House Prices and Consumption: A New Instrumental Variables Approach *

James Graham †

November 15, 2018

Abstract

Fluctuations in house prices can generate large movements in household expenditure. However, empirical work exploring this relationship must deal with the endogeneity problems associated with using house prices as a regressor. A popular instrumental variables strategy exploits cross-sectional variation in house prices as predicted by the local housing supply elasticities of Saiz (2010). As an alternative, I introduce a Bartik instrument for house prices that consists of the interaction between the pre-existing local supply of housing characteristics and broad changes in the relative demand for those characteristics. I show that the instrument is a strong predictor of house price growth in both the cross-section and through time. I then use household panel data on non-durable expenditures to estimate the elasticity of consumption with respect to local house price growth. I report precise estimates in the range of 0.10 to 0.15, which correspond to marginal propensities to consume out of housing wealth of 1.2 to 1.8 cents in the dollar. These estimates are robust to controls for aggregate fluctuations, local business cycles, and local industry and demographic composition. In contrast, estimates I show that the traditional housing supply elasticity instrument produces inconsistent estimates when confronted with these same controls. Thus, the Bartik instrument succeeds in generating plausibly exogenous variation in house prices when housing supply elasticity instruments may fail. When decomposing variation in the Bartik instrument, I find that the identified consumption response to house prices is largely driven by times and locations where house prices varied the most: during the 2008 recession and in the Western US.

*First draft: June 2018. I would like to thank Simon Gilchrist, Virgiliu Midrigan, Mark Gertler, Tim Cogley, Gianluca Violante, Martin Rotemberg, Mike Gilraine, and Philipp Schnabl for many productive discussions. Many thanks to Skylar Olsen at Zillow Research for her help in understanding the data. Thanks also to seminar participants at London Business School, New York University, Yale, Brown, Washington University of St. Lousis, George Washington University, The Federal Reserve Bank of St. Louis, and the Reserve Bank of New Zealand for many useful comments.

†New York University. Email: james.graham@nyu.edu

1. Introduction

Fluctuations in house prices are thought to have a significant impact on the macroeconomy, especially through their effect on consumption expenditures. The household balance sheet channel is a commonly cited explanation for this effect, whereby movements in house prices induce wealth effects and changes in the value of collateral. The primary difficulty in trying to identify these effects empirically is that house prices are endogenous equilibrium objects. Instrumental variables strategies are one way to isolate potentially exogenous variation in these prices. To that end, this paper introduces a new Baritk instrument to study household consumption responses to local house price movements.

In recent literature, popular instrumental variables strategies have employed the housing supply elasticities of Saiz (2010), and the residential land use regulation indexes presented in Gyourko, Saiz, et al. (2008). These instruments use cross-city variation in the difficulty of building new houses to predict variation in house prices. Intuitively, relatively inelastic cities build few houses in response to increased demand, which results in higher house prices instead. Housing supply elasticity instruments have been shown to predict house price growth across cities (Mian and Sufi (2009)), which accounts for their popularity in studying the effect of local house price movements on consumption (see Mian, Rao, et al. (2013), Kaplan, Mitman, et al. (2016), Aladangady (2017)).

However these instruments have several drawbacks. First, they may not be exogenous to unobserved determinants of consumption demand. For example, Davidoff (2016) argues that inelastic cities tend to have more desirable amenities, which attracts highly skilled and also immigrant workers whose characteristics and income growth are likely to be correlated with consumption demand.[1] Second, the instruments provide only cross-sectional information, which limits their use for studying cyclical or time-series fluctuations in house prices and consumption.[2] Third, the instruments are measured for a set of cities (MSAs), which limits their use in studying house price changes for smaller units of geography.[3]

In order to address these issues, I introduce an alternative empirical strategy following the literature on Bartik instruments (Bartik (1991), and for recent surveys of the literature see Goldsmith-Pinkham et al. (2018) and Borusyak et al. (2018)). I instrument for local house price growth using the interaction of the local share of houses possessing particular characteristics with regional growth in the marginal prices of those characteristics. These include features that capture the overall quality of the house, such as age, number of bedrooms, or number of bathrooms. To the extent that the housing stock in each location is more concentrated in particular characteristics, the more exposed is that location to shocks to their relative prices. For example, suppose San Francisco consists mostly of two-bedroom houses built prior to the 1940s, while Las Vegas has mostly four-bedroom houses built in the early 2000s, then an increase in the price of larger and newer houses in the Western US would result in relatively faster house price appreciation in Las Vegas.

To construct the Bartik instrument, I make use of new housing data from the Zillow Transaction and Assessment Dataset (ZTRAX).[4] ZTRAX provides detailed information on millions of individual housing transactions from across the US. The geographic and property information in the dataset is used to compute the house characteristic shares for each location. Then using the sales prices associated

[1] Inelastic cities are frequently coastal, as can be seen in Table I in Saiz (2010). Inelastic cities also include many of the so-called "superstar" cities which report high housing demand growth, such as Los Angeles, San Francisco, and New York (see Gyourko, Mayer, et al. (2013)). Saiz (2010) also reports that inelastic cities have higher income levels, greater production of patents, and more tourist visits.

[2] The seminal paper of Mian, Rao, et al. (2013) is limited to studying a single cross-section of consumption and house price growth, covering the period from 2006 to 2009. Kaplan, Mitman, et al. (2016) consider other sample periods during the 2000s housing boom and bust, but also limit themselves to single cross-sectional regressions.

[3] Prior studies of the effects of house price changes at the zip code-level have had to resort to OLS estimation. For example, see Mian, Rao, et al. (2013) and Stroebel et al. (2014)

[4] Data provided by Zillow through the Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at http://www.zillow.com/ztrax. The results and opinions are those of the author and do not reflect the position of Zillow Group.

with each property transaction, I estimate hedonic pricing regressions with time-varying coefficients for houses sold in each of the four US Census regions. These coefficients provide estimates of the marginal prices of each house characteristic. I show that there is significant cross-sectional variation in the local shares and that the marginal prices of house characteristics are indeed time-varying. These results help the instrument generate both cross sectional and time series variation in house prices. Finally, note that because ZTRAX covers virtually all house sales in the US, it is straightforward to construct instruments for house prices at different levels of geographic aggregation. Although the main results in the paper use instruments constructed at the county-level, I demonstrate the use of a zip code-level instrument in results reported in the Appendix.

To study the effect of house prices on consumption, much of the literature uses geographically aggregated consumption measures. In this paper, I follow recent work by Aladangady (2017) which makes use of household-level panel data on consumption expenditures.[5] While Aladangady (2017) uses short-period panel data from the Consumption Expenditure Survey (CEX), I use the Nielsen Homescan Consumer Panel data which tracks individual households over longer periods of time.[6] The Consumer Panel data includes expenditures on mostly non-durable goods. Most importantly, it reports zip code and county information for each household, which enables me to match them to local house prices. I can then use this data to estimate household-level consumption elasticities with respect to local house price movements.

In first stage IV regressions, I show that the Bartik instrument is a strong predictor of local house price growth. However, the most important feature of the instrument is its exogeneity with respect to unobserved determinants of household consumption. Recent work by Goldsmith-Pinkham et al. (2018) shows that identification using Bartik instruments relies on exogeneity of the local shares, rather than exogeneity of the aggregated shocks.[7] In this case, the local shares of house characteristics must be uncorrelated with shocks to household consumption growth. The first argument in favor of this assumption is that the distribution of housing characteristics changes very slowly over time, since new construction is a small proportion of the total housing stock. This suggests that the composition of local housing is predetermined when shocks to consumption are realized.

However, it is possible that exposure of the local housing stock to variation in the relative demand for different house characteristics is correlated with households' exposure to consumption shocks. This could occur for a variety of reasons, but the most likely reflects the same problem of demographic composition that affects the exogeneity of the housing supply elasticity instruments (Davidoff (2016)). Specifically, the local shares may be correlated with the demographic makeup of the households who choose to live there. If demographic factors predict exposure to consumption shocks, then the local housing shares, and thus the Bartik instrument, will also be correlated with those shocks. Although it is not possible to test for endogeneity of this sort directly, I check the stability of the estimated consumption elasticities after conditioning on a range of observable local demographic controls.[8] These controls include factors that may predict differential consumption exposures, such as average age, education, immigration status, homeownership, family size, and so on. Since conditioning on these variables leads to little change in estimated consumption responses, the Bartik instrument is unlikely to be unduly influenced by unobserved demographic factors.

[5] Campbell et al. (2007) create a synthetic panel data set using repeated cross-sections of households in the UK matched to local house price movements.

[6] All results are the researcher's own analysis calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

[7] When exogeneity of the local shares cannot be assumed, Borusyak et al. (2018) provide an identification condition if exogeneity of the aggregate shocks can be assumed.

[8] The specific form of these controls follows advice in Goldsmith-Pinkham et al. (2018). See Section 6.2 for details.

The main empirical results of the paper report estimates of the elasticity of non-durable consumption expenditures with respect to local house prices. I directly compare IV estimates using the Bartik instrument and the housing supply elasticity instrument of Saiz (2010). I modify the latter instrument for use in a panel data context by interacting the cross-city elasticities with regional house price growth.[9] I show that the Bartik instrument consistently produces precise estimates of the consumption elasticity in the range of 0.1 to 0.15. This corresponds to marginal propensities to consume non-durable goods out of housing wealth of approximately 1.2 to 1.8 cents in the dollar. I show that these estimates are robust to the inclusion of controls for a variety of possible confounding factors: time fixed effects, local economic activity, local industry composition, as well as the local demographic factors discussed above. In contrast, the housing supply elasticity instrument produces consumption elasticities that are very sensitive to the inclusion of these controls, and that are often statistically insignificant from zero. The poor performance of the housing supply elasticity instrument reflects the lack of times series variation in the instrument, despite the interaction with regional house price growth. This is due to the fact that a shock to housing demand is simply scaled by local housing supply elasticity. Thus, the instrument cannot, for example, simultaneously generate both increasing and decreasing local house prices within a region.

In contrast, the Bartik instrument performs well because it draws on multiple sources of cross sectional and time series house price variation. To explore this in more detail, I present results of an over-identified instrument decomposition due to Rotemberg (1983) and recently applied to Bartik instruments by Goldsmith-Pinkham et al. (2018). The method decomposes an IV estimate using the Bartik instrument into a set of just-identified IV estimates associated with the combination of each of the local housing characteristic shares and the region-by-time growth rates of the marginal price of those characteristics. The set of just-identified estimates is associated with a set of Rotemberg weights, which account for the contribution of the shares and growth rates to the overall IV estimate. I show that variation in the Bartik instrument that explains consumption responses to house prices occurs in virtually all years and regions in the sample. Nevertheless, I find that variation in the instrument is largely associated with price variation in the West of the country and occurs mostly during the 2000s housing bust and subsequent recovery. For example, around 50 percent of the contribution to consumption fluctuations occurs in 2008 and 2009 alone, while nearly 30 percent is associated with the years from 2013 to 2015.

Finally, I discuss economic mechanisms that could generate the pattern of local house price movements described by the Bartik instrument. These mechanisms must account for broad-based changes in the relative demand for different house characteristics. For example, what could cause the price of older houses to rise relative to the price of younger houses? I argue that these relative price movements can be accounted for by changes in the purchasing power of the marginal buyers of houses with different characteristics. Consider, for example, recent work by Landvoigt et al. (2015) on the housing market of San Diego during the 2000s housing boom. They show that initially low-priced houses appreciated much faster during the 2000s than did initially high-priced houses. They attribute this to an increase in credit supply in the early 2000s, which had a disproportionate effect on lower income households who increased their demand for the low-priced houses for which they were marginal buyers. Revisiting the housing market of San Diego, I show that initially low-priced houses were significantly older and smaller than initially high-priced houses. Thus, changes in the purchasing power of the marginal buyers for these houses can induce relative increases in demand for these house characteristics, consistent with the house price variation generated by the Bartik instrument.

The remainder of the paper proceeds as follows. Section 2 discusses related literature, Section presents the data sources used in the analysis, Section 4 discusses the empirical strategy for estimating the consumption relationship to house prices, Section 5 explains the construction of the Bartik instrument, Section 6 presents the empirical results, Section 8 presents several robustness exercises, and

[9] In results reported in the Appendix, I also try a variant of the instrument interacted with national house price growth. In a similar instrument construction, Chaney et al. (2012) and Aladangady (2017) interact the housing supply elasticities with national real interest rates.

Section 9 concludes.

2. Related Literature

In the recent literature studying the relationship between consumption and house prices, a common strategy to address endogeneity concerns is to use an exogenous instrument for house price movements. By far the most popular approach uses cross-sectional housing supply elasticities produced by Saiz (2010). These are constructed for a number of cities using satellite-generated data on the difficulty of building on a particular topography as a proxy for housing supply constraints. Gyourko, Saiz, et al. (2008) present a related instrument via the Wharton residential land use regulation index. The Saiz (2010) instruments is assumed to be exogenous to other determinants of housing and consumption demand because it is generated by features of a city's geography, which is unchanging over the business cycle. They generate cross-sectional variation in house prices because more elastic cities meet higher housing demand with increased housing construction, while more inelastic cities build less and instead experience more rapid rises in house prices. Indeed, Mian and Sufi (2009) show that the instruments are strong predictors of house price movements during the 2000s housing boom and bust.

Mian, Rao, et al. (2013) were the first to use these housing supply elasticity instruments to explore the response of consumption to housing wealth. Using county-level data on the 2006 to 2009 growth in various consumption measures, their IV estimates imply elasticities with respect to house prices in the range of 0.13 to 0.21. They also report OLS-estimated results for non-durable consumption, with an elasticity equivalent to 0.09.[10]

Kaplan, Mitman, et al. (2016) replicate the findings of Mian, Rao, et al. (2013) using publicly available data sources. They make use of data from the Kilts-Nielsen Retail Scanner Dataset, which is similar to the Conusmer Panel data used in the current paper. The data sets differ in that the Scanner data provides store-level sales while the Consumer Panel provides household-level purchases. As in the Consumer Panel, the Scanner data reports mostly non-durables consumption expenditure. Following the same empirical model specifications as Mian, Rao, et al. (2013), Kaplan, Mitman, et al. (2016) report IV-estimated average elasticities of non-durable consumption of between 0.05 and 0.11.[11] The results for non-durable consumption using the Bartik instrument in the current paper are similar to those produced by Mian, Rao, et al. (2013) and Kaplan, Mitman, et al. (2016). I report IV-estimated consumption elasticities of between 0.1 and 0.15.

The purely cross-sectional Saiz (2010) instrument has been extended by Chaney et al. (2012) to study the effect of house prices on corporate investment, and by Aladangady (2017) to investigate the effect of house prices on consumption. The extended instrument is constructed by interacting the supply elasticity and long-term real interest rates. This provides both cross-sectional and time series variation in house prices since changes in real interest rates affect housing demand through the cost of mortgages, and the local house prices vary with how this demand channel interacts with local housing supply. Using micro-panel data from the CEX on total household consumption, Aladangady (2017) finds an IV-estimated MPC out of housing wealth for homeowners in the range of 3.4 to 5.8 cents in the dollar. These estimates are equivalent to elasticities with respect to house prices of 0.14 to 0.24.[12]. The Bartik instrument presented in this paper has similarities to these instruments in that it consists of the

[10] Note that Mian, Rao, et al. (2013) report elasticities of consumption with respect to housing net worth shocks: $\Delta logP_t \times \frac{H_{t-1}}{NW_{t-1}}$. The estimates for total consumption are 0.5 to 0.8, and the estimate for non-durable consumption is 0.34. To convert to elasticities with respect to house prices, I scale by the average housing wealth to total wealth ratio, which Mian, Rao, et al. (2013) report is 0.26.

[11] I use the same re-scaling as used for Mian, Rao, et al. (2013), above. The reported elasticities with respect to housing wealth shocks are in the range of 0.19 to 0.41.

[12] Aladangady (2017) uses a measure of changes to housing net worth given by $\Delta w_{i,t+1} = \Delta \log(P_{h,t+1})\frac{P_{h,t}}{C_{i,t}}$. To convert to an elasticity out of house prices, I multiply the estimated MPC by the ratio of average house prices to average consumption, which Aladangady (2017) reports as $\frac{\$156,653}{\$37,873} = 4.14$

interaction between broad changes in the demand for housing (i.e. through changes in relative demand for house characteristics) and the pre-existing cross-section of housing supply (i.e. the stock of housing characteristics available).

Guren et al. (2018) develop what they call a 'sensitivity' instrument for house prices.[13] The instrument reflects cities' differential exposure to regional house price cycles, where higher sensitivity cities experience greater price movements following regional demand shocks. As long as the cross-sectional sensitivity instrument is uncorrelated with the cross-sectional sensitivity of cities' consumption to other regional shocks, then the instrument provides exogenous variation in city-level house prices. Guren et al. (2018) use CBSA-level retail employment as a proxy for consumption, and report average IV-estimated elasticities for the period 1990-2015 of 0.06. Their procedure allows them to estimate time-varying elasticities, and they show that the elasticity may have been as high as 0.2 in the early 1990s, but has been much lower and relatively constant since the mid-1990s.

The Bartik instrument presented in this paper also has similarities to the sensitivity instrument in Guren et al. (2018). For both instruments, exogenous variation in local house prices is identified from the local exposure or sensitivity to regional house price movements. Guren et al. (2018) interpret the variation in their instrument as reflecting cross-sectional variation in current or expected housing supplies. The Bartik instrument is more explicit about the relationship to local housing supplies, since it is directly constructed from the relative supplies of different house characteristics. Thus, local house prices are more exposed to regional price cycles the more concentrated is the distribution of housing in those characteristics most subject to changes in relative demand.

In addition to estimating average consumption responses to house prices, recent papers in the literature test for the presence of the collateral channel of house prices directly. Households with higher initial debt burdens should be more sensitive to house price changes since they are more likely to be borrowing constrained, and price changes can tighten or relax borrowing constraints through changes in the value of housing collateral. Mian, Rao, et al. (2013) compare estimated MPCs across zip codes with different housing loan to value (LTV) ratios. They find that zip codes with LTVs above 90 percent have MPCs three times larger than zip codes with LTVs of 30 percent. Aladangady (2017) conducts a similar exercise, testing for differences in MPCs among households with LTV ratios above and below 80 percent. He finds MPCs almost twice as large for households with high LTV ratios. In this paper, I test for the effect of the collateral channel by exploiting geographic heterogeneity in LTV ratios in the ZTRAX data prior to the recent housing bust. I stratify households by their zip code's average LTV ratio for mortgage originations from 2004 to 2006. I find that households in zip codes with initial average LTV ratios above 0.8 are 65 to 75 percent more responsive to house prices than other households, which corresponds to MPCs as high as 3 cents in the dollar.

