House Prices and Consumption: A New Instrumental Variables Approach

ONLINE APPENDIX

James Graham and Christos A. Makridis *

November 7, 2021

*James Graham: University of Sydney, james.a.graham@sydney.edu.au. Christos: Arizona State University and Stanford University, cmakridi@stanford.edu

A Supplement to Data and Measurement

A.1 Data Dictionary

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/.

A.2 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. We aim to work with a consistent data set containing typical property transactions conducted by residential owner-occupiers. To this end, we carry out the following cleaning procedure.

We 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. We restrict properties to those that are are non-commercial, and that are either single family residences or owner-occupied properties as long as they are not apartments or vacation homes. These property types include: Single Family Residences, Townhouses, Cluster Homes, Condominiums, Cooperatives, Row Houses, Bungalows, Patio Homes, Garden Homes, and Landominiums. This latter restriction includes 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. We also focus only on property transactions with non-zero sales prices, thereby removing all mortgages, mortgage refinancing, and transfers or gifts. We exclude transactions that may have been subject to 'house flipping', thereby distorting the market value of the house. To do this, we remove any house sale that occurs within 180 days of a prior sale of the same house. Additionally, we remove transactions with a sale price of less than $10,000 as well as those with no reported transaction date. We 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.[1] For these states, property deeds and assessment records may still be reported to the ZTRAX database. We collect data on housing characteristics for these states, but we cannot use the transaction data on sales prices.[2] Instead, for these states we use publicly available, geographically aggregated Zillow house price indexes. After data cleaning, there are 55 million individual transactions available between 1994 and 2016.

A.3 Consumer Panel Data

Table A.5 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.6 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, we report the proportion of households that have moved in the

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

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

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, we 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.

While home ownership is not directly reported in the Consumer Panel, we follow Stroebel and Vavra (2019) who 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, and does not change significantly over the sample. From 2004 to 2015, the home ownership rate for the US as a whole fell from 69 percent to 64 percent.[3] The second panel of Figure A.1 reports the age profile of home ownership, which reveals that implied home ownership rates are overstated by between 15 and 30 percentage points for young households relative to data from the SCF. Implied home ownership rates for older households are very similar to those reported in the SCF. For most of the empirical results, we make use of the sample of implied home owners only.

Table A.5: Household summary statistics, Neilsen Consumer Panel

	Income	Expenditure	Age	FamilySize
Mean	54205.91	6106.28	54.05	2.49
Median	51403.83	5171.33	53.00	2.00
StdDev	34420.40	3928.87	14.41	1.42

Notes: Means, medians, and standard deviations computed using Consumer Panel survey weights. Income is the households real 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 real 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 ZTRAX, FRED.

Table A.7 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 A.1 presents the age profile of CPI-deflated consumption expenditure and the home-ownership 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.

[3]Home ownership rates for the United States are from FRED (code: USHOWN).

Table A.6: Demographics, Nielsen Consumer Panel

	College	Non-Employed	Homeownership	Moved Zip	Moved County	Moved State
Nielsen	0.418	0.196	0.752	0.030	0.017	0.008
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 and Vavra (2019), 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.7: Number of panelists, Nielsen Consumer Panel

Year	Number Panelists	Remain, 2 years	Remain, 3 years	Remain, 5 years
2004	34260	0.91	0.74	0.59
2005	36769	0.82	0.73	0.57
2006	33882	0.89	0.77	0.58
2007	52931	0.90	0.75	0.54
2008	57051	0.85	0.69	0.52
2009	55252	0.83	0.69	0.54
2010	54978	0.85	0.72	0.56
2011	56534	0.87	0.75	0.59
2012	55090	0.87	0.75	0.59
2013	55637	0.86	0.75	
2014	56035	0.87	0.76	
2015	55875	0.88		
2016	50828			

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 A.2 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 of possible selection effects arising from panelist attrition, we 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 A.6: Consumption and Homeownership Over the Life Cycle

