure plots the elasticity of retail employment per capita to house prices at the CBSA level for rolling 10-year sample periods as in Figure 2 in blue and a version with all regressions weighted by 2000 population in red (dashed). Each point indicates the elasticity for a 10-year sample period that begins in the quarter stated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. Confidence intervals are similar to Figure 2 and are not shown so that the comparison between the two specifications is clearer.

is somewhat higher in the 10-year windows beginning in the 1980s and early 1990s and somewhat lower for the 10-year windows starting in the 2000s.

Figure 16 shows results that estimate γ_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 γ_i across windows (which may be partly real and partly due to sampling error). Since 1990, this specification yields a declining elasticity, in line with our baseline analysis. The elasticity is somewhat lower both in the 1980s and in the 2000s than in our baseline analysis.

Figure 17 shows results using the log change in retail employment rather than the log change in retail employment per capita as the dependent variable. The time series pattern is similar to the per-capita specification, indicating that the declining elasticity since the 1990s does not arise from changes in migration responses over time. However, the level is several percentage points higher, indicating that population flows toward booming CBSAs, which is consistent with the literature (e.g., Blanchard and Katz, 1992).

Figure 18 shows results dropping the “sand states” of California, Arizona, Nevada, and Florida from the analysis. Critics of the Saiz instrument such as Davidoff (2013) often argue that much of the identification is driven by these states. This figure shows that the declining pattern of




Figure 16: Elasticity of Retail Employment Per Capita to House Prices: Single Gamma

Note: The figure plots the elasticity of retail employment per capita to house prices at the CBSA level for rolling 10-year sample periods as in Figure 2 in blue. The specification in red is the same except that γ_i is estimated for each CBSA by equation (3) 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 that begins in the quarter stated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. Confidence intervals are similar to Figure 2 and are not shown so that the comparison between the two specifications is clearer.




Figure 17: Elasticity of Retail Employment (Not Per Capita) to House Prices

Note: The figure plots the elasticity of retail employment per capita to house prices at the CBSA level for rolling 10-year sample periods as in Figure 2 except with retail employment rather than retail employment per capita as the dependent variable. Each point indicates the elasticity for a 10-year sample period that begins in the quarter stated 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 region-time.




Figure 18: Elasticity of Retail Employment Per Capita to House Prices: No Sand States

Note: The figure plots the elasticity of retail employment per capita to house prices at the CBSA level for rolling 10-year sample periods as in Figure 2 in blue and a version dropping the “sand states” of California, Nevada, Arizona, and Florida in red (dashed). Each point indicates the elasticity for a 10-year sample period that begins in the quarter stated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. Confidence intervals are similar to Figure 2 and are not shown so that the comparison between the two specifications is clearer.

elasticities since the 1990s is not affected by these states. Indeed, the quantitative results since the 1990s are quite similar whether or not one includes these states.

Figure 19 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. This has 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 a three-year buffer around the 10-year window in constructing the instrument using equation (3) and in constructing the controls for differential city-level exposure to regional retail employment, real 30-year mortgage rates, and the Gilchrist-Zakrajek excess bond premium using equation (4). This has a slight effect on the results for a few 10-year windows in the early 1990s, but in general the difference is not significant economically or statistically. The fourth specification uses an alternative specification for the differential city-level exposure to regional retail employment. Specifically, the “alternate cyclical control” creates the control for differential city-level exposure to retail employment $a_i \Delta Y_{r,t}$ using a regression of the form $\Delta y_{i,r,t} = \psi_i + \xi_{r,t} + \alpha_i \Delta Y_{r,t} + \omega_i \Delta x_{i,r,t} + \zeta_i \Delta X_{r,t} + \varepsilon_{i,r,t}$ rather than omitting the Δx and ΔX terms. This has essentially no effect on the




Figure 19: Elasticity of Retail Employment Per Capita to House Prices: Misc. Robustness Tests