3. Data

3.1. Housing Data

House price data come from the Zillow Transaction and Assessment Dataset (ZTRAX), made available by Zillow Research.[14] The full ZTRAX dataset contains more than 370 million public records from across the US and includes information on deed transfers, mortgages, property characteristics, and geographic information for residential and commercial properties.

I restrict the data to observations on arm's-length, non-foreclosed sales of residential properties made by owner-occupiers. I exclude data from Rhode Island, Tennessee, and Vermont entirely due to missing data problems. Additionally, several states are missing price data for a large number of transactions.[15]

[13]This follows a similar empirical strategy introduced by Palmer (2015).

[14]The conclusions drawn from the ZTRAX dataset are those of the researcher and do not reflect the views of Zillow. Zillow is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

[15]Large numbers of observations are missing price data in Alaska, Idaho, Indiana, Kansas, Maine, Mississippi, Montana,

When working with house prices, I exclude all observations with missing data or where the sale price is less than \$10,000. However, when a transaction price is missing, the property's house characteristics can still be used in the construction of the Bartik instrument. Specifically, I use this data to compute the share of houses with different characteristics in each location. These characteristics include the age of houses, the number of bedrooms, the number of bathrooms, the floor size, and the property lot size. Other house characteristics are are reported in ZTRAX, however many of these fields are not broadly populated, and are not used in the analysis. After filtering, I work with 55 million observations between 1994 and 2016. The data filtering procedure is described in detail in the Appendix.

3.2. Consumption Data

Household-level consumption data come from the Nielsen Consumer Panel data. The panel runs from 2004 to 2016 and contains between 40,000 and 60,000 households each year. Households report, via an in-home scanning device, the price paid for and quantity purchased of all goods bought during their time in the survey. I aggregate these purchases into household-level annual expenditure. The data reports on approximately 1.5 million unique goods, which account for approximately 30 percent of all household consumption categories (Marketing (2016)). These goods are largely non-durables from the following categories: health and beauty, dry grocery, frozen foods, dairy, deli, packaged, meat, fresh produce, non-food grocery, alcohol, general merchandise. Importantly, the data reports the state, county, and zip code in which each household lives. Each household can then be linked to a measure of local house prices. And with this information, I can estimate household-level consumption responses to local house price changes.

Although the Consumer Panel reports typical demographic information associated with each household, homeownership status is not one of these variables. In order to infer homeownership status, I follow the procedure in Stroebel et al. (2014) who also use the Consumer Panel data. Households report whether they live in a one-, two-, or three-family dwelling, and also whether the house is a condo or co-op. I assume that single-family, non-condo/co-op residences are inhabited by homeowners, while remaining households are assumed to be renters. The average proportion of households living in single-family homes is 75 percent, and does not change significantly across sample years. For comparison, the national homeownership rate fell from 69 percent in 2004 to 64 percent in 2015.[16] Because only homeowners should experience the wealth and collateral effects of house prices on consumption (see Buiter (2008)), I restrict the empirical analysis to the sample of inferred homeowners. Households in the data are occasionally observed to move across geographies, although this is less common than is the case in the general population (see Table A.2 in the Appendix). Because consumption patterns are likely to differ for movers and non-movers, I restrict the sample to those households that never move during their time in the panel.

Summary statistics for the demographic composition and consumption patterns of households in the Consumer Panel are reported in the Appendix.

3.3. Additional Data Sources

County-level house price indexes come from the Federal Housing Finance Agency, city (MSA) and zip code house price indexes come from Zillow's publicly available ZHVI All Home price indexes.[17]

New Mexico, Texas, Utah, and Wyoming due to either non-mandatory disclosure of or outright prohibitions on the reporting of transactions prices. See http://www.zillowgroup.com/news/chronicles-of-data-collection-ii-non-disclosure-states/ for more details.

[16] Homeownership rates for the United States are from FRED (code: USHOWN).

[17] Although the ZTRAX data is a rich source of individual housing transactions, zip code, county, and city data have varying degrees of completeness, which presents difficulties in constructing broad and consistent house price indexes. I instead rely on the published house price indexes of the FHFA and Zillow.

To account for changes in the general level of prices, I use the CPI for all urban consumers from FRED. Average after tax income by zip code is computed from the IRS Statistics of Income (SOI), using the adjusted gross income variable less total tax payments. County-level unemployment data is collected from the BLS Local Area Unemployment statistics. Zip code and county-level demographic information is computed from the 2000 Decennial Census. I also use annual county-level employment by industry from Country Business Patterns data. I aggregate employment using the 6 digit NAICS codes into broad categories for construction (NAICS: 23), retail trade (NAICS: 44, 45), and finance/insurance/real estate (NAICS: 52, 53). A detailed list of all data sources is reported in the Appendix.

4. Theory and Empirical Approach

Aggregate house prices and consumption are known to co-move over the business cycle. But because house prices are endogenous equilibrium objects, it is difficult to distinguish the effect of price movements on consumption from the effect of underlying macroeconomic shocks. Nevertheless, economic theory suggests direct a relationship from house prices to household consumption through wealth and collateral effects on household balance sheets.

Owner occupied housing is a large component of household wealth, and its market value rises and falls with changes in house prices. Friedman's Permanent Income Hypothesis model of household consumption shows that an increase in permanent wealth leads to an increase in current expenditures (Friedman, 1957). However, housing wealth may be unlike other components of household wealth (Buiter, 2008). Owner-occupiers forgo rental income from leasing their homes, but also save on rental payments for housing services. If houses are priced according to the future stream of rental rates, then an increase in house prices reflects rising rental costs. Higher house values would mean that homeowners enjoy an increase in wealth on paper but also face higher implicit housing service costs. These offsetting effects imply that changes in house prices may have no net wealth effect on consumption. Despite this, many empirical papers in the recent literature find some support for a general wealth effect channel (Case et al. (2005), Campbell et al. (2007), Mian, Rao, et al. (2013), Kaplan, Mitman, et al. (2016), Aladangady (2017)).

While wealth effects may be zero in infinite horizon models, life-cycle models suggest that they may differ according to age (Campbell et al., 2007). For young households, future changes in income, family size, and preferences suggest a rising profile of housing consumption over their life time. In contrast, older households may expect to decrease housing consumption in the near future. Thus, rising house prices reflect a decrease in the present discounted value of wealth for young households since they expect to purchasing more housing in the future. For older households, an expected decline in housing consumption means that higher house prices reflect an increase in the present discounted value of wealth. Although Campbell et al. (2007) find empirical support for this age-dependent wealth channel, Attanasio et al. (2009) present contradictory evidence.[18]

The effect of a change in wealth on consumption can be non-linear. Consumption functions are concave in current wealth when households are faced with uncertainty and either convex marginal utility (Carroll and Kimball, 1996) or borrowing constraints (Carroll, 2001). Then, households can have very large marginal propensities to consume (MPCs) out of shocks to wealth.[19] Borrowing constraints are especially relevant in the context of housing wealth since houses act as collateral for mortgage borrowing. For example, rising house prices can relax borrowing constraints by reducing current loan-to-value ratios, which allows households to refinance mortgages or make use of home equity lines of credit (Chen et al., 2013). Thus, we expect that the expenditure of households with little wealth or who are close to relevant borrowing constraints is especially sensitive to changes in house prices. Recent empirical work

[18] Mian and Sufi (2011) also find little evidence for differences across age in borrowing repsponses to house prices.

[19] Large MPCs are even possible for high net-worth households if they are otherwise borrowing constrained. Consider, for example, the wealthy hand-to-mouth households in Kaplan, Violante, et al. (2014).

provides evidence that this collateral channel of house prices has a strong effect on borrowing (Mian and Sufi, 2011; Cloyne et al., 2017) and consumption (Mian, Rao, et al., 2013; Aladangady, 2017).

The empirical literature has employed two reduced form model specifications for investigating the effect of house prices on consumption. One model estimates the elasticity of consumption with respect to house prices, and the other estimates the MPC out of housing wealth. While the former approach requires only data on consumption and house prices, the latter also requires information about the composition of wealth. Because the household panel data I use this paper lacks information on household balance sheets, I opt to estimate the elasticity of consumption with respect to house prices. The literature has also employed various combinations of zip code-, county-, and city-level house prices when investigating the effect of prices on consumption. I show that it is possible to estimate effects for house prices at any level of geography using the Bartik instrument introduced in this paper. However, to conserve space I only report results for county-level house prices in the body of the paper.[20].

I primarily work with a simple linear relationship between consumption growth and house price growth. I later test for non-linearity in the relationship due to the influence of collateral constraints. The reduced form model is:

$$\Delta c_{i,g,t} = \beta_1 \Delta p_{g,t} + \beta_2 x_{i,t} + \beta_3 y_{g,t} + \alpha_{cbsa} + \alpha_t + \varepsilon_{i,g,t}, \tag{1}$$

where i denotes an individual household, g denotes the location of that household (e.g. zip code, county, city), and t denotes the year of observation. $\Delta c_{i,g,t}$ is the annual log-change in real household consumption expenditure, $\Delta p_{g,t}$ is the annual log-change in real local house prices, and β_1 is the elasticity of consumption with respect to local house prices. Household controls $x_{i,t}$ include a typical set of demographic characteristics, including real income growth. Local controls $y_{g,t}$ include business cycle variables, industry employment composition, and local demographic variables. City-level (CBSA) fixed effects α_{cbsa} are included since county-level house prices tend to be strongly correlated within CBSAs. Time fixed effects α_t are included when controlling for common movements in house prices and consumption. See Appendix F for a full description of control variables.

4.1. *Identification*

Despite the inclusion of household and local controls, OLS estimates of the consumption elasticity in Equation (1) may be biased by endogeneity or measurement error in house prices. Endogeneity problems arise if unobserved shocks in the error term $\varepsilon_{i,t}$ are correlated with both consumption and house price growth. For example, a positive shock to local labor productivity increases household income which simultaneously increases consumption and demand for housing, which increases house prices. Similar effects could be generated by decreases in local property taxes, or increases in migration. These unobserved shocks would lead to an upward bias in OLS estimates of the consumption elasticity. On the other hand, mis-measurement of house prices could lead to a downward bias in OLS estimates. An additional problem for the interpretation of Equation 1 is reverse causality: increases in household consumption can lead to increases in local income which increase house prices.

I address these problems by estimating Equation 1 with an instrumental variables strategy. Consider an instrument $B_{g,t}$ for house price growth $\Delta p_{g,t}$. The full empirical model is then:

$$\Delta c_{i,g,t} = \beta_1 \Delta p_{g,t} + \beta_2 x_{i,t} + \beta_3 y_{g,t} + \alpha_{cbsa} + \alpha_t + \varepsilon_{i,g,t}, \tag{2}$$

$$\Delta p_{g,t} = \gamma_1 B_{g,t} + \gamma_2 x_{i,t} + \gamma_3 y_{g,t} + \tilde{\alpha}_{cbsa} + \tilde{\alpha}_t + \eta_{i,g,t}, \tag{3}$$

$$\text{cov}(B_{g,t}, \varepsilon_{i,g,t}) = 0, \tag{4}$$

where Equation 3 is the first stage regression, and 4 is the exclusion restriction for the instrumental variable. Notice that the first stage regression includes all of the household controls included in the second stage, although in practice these have no effect on the first stage estimates.

[20] Estimates using zip code- and city-level house prices are reported in the Appendix.

The instrumental variable $B_{g,t}$ follows the Bartik instrument construction (Bartik, 1991). In the canonical setting, local employment growth is instrumented with a variable that consists of the interaction between local industry employment shares and national industry employment growth. Here, the Bartik instrument for house prices is constructed from the interaction between the local share of houses with particular characteristics (e.g. age and size), and broad changes in the marginal prices of house with those characteristics. I assume that house characteristics such as age, size, or number of bedrooms are a measure of house quality (see Section 7), so I interchangeably refer to characteristics and qualities. I interpret the local housing shares as the supply of particular house qualities, and changes in the marginal prices reflect shocks to the relative demand for those qualities. So for example if the housing stock in San Francisco consists mostly of two-bedroom houses built prior to the 1940s, while Nevada is mostly four bedroom houses built in the early 2000s, then an increase in the demand for larger and newer homes would result in relatively faster house price appreciation in Nevada than San Francisco.

The identifying assumption for estimating the consumption elasticity β_1, expressed in Equation (4), states that the Bartik instrument must be uncorrelated with unobserved shocks to household consumption growth. Notice that the inclusion of time fixed effects rules out any pure time-series correlation between the two variables. Thus, identification depends crucially on the cross sectional relationship between the instrument and consumption shocks. Indeed, Goldsmith-Pinkham et al. (2018) show that in general Bartik instruments are equivalent to using the local shares as instruments directly, while the aggregated shocks contribute only to the relevance of the instrument.[21][22] Because identification depends on the exogeneity of the local shares, it must be that the composition of the housing stock in location g is uncorrelated with the exposure of households in g to unobserved shocks.[23] Note that the identification condition would be satisfied if shocks to household consumption were idiosyncratic and uncorrelated across households within locations.

The first argument for exogeneity of the local shares is that they are pre-determined at the time that shocks to household consumption are realized. This is true mechanically because I compute the shares from data prior to the sample period for the main empirical analysis, as is in keeping with the use of Bartik instruments in the literature (see Section 5.1). However, it is also the case that local shares change very slowly over time since new housing construction is a small proportion of the existing housing stock. This is verified in Tables I and II, which report the within-county correlation of shares of houses built in various year groups and shares of houses with different numbers of bedrooms. The correlations are extremely high even after 10 and 15 years: 94 percent of the 2015 housing age distribution can be predicted from the 2000 housing age distribution. This means that the composition of the housing stock does not respond to shocks that affect household consumption in those locations.

The second factor that is important for exogeneity of the local shares is that they be uncorrelated with the demographic composition of the households in each location. We might worry that household types and the housing stock are correlated across locations since different households may choose where to live based on the different types of houses available. Exogeneity of the instrument would then be violated if, for example, highly educated households cluster in large cities where houses are older and smaller, and these households face particular consumption shocks. This kind of correlation is precisely what Davidoff (2016) argues is a problem for the exogeneity of housing supply elasticity instruments.

As a simple test of the identifying assumption, Table III reports the cross-county correlation between the share of houses built in different decades and a variety of demographic characteristics from the 2000 Census. Although no correlation is greater than 0.5 in absolute magnitude and the mean correlation

[21] This interpretation comes from a representation of the Bartik instrument as an over-identified GMM estimator, with the shares as instruments and the associated shocks as a particular weighting matrix.

[22] Alternatively, Borusyak et al. (2018) show that Bartik instruments are equivalent to just-identified instrumental variables regressions at the level of aggregation of the shocks rather than that at the local level at which the regression outcomes are observed. In that case, the local shares may be endogenous, but identification still obtains if the shocks are exogenous.

[23] Notice that identification here is similar to the identification conditions for the regional sensitivity house price instrument in Guren et al. (2018).

across all house characteristics and demographics is 0.015, I find that counties with a higher proportion of new houses have: higher homeownership rates, higher rates of college educated households, more white households, fewer black households, and fewer immigrant households. Given that some correlation is present, I test the plausibility of the exclusion restriction on the IV-estimated elasticities by reporting estimation results that include demographic characteristics as control variables (see Section 6.2). I show that the inclusion of these controls has little effect on the IV-estimated consumption elasticities, suggesting that demographic factors do not pose a significant problem for identification.

Table I
Correlation of Age Shares

	2000	2005	2010	2015
2000	1.00			
2005	0.98	1.00		
2010	0.95	0.98	1.00	
2015	0.94	0.98	0.95	1.00

Table II
Correlation of Bedroom Shares

	2005	2010	2015
2005	1.000		
2010	0.998	1.000	
2015	0.997	0.997	1.000

Notes: Correlation of county shares of houses built by age and number of bedrooms, computed across 5-year intervals. Data taken from the 2000 Census, which reports data for all dwellings, and the 5-year American Community Survey in 2005, 2010, and 2015, all of which report data for occupied dwellings only. 775 counties available across all four surveys. Data on bedrooms is only available from the ACS in 2005-2015. Age-groups chosen for consistency across surveys and include: built prior to 1939, built from 1940 to 1959, built from 1960 to 1979, and built from 1980 to 1999. Bedroom groupings are: 0 bedrooms, 1 bedroom, 2 or 3 bedrooms, 4 or more bedrooms.
Sources: ACS.

Table III
Correlations of Local Characteristic Shares and Local Demographics

	Years built							
	pre-1940	1940-1949	1950-1959	1960-1969	1970-1979	1980-1989	1990-1999	2000-2005
Frac. Homeowner	-0.40	-0.21	-0.15	-0.06	0.23	0.06	0.36	0.29
Frac. College+	-0.24	-0.13	-0.03	0.08	0.16	0.26	0.23	0.00
Frac. White	-0.19	-0.27	-0.24	-0.16	0.14	0.00	0.28	0.27
Frac. Black	0.27	0.24	0.17	0.09	-0.28	-0.11	-0.27	-0.21
Frac. Hispanic	-0.13	0.04	0.11	0.04	0.07	0.15	-0.02	-0.01
Frac. Foreign born	0.08	0.12	0.19	0.13	0.04	0.11	-0.18	-0.24
Median Age	0.05	-0.05	-0.04	0.03	0.10	0.06	-0.08	-0.10
Mean HH. Size	-0.24	-0.01	0.08	0.05	0.13	0.09	0.14	0.07
Mean Commute	0.18	0.02	0.02	-0.02	-0.12	0.06	-0.11	-0.12

Notes: Correlation between county shares for housing characteristics and county demographics from the 2000 Census. Correlations computed for 1674 counties, weighted by Census population counts.
Sources: Author's calculations using 2000 Census, ZTRAX.

5. Construction of the Bartik Instrument

Following the discussion in Goldsmith-Pinkham et al. (2018), a useful way to understand the construction of the Bartik instrument is to consider an inner product representation of the growth in house prices $\Delta p_{g,t}$. I decompose house prices into the interaction between local housing shares and local quality price growth rates $\Delta q_{g,c,t}$:

$$\Delta p_{g,t} = \sum_c \lambda_{g,c,t} \Delta q_{g,c,t}, \tag{5}$$

where $\lambda_{g,c,t}$ is the share of houses in location g with characteristic c at time t, and $\Delta q_{g,c,t}$ is quality price growth for houses with characteristic c in location g at time t. House price growth is then given by changes in quality prices weighted by the proportion of those qualities in a particular location. Suppose, for example, that there is one location and one time period, that there are only two housing types (small and large), that the share of small houses in a location is λ_S, and that price growth for each type is $\Delta q_{g,S}$ and $\Delta q_{g,L}$. Then overall house price growth is $\Delta p_g = \lambda_S \Delta q_{g,S} + (1 - \lambda_S)\Delta q_{g,L}$.