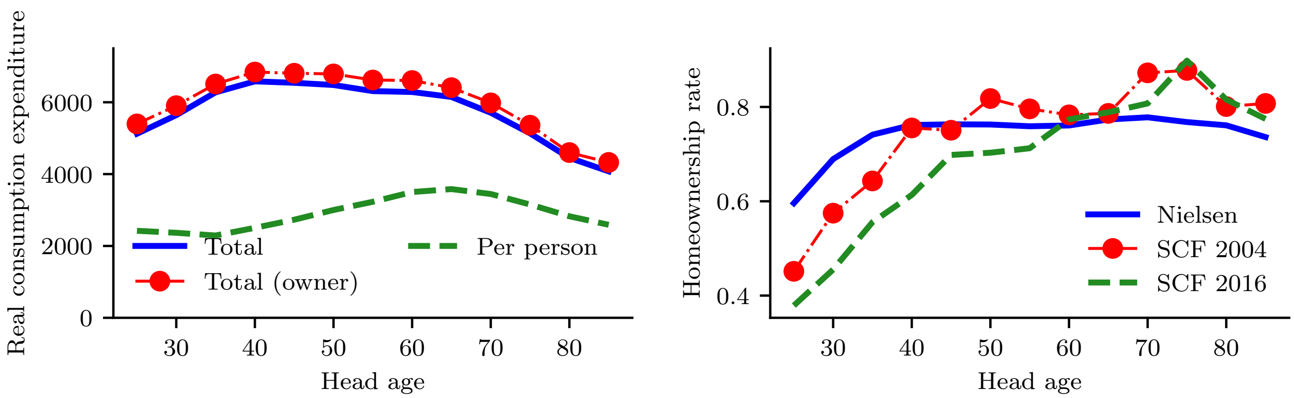

Notes: Consumption and homeownership rates in the Nielsen data are pooled across all years by age group. The left panel plots total household consumption for all households (blue, solid line), total household consumption for all (implied) homeowners (red, dash-circle), and total household consumption 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 A.7: Per Capita Non-Durables Consumption Growth

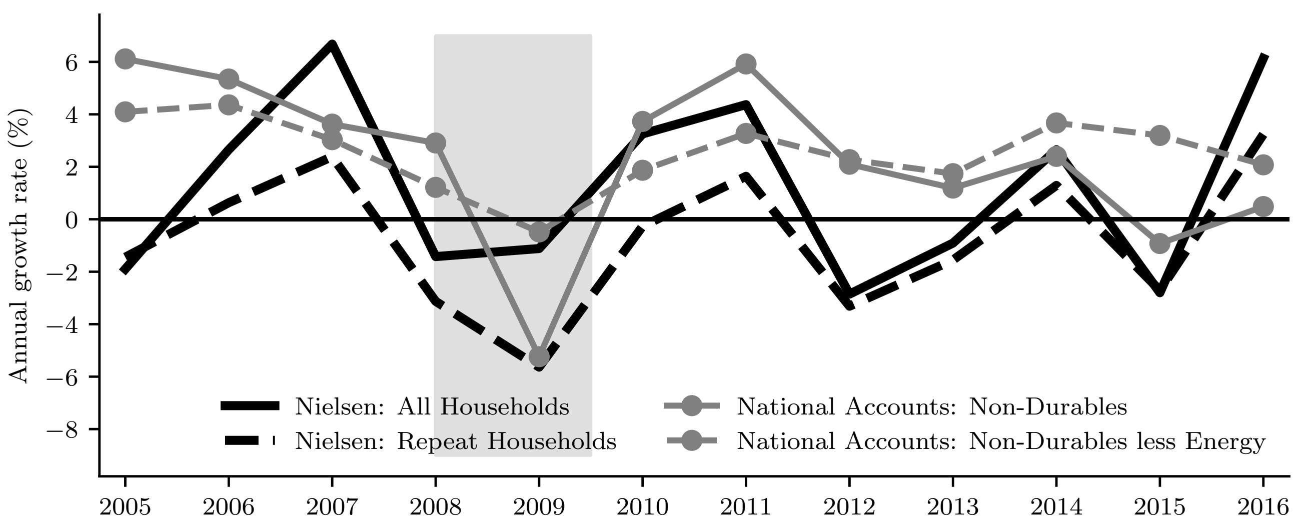

Notes: Annual nominal non-durable consumption growth per capita in the Nielsen Consumer Panel and national accounts data. The solid blue line is the growth rate in the survey-weighted average of total consumption-to-household size. The dashed blue line is the growth rate in the survey-weighted average 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.

A.4 Control Variables in 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, we 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, we 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: County fixed effects are included in all specifications. Some specifications include year fixed effects. We alternatively experimented with CBSA fixed effects, but this had no material effect on our results.

Local business cycle controls: County-level pre-tax 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, we take the total number of employees in a given industry, and divide by total employment in that county. We use employment shares for the following industries: 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, Sorkin and Swift (2020).

B Details of Bartik-like Instrument Construction

In order to compute the local housing characteristic shares, we use data on unique houses reported in ZTRAX. Because the sample period for the main empirical analysis is 2005-2016, we construct the local shares for a pre-sample period: 1994-2005. We 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, we 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.

We 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, Panel A in Figure B.3 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.[4] We 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 under-sample 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.

Panel B in Figure B.3 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, we exclude houses reporting zero bedrooms from the analysis.[5] Additionally, Section C considers a version of the Bartik-like instrument using housing age as the only house characteristic.

Figure B.8: Local House Characteristic Shares in ZTRAX and 2000 Census