Note: The figure plots the elasticity of retail employment per capita to house prices at the CBSA level for rolling 10-year sample periods as in Figure 2 in dark blue and with seven other specifications that do not substantially affect the results. We use an instrumental variables estimator that is described in Section 3. Confidence intervals are similar to Figure 2 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 (3) and for the regressions as in equation (4) 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 “alternate cyclicity control” creates the control for differential city-level exposure to retail employment using a regression of the form $\Delta y_{i,r,t} = \psi_i + \xi_{r,t} + \alpha_i \Delta Y_{r,t} + \omega_i \Delta x_{i,r,t} + \zeta_i \Delta X_{r,t} + \varepsilon_{i,r,t}$ rather than omitting the Δx and ΔX terms. 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.

results. 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. This too has essentially no impact on the results.

Finally, Figure 20 shows results beginning in 2000 for 5-year windows rather than 10-year windows. We focus on the post-2000 period for the specification using 5-year windows because the standard errors for this specification (which are generally larger than in our baseline due to the shorter window) are smaller post-2000 due to the large movements in house prices. This finer resolution does not yield evidence that the elasticity rose during the bust. One notable feature of this plot is that the elasticity falls close to zero at the end of the sample, although the standard




Figure 20: 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 house prices at the CBSA level as in Figure 2 except for rolling 5-year windows instead of 10-year windows. The figure starts with the 2000-2005 window due to wide standard errors in the 5-year specification prior to 2000. Each point indicates the elasticity for a 5-year sample period that begins in the quarter stated 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 region-time.

errors widen significantly in this period.

C.2.3 Single Cross-Section Results

How do our results compare with the repeated cross-section approach used by most of the literature? Figure 21 shows the point estimates and standard errors for repeated 3-year cross-sections. We follow the same empirical approach but use a single cross section and replace the region-time fixed effect in equation (1) with a region fixed effect. We also demean all variables including controls once for the whole period from 1976-2015.

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 three 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 effect to be greater from 1990 to 2003 than it was in the 2000s.




Figure 21: Repeated Cross-Sections: 3 Year Differences

Note: The figure plots the elasticity of retail employment per capita to 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 1976-2015 sample. 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 region-time.

C.2.4 Alternate Data Sets

Our results use the QCEW and Freddie Mac House Price Indices. In this section, we present results using alternate data sources to show that our results are broadly robust to the data used.

Figure 22 uses County Business Pattern employment data rather than QCEW data for retail employment. The CBP is only available annually, and so we report annual results rather than quarterly. As a consequence, the standard errors are wider. Nonetheless, the same general time pattern we observe with the QCEW are evident with the CBP.

Figure 23 presents results using the FHFA purchase-only house price index rather than the Freddie Mac indices. This index only uses purchases, while the Freddie mac index uses both appraisals and purchases to be available for more cities for a longer time period. However, this FHFA index is only available for 72 CBSAs starting in 1992.




Figure 22: Alternate Data: County Business Patterns Employment Data

Note: The figure plots the elasticity of retail employment per capita to house prices at the CBSA level for rolling 10-year sample periods as in Figure 2 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 that begins in the year stated 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 region-time.




Figure 23: Alternate Data: FHFA Purchase Only House Price Indices

Note: The figure plots the elasticity of retail employment per capita to house prices at the CBSA level for rolling 10-year sample periods as in Figure 2 using the FHFA purchase only house price index rather than the Freddie Mac house price index. This index only includes transactions as opposed to the Freddie Mac index which uses both transactions and appraisals. The index starts in 1992 for 72 MSAs, so the graphs only start with the 1992-2001 ten year window. Each point indicates the elasticity for a 10-year sample period that begins in the quarter stated 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 region-time.

One concern is that the Freddie Mac house price index we use in our baseline analysis is “smoothed” due to the behavior of appraisers relative to a transactions-based house price index. This would lead to 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. Such bias is not, however, evident in the figure: the estimates are quantitatively similar to our baseline estimates.