Quality prices can then be decomposed as:

$$q_{g,c,t} = q_g + q_{c,t} + \tilde{q}_{g,c,t}, \tag{6}$$

where q_g is a location fixed effect, $q_{c,t}$ is a characteristic-time component, and $\tilde{q}_{g,c,t}$ is an idiosyncratic location-characteristic-time component. Willingness to pay for housing qualities over time depends on location, the qualities themselves, and interactions between the two. For example, poor rural areas are less able to pay for any given characteristic yielding a low value of q_g. Large houses are relative luxuries, meaning that $q_{c,t}$ is high for large houses when income is high. But since rural areas already have a lot of space they put less of a premium on large houses so that $\tilde{q}_{g,c,t}$ is relatively low for large houses in rural areas when income is high.

The location and idiosyncratic components of quality prices are likely to be correlated with local shocks to household consumption growth. Thus, we form the Bartik instrument using only the characteristic-time component $\Delta q_{c,t}$ of quality prices. I also restrict the local shares to an initial period: $\lambda_{g,c} = \lambda_{g,c,0}$. The instrument is then expressed as:

$$B_{g,t} = \sum_{c} \lambda_{g,c} \Delta q_{c,t}, \tag{7}$$

where $B_{g,t}$ denotes the Bartik instrument for house prices.

In practice, I modify Equation (7) to allow for separate house characteristics c, each with mutually exclusive categories i. I do this because houses consist of bundles of characteristics, such as house age and number of bedrooms. Recall that the canonical Bartik instrument uses industries as the only categorical variable, so the industry shares sum to one in each location. For the the housing Bartik instrument, the share of houses in category i with characteristic c is denoted $\lambda^i_{g,c}$, where $\sum_i \lambda^i_{g,c} = 1$ for each c in each location. Equation (7) can be rewritten as:

$$B_{g,t} = \sum_{c} \sum_{i} \lambda^i_{g,c} \Delta q^i_{c,t}, \tag{8}$$

where $\lambda^i_{g,c}$ are the shares of houses in category i for characteristic c in each location g.[24]

5.1. Local Housing Characteristic Shares

I compute the local shares of housing characteristics using ZTRAX data on unique houses sold between 1994 and 2005. I divide the data associated with each house characteristic into several categories. Building age is split into decadal bins: {pre-1939, 1940-1949, 1950-1959, 1960-1969, 1970-1979, 1980-1989, 1990-1999, 2000-2005}.[25] The number of bedrooms is split into the categories {1, 2, 3, 4, 5+}. The number of bathrooms is split into the categories {0, 1, 2, 3, 4+}, where half-bathrooms are rounded down to the nearest whole-number category. The Appendix tests for the reliability of the local housing shares computed from ZTRAX data.[26].

Because the local housing shares provide identifying information for the instrument, there must be sufficient cross-sectional variation in these shares. Figure I presents the distribution of housing age across counties in the US. For ease of presentation, I report the proportion of houses in each county built prior to 1960, between 1960 and 1990, and between 1990 and 2005. The figure suggests that there is significant cross-county variation in house age. For example, counties in the North East and Mid West have particularly high proportions of houses built prior to 1960. Counties in the South (e.g. Texas) and also in parts of the West (e.g. Nevada and Arizona) have a relatively large proportion of houses built in the latter half of the twentieth century. Importantly, there is variation in the housing age distribution even within regions, notably in the Western US where inland counties have much newer housing stocks than the cities in the coastal states. Figures IX and VIII in the Appendix present the cross-city (CBSA) and cross-zip code distributions of housing age. There is significant variation in these geographies also, which suggests that house price instruments can be constructed at all levels of geographic aggregation.

Because Figures I shows higher (lower) concentrations of new (old) housing in states where new construction occurred in recent years, it may be that the local shares are correlated with housing supply

[24]Equation (11) seems to suggest that the various characteristics are unweighted in the relation to house prices: since the category shares for each characteristic sum to 1, then house age is no more or less important than house size. However, the relative importance of each characteristic is embodied in the quality prices $\Delta q^i_{c,t}$. Equivalently, I could normalize the quality prices and apply a characteristic-specific weighting factor to each of them.

[25]This categorization broadly corresponds to the categories reported in the 2000 Census and subsequent American Community Surveys.

[26]I find that building age is well-measured relative to data in the American Community Survey, although data on the number of bedrooms may be less reliable. Section 8 considers a version of the Bartik instrument using only the local shares for housing age

Table IV
Correlations of Local Characteristic Shares and Housing Supply

Built	Saiz	Wharton	Beds	Saiz	Wharton	Baths	Saiz	Wharton
pre-1940	-0.03	0.10	One	-0.16	0.16	Zero	0.06	-0.08
1940-1949	-0.15	0.08	Two	-0.20	0.24	One	-0.01	0.07
1950-1959	-0.27	0.12	Three	0.21	-0.14	Two	-0.10	0.09
1960-1969	-0.19	0.07	Four	-0.18	0.06	Three	-0.12	0.09
1970-1979	-0.22	0.08	Five+	-0.36	0.30	Four+	-0.20	0.04
1980-1989	-0.22	0.11						
1990-1999	0.22	-0.13						
2000-2005	0.37	-0.25						

Notes: Correlation between local shares for housing characteristics and measures of local housing supply. Correlation computed for cities (CBSAs), using 240 observations available for all three sets of instruments, weighted by population as at the 2000 Census.
Sources: Author's calculations using 2000 Census, Gyourko, Saiz, et al. (2008), Saiz (2010), ZTRAX.

elasticities. Table IV reports the population-weighted correlations between the local housing characteristic shares used in the Bartik instrument, and the housing supply elasticities from Saiz (2010) and the Wharton residential land use regulation indexes in Gyourko, Saiz, et al. (2008). The local shares are only weakly correlated with the two measures, suggesting that the Bartik instrument provides largely independent variation to that generated by housing supply elasticity instruments. Interestingly, however, the share of houses built in the decades to 1999 and 2005 are (weakly) positively correlated with housing supply elasticities, while a (weak) negative correlation is reported for the shares built in other decades. This is consistent with economic intuition that high elasticity locations should have built relatively more houses during the recent house price boom.

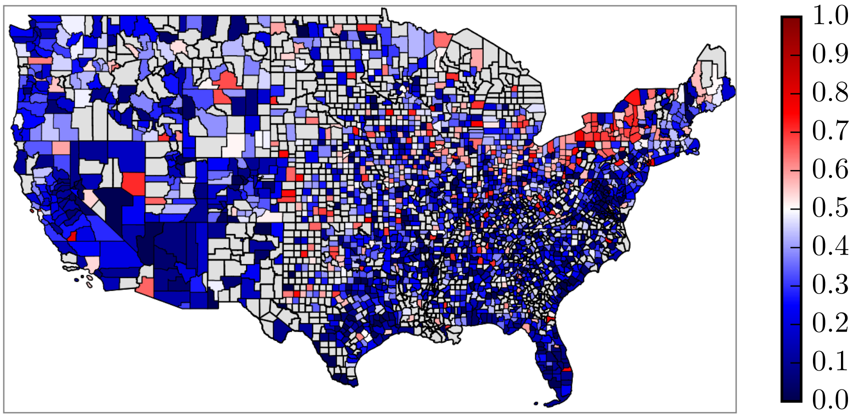

(a) Proportion Built Prior to 1960

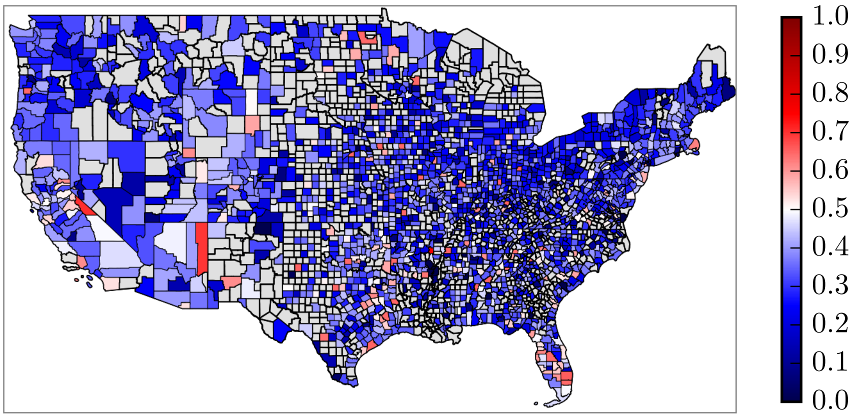

(b) Proportion Built 1960 to 1990

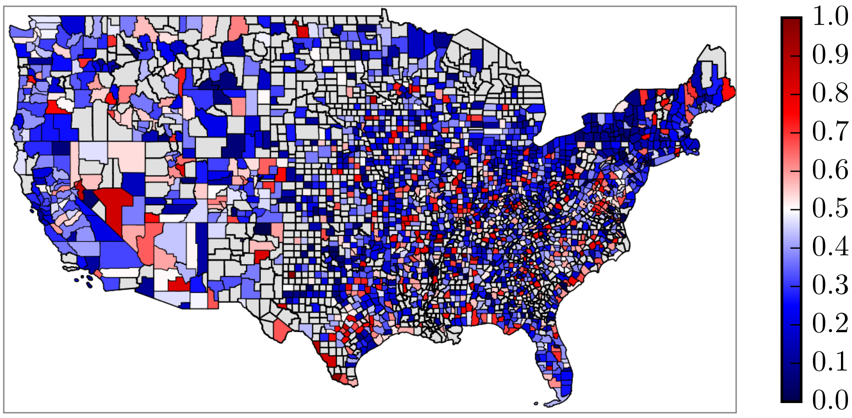

(c) Proportion Built After 1990

Figure I
Distribution of Housing Age Across Counties

Notes: Maps show the within-county proportion of all unique houses sold between 1994 and 2005 that were built before 1960, between 1960 and 1990, and between 1990 and 2005. Maps show 1283 counties for which there are at least 100 individual transactions of unique houses from 1994 to 2005.
Sources: ZTRAX.

5.2. *Housing Quality Prices*

I estimate housing quality prices using simple hedonic price regressions. These regressions condition on the same set of house characteristics and categories as used to compute the local housing shares, above. The hedonic pricing regression model is:

$$
\begin{aligned}
p_{i,g,t} = {} & \alpha_g + q'_{b,t}\mathbb{1}(b \in B) + q'_{h,t}\mathbb{1}(h \in H) + q'_{d,t}\mathbb{1}(d \in D) \\
& + \beta_t^f f_i + \beta_t^l l_i + \eta_{i,g,t}
\end{aligned} \tag{9}
$$

where $p_{i,g,t}$ is the log of the real house price for property i in location g, α_g is a county-specific fixed effect, $\mathbb{1}(b \in B)$ is a dummy variable for the number of bedrooms, $\mathbb{1}(h \in H)$ is a dummy variable for the number of bathrooms, $\mathbb{1}(d \in D)$ is a dummy variable for the decade in which the house was built. I allow for time-varying coefficients in these regressions so as to capture the characteristic-time component $q_{c,t}$ of quality prices in Equation (6).

The additional variables f_i and l_i are the log of floor size in square feet and the log of property lot size in square feet. These house size variables are not included in the Bartik instrument, since size is a continuous variable and does not have an obvious natural categorization from which to compute local shares. However, conditioning on these variables means that the other regression coefficients can be interpreted as the marginal price of characteristics holding size constant. For example, I find that the coefficients on the number of bedrooms are typically negative and monotonically decreasing, suggesting an additional bedroom crowds out other living space in a house of otherwise fixed size.

Equation (9) is estimated separately for each Census region in the US: Mid-West, North-East, South, and West. Grouping houses by region, rather than for the entire country, produces more variation in the instrument. This construction is also similar to the regional house price sensitivity instrument introduced in Guren et al. (2018). Their instrument relies on the fact that there have been significantly different regional house price cycles over the last 30 years. As discussed in Section 8, I also consider a version of the Bartik instrument which estimates Equation (9) at the national level.

In keeping with the construction of Bartik instruments in the literature, the hedonic pricing regressions are estimated using a leave-one-out procedure. For the instrument in location g, I drop all of the observations for that location and then estimate the coefficients in (9). This ensures that housing market movements in a particular location do not dominate the variation in its own instrument. In practice, counties are small relative to the surrounding region so the leave-one-out procedure has virtually no effect on either the instruments or the estimated consumption elasticities. The hedonic regressions explain a significant proportion of cross-sectional variation in house prices, with a median R-squared statistic of 0.6.

Figure II shows quality price growth for houses constructed in different decades. The horizontal axis plots the decade in which a house was built and the vertical axis reports housing quality price growth. Note that all values are reported relative to a house built prior to 1940 with one reported bedroom and zero reported bathrooms. These figures show, for example, that in the South between 2006 and 2009, houses built in the 1950s experienced a relative fall in price of more than 20 percent. During the same period, houses in the Mid West built in the 1980s experienced relative appreciation of around 12 percent. Altogether these quality price changes describe the patterns of price variation for houses with different features. To the extent that a location has a large proportion of houses with characteristics that appreciated by relatively more, that location should experience faster overall house price growth.

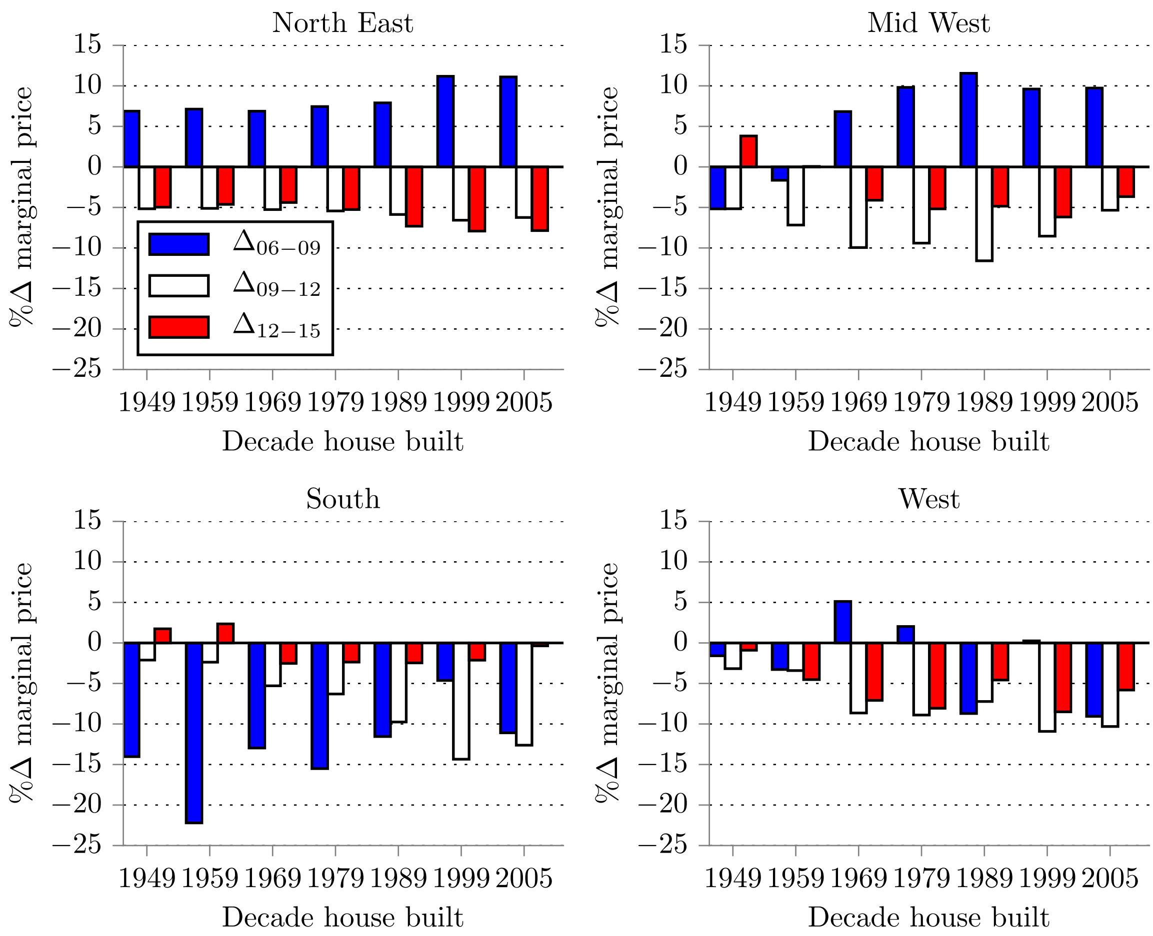

Figure II
Change in Marginal House Prices By Housing Age

Change in marginal house prices according to decade of house construction. Constructed from the coefficients estimated in Equation 9. Growth rates are interpreted as the marginal price changes for a house with the given characteristic relative to a house built prior to 1939, with one bedroom, and zero bathrooms. Growth rates are calculated for 2006-2009, 2009-2012, and 2012-2015.
Source: Author's calculations using ZTRAX.

6. The Elasticity of Consumption with Respect to House Prices

6.1. Strength of the Instrument

I first evaluate the strength of the Bartik instrument by estimating the first stage regression, Equation (3). Figure III presents a simple binned scatter plot of the residualized instrument against residualized house price growth, after projecting out the exogenous variables.[27] The figure shows that despite the inclusion of a large number of control variables, there remains a tight relationship between the instrument and house prices.

More formally, Table V presents the results of the first stage regressions under various specifications (these follow the specifications presented in Section 6.2). Column (1) reports results conditioning on household variables and CBSA fixed effects, column (2) adds time fixed effects, column (3) adds local

[27] I use the specification that includes all household, local, industry, and demographic controls as well as CBSA and time fixed effects. See Section 6.2 for details.