C.2.5 Alternate Versions of the Saiz Instrument

In Figure 5 in the main text we compare the standard national Saiz instrument to our sensitivity instrument. In this section, we present rolling window estimates for two other permutations of the Saiz instrument. Both generate large standard errors relative to our baseline sensitivity instrument. Figure 24 constructs a “regional Saiz” instrument by interacting each period’s regional house price index change with Saiz’s estimated CBSA housing supply elasticity with region-time fixed effects. Figure 25 presents results for the national instrument using Saiz’s (2010) land unavailability measure instead of his elasticity variable. This version generates even larger standard errors than the baseline in Figure 5 but a similar time pattern.

C.2.6 Accounting for Sampling Error in the γ_i s

We use two-way clustered standard errors by CBSA and region-time in most of our figures and our pooled analysis in Table 1. These standard errors do not account for sampling error in the estimation of the γ_i s. To evaluate whether accounting for sampling error in the γ_i s would significantly change our standard errors, Figure 26 compares a specification that uses parametric standard errors clustering only by time only with bootstrapped standard errors that block-bootstrap by time and account for the sampling error in γ_i .³³ To minimize computational burden, we use a specification without controls. The Figure shows that accounting for sampling error in γ_i does not substantially widen the standard errors.

C.3 Robustness of the Pooled Results

This section presents results on the robustness of the pooled specifications in Table 1. Table 5 presents the same results weighting by 2000 population. Table 6 presents the same results without

³³One can only block-bootstrap on one dimension. We choose time because block-bootstrapping by CBSA would involve drawing CBSAs randomly and each CBSA would have the same γ_i that it would in the baseline procedure under such a sampling scheme. Any sampling error in γ_i would thus come from the time dimension.




Figure 24: Housing Wealth Effect: Sensitivity Instrument vs. Regional Saiz Instrument

Note: The red dashed line plots the point estimates of the housing wealth effect based on 10-year rolling windows using our instrument (same as in Figure 2). The light red dashed lines plot the upper and lower bounds of 95% confidence intervals for these estimates. The dark blue line plots the point estimates of the housing wealth effect based on 10-year rolling windows estimated using an instrument based on the estimated housing supply elasticity of Saiz (2010) interacted with the regional annual log change in house prices with region-time fixed effects. The lighter blue lines plot the upper and lower bounds of 95% confidence intervals for the estimates based on the Saiz instrument. Standard errors are two-way clustered by region-time and CBSA.




Figure 25: Housing Wealth Effect: Sensitivity Instrument vs. Saiz Unavailability Instrument

Note: The red dashed line plots the point estimates of the housing wealth effect based on 10-year rolling windows using our instrument (same as in Figure 2). The light red dashed lines plot the upper and lower bounds of 95% confidence intervals for these estimates. The dark blue line plots the point estimates of the housing wealth effect based on 10-year rolling windows estimated using an instrument based on the land unavailability of Saiz (2010) interacted with the national annual log change in house prices. The lighter blue lines plot the upper and lower bounds of 95% confidence intervals for the estimates based on the Saiz instrument. The lighter blue lines plot the upper and lower bounds of 95% confidence intervals. Standard errors are two-way clustered by region-time and CBSA.




Figure 26: Accounting for Sampling Error in the γ_{is}

Note: The dark blue line plots the elasticity of retail employment per capita to house prices at the CBSA level for rolling 10-year sample periods. The light blue lines indicate 95% confidence intervals clustered by time. The red dashed line is identical but it block bootstraps by time and accounts for sampling error in the estimation of the γ_{is} . The light red dashed lines indicate 95% confidence intervals. The point estimates differ because the red dashed line indicates the mean point estimate over all of the bootstrap samples. Each point indicates the elasticity for a 10-year sample period that begins in the quarter stated on the horizontal axis. We use an instrumental variables estimator that is described in Section 3. This specification includes no controls to reduce the computational burden.

controls. The point estimates for the weighted regression are somewhat lower, and the point estimates for the version without controls are somewhat higher. In no specification do we find any strong evidence of a boom-bust asymmetry.