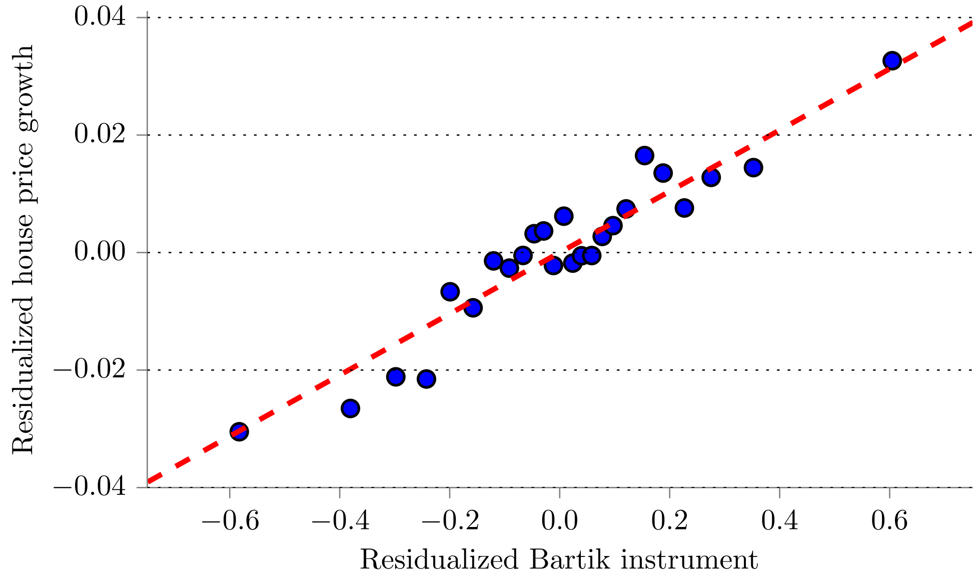

Figure III
Effect of Bartik Instrument on House Price Growth

Notes: Residualized Bartik instrument and county house price growth, representing the first stage regression. The residualized variables are constructed using the same household-level data and include the full set of controls as in the IV estimation of the consumption elasticities. See Appendix F for a list of control variables. The value of the Bartik instrument is split into equal sized bins, where the mean of the instrument and house prices is computed for observations falling within each bin. The red dashed line plots the first stage regression coefficient on the Bartik instrument. The slope is 0.050 as reported in column (5) of Table V.
Sources: Author's calculations using data from BLS, CBP, 2000 Census, FHFA, IRS SOI, Nielsen, ZTRAX.

business cycle controls, column (4) adds local industry employment shares, and column (5) adds local demographic controls. I find that the first stage relationship between the instrument and house prices is strong at standard levels of statistical significance and consistently estimated across specifications. Moreover, the first stage F-statistics and R-squared statistics are large in every case. These results suggest that the relationship between the instrument and house prices is not simply an artifact of common time series variation in house prices, nor is it only due to local business cycle variation, local industry composition, or demographics across locations.

6.2. *Estimation Using Bartik and Housing Supply Elasticity Instruments*

I estimate the elasticity of consumption with respect to house prices using the Bartik instrumental variables strategy. All estimates use household panel data for homeowners, county-level house prices, and cover the sample 2005 to 2016. I compare these results to estimates using the housing supply elasticity instrument of Saiz (2010), which I refer to throughout as the Saiz instrument.

The Saiz instrument is often used in cross-sectional settings, such as the 2006 to 2009 long-difference specification in Mian, Rao, et al. (2013). In this paper I use household level panel data with annual observations on consumption growth. To construct an appropriate Saiz instrument for this setting, I follow Guren et al. (2018) in interacting the cross-sectional housing supply elasticity instruments with regional house price growth. The use of regional house price growth means the Saiz instrument is comparable to the Bartik instrument, since they make use of the same source of time series variation

Table V
Effect of Bartik Instrument on County House Prices

	Annual county real house price growth				
	(1)	(2)	(3)	(4)	(5)
$B_{county,t}$	0.093***	0.059***	0.058***	0.060***	0.052***
	(0.008)	(0.010)	(0.010)	(0.010)	(0.008)
Observations					
Total	256,824	256,824	230,865	227,490	227,490
Households	54,551	54,551	51,408	51,000	51,000
CBSAs	216	216	216	216	216
Counties	571	571	570	563	563
Controls					
Household	✓	✓	✓	✓	✓
CBSA FE	✓	✓	✓	✓	✓
Time FE	×	✓	✓	✓	✓
Local	×	×	✓	✓	✓
Industry	×	×	×	✓	✓
Demographic	×	×	×	×	✓
F-statistic	127.25	36.02	37.65	38.33	43.97
Adj. R^2	0.34	0.62	0.63	0.64	0.76

Notes: Models estimate the effect of the Bartik instrument on county house prices, as in the first stage regression of Equation 3. The Bartik instrument is constructed as described in Section 5. The first stage regressions use the same household-level data and include the same household, local, industry, and demographic controls as in the full IV estimation of the consumption elasticities. See Section 6 for details. Both standard errors (reported in parentheses) and first stage F-statistic are clustered at the CBSA level. * significant at 10%, ** significant at 5%, *** significant at 1%.
Sources: Author's calculations using data from BLS, CBP, 2000 Census, FHFA, IRS SOI, Nielsen, ZTRAX.

in house prices.[28] This interaction between the Saiz instrument and a source of time series variation is similar to the approach in Chaney et al. (2012) and Aladangady (2017), who both use the cross-sectional elasticities interacted with national real interest rates.

The estimation results are reported in Table VI. It compares the IV-estimated elasticity of consumption with respect to house prices using the Bartik instrument to IV estimates using the Saiz instrument. The table consists of four pairwise comparisons, reflecting empirical specifications with different sets of controls. These control variables are described in detail in Appendix F. Columns (1) and (2) compare estimates using only household-level covariates and city (CBSA) fixed effects. The Bartik and Saiz estimates are 0.095 and 0.120, respectively. Both instruments produce precise estimates of the consumption elasticity, although they are statistically insignificant from each other.

Columns (3) and (4) add time fixed effects to the model. The fixed effects remove the influence of common movements in consumption, which are likely to have been significant during the late 2000s recession and recovery period. The Bartik estimate changes very little, to 0.106 and remains precisely estimated. In contrast, the Saiz estimate falls to -0.055 and is no longer distinct from zero. This suggests that much of the variation in the Saiz instrument that predicts household consumption is simply due to the co-movement between national house prices and the aggregate economy.

Columns (5) and (6) additionally control for local business cycle variation due to local income growth and local unemployment growth. Local house prices may move differently from national house prices because of idiosyncratic local economic activity, which would confound estimates of consumption elasticities. I find that estimates using the Bartik instrument are largely unaffected by these controls, with a reported elasticity of 0.113. The estimate using the Saiz instrument is slightly more negative at -0.077, but remains insignificantly different from zero.

Columns (7) and (8) control for the local industry composition of employment. I include local industry shares for agriculture, construction, manufacturing, retail trade, and finance, real estate, and insurance (FIRE). The inclusion of construction and FIRE industry shares ensures that the fall in household expenditure is not due to the loss of current or expected income associated with exposure to those industries most heavily affected by the housing boom and bust. The inclusion of the retail trade share, a proxy for the share of non-tradable employment, dampens general equilibrium feedback effects from house prices to local demand. Again, the elasticity estimate using the Bartik instrument is barely changed at 0.110. The estimate using the Saiz instrument is again negative and insignificant at -0.086.

[28] I construct regional house prices by weighting county level house prices within a region by county population as at the 2000 Census. Table **??** in the Appendix presents estimates using national house price growth.

Table VI
Consumption Response to House Prices Using Bartik and Saiz Instruments

	Annual household real non-durable consumption growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta p_{county,t}$	0.095^{***}	0.120^{***}	0.106^{***}	-0.055	0.113^{***}	-0.077	0.110^{***}	-0.086
	(0.017)	(0.022)	(0.031)	(0.076)	(0.032)	(0.094)	(0.031)	(0.101)
Instrument	Bartik	Saiz	Bartik	Saiz	Bartik	Saiz	Bartik	Saiz
Observations								
Total	256,824	256,824	256,824	256,824	230,865	230,865	227,490	227,490
Households	54,551	54,551	54,551	54,551	51,408	51,408	51,000	51,000
CBSAs	216	216	216	216	216	216	216	216
Counties	571	571	571	571	570	570	563	563
Controls								
Household	✓	✓	✓	✓	✓	✓	✓	✓
CBSA FE	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	×	×	✓	✓	✓	✓	✓	✓
Local	×	×	×	×	✓	✓	✓	✓
Industry	×	×	×	×	×	×	✓	✓
F-statistic	127.25	70.54	36.02	5.52	37.65	4.87	38.33	4.4

Notes: Models estimated using county prices and household consumption growth for homeowners. House prices are instrumented, alternately, using the Bartik and Saiz instruments. The Bartik instrument uses county-level house characteristic shares, as described in Section 5. The Saiz instrument is the interaction of CBSA-level housing supply elasticities and regional house price growth. Control variables are described in Section F. Both standard errors (reported in parentheses) and first stage F-statistic are clustered at the CBSA level. * significant at 10%, ** significant at 5%, *** significant at 1%.
Sources: Author's calculations using data from BLS, CBP, 2000 Census, FHFA, IRS SOI, Nielsen, Saiz (2010), ZTRAX.

Identification for the Bartik instrument depends on the exogeneity of the local housing shares. This means that the shares must be uncorrelated with other features of the locations from which they are drawn. As discussed in Section 4.1, the most obvious possible confounding factors reflect the demographic makeup of households across counties. To test the plausibility of the exogeneity of the Bartik shares, Goldsmith-Pinkham et al. (2018) suggests conditioning on local demographics measured in the same period as the shares and interacting these with time dummies. The interaction with time dummies means that the demographic controls do not become irrelevant over the course of the sample. Note that this significantly increases the number of additional controls: T additional controls for each demographic variable. I include a broad set of variables computed from the 2000 Census, all of which are described in Appendix F. Many of these variables, such as education and immigrant exposure, are directly motivated by the variables that Davidoff (2016) suggests are correlated with the Saiz housing supply elasticities.

Table VII reports the results of conditioning on these demographic variables. While columns (1) and (2) reproduce columns (1) and (2) from Table VI, columns (3) and (4) report results with the additional demographic controls. The estimated consumption elasticity is slightly larger under this specification, at 0.146, although it is not significantly different from the estimates in previous specifications. The estimate using the Saiz instrument is now large and negative at -0.834, although imprecisely estimated.

Across various empirical specifications, the Bartik instrument consistently estimates a consumption elasticity in the range of 0.10 to 0.15. For context, consider that in 2009, the deepest part of the recession, national house prices fell by approximately 9.5 percent while aggregate non-durable consumption expenditures fell by approximately 4.25 percent. According to the elasticities reported here, the fall in national house prices predicts a 0.9 to 1.4 percent decline in non-durables expenditure, which is around 20 to 30 percent of the overall decline.

It is also possible to express these elasticities in terms of marginal propensities to consume out of housing wealth.[29] The marginal propensity to consume non-durables out of housing wealth is 1.21 to 1.81 cents in the dollar. Note that this compares to other recent MPC estimates for non-durable consumption of 1.6 cents (Mian, Rao, et al. (2013)), and estimates for total consumption of 2.8 cents (Guren et al. (2018)), 4.7 cents (Aladangady (2017)), and 5.4 (Mian, Rao, et al. (2013)).

Note that these results reflect average consumption responses to house price movements. In the next section I consider possible heterogeneity in consumption responses due to the effect of collateral constraints and possible asymmetries due to the boom and bust in house prices.

[29] I follow a convention in the literature that makes use of the fact that the MPC is equal to the elasticity of consumption multiplied by the consumption-to-housing wealth ratio. I take consumption to be aggregate expenditure on non-durable goods (FRED code: PCND) and housing wealth is the market value of owner-occupied real estate (FRED code: HOOREVLMHMV). The average ratio from 2000 to 2016 is 0.12. Note that the ratio for aggregate total expenditure (less housing and utilities) is 0.45.

Table VII
Consumption Response to House Prices, Demographic Controls

	Annual household real non-durable consumption growth			
	(1)	(2)	(3)	(4)
$\Delta p_{county,t}$	0.095^{***}	0.120^{***}	0.146^{***}	-0.834
	(0.017)	(0.022)	(0.041)	(1.715)
Instrument	Bartik	Saiz	Bartik	Saiz
Observations				
Total	256,824	256,824	227,490	227,490
Households	54,551	54,551	51,000	51,000
CBSAs	216	216	216	216
Counties	571	571	563	563
Controls				
Household	✓	✓	✓	✓
CBSA FE	✓	✓	✓	✓
Time FE	×	×	✓	✓
Local	×	×	✓	✓
Industry	×	×	✓	✓
Demographic	×	×	✓	✓
F-statistic	127.25	70.54	43.97	0.25

Notes: Models estimated using county prices and household consumption growth for homeowners. House prices are instrumented, alternately, using the Bartik and Saiz instruments. The Bartik instrument uses county-level house characteristic shares, as described in Section 5. The Saiz instrument is the interaction of CBSA-level housing supply elasticities and regional house price growth. Control variables are described in Section F. Both standard errors (reported in parentheses) and first stage F-statistic are clustered at the CBSA level. * significant at 10%, ** significant at 5%, *** significant at 1%.
Sources: Author's calculations using data from BLS, CBP, 2000 Census, FHFA, IRS SOI, Nielsen, Saiz (2010), ZTRAX.

6.3. *Heterogeneity in Consumption Elasticities*

The previous section reports IV estimates of the average consumption elasticity with respect to house prices. However, as discussed elsewhere in the literature, there are theoretical reasons to think that there is heterogeneity in consumption responses to house prices.

One important source of heterogeneity in consumption responses is the collateral channel of house price movements. The collateral channel predicts large differences in the response of households who are and are not borrowing constrained. A change in house prices can affect the terms of financing for households with mortgage debt that is collateralized against the value of their house. For example, when house prices rise, households with previously high loan-to-value (LTV) ratios can borrow more, and perhaps borrow more cheaply (see Cloyne et al. (2017)), at the new, lower LTV ratio. The increase in house prices can then relax borrowing constraints for the household, which induces larger changes in consumption expenditure than in comparison to households who were not constrained. In order to test for the heterogeneity in consumption responses implied by this collateral channel, I compare estimated consumption elasticities for households likely to be more or less borrowing constrained in Table VIII.

The Consumer Panel does not report household asset or debt positions, so we cannot directly observe borrowing constraints. However, the transactions data in ZTRAX reports both house prices and mortgage sizes at origination. From this data I compute mean LTV ratios for mortgages at origination by zip code during the 2004 to 2006 housing boom. I assume that average LTV ratios at mortgage origination during this time are a good proxy for overall LTV ratios because many households bought houses or refinanced mortgages during the boom period. I split the sample of homeowner households into those living in zip codes with an average LTV at origination above 0.8 (high LTV) and below 0.8 (low LTV). New mortgages with LTV ratios above 0.8 have more stringent borrowing requirements if insured by GSEs, and are typically required by lenders to have additional private mortgage insurance. This suggests that households with LTVs in this range are more likely to be close to relevant borrowing constraints than households with mortgages below 0.8.

Columns (1) and (2) of Table VIII compare consumption elasticities estimated separately for the high- and low-LTV sample. Column (3) presents estimates using the pooled sample, but includes an interaction between house price growth and an indicator variable for households in high-LTV zip codes.[30] All specifications include the full set additional controls discussed in the previous section. The IV-estimated elasticity for high LTV households (0.261) is larger than that for low LTV households (0.161), although these estimates are not statistically distinct. However, the pooled sample estimates suggest that the consumption response of high-LTV households is significantly larger than that of low LTV households. These results are consistent with previous findings in the literature, which show large and significant effects of the collateral channel (see Mian, Rao, et al. (2013), Aladangady (2017)).

Several papers in the recent literature have also investigated whether the consumption response to house prices varies over time. Aladangady (2017) estimates the marginal propensity to consume out of housing wealth shocks using the Consumer Expenditure Survey over the period 1986 to 2008. He considers whether MPC might have been different in the 2000s, perhaps due to the ease of access to mortgage mortgage credit in this period. He finds that the MPC is slightly lower during the 2000s, although this difference from the benchmark result is statistically insignificant. Guren et al. (2018) use an estimation method that explicitly allows for time variation in the estimated consumption elasticities. They report virtually no variation in the elasticity from 1995 to 2005, but show that the consumption elasticity may have been almost twice as large in the early 1990s.

In order to investigate time-variation in consumption elasticities using the Bartik instrument, Table IX reports estimates that include a indicator variable for the 2006-2009 period interacted with house price growth. I focus on these three years since this is sample period used in Mian, Rao, et al. (2013),

[30] For the pooled regression, the interaction term is instrumented via the interaction of the Bartik instrument and the indicator variable for high-LTV zip codes.

but also because these years experienced some of the fastest declines in house prices. Column (1) shows reports a slightly negative, but statistically insignificant, coefficient on the 2006 to 2009 period. Column (2) adds interaction terms for households that live in high LTV zip codes. Consistent with the results reported in Table VIII, I find that high LTV households respond more strongly to house price movements. But as was the case in column (1), I find no significant difference in the consumption elasticity during the housing bust period.

Table VIII
Consumption Response to House Prices for High LTV Households

	Annual household real non-durable consumption growth		
	(1)	(2)	(3)
$\Delta p_{county,t}$	0.261^{***}	0.161^{***}	0.141^{***}
	(0.088)	(0.061)	(0.042)
$\Delta p_{county,t} \times \mathbb{1}(HighLTV)$			0.109^{***}
			(0.037)
Sample	High LTV	Low LTV	Pooled
Observations			
Total	138,420	85,761	261,405
Households	32,443	20,281	61,914
CBSAs	443	334	630
Counties	821	588	1,174
Zip codes	5,396	3,847	11,773
Controls			
Household	✓	✓	✓
CBSA FE	✓	✓	✓
Time FE	✓	✓	✓
Local	✓	✓	✓
Industry	✓	✓	✓
Demographic	✓	✓	✓

Notes: Models estimated using county prices and household consumption growth for homeowners in high LTV zip codes, low LTV zip codes, and the pooled sample of households. High LTV zip codes had mean LTV ratios greater than 0.8 for mortgages at origination between 2004 and 2006. Low LTV zip codes had ratios of 0.8 and below. IV estimates of consumption elasticities reported, where house prices are instrumented using the Bartik instrument described in Section 5. Control variables are described in Appendix F Both standard errors (reported in parentheses) and first stage F-statistic are clustered at the CBSA level. * significant at 10%, ** significant at 5%, *** significant at 1%.
Sources: Author's calculations using data from BLS, CBP, FHFA, IRS SOI, Nielsen, ZTRAX.

Table IX
Consumption Response to House Prices During the Bust

Annual household real non-durable consumption growth		
	(1)	(2)
$\Delta p_{county,t}$	0.179^{**}	0.141^{*}
	(0.078)	(0.076)
$\Delta p_{county,t} \times \mathbb{1}(2006-2009)$	-0.003	-0.002
	(0.118)	(0.114)
$\Delta p_{county,t} \times \mathbb{1}(HighLTV)$		0.103^{***}
		(0.034)
$\Delta p_{county,t} \times \mathbb{1}(HighLTV) \times \mathbb{1}(2006-2009)$		0.012
		(0.066)
Observations		
Total	261,405	261,405
Households	61,917	61,917
CBSAs	631	631
Counties	1,175	1,175
Zip codes	11,776	11,776
Controls		
Household	✓	✓
CBSA FE	✓	✓
Time FE	✓	✓
Local	✓	✓
Industry	✓	✓
Demographic	✓	✓

Notes: Models estimated using county prices and household consumption growth for homeowners. Models include dummy variables for the 2006-2009 period, for households in high LTV zip codes, and the interaction between the two. High LTV zip codes had mean LTV ratios greater than 0.8 for mortgages at origination between 2004 and 2006. IV estimates of consumption elasticities reported, where house prices are instrumented using the Bartik instrument described in Section **??**. Control variables are described in Appendix **??** Both standard errors (reported in parentheses) and first stage F-statistic are clustered at the CBSA level. * significant at 10%, ** significant at 5%, *** significant at 1%.
Sources: Author's calculations using data from BLS, CBP, FHFA, IRS SOI, Nielsen, ZTRAX.

7. The Bartik Instrument for House Prices: What, When, Where, and How

The Bartik instrument strongly predicts house prices and provides consistent and precise IV estimates of the consumption response to house prices. But how does the Bartik instrument work so well in this context when the Saiz instrument does not? Additionally, given the extreme variation in recent house price movements across time and space, from when and where does the instrument draw its variation? Finally, what economic mechanisms might explain how the instrument generates such variation in house prices?

7.1. Decomposing the Bartik Instrument

At a glance, it seems that the performance of the Bartik is due to the multiple sources of house price variation on which it draws for each location. Specifically, a separate housing quality price is associated with each house characteristic to which locations are exposed. And as discussed in Section , there is significant variation in housing quality prices across characteristics, time, and regions. But notice that even if the housing stock in a given location is entirely uniform, because the instrument includes three sets of characteristics – age, bedrooms, and bathrooms – it provides three sources of variation in house prices for that location. In contrast, for the Saiz instrument locations lie along the distribution of housing supply elasticities, but a shock to housing demand (i.e. movement in regional house price growth) simply scales the local house price response up or down. Thus, within a region the instrument cannot, for example, generate both increasing and decreasing local house prices while regional prices are increasing.

Mechanically, it is obvious that the Bartik instrument generates more varied house price movements than is possible for the Saiz instrument. On the other hand, a drawback of Bartik instruments is the "black box" nature of their statistical mechanisms. Because the instrument consists of both a set of local housing shares and a number of housing quality prices, it is not clear which of these shares and prices contributes to the variance of the instrument. Moreover, since identification depends on the exogeneity of the local housing shares, it would be useful to know whether the housing characteristics that drive most of the variation in the instrument are plausibly exogenous.

In order to explore Bartik instruments more carefully, Goldsmith-Pinkham et al. (2018) present a useful decomposition of the instrument. They first show that the IV estimator using a Bartik instrument can be recast as an over-identified GMM estimator, where the local shares are treated as a set of individual instruments under a particular weighting matrix. The IV estimator can then be decomposed into a set of estimators using each of the local shares, and a set of "Rotemberg" weights associated with each of these estimates (see also Rotemberg (1983)). Note that each of the local shares provides a just-identified estimate of the parameter of interest. And each of these just-identified estimates along with their associated Rotemberg weights provide the contribution of the associated local share to the overall Bartik estimate.

Returning to the application to house prices and consumption, consider a simplification of the the reduced form and first stage Equations 2 and 3:

$$\Delta c_{l,t} = \beta \Delta p_{l,t} + \varepsilon_{l,t}$$
$$\Delta p_{l,t} = \gamma B_{l,t} + \eta_{l,t}.$$

For ease of exposition, suppose only one household is observed in each location, and that there is only one time period. Suppose, also, that consumption growth, house prices, and the instrument have all been residualized with respect to the various control variables.[31] Recall that $\Delta c_{l,t}$ consumption growth, $\Delta p_{l,t}$

[31] Although tedious, notation can be extended to accommodate the more general cases with H_l households per location l, T time periods, and an explicit vector of control variables.

is house price growth, $B_{l,t}$ is the Bartik instrument, and β is the elasticity of consumption with respect to house prices. γ is the effect of the instrument on house prices from the first stage regression.

Let L denote the number of locations, and K the total number of house characteristics used in the instrument. Then C is the $L \times 1$ vector stacking $\Delta c_{l,t}$, P is the $L \times 1$ vector stacking $\Delta p_{l,t}$, and B is the $L \times 1$ vector stacking the instrument $B_{l,t}$. Recall that the instrument is constructed via $B_{l,t} = \sum_k \lambda_{l,k} \Delta q_{k,t}$, where $\lambda_{l,k}$ are the local housing characteristic shares for each location l and characteristic k, and $\Delta q_{k,t}$ is the growth in housing quality prices for characteristic k. Let Λ denote the $L \times K$ matrix of local housing characteristic shares, and Q is the $K \times 1$ vector of stacked quality price growth rates. Then the stacked vector of Bartik instruments is $B = \Lambda Q$. The IV estimator of the consumption elasticity using the Bartik instrument has the familiar form:

$$\beta^{bartik} = \frac{B'C}{B'P} = \frac{Q'\Lambda'C}{Q'\Lambda'P} \tag{10}$$

Following Goldsmith-Pinkham et al. (2018), the Bartik estimate can then be decomposed into the just-identified estimates β_k^{bartik} and the associated Rotemberg weights α_k. The Bartik estimate of the consumption elasticity is the Rotemberg-weighted average of the just-identified estimates: $\beta^{bartik} = \sum_k^{K=1} \alpha_k \beta_k^{bartik}$, where the Rotemberg weights α_k sum to one. Goldsmith-Pinkham et al. (2018) notes that individual Rotemberg weights α_k may be negative, which means that the overidentified IV estimate using the full Bartik instrument β^{bartik} can be outside of the range of the individual estimates β_k^{bartik}. The just-identified estimates are given by:

$$\beta_k^{bartik} = \frac{\Lambda_k' C}{\Lambda_k' P},$$

where Λ_k' is the k^{th} column of Λ. And the Rotemberg weights are given by:

$$\alpha_k = \frac{\Delta q_{k,t} \Lambda_k' P}{\sum_{k=1}^K \Delta q_{k,t} \Lambda_k' P} = \frac{\hat{\gamma} \Delta q_{k,t} \Lambda_k' P}{\hat{\gamma} B' P} = \frac{P_k^{bartik'} P}{P^{bartik'} P},$$

where the second equality follows from the definition of the Bartik instrument and $\hat{\gamma}$ is the estimated first stage coefficient used to predict house prices with the instrument. Then P^{bartik} are the fitted values for house price growth from the first stage, and P_k^{bartik} are the fitted values from the first stage but using only the k^{th} component of the Bartik instrument for prediction.

Summary statistics from the decomposition of the Bartik instrument are reported in Table X. I present the decomposition of the instrument using all households, time periods, and locations. Panel A summarizes the individual estimates and Rotemberg weights. Panel B explores the correlations between these, housing quality price growth, and the variance of the local housing shares. The Rotemberg weights and housing quality price growth are negatively correlated, while the weights and the variance of the local housing shares are positively correlated. Because the over-identified IV estimator places more weight on instruments that vary more, it is unsurprising that it draws heavily on the dramatic swings in housing quality prices that occur throughout the sample period. Additionally, dispersion in the local housing shares implies variation in the cross-sectional exposure to quality price movements, which also helps generate variation in the instrument.

Panel C reports the house characteristics with the largest share of absolute Rotemberg weights, by region and year. The top five characteristics account for 25 percent of the Rotemberg weights, suggesting skewness in the influence of the just-identified estimates. The benchmark house (built prior to 1940 with one bedroom and zero bathrooms) is associated with much of the Rotemberg weight, while the next most influential characteristics are associated with two bathroomed houses, three bedroomed houses, and houses built in the 1990s and 2000s. Strikingly, virtually all of the weight is associated with the Western region, and is largely concentrated in the bust years of 2008 and 2009, but also the recovery years of 2013 to 2015. This is perhaps unsurprising given that Western parts of the country such as Arizona, California, and Nevada had some of the largest house price fluctuations in the entire country during this period. These results are emphasized in Figure IV, which shows that much of the Rotemberg weight occurs in years featuring rapid house price movements: 2005 (end of the boom), 2008 and 2009 (deepest part of the bust), and 2013 and 2014 (fastest part of the recovery).

Table X
Summary of Rotemberg Weights for the Bartik Instrument

Panel A: Summary of β_k and α_k

	Wgt. Mean	Mean	Median	25th Perc.	75th Perc.	Share Negative
β_k	0.155	-0.303	0.064	-0.348	0.304	0.306
α_k	–	0.001	0	-0.001	0.003	0.416

Panel B: Correlations

	α_k	Δq_k	β_k	var(λ_k)
α_k	1			
Δq_k	-0.248	1		
β_k	0.009	-0.007	1	
var(λ_k)	0.356	0.003	0.018	1

Panel C: Top 20 house characteristics by share of absolute Rotemberg weight

Characteristic	Year	Region	α_k	$\frac{\lvert\alpha_k\rvert}{\sum\lvert\alpha_k\rvert}$	Δq_k	β_k
Benchmark house	2008	West	-0.6	0.1	-1.68	0.13
Benchmark house	2009	West	-0.28	0.05	-0.87	0.16
Benchmark house	2013	West	-0.25	0.04	0.96	-0.05
Bathrooms: 2	2008	West	0.22	0.04	-1.69	0.08
Benchmark house	2014	West	-0.17	0.03	0.62	0.05
Bedrooms: 3	2008	West	0.14	0.02	-1.8	0.14
Bathrooms: 2	2009	West	0.11	0.02	-0.94	0.12
Benchmark house	2008	South	0.11	0.02	-0.72	-0.29
Decade to 2005	2008	West	0.1	0.02	-1.74	0.14
Decade to 1999	2008	West	0.09	0.01	-1.68	0.18
Bathrooms: 2	2013	West	0.09	0.01	0.93	-0.05
Bedrooms: 4	2008	West	0.09	0.01	-1.84	0.14
Decade to 1989	2008	West	0.07	0.01	-1.74	0.11
Benchmark house	2015	West	-0.07	0.01	0.53	0.16
Bedrooms: 3	2009	West	0.07	0.01	-0.86	0.15
Bathrooms: 1	2008	West	0.06	0.01	-1.71	0.19
Benchmark house	2014	NorthEast	-0.06	0.01	-0.29	0.21
Benchmark house	2005	South	0.06	0.01	0.35	-0.14
Decade to 1999	2009	West	0.06	0.01	-0.8	0.05
Bedrooms: 3	2013	West	0.06	0.01	0.96	-0.05

Notes: Throughout the table β_k are the characteristic-specific estimates, α_k are the Rotemberg weights, Δq_k are the housing quality price growth rates for characteristic k, and var(λ_k) is the cross-sectional variance of the local housing characteric shares.

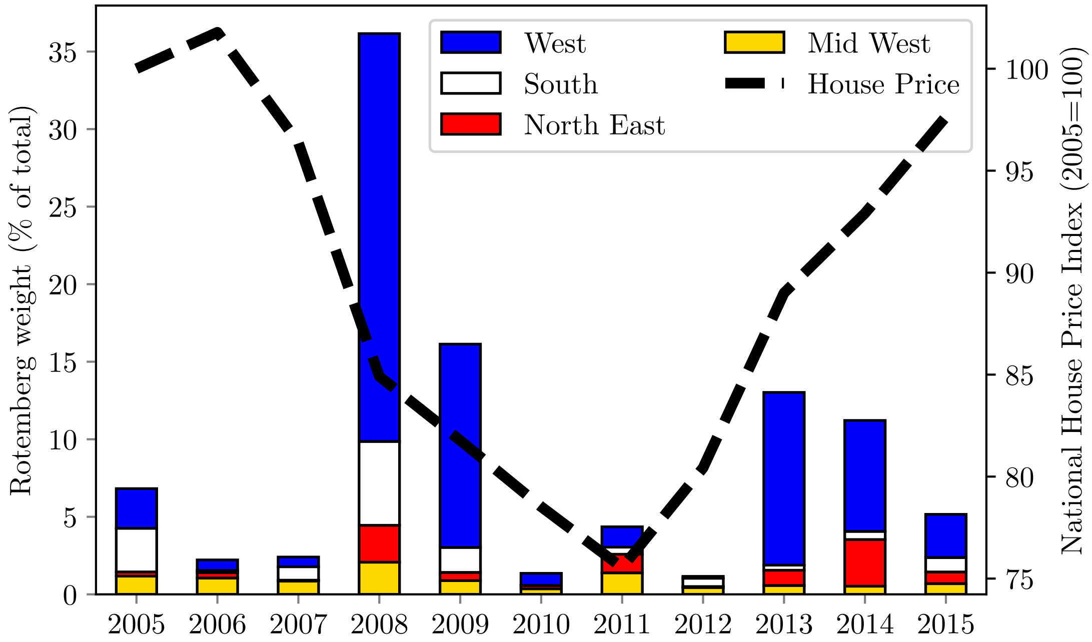

Figure IV
Rotemberg Weights by Region and Year

Notes: Sums of the share of the absolute Rotemberg weights for each region and year. The dashed black line is the S&P/Case-Shiller National House Price Index.
Sources: Author's calculations, FRED.

7.2. *Economic Intuition for the Bartik Instrument*

Beyond understanding the statistical sources of variation in the Bartik instrument, it is also useful to consider potential economic mechanisms underlying this variation. In particular, why might the marginal prices of certain property characteristics rise and fall relative to others?

Note that housing markets are segmented. Piazzesi et al. (2015) show that buyers do not consider all houses currently on the market, instead they tend to search along several dimensions, including geography, price, and house size. As a result, the average price of houses with different qualities can have very different price dynamics. For example, Landvoigt et al. (2015) show that in San Diego, initially low priced (i.e. low quality) houses appreciated much faster during the 2000s housing boom than did initially high priced houses. They reproduce these differential changes in house prices with a heterogeneous households model in which different groups are the marginal buyers for different house qualities. They suggest that lower income households were more affected by the increase in credit supply during the housing boom, which led to an increase in demand for houses where low income households were marginal buyers.

Confirming the findings in Landvoigt et al. (2015), the upper left panel of Figure V shows house price appreciation for houses that were in the bottom, middle, and top third of San Diego houses by sale price in 2000 using data from ZTRAX.[32] House prices in the lowest tier were much more volatile price dynamics than in the other tiers. The remaining three panels of the figure then report the distribution of housing characteristics across these house price tiers. I find that the lowest tier of houses are older and smaller, according to building size and number of bedrooms. Following Landvoigt et al. (2015), as low

[32] I compute the median prices in each tier for houses that were initially sold in 2000 and the re-sold in each year after that.

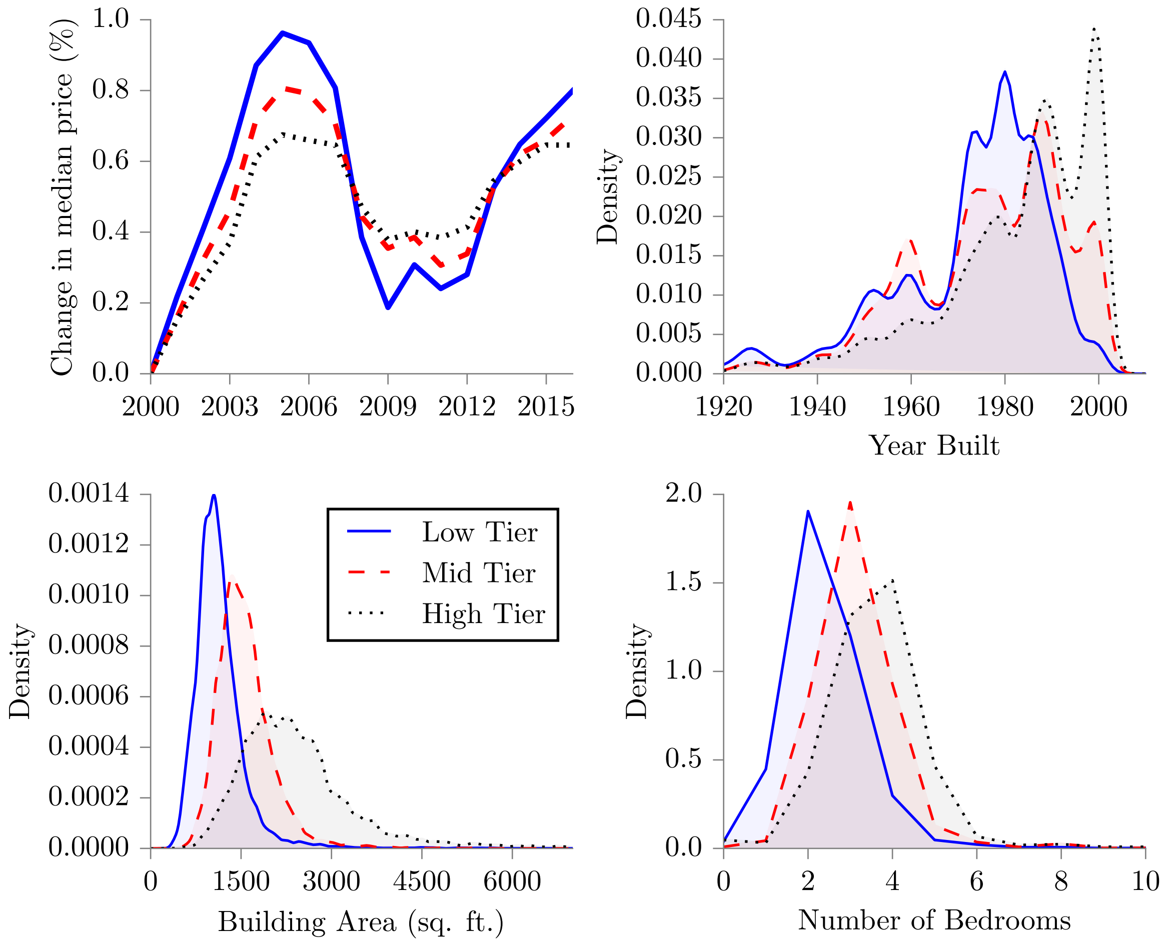

Figure V
Prices and Characteristics of Houses in San Diego by 2000 Price Tier

Notes: Individual houses in San Diego sorted into three tiers based on sale price in 2000. The median price for each tier is computed using houses sold repeatedly after 2000. Kernel densities estimated using a Gaussian kernel.
Sources: Author's calculations using ZTRAX.

income households bid up the price of low quality houses during the housing boom, older and smaller houses should have appreciated more quickly since these house characteristics are disproportionately represented amongst the lowest quality houses. The results of the previous section provide some justification for this intuition, since benchmark houses (i.e. the smallest and oldest houses) drive a lot of the variation in the Bartik instrument (see Table X).