D Retail Employment vs. Consumption

In this section, we analyze the relationship between consumption and retail employment using aggregate and city-level data. We should emphasize that the findings and analysis below are by no means original or surprising in the context of the measurement literature. Indeed, retail employment has long been used to impute consumption in government and private sector attempts to construct regional consumption measures (e.g., BEA’s Regional PCE measures and the private sector “Survey of Buying Power” both make use of retail employment data).

Let us start with the aggregate evidence. Figure 27 presents aggregate statistics on retail employment and real personal consumption expenditures, and shows that the two series comove strongly. Since our analysis allows for time trends, it will not be affected by long-term trends such as shifts in retailing toward online retailers or big-box stores. However, this graph also suggests

Table 5: 1990-2015 Elasticity of Retail Employment Per Capita to House Prices: Weighted

	(1)	(2)	(3)
$\Delta \log (P)$	0.050*** (0.019)		
$\Delta \log (P) -$		0.081*** (0.027)	
$\Delta \log (P) +$		0.035 (0.026)	
P Test for Equality		0.234	
$\Delta \log (P)$			0.059*** (0.020)
$\Delta \log (P)^2$			0.054 (0.064)

Note: For these estimates, we first construct our instrument for each quarter by estimating the γ_i 's in equation (3) 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 region-time and CBSA level. All regressions are weighted by 2000 population.

Table 6: 1990-2015 Elasticity of Retail Employment Per Capita to House Prices: No Controls

	(1)	(2)	(3)
$\Delta \log (P)$	0.082*** (0.013)		
$\Delta \log (P) -$		0.097*** (0.018)	
$\Delta \log (P) +$		0.074*** (0.017)	
P Test for Equality		0.358	
$\Delta \log (P)$			0.091*** (0.0153)
$\Delta \log (P)^2$			0.067 (0.038)

Note: For these estimates, we first construct our instrument for each quarter by estimating the γ_i 's in equation (3) 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 region-time and CBSA level. All regressions do not include the standard controls in our baseline specification.




Figure 27: 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 trend growth.

that these trends have not led the growth rates in these two series to drift apart.³⁴

We next study 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:

$$\Delta \log \bar{C}_{i,t} = \xi_i + \zeta_t + \beta \Delta \log \bar{Y}_{i,t} + \varepsilon_{i,t}, \tag{10}$$

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 ξ_i and ζ_t are city and time fixed effects respectively. We estimate the equation over the sample period 1986-2014. To construct real consumption,

³⁴At an aggregate level, retail employment actually does a much better job capturing time series variation in non-durable PCE than the CEX (see Heathcoate et al, 2009).

Table 7: City-Level Consumption vs. Retail Employment Regressions

	1	2	3	4
	OLS	IV	OLS	IV
	Total Cons	Total Cons	Ex Imputed Rent	Ex Imputed Rent
Retail Emp.	0.460**	0.940**	0.521*	0.969**
Per Capita Growth	(0.179)	(0.314)	(0.230)	(0.400)
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 (10). 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.

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 7 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 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). The IV coefficient is a ratio of the covariances of retail employment with house prices, and consumer expenditures with house prices. The IV specifications in columns 2 and 4 yield larger effects. 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. The fact that these IV estimates are higher than their OLS counterparts reflects the role of sampling error, as we discuss above.

These regressions (and our main specification) will not be affected by long-term trends such as shifts in retailing toward online retailers or big-box stores, since our regressions are run in first-differences, so a trend divergence would be reflected in the constant term. However, the comparison between the aggregate retail employment and consumption data in Figure 27 suggest that these trends are quantitatively small.

E Related Literature

Previous work has studied the effect of house prices on both consumption and employment, and has reported the results both in terms of elasticities and the cent-for-cent “marginal propensity to consume out of housing wealth” (MPCH), the increase in consumption in response to a dollar increase in home values. Our full-sample estimate for the sample period 1982-2015 is an elasticity of 0.06, which implies a MPCH of roughly 2.76 cents for each dollar of extra housing wealth. Aladangady (2017) estimates a dollar-for-dollar MPCH of 4.7 cents for owners and zero for renters, using an instrument constructed by interacting the Saiz instrument with the real interest rate. Multiplying by a long-run homeownership rate of 65 percent implies an overall MPCH of roughly 3.1 cents. Mian, Rao, and Sufi (2013) estimate an MPCH out of housing wealth of 5.4 cents for total consumption, but 2.3 cents of that comes from automobiles and only 1.6 cents for non-durable goods.

Several other papers report the response of log employment to a housing net worth shock, defined as $\Delta \log P_{2006,2009} * H_{2006}/NW_{2006}$. 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 instrumenting the change in net worth with Saiz’s (2010) estimate of the housing supply elasticity. To convert this to a housing wealth elasticity, we must divide by the ratio of housing net worth to total net worth, which is between three and four (Berger et al., 2017). This implies a range for the elasticity of retail employment to house prices of between 0.09 and 0.16. Kaplan, Mitman, and Violante (2016) estimate the elasticity of non-durable consumption with respect to this shock 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. Recall that our pooled elasticity with respect to house prices is 0.06, which is slightly below this range.

Several other studies use OLS. Case et al. (2005) find an elasticity with respect to house prices across countries of 0.11 to 0.17 and across states of 0.05 to 0.09. In a follow up paper, Case et al. (2013) find an elasticity with respect to house prices of between 0.04 and 0.09 with OLS and 0.08 and 0.17 when they instrument with lags of income and wealth. Bostic et al. (2009) find an elasticity with respect to house prices of 0.06. Carroll et al. (2011) finds MPCH of 2 cents in the short run and 9 cents in the long run. Campbell and Cocco (2007) find elasticities above one for selected groups of old homeowners but no response for young renters in British data.

A few studies estimate the elasticity of borrowing with respect to house prices, and find much larger responses. Mian and Sufi (2011) estimate this elasticity to be between 0.5 and 0.6 in U.S. data from 2002-2006 using the Saiz instrument. Cloyne et al. (2017) 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. However, one might expect that home equity borrowing might respond more to house price changes than consumption because households partially offset this by saving more or borrowing less from other sources.

There is little evidence on how housing wealth effects have changed over time. Case et al. (2013) show that their OLS effect is larger post 1986, when the Tax Reform Act made HELOCs more tax-advantageous. Aladangady (2017) presents estimates for post-2002, but the estimates are imprecise (the standard error is roughly equivalent in size to the magnitude of the coefficient). It has sometimes been argued that the contrast between Case et al. (2005), which uses a sample from 1982 to 1999, and Case et al. (2013), which uses a sample from 1975 to 2012, shows that the housing wealth effect is larger for the latter period (when the 2000-2012 period is added). While this is the case, there are enough econometric differences across these papers that a comparison is not appropriate.

The evidence is also mixed on whether there is an asymmetry in responses to house price increases versus decreases. Case et al. (2005) found an asymmetry, but Case et al. (2013) rejected this initial finding with additional years of data and find no asymmetry. 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 in the recession: he reports that the elasticity for durables rose from 0.08 in 2000-2006 to 0.31 in 2007-2009. However, he does not find an asymmetry for consumption overall, which is consistent our empirical findings.