Variation in the Bartik instrument may also be driven by changes in the purchasing power of other types of marginal house buyers, such as residential property investors. Several recent papers suggest that investors were especially active in housing markets with the largest swings in house prices (Haughwout et al. (2011), Mian and Sufi (2018)). These investors appear to have been especially sensitive to the availability of mortgage credit, and may have contributed to both the rise and fall in house prices. If these buyers invest in houses with particular features, they will influence the price dynamics of those houses for which they act as marginal buyers. It is also possible that the prices of particular houses in some housing markets are driven by variation in demand for luxury apartments or demand from out-of-town or foreign investors (Ait-Sahalia et al. (2004), Chinco et al. (2015), Favilukis et al. (2017)).

8. Robustness

In this section I consider several variations on the Bartik instrument, which explore the robustness of the instrument to various model specifications. All tables and figures are reported in Appendix G.

Table A.8 provides IV-estimates of the consumption elasticity with respect to zip code house prices. Columns (1) to (5) explore the same empirical model specifications as for the benchmark estimates. I find elasticity estimates in the range of 0.11 to 0.14, which is very similar to the range of estimates using county level house prices.

Table A.4 replicates the main results in Table VI using the same sample across all four empirical specifications. This ensures that changes in the sample size do no affect the results. I find negligible differences in the consumption elasticity estimates using the common sample.

Table A.5 compares regional and national variants of the Bartik and Saiz instruments. The regional Bartik instrument is that used for the benchmark results. The national version of the instrument is constructed from housing quality prices estimated from the hedonic regression in Equation 9, but using housing transactions from the entire country. The regional and national Saiz instruments are simply constructed using the interaction of local housing supply elasticities with regional and national house prices, respectively. Columns (1)-(4) show that estimates using the four instruments with only household level controls and city (CBSA) fixed effects produce very similar results. However, the inclusion of time fixed affects and other local, industry, and demographic controls significantly impact estimates using the national Bartik instrument, as well both the regional and national Saiz instruments (see columns (6)-(8)). This highlights the difficulty of producing exogenous instruments for house prices that rely on interactions between cross-sectional exposures and national shocks to housing demand.

Next, I consider two alternative constructions of the Bartik instrument. The first of these only uses house age as a characteristic for building the instrument. One reason to do this is due to a concern that the number of bedrooms or bathrooms may not be exogenous to other determinants of household consumption. For example, a local consumption boom might be associated with an increased desire to renovate houses, which might include adding bedrooms or bathrooms to existing houses. To construct this instrument, I estimate the housing quality prices from the hedonic price regression (9) as before. This ensures that the estimated marginal prices of house age are conditional on the other house characteristics. However, I then construct the Bartik instrument using only the local shares and quality prices for house age.

The second alternative takes the benchmark instrument but adds local exposures to house floor size and property lot size. One reason to do this is that house price variation might be driven by fluctuations in the value of land. Thus, the exclusion of these size variables from the benchmark instrument may throw out useful variation in house prices.[33] To construct the instrument, I first compute the median log-floor size and median log-lot size for each location. Equation (9) is again estimated to form the quality prices, as before. But now the hedonic price regression coefficients on the house size variables enter the Bartik instrument. I rewrite the instrument as:

$$B_{g,t} = \sum_c \sum_i \lambda^i_{g,c} \Delta q^i_{c,t} + \mu_{g,f} \Delta q^i_{f,t} + \mu_{g,l} \Delta q^i_{l,t}, \tag{11}$$

where $\mu_{g,f}$ and $\mu_{g,l}$ are the local median log-floor size and log-lot size. This version of the instrument is not a standard Bartik instrument, since the additional size variables do not reflect the typical shift-share construction. However, a similar intuition is retained in that locations with larger houses on average should experience faster house price appreciation when the marginal price per square foot of floor or lot increases.

[33] On the other hand, one reason not to include size variables in the instrument is due to a concern house price variation associated with fluctuations in land prices is likely to be strongly correlated with both local and aggregate business cycle variation that drives household consumption. Including these variables in the instrument risks re-introducing endogenous variation in house prices.

Table A.6 compares estimates of the consumption elasticity using the benchmark Bartik instrument, the instrument using housing age only, and the size-augmented instrument. Columns (1)-(3) report IV-estimates using household controls and CBSA fixed effects only. The benchmark and age-only instruments produce virtually identical estimates (0.128), while the size-augmented instrument estimate is slightly larger (0.160), although not statistically distinct. Columns (4)-(6) include time fixed effects, local controls, industry employment shares, and local demographic variables. Estimates for all three estimates are slightly larger than without controls (0.143, 0.144, and 0.167) but all remain precisely estimated.

Out of interest, Table A.7 reports first stage regressions using the size-augmented instrument, and Figure XI plots a binned scatter plot of the residualized instrument against residualized house prices. The strength of the instrument as determined by the F-statistic is similar to the benchmark instrument under all specifications. But notice that the first stage regression coefficient is much larger at 0.2-0.5 than the first stage coefficients for the benchmark instrument, at0.05-0.09. This reflects the fact that fluctuations in land prices are also a significant source of house price fluctuations, which is being picked up by the size-augmented instrument. Research contexts that require more sensitive instruments for house prices might consider this specification of the Bartik instrument.

9. Conclusion

This paper introduces a new instrumental variables strategy for estimating the effect of house prices on consumption. I follow the literature on Bartik instruments, which adopts a shift-share approach to constructing instruments. The house price instrument presented here consists of the interaction between the local share of different housing characteristics and changes in the marginal prices of those characteristics. When the relative demand for particular house characteristics is rising, locations with a high proportion of the housing stock possessing those characteristics will experience faster house price appreciation.

I argue that the Bartik instrument is plausibly exogenous to a range of possible confounding factors including aggregate time series variation, local business cycles, local industry composition, and local demographic factors. Under a variety of empirical specifications, I find that the elasticity of consumption with respect to house prices falls in a range between 0.1 and 0.15.

The most popular instrumental variables strategy in the existing literature uses local housing supply elasticities produced by Saiz (2010). I show that relative to the Bartik instrument, the Saiz instrument performs poorly. The main reason for this is the lack of time series variation in the instrument. This is the case even though I use a variant of the instrument that is interacted with regional house price growth in order to provide time series variation.

In contrast, I show that the Bartik instrument provides both cross-sectional and time series variation. The local housing characteristic shares that form the basis of the instrument display significant heterogeneity at various levels of geography. Indeed, this is what enables the instrument to be constructed for smaller units of geographic aggregation, such as at the zip code level. In the time series dimension, I show that there is significant variation in the growth rates of the housing quality prices. This suggests that there are indeed time-varying shocks to the relative demand for different house characteristics.

I also decompose the instrument into a set of just-identified IV estimates of the consumption elasticity in order to examine variation in the instrument in more detail. I find that over 50 percent of variance in the instrument that contributes to identifying the consumption response to house prices is due to 2008 and 2009 alone. Another 30 percent is due to the recovery years of 2013 and 2015. Additionally, much of the identifying variation in the instrument is due to house prices changes in the West of the US. These findings are consistent with the fact that house prices varied the most over the sample period in the bust and recovery and also in the western parts of the US. It is unsurprising that the instrument places so much emphasis on times and places with rapid house price changes: this is precisely when and where house prices are most likely to affect household balance sheets, and thus to influence household consumption decisions.

References

Ait-Sahalia, Yacine, Jonathan A Parker, and Motohiro Yogo, "Luxury goods and the equity premium", *The Journal of Finance* 59 (2004), 2959–3004.

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

Attanasio, Orazio P, Laura Blow, Robert Hamilton, and Andrew Leicester, "Booms and busts: Consumption, house prices and expectations", *Economica* 76 (2009), 20–50.

Bartik, Timothy J, "Who benefits from state and local economic development policies?" (1991).

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, *Quasi-experimental Shift-share Research Designs*, tech. rep., NBER Working Paper, 2018.

Buiter, Willem H, *Housing wealth isn't wealth*, tech. rep., NBER Working Paper, 2008.

Campbell, John Y and Joao F Cocco, "How do house prices affect consumption? Evidence from micro data", *Journal of monetary Economics* 54 (2007), 591–621.

Carroll, Christopher D, "A theory of the consumption function, with and without liquidity constraints", *Journal of Economic perspectives* 15 (2001), 23–45.

Carroll, Christopher D and Miles S Kimball, "On the concavity of the consumption function", *Econometrica: Journal of the Econometric Society* (1996), 981–992.

Case, Karl E, John M Quigley, and Robert J Shiller, "Comparing wealth effects: the stock market versus the housing market", *Advances in macroeconomics* 5 (2005).

Chaney, Thomas, David Sraer, and David Thesmar, "The collateral channel: How real estate shocks affect corporate investment", *American Economic Review* 102 (2012), 2381–2409.

Chen, Hui, Michael Michaux, and Nikolai Roussanov, *Houses as ATMs? Mortgage refinancing and macroeconomic uncertainty*, tech. rep., NBER Working Paper, 2013.

Chinco, Alex and Christopher Mayer, "Misinformed speculators and mispricing in the housing market", *The Review of Financial Studies* 29 (2015), 486–522.

Cloyne, James, Kilian Huber, Ethan Ilzetzki, and Henrik Kleven, *The effect of house prices on household borrowing: A new approach*, tech. rep., NBER Working Paper, 2017.

Davidoff, Thomas, "Supply constraints are not valid instrumental variables for home prices because they are correlated with many demand factors", *Critical Finance Review* 5 (2016), 177–206.

Favilukis, Jack Y and Stijn Van Nieuwerburgh, "Out-of-town Home Buyers and City Welfare" (2017).

Friedman, Milton, *Theory of the consumption function*, Princeton university press, 1957.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, *Bartik Instruments: What, When, Why, and How*, tech. rep., NBER Working Paper, 2018.

Guren, Adam M, Alisdair McKay, Emi Nakamura, and Jón Steinsson, *Housing Wealth Effects: The Long View*, tech. rep., NBER Working Paper, 2018.

Gyourko, Joseph, Christopher Mayer, and Todd Sinai, "Superstar cities", *American Economic Journal: Economic Policy* 5 (2013), 167–99.

Gyourko, Joseph, Albert Saiz, and Anita Summers, "A new measure of the local regulatory environment for housing markets: The Wharton Residential Land Use Regulatory Index", *Urban Studies* 45 (2008), 693–729.

Haughwout, Andrew, Donghoon Lee, Joseph S Tracy, and Wilbert Van der Klaauw, "Real estate investors, the leverage cycle, and the housing market crisis" (2011).

Kaplan, Greg, Kurt Mitman, and Giovanni L Violante, *Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data*, tech. rep., NBER Working Paper, 2016.

Kaplan, Greg, Giovanni L Violante, and Justin Weidner, *The wealthy hand-to-mouth*, tech. rep., NBER Working Paper, 2014.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider, “The housing market(s) of San Diego”, *The American Economic Review* 105 (2015), 1371–1407.
Marketing, Kilts Center for, *Consumer Panel Dataset Manual*, 2016.
Mian, Atif, Kamalesh Rao, and Amir Sufi, “Household Balance Sheets, Consumption, and the Economic Slump”, *Quarterly Journal of Economics* 128 (2013).
Mian, Atif and Amir Sufi, *Credit supply and housing speculation*, tech. rep., NBER Working Paper, 2018.
— “House prices, home equity-based borrowing, and the US household leverage crisis”, *American Economic Review* 101 (2011), 2132–56.
— “The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis”, *The Quarterly Journal of Economics* 124 (2009), 1449–1496.
Palmer, Christopher, “Why did so many subprime borrowers default during the crisis: Loose credit or plummeting prices?”, *Working Paper* (2015).
Piazzesi, Monika, Martin Schneider, and Johannes Stroebel, *Segmented housing search*, tech. rep., NBER Working Paper, 2015.
Rotemberg, Julio, “Instrument variable estimation of misspecified models” (1983).
Saiz, Albert, “The geographic determinants of housing supply”, *The Quarterly Journal of Economics* 125 (2010), 1253–1296.
Stroebel, Johannes and Joseph Vavra, *House prices, local demand, and retail prices*, tech. rep., NBER Working Paper, 2014.

Appendices

A. *Data Sources*

This section documents the sources of data used in the paper.

- Panel consumption data comes from the Nielsen Consumer Panel Data survey made available by the Kilts Center at Chicago Booth. This data is proprietary and is typically available only by institutional subscription. See the Kilts Center website for more information regarding access: https://research.chicagobooth.edu/nielsen/.
- The individual housing transaction data comes from Zillow's Assessment and Transaction Database (ZTRAX). This data is proprietary, but is available from Zillow by request. For information regarding access, contact see http://www.zillow.com/ztrax.
- Annual county house price indexes are publicly available from the Federal Housing and Finance Agency at https://www.fhfa.gov/DataTools/Downloads/Pages/House-Price-Index-Datasets.aspx.
- Additional house price indexes for zip codes, counties, and metropolitan areas (CBSA) are publicly available from Zillow at https://www.zillow.com/research/data/.
- The consumption price index is the monthly seasonally adjusted CPI-U for all items. This is available from FRED at https://fred.stlouisfed.org/, using code CPIAUCSL.
- Zip code level income is retrieved from the IRS Statement of Income (SOI) statistics at https://www.irs.gov/statistics/.
- County unemployment data is from the Bureau of Labor Statistics, available at https://www.bls.gov/lau/data.htm. Python code to clean this data is available at Github: https://github.com/jagman88/Clean-BLS-County-Level-Employment-Data.
- Zip code and county level demographic characteristics are retrieved from 2000 Census, available at https://factfinder.census.gov/.
- County employment by industry is in the County Business Patterns data, available at https://www.naics.com/business-lists/counts-by-naics-code/.
- Zip code, FIPS (county) code, and metropolitan area (CBSA) cross-walk information is retrieved from the Department of Housing and Urban Development at https://www.huduser.gov/portal/datasets/usps_crosswalk.html.
- Cartographic boundary files (i.e. TIGER shape files) used in the construction of maps are available from the Census Bureau at https://www.census.gov/geo/maps-data/.
- Additional figures use data from the Survey of Consumer Finances, available at https://www.federalreserve.gov/econres/scfindex.htm. Other data comes from the Current Population Survey, available via IPUMBS at https://cps.ipums.org/cps/.

B. *ZTRAX House Price Data*

Each transaction in ZTRAX contains information on the characteristics of the property and the sale including date of sale, property type, sale type, buyer type, and so on. I aim to work with a consistent data set containing typical property transactions conducted by residential owner-occupiers. To this end, I carry out the following cleaning procedure.

I restrict the data to housing transactions made at arm's-length and when not sold due to foreclosure. This removes all distressed sales, and all transactions with builders, developers, or real estate agents on either side of the transaction. I restrict properties to those that are are non-commercial, and that are either single family residences or owner-occupied properties. This latter restriction allows me to include properties that are apartments, as long as they are owner occupied (i.e. not sold by a landlord). This is important in cities like New York where a significant proportion of the owner-occupied housing stock consists of apartments. I also focus only on property transactions with non-zero sales prices, thereby removing all mortgages, mortgage refinancing, and transfers or gifts. I exclude transactions that may have been subject to 'house flipping', thereby distorting the market value of the house. To do this, I remove any house sale that occurs within 180 days of a prior sale of the same house. Additionally, I remove transactions with a sale price of less than $10,000 as well as those with no reported transaction date. I exclude houses with no recorded build year (i.e. no known age of the building), no reported floor size, and no reported zip code.

The ZTRAX data is held in state-level files, each of which contains the entire set of property characteristics and transactions for that state. Three states – Rhode Island, Tennessee, and Vermont – have various missing data in the ZTRAX database, and are excluded from the analysis. For several other states, non-mandatory disclosure and outright prohibitions on the reporting of transactions prices mean that a very large proportion of transactions feature sales with prices reported as zero or missing.[34] For these states, property deeds and assessment records may still be reported to the ZTRAX database. I collect data on housing characteristics for these states, but I cannot use the transaction data on sales prices.[35] Instead, for these states I use publicly available, geographically aggregated Zillow house price indexes. After data cleaning, there are 55 million individual transactions available between 1994 and 2016.

C. Consumer Panel Data

Table A.1 reports household summary statistics from the Consumer Panel. Notice that average consumption is much lower than average income, which is because only non-durable expenditure is surveyed. Table A.2 reports several demographic summary statistics. In comparison with data from the Current Population Survey (CPS) over the same sample period, the Consumer Panel has a similar proportion of households whose heads have attended college, are not in employment, and are homeowners. Additionally, I report the proportion of households that have moved in the past year across zip codes, counties, or states. Relative to the CPS, households are similarly likely to have moved across states, about half as likely to have moved across counties (not including cross-state moves), and less than a third as likely to have moved across zip codes (not including cross-county or cross-state moves). Since households are less likely to move than typical households in the population, they may experience greater consumption sensitivity with respect to house prices than the typical household in the population. For this reason, in the empirical analysis I restrict households to those that do not move during the sample, and so all results should be interpreted as consumption responses to a house price change for non-moving households.

Note that homeownership is not directly reported in the Consumer Panel. In order to compute homeownership rates, I follow Stroebel et al. (2014) who, also using the Consumer Panel Data, infer ownership status household residence type. Households report whether they live in a one-, two-, or three-family dwelling, and also whether the house is a condo or co-op. Single-family, non-condo/co-op residences are assumed to be inhabited by homeowners, with remaining households assumed to be renters. The average sample weighted-proportion of households living in single-family homes is 0.75,

[34]See http://www.zillowgroup.com/news/chronicles-of-data-collection-ii-non-disclosure-states/ for more details.

[35]The states with large numbers of missing transaction data are: Alaska, Idaho, Indiana, Kansas, Maine, Mississippi, Montana, New Mexico, Texas, Utah, and Wyoming

and does not change significantly over the sample. From 2004 to 2015, the homeownership rate for the US as a whole fell from 69 percent to 64 percent.[36] The second panel of Figure VI in the Appendix reports the age profile of homeownership, which reveals that implied homeownership rates are overstated by between 15 and 30 percentage points for young households relative to data from the SCF. Implied homeownership rates for older households are very similar to those reported in the SCF. For most of the empirical results, I make use of the sample of implied homeowners only.

Table A.1
Household Summary Statistics

	Income	Expenditure	Age	Family Size
Mean	\$68,141	\$7,489	54	2.6
Median	\$59,999	\$6,317	52	2.0
StdDev	\$42,330	\$4,896	15	1.5

Notes: Means, medians, and standard deviations computed using Consumer Panel survey weights. Income is the households income two years prior to the panel year, and is recorded categorically. Income statistics are computed using the upper bound of each category. Expenditure is total nominal household consumption expenditure within the panel year. Age is computed using the male household head, or the female household head if no male head is reported. Family size is the number of family members reported to live in the household.
Sources: Authors calculations using Nielsen Consumer Panel.