Table 8: 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

F Theory Appendix

F.1 Income process

We use an income process that captures some features of earnings dynamics reported in Guvenen et al. (2016), hereafter GKOS. Specifically we model log annual income as $\log y = \ell + z + \xi$, where ℓ is a deterministic life-cycle component, z is a persistent shock that follows an AR(1) process, and ξ is a transitory shock. 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. ξ is zero with some probability or is equal to $\log(1 - x)$ where x is drawn from an exponential distribution on the interval $(0, 1)$. We use data from the 2002 March CPS on hours worked last year to estimate the parameter of the exponential distribution to be 2.25 and the probability that ξ 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 8 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 as follows: the first component has a mixture weight of 0.984 and a standard deviation of 0.071 and the second mixture component has a weight of 0.016 and a standard deviation of 1.60. The data used by GKOS refers to 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 $0.135 + 0.062(Y/\bar{Y})$.

F.2 Internal Calibration

This appendix describes the procedure to calibrate the discount factor, β ; the strength of the preference for housing, ω ; the strength of the bequest motive, B_0 ; the degree to which a bequest is a luxury, B_1 ; the rent-price ratio, δ ; the mortgage origination cost, ψ^M ; and the transaction cost for selling a house, ψ^{Sell} . 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 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 the median income among 40 year-olds is 1.0. To compute life-cycle profiles for these variables 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 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. Table 9 lists the reasons for moving responses, whether we included or excluded these moves, and the frequency of these responses.

Table 9: 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

F.3 Value Functions and Model Solution

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

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

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, ℓ is liquid cash on hand, and I is an indicator for the mortgage being high-LTV at origination. This maximization problem is subject to the LTV constraint:

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

where w is defined as:

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

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:

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

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:

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

subject to:

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

A renter solves:

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

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

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

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 ϵ and incurs a utility cost ϕ . 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 (ℓ, h, m, I, 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 (ℓ, h, m, I, 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 (ℓ, h, m, I, 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 (ℓ, h, m, I, z, a) . To find the decision rules, we cannot avoid solving the problem for the specific combinations of (ℓ, h, m, I, 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.

F.4 Constructing the Distribution of Idiosyncratic States from SCF Data

Our model has five state variables: cash on hand, mortgage debt, housing position, persistent income, and age. We create analogous variables using SCF data using each wave of the SCF data from 1983 to 2016. We equate cash on hand to liquid assets plus annual wage income, where liquid assets are defined as in Section F.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 pensions. We remove taxes from income in the same manner as described in Section F.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). The housing position is the value of real estate deflated by the relative price of housing. 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 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




Figure 28: Linearity of Housing Wealth Effect

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

mass of the SCF observations to the grid points in a manner that preserves the means of variables.

F.5 Linearity and Interaction Effects

Figure 28 shows model estimates of the housing wealth effect for both positive and negative changes larger than the 10% changes we use in our main analysis. The figure shows that the housing wealth effect 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 a common shock to the demand for housing will be absorbed by the time fixed effects except for its differential implications for home prices depending on the elasticity of housing supply in a city. This argument allows us to remain agnostic about the shocks driving home prices. The argument assumes that the second-order interaction of home prices and the demand shock does not have important consequences for consumption. To assess the validity of the argument we take a stand on the nature of the shock and suppose it is a shock to the preference for housing, ω . Figure 29 shows that the housing wealth effect does not change meaningfully as we change ω . This implies that the cross derivative $d^2C/(dpd\omega)$ is small.




Figure 29: Housing Wealth Effect For Alternative ω

Note: The figure shows the housing wealth effect 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 ω on the consumption decision rule. In the long run, a larger ω will lead households to own more housing and this will raise the housing wealth effect.

F.6 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 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 effect, 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 $\theta = 0.70$.³⁵

Figure 30 shows that the housing wealth effect is barely changed by shifts in the LTV constraint. The intuition Table 4 is useful to understand these results: a large part of the housing wealth effect comes from households who are far from the LTV constraint and their behavior is little affected by the details of the constraint. We should 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

³⁵Some analyses of changing credit conditions (e.g., Guerrieri and Lorenzoni, 2015) take the distribution of individual states from a model simulation. 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. Therefore we do not need to take a stand on how credit constraints changed and can compare different constant levels of the constraint.




Figure 30: Impact of Borrowing Constraints on the Housing Wealth Effect

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

consumption. This is an important difference from other analyses that typically focus on more on how the level of consumption reacts to changes in borrowing constraints.^{36 37}

F.7 Interest Rate Changes

Figure 31 shows results with alternate values for the real mortgage interest rate. Recall that we take the distribution of state variables from the SCF as given, so the new mortgage interest rate only affects decisions going forward not the distribution of mortgage debt that households have at the beginning of each period. 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 effect 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 effect coming from higher home values and leverage.

³⁶For example, Table 2 of Favilukis et al. (2017) shows that relaxing the collateral constraint from a 25 percent downpayment require to a 1 percent requirement raises the consumption share of GDP by two percentage points.

³⁷While past constraints are irrelevant conditional on the current states, future constraints are not. We are assuming that households always expect the value of the constraint to remain constant going forward.




Figure 31: Sensitivity to Interest Rates

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

F.8 Rental Cost of Housing

We assume that the rent is a constant fraction of the home price. 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 32 shows that the level of the housing wealth effect rises but its time series remains unchanged if we make the polar opposite assumption that rents remain steady when home prices change. In our baseline, renters are more-or-less unresponsive to home price changes. Given that rents are changing along with home prices, this may seem surprising. There is, however, a counter-vailing force, which is that an increase in home prices causes some renters to delay purchasing homes and these individuals consume more as a result. This force is still present under the alternative assumption that rents remain constant when home prices change. In this case we find that renters actually consume more when home prices increase, and this is the main source of the difference in Figure 32.




Figure 32: Sensitivity to Rent-Price Ratio

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

F.9 Short-Term Debt

Figure 33 shows the elasticity by LTV in 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 they are “margin called” when house prices fall. The model consequently generates an increase in the wealth effect in the Great Recession as well as a substantial boom-bust asymmetry.

F.10 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. 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-elasticity city. At higher prices, this city will demand less housing than the more elastic city due to a movement along the supply curve. We assume that the price of housing rises by 10 percent more in the less elastic city




Figure 33: Elasticity by LTV for Short-Term Debt Model

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

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 in which we focus on 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 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 34 shows the housing wealth effect 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 effect is very similar to in our baseline case, though slightly lower, because the less elastic city is no longer substituting out of housing in the short-run. This arises 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).




Figure 34: No Short-Run Housing Adjustment.

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

F.11 Accounting for the Evolution of Household Balance Sheets

Our baseline analysis takes the distribution of individual states from the data. 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. 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 effect.

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 $t+1, \dots, t+K$. We begin a simulation with each wave of the SCF (i.e. 1983, 1986, ...) and simulate four years of data for each using the observed evolution of home prices.

Figure 35 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




Figure 35: 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.

fails is during the housing boom. During those years, the increases in home values would push LTV 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 the period. The model does not predict this increase in mortgage debt and 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. One interpretation of free 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 expectations of capital gains to be nonzero. Specifically, we assume the one-year expected capital gain rises from 0 to 2 percent in a linear fashion from 2004 to 2007. That is, in 2004 people expect home prices to be constant going forward and in 2007 they expect a 2 percent appreciation in 2008 followed by constant prices thereafter. Expected capital gains increase leverage for two reasons: first, homeowners feel richer and increase consumption due to a wealth effect and, second, the expected return on housing lowers the user cost of housing and prompts an




Figure 36: Simulated LTV Distribution

Note: The figure repeats the simulation from Figure 35 (baseline) and also shows results for the boom parameters and the boom parameters with expected capital gains for the period 2004-2007.

increase in the demand for housing financed with mortgage debt.

Figure 36 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 wealth effect (which should not be surprising given our results on changes in the credit constraint and no short-run adjustment earlier in the appendix). Indeed, in 2007, the baseline parameters lead to a housing wealth effect of 0.107, the boom parameters lead to 0.108, and the boom parameters with the expected capital gain leads to 0.115.