Table A.2
Household Demographics

	College	Non-Employed	Home Owners	Moved across: Zip	County	State
Nielsen	0.420	0.190	0.749	0.021	0.012	0.012
CPS	0.421	0.242	0.686	0.078	0.023	0.017

Notes: Computed using survey-weighted averages in the Nielsen Consumer Panel and the Current Population Survey (CPS) for the period 2004-2015. In the Nielsen data: college is computed using the male household head, or the female household head if no male head is reported; non-employed is computed using only households with a male head aged between 18 and 65; homeownership is computed following Stroebel et al. (2014), where a household is considered a homeowner if they report living in either a one-family house or a one-family condo or co-op. In the CPS, college is computed for the whole population, and non-employed is computed for men aged between 18 and 65. The final three columns report average proportion of households moving across zip codes (within counties in the CPS), moving across counties, and across states.
Sources: Authors calculations using Nielsen Consumer Panel, CPS.

Table A.3 reports the number of panelists in each year, as well as the proportion of panelists remaining in the panel 2, 3 and 5 years after observing them in a given year. From 2006 to 2007 the size of the panel increases substantially, from 37,786 to 63,350. Attrition rates in both the short and medium term do not vary substantially over time and appear to be relatively low. The average proportion of panelists remaining after 2, 3, and 5 years is 0.81, 0.68, and 0.52, respectively.

Figure VI presents the age profile of CPI-deflated consumption expenditure and the homeownership rate. Total household expenditure follows a well-known hump-shaped pattern over the life-cycle. Consumption expenditure for homeowners does not differ markedly from the average household. Household expenditure per person also follows a hump shape, although the initial rise in expenditure occurs later than for total household expenditures.

[36]Homeownership rates for the United States are from FRED (code: USHOWN).

Table A.3
Number of Panelists in Nielsen Consumer Panel

Year	Number Panelists	Remain, 2 years	Remain, 3 years	Remain, 5 years
2004	39577	0.79	0.64	0.51
2005	38863	0.78	0.69	0.53
2006	37786	0.85	0.73	0.55
2007	63350	0.79	0.66	0.47
2008	61440	0.80	0.65	0.48
2009	60506	0.77	0.64	0.50
2010	60658	0.78	0.67	0.52
2011	62092	0.82	0.71	0.55
2012	60538	0.82	0.70	0.55
2013	61097	0.82	0.70	–
2014	61557	0.82	0.71	–
2015	61380	0.83	–	–
2016	63150	–	–	–

Notes: The first column reports the number of unique panelists per year. The remaining columns report the proportion of unique panelists remaining in the panel for 2, 3, and 5 years.
Sources: Authors calculations using Nielsen Consumer Panel.

Figure VII shows the annual growth rate of nominal non-durable consumption per capita for the Nielsen data and for data taken from the personal consumption expenditures section of NIPA. Growth rates are computed from the Consumer Panel data first by computing the growth rate in the survey-weighted average of total consumption-to-household size for all households in the panel. Because there may be selection in panelist attrition, I also compute the growth rate in the survey-weighted average of total consumption-to-household size for households that remain in the panel for each pair of consecutive years. For national accounts data, growth rates are computed as non-durable personal consumption expenditures-to-population, and non-durable personal consumption less energy expenditures-to-population. The patterns of growth rates in non-durable consumption for the Consumer Panel and national accounts data are similar, with the notable exceptions of 2005 and 2012.

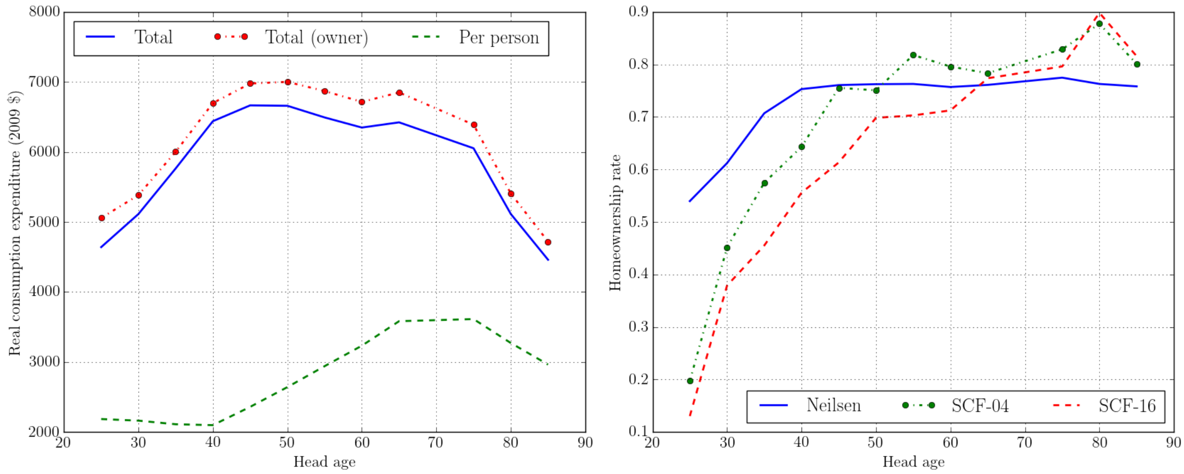

Figure VI
Consumption and Homeownership Over the Life Cycle

Notes: Consumption and homeownership rates in the Nielsen data are pooled across all years by age group. The left panel plots total household consumption for all households (blue, solid line), total household consumption for all (implied) homeowners (red, dash-circle), and total household consumption normalized by the household size (green, dashed line). Consumption values are reported in real, 2009 dollars. The right panel plots homeownership rates in the Nielsen data (blue, solid line), the 2004 SCF (red, dash-circle), and the 2016 SCF (green, dashed line).
Source: Author's calculations using Nielsen Consumer Panel, Survey of Consumer Finances.

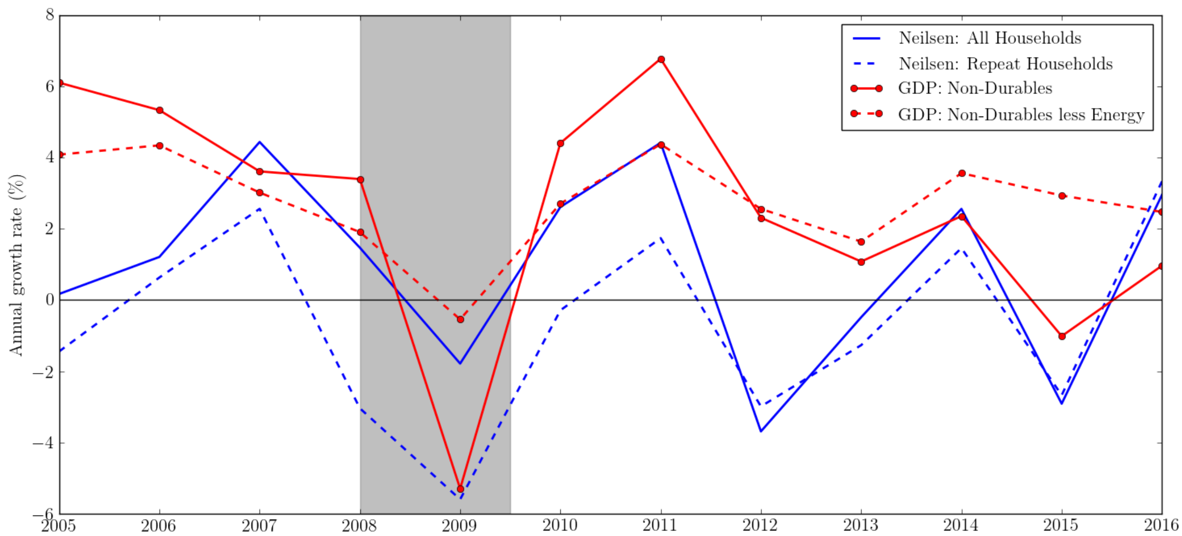

Figure VII
Per Capita Non-Durables Consumption Growth

Notes: Annual nominal non-durable consumption growth per capita in the Nielsen Consumer Panel and national accounts data. The solid blue line is the growth rate in the survey-weighted average of total consumption-to-household size. The dashed blue line is the growth rate in the survey-weighted average of total consumption-to-household size for households that remain in the panel for consecutive years. The solid red line is the growth rate in non-durable personal consumption expenditures-to-population. The dashed red line is the similar, but using non-durable personal consumption expenditure for all goods minus non-durable personal consumption expenditure for gasoline and other energy goods. Shaded area represents recession dates.
Source: Author's calculations using Nielsen Consumer Panel, NIPA via FRED.

D. Bartik Instrument Characteristics and Quality Prices

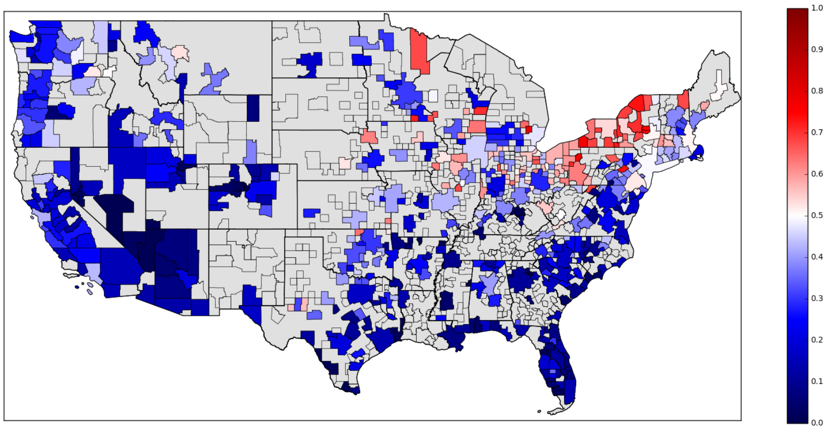

(a) Proportion Built Prior to 1960

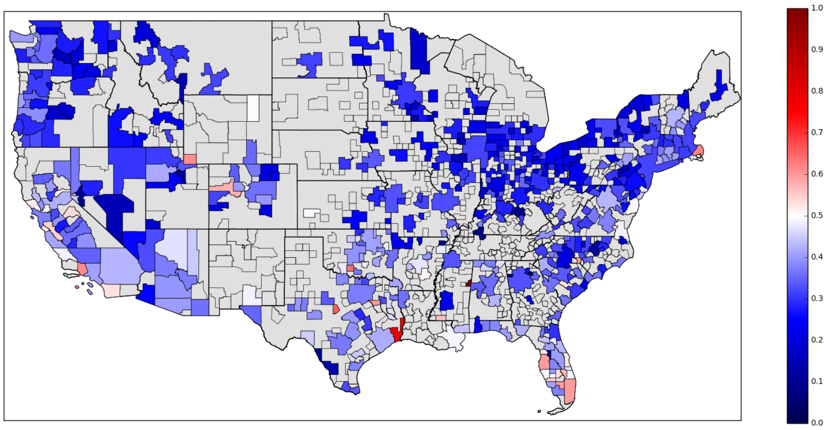

(b) Proportion Built 1960 to 1990

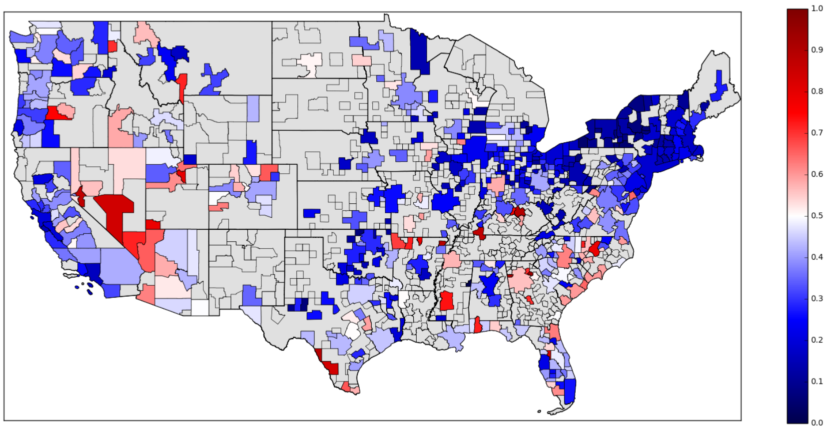

(c) Proportion Built After 1990

Figure VIII
Distribution of Housing Age Across Cities

Notes: Maps show the proportion of houses built in three different year groups for each CBSA. Maps show 429 CBSAs for which there are at least 500 individual transactions of unique houses from 1994 to 2005. These CBSAs contained 80 percent of the US population in the 2000 Census.
Sources: Author's calculations using ZTRAX.

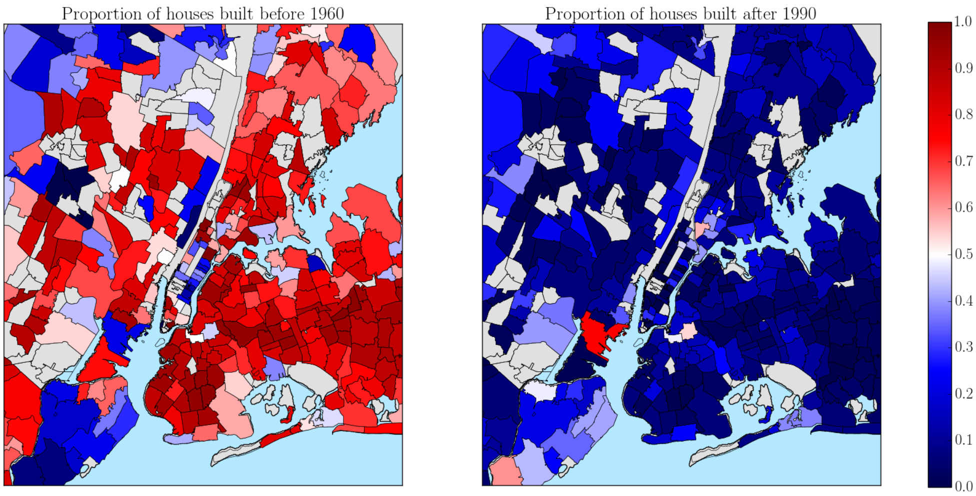

(a) New York

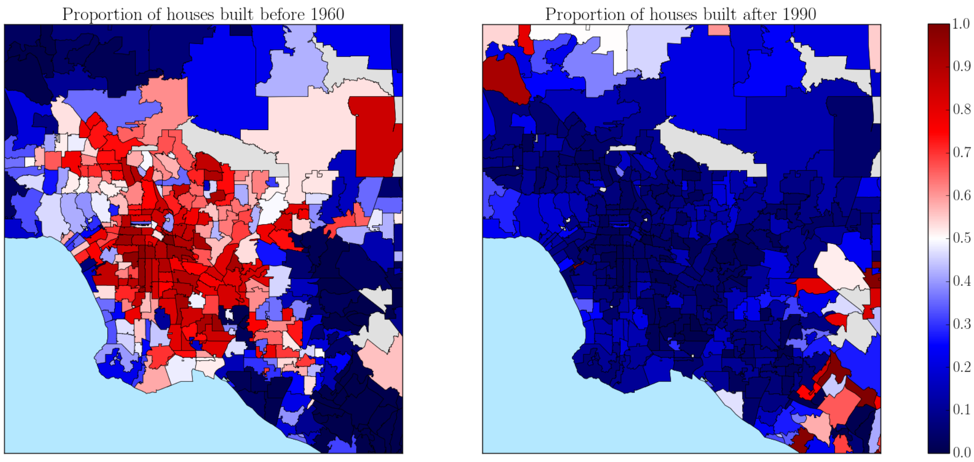

(b) Los Angeles

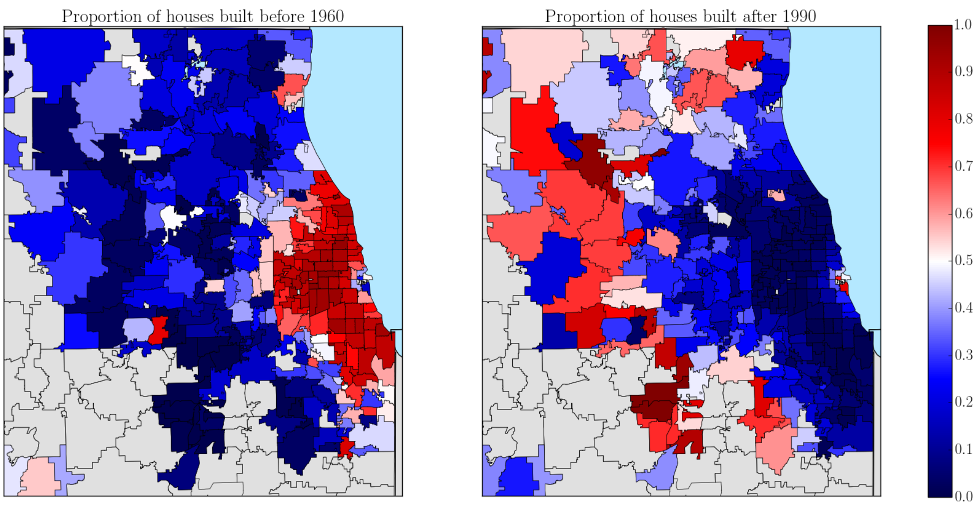

(c) Chicago

Figure IX
Distribution of Housing Age Across Zip Codes

Notes: Maps show the proportion of houses built in two different year groups for each zip code within a city. Maps show zip codes for which there are at least 50 individual transactions of unique houses from 1994 to 2005.
Sources: Author's calculations using ZTRAX.

E. Details of Bartik Instrument Construction

In order to compute the local housing characteristic shares, I use data on unique houses reported in ZTRAX. Because the sample period for the main empirical analysis is 2005-2016, I construct the local shares for a pre-sample period: 1994-2005. I include 2005, because housing data for some locations is not available in ZTRAX prior to 2005. However, the results of the analysis are not quantitatively affected by excluding these locations and ending the pre-period in 2004.

The set of housing characteristics used to construct instruments are: house age, building floor size, property lot size, number of bedrooms, and number of bathrooms. There are many other housing characteristics described in ZTRAX, however many of the fields containing this information are not broadly populated. Moreover, several important fields, such as total number of rooms, are not reported consistently across the data set. For example, in an unreported exercise, I found that the average number of rooms computed from ZTRAX was extremely inconsistent with the average number of rooms computed from the 2000 Census. One reason for this is that a 'room' is not easily defined, leading to variation in reports from assessors. Other variables, such as floor size, number of bedrooms, number of bathrooms, or property age are much better defined, and so likely to reflect higher quality data.

I assume that the houses transactions recorded in ZTRAX reflect a random sample of the existing housing stock. However, there could be a selection bias in this measure if, for example, lower quality houses tend to sell less often (i.e. a classic 'lemons' problem). In order to investigate whether selection bias is a problem, Figure X compares the proportion of the housing stock built during different periods of the 20th century for each county according to the data from the 2005 American Community Survey and the data derived from transactions in ZTRAX.[37] I present population weighted scatter plots against the 45-degree line reflecting perfect correlation between the two measures. For most year groups, the data like close to the 45-degree line, indicating that the ZTRAX data does not generally over- or undersample housing age. Although the fraction of houses built in the 2000s is somewhat overstated in the ZTRAX data, this is likely attributable to the fact that a higher proportion of all new houses are sold at any given time than the proportion of old houses sold.

Figure Xb reports a similar exercise but for number of bedrooms. There appears to be systematic mis-reporting of the share of houses with zero bedrooms, although the proportion of houses with 2-3 or 4 or more bedrooms appears to reasonable. For this reason, I exclude houses reporting zero bedrooms from the analysis.[38] Additionally, Section 8 considers a version of the Bartik instrument using housing age as the only house characteristic.

[37]The year groups are selected to correspond to the categories reported in the ACS.

[38]This is approximately 16 percent of the sample. Despite the apparent measurement error, the main results are unaffected if include these zero bedroom houses.

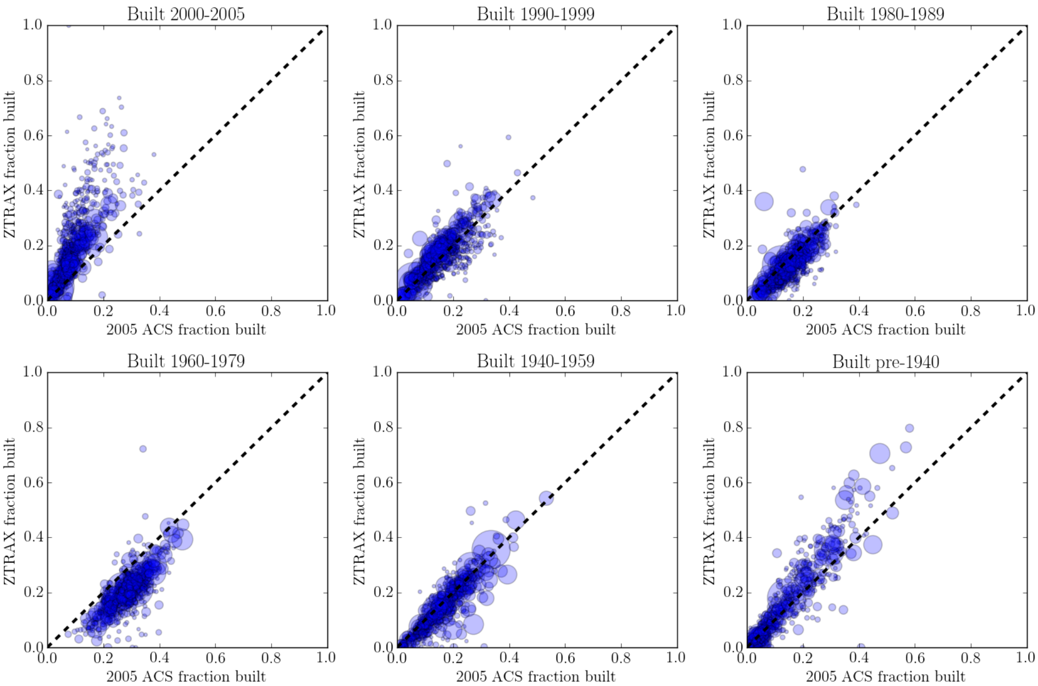

(a) Share of Houses by Year Built

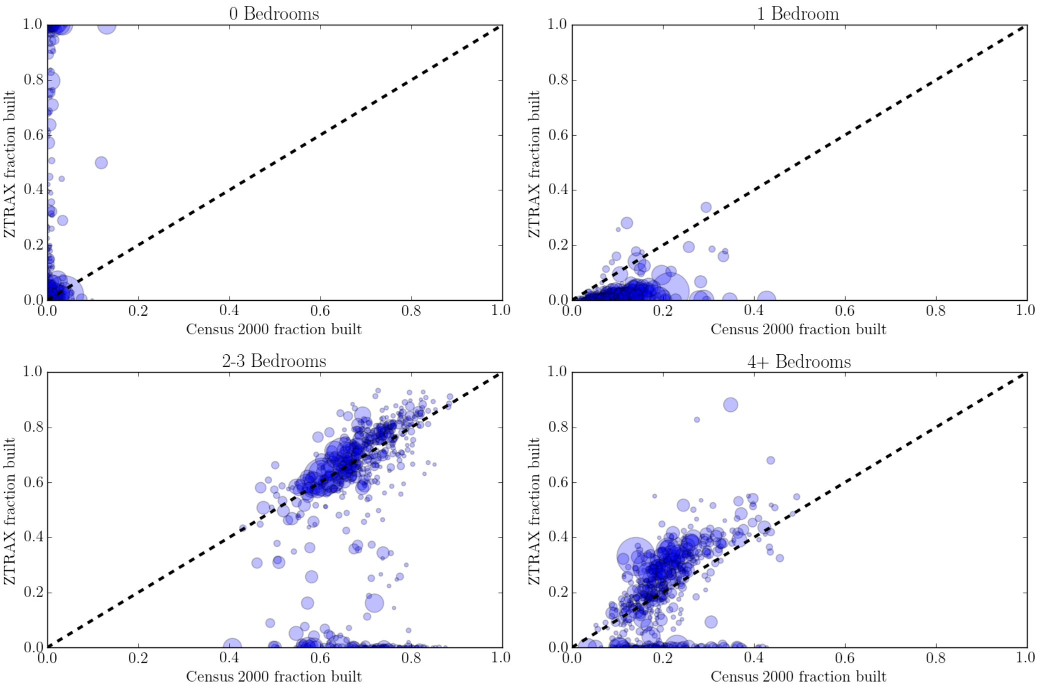

(b) Share of Houses by Number of Bedrooms

Figure X
Local House Characteristic Shares in ZTRAX and 2005 ACS

Notes: County share of housing stock by year built and number of bedrooms. Shares computed from the 2005 American Community Survey and ZTRAX data for unique houses sold between 1994 and 2005. Note that the ACS reports data for occupied houses only, while the ZTRAX data is drawn from all houses sold. Each blue circle is an observation for a county, weighted by the relative size of the housing stock as reported in the ACS. The black dashed line is the 45 degree line.
Sources: ACS, ZTRAX

F. Control Variables Included in Various Regression Specifications

Household controls: All household controls are reported in the Consumer Panel. The controls refer either to the head of household, or apply to the household as a whole. When a household reports two household heads, I use information from the head male. Controls include: real household income growth, age, the square of age, the change in household size, an indicator variable for the presence of children, marital status, race, an indicator for Hispanic or Latino origin, occupation, education. Household income is reported as for the year two years prior to the current panel date. Income is reported as a categorical variable. In order to construct income growth, I record current income as the value assigned to the upper boundary of the current income category. Income is then deflated by the CPI, before the annual growth rate is computed.

Fixed effects: city-level (i.e. CBSA) fixed effects are included in all specifications. Some specifications include year-fixed effects.

Local business cycle controls: Zip code-level real income growth from the IRS SOI data, and county-level unemployment growth from BLS data. These data are reported annually.

Local industry composition controls: All industry controls are annual time series from the County Business Patterns survey. For each county, I take the total number of employees in a given industry, and divide by total employment in that county. I use employment shares for the following industries: agriculture, construction, manufacturing, retail trade, and finance/real estate/insurance (FIRE).

Local demographic controls: All demographic controls are county-level observations from the 2000 Census (i.e. a single cross-section of observations). The demographic controls reported as a proportion of the local population are: race=white, race=black, Hispanic ethnicity, foreign-born, those with at least some college education, homeowners. Other demographic controls are: median age, mean household size, mean travel time to work. Each demographic variable is interacted with year-dummy variables, as suggested by Goldsmith-Pinkham et al. (2018).

G. Robustness to Variations in the Bartik Instrument

Table A.4
Consumption Response to House Prices Using Bartik and Saiz Instruments, Common Sample

	Annual household real non-durable consumption growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta p_{county,t}$	0.070^{***}	0.094^{***}	0.105^{***}	-0.077	0.111^{***}	-0.081	0.110^{***}	-0.086
	(0.017)	(0.022)	(0.030)	(0.094)	(0.032)	(0.096)	(0.031)	(0.101)
Instrument	Bartik	Saiz	Bartik	Saiz	Bartik	Saiz	Bartik	Saiz
Observations								
Total	227,490	227,490	227,490	227,490	227,490	227,490	227,490	227,490
Households	51,000	51,000	51,000	51,000	51,000	51,000	51,000	51,000
CBSAs	216	216	216	216	216	216	216	216
Counties	563	563	563	563	563	563	563	563
Controls								
Household	✓	✓	✓	✓	✓	✓	✓	✓
CBSA FE	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	×	×	✓	✓	✓	✓	✓	✓
Local	×	×	×	×	✓	✓	✓	✓
Industry	×	×	×	×	×	×	✓	✓
F-statistic	114.1	62.99	35.67	4.52	38.15	4.66	38.33	4.4

Notes: Models estimated using county prices and household consumption growth for homeowners. All specifications in this table use the same sample of households, counties, and time periods. House prices are instrumented, alternately, using the Bartik and Saiz instruments. The Bartik instrument uses county-level house characteristic shares, as described in Section 5. The Saiz instrument is the interaction of CBSA-level housing supply elasticities and regional house price growth. Control variables are described in Section F. Both standard errors (reported in parentheses) and first stage F-statistic are clustered at the CBSA level. * significant at 10%, ** significant at 5%, *** significant at 1%.
Sources: Author's calculations using data from BLS, CBP, 2000 Census, FHFA, IRS SOI, Nielsen, Saiz (2010), ZTRAX.

Table A.5

Consumption Response to House Prices Using National and Regional Instruments

	Annual household real non-durable consumption growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta p_{county,t}$	0.095***	0.130***	0.120***	0.156***	0.146***	0.265	−0.834	0.041
	(0.017)	(0.019)	(0.022)	(0.020)	(0.041)	(0.850)	(1.715)	(0.069)
Instrument	Bartik (Region)	Bartik (Nation)	Saiz (Region)	Saiz (Nation)	Bartik (Region)	Bartik (Nation)	Saiz (Region)	Saiz (Nation)
Observations								
Total	256,824	256,824	256,824	256,824	227,490	227,490	227,490	227,490
Households	54,551	54,551	54,551	54,551	51,000	51,000	51,000	51,000
CBSAs	216	216	216	216	216	216	216	216
Counties	571	571	571	571	563	563	563	563
Controls								
Household	✓	✓	✓	✓	✓	✓	✓	✓
CBSA FE	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	×	×	×	×	✓	✓	✓	✓
Local	×	×	×	×	✓	✓	✓	✓
Industry	×	×	×	×	✓	✓	✓	✓
Demographic	×	×	×	×	✓	✓	✓	✓
F-statistic	127.25	86.37	70.54	87.19	43.97	0.18	0.25	17.65

Notes: Models estimated using county prices and household consumption growth for homeowners. House prices are instrumented, alternately, using the Bartik and Saiz instruments. Separate estimates presented using Bartik and Saiz instruments constructed using national and regional variation in house prices. Control variables are described in Section F. Both standard errors (reported in parentheses) and first stage F-statistic are clustered at the CBSA level. * significant at 10%, ** significant at 5%, *** significant at 1%.
Sources: Author's calculations using data from BLS, CBP, 2000 Census, FHFA, IRS SOI, Nielsen, Saiz (2010), ZTRAX.

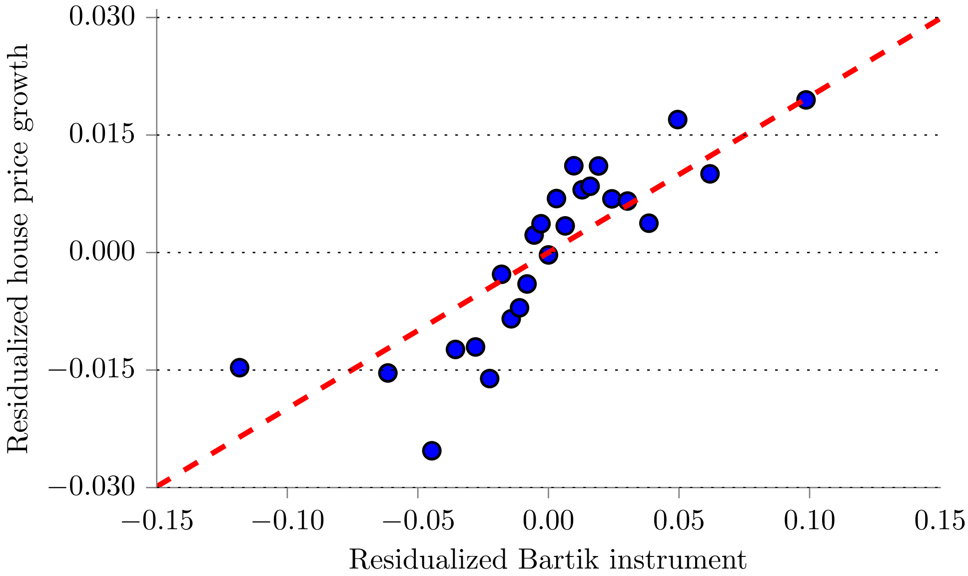

Figure XI
Effect of Extended Bartik Instrument on House Price Growth

Notes: Residualized extended Bartik instrument and county house price growth, representing the first stage TSLS regression. The residualized variables are constructed using the full set of controls used the IV estimation of the consumption elasticities. The value of the Bartik instrument is split into equal sized bins, where the mean of the instrument and house prices is computed for observations falling within each bin. The red dashed line plots the first stage regression coefficient on the Bartik instrument. The slope is 0.050 as reported in column (5) of Table V.
Sources: Author's calculations using data from BLS, CBP, 2000 Census, FHFA, IRS SOI, Nielsen, ZTRAX.

Table A.6
Consumption Response to House Prices Using Alternative Bartik Instruments

	Annual household real non-durable consumption growth					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta p_{county,t}$	0.095***	0.097***	0.150***	0.146***	0.148***	0.155**
	(0.017)	(0.018)	(0.016)	(0.041)	(0.041)	(0.061)
Instrument	Bartik (Benchmark)	Bartik (Age only)	Bartik (Add size)	Bartik (Benchmark)	Bartik (Age only)	Bartik (Add size)
Observations						
Total	256,824	256,824	256,824	227,490	227,490	227,490
Households	54,551	54,551	54,551	51,000	51,000	51,000
CBSAs	216	216	216	216	216	216
Counties	571	571	571	563	563	563
Controls						
Household	✓	✓	✓	✓	✓	✓
CBSA FE	✓	✓	✓	✓	✓	✓
Time FE	×	×	×	✓	✓	✓
Local	×	×	×	✓	✓	✓
Industry	×	×	×	✓	✓	✓
Demographic	×	×	×	✓	✓	✓
F-statistic	127.25	137.87	154.48	43.97	43.37	28.88

Notes: Models estimated using county prices and household consumption growth for homeowners. House prices are instrumented using: the benchmark Bartik instrument as described in Section 5; the Bartik instrument using house age as the only characteristic as in Section 8; and the Bartik instrument with added characteristics for building floor size and property lot size as in Section 8. Control variables are described in Section F. Both standard errors (reported in parentheses) and first stage F-statistic are clustered at the CBSA level. * significant at 10%, ** significant at 5%, *** significant at 1%.
Sources: Author's calculations using data from BLS, CBP, 2000 Census, FHFA, IRS SOI, Nielsen, Saiz (2010), ZTRAX.

Table A.7
Effect of Extended Bartik Instrument on County House Prices

	Annual county real house price growth				
	(1)				
$B_{county,t}$	0.530^{***}	0.325^{***}	0.302^{***}	0.307^{***}	0.199^{***}
	(0.043)	(0.063)	(0.060)	(0.061)	(0.037)
Observations					
Total	256,824	256,824	230,865	227,490	227,490
Households	54,551	54,551	51,408	51,000	51,000
CBSAs	216	216	216	216	216
Counties	571	571	570	563	563
Controls					
Household	✓	✓	✓	✓	✓
CBSA FE	✓	✓	✓	✓	✓
Time FE	×	✓	✓	✓	✓
Local	×	×	✓	✓	✓
Industry	×	×	×	✓	✓
Demographic	×	×	×	×	✓
F-statistic	154.48	26.45	25.46	25.67	28.88
Adj. R^2	0.46	0.61	0.62	0.62	0.74

Notes: Models estimate the effect of the extended Bartik instrument on county house prices, as in the first stage regression of Equation 3. The extended Bartik instrument is constructed as described by Equation (11). The first stage regressions use the same household-level data and include the same household, local, industry, and demographic controls as in the full IV estimation of the consumption elasticities. See Section efsec: Construction of the Bartik Instrument for details. Both standard errors (reported in parentheses) and first stage F-statistic are clustered at the CBSA level. * significant at 10%, ** significant at 5%, *** significant at 1%.
Sources: Author's calculations using data from BLS, CBP, 2000 Census, FHFA, IRS SOI, Nielsen, ZTRAX.

Table A.8
Consumption Response to Zip Code House Prices

	Annual household real non-durable consumption growth				
	(1)	(2)	(3)	(4)	(5)
$\Delta p_{zipcode,t}$	0.078^{***}	0.091^{***}	0.102^{***}	0.102^{***}	0.113^{***}
	(0.011)	(0.027)	(0.029)	(0.029)	(0.035)
Instrument	Bartik	Bartik	Bartik	Bartik	Bartik
Observations					
Total	223,726	223,726	201,254	197,925	194,560
Households	50,446	50,446	47,601	47,131	46,124
CBSAs	466	466	466	463	460
Zip codes	7,870	7,870	7,772	7,739	7,509
Controls					
Household	✓	✓	✓	✓	✓
CBSA FE	✓	✓	✓	✓	✓
Time FE	×	✓	✓	✓	✓
Local	×	×	✓	✓	✓
Industry	×	×	×	✓	✓
Demographic	×	×	×	×	✓
F-statistic	247.16	58.58	61.82	62.92	66.18

Notes: Models estimated using zip code prices and household consumption growth for homeowners. IV estimates reported where house prices are instrumented using the Bartik instrument described in Section 5, but constructed at the zip code level. Control variables are described in Section F. Standard errors (reported in parentheses) and first stage F-statistic clustered at the CBSA level. * significant at 10%, ** significant at 5%, *** significant at 1%.
Sources: Author's calculations using data from BLS, CBP, FHFA, IRS SOI, Nielsen, Zillow, ZTRAX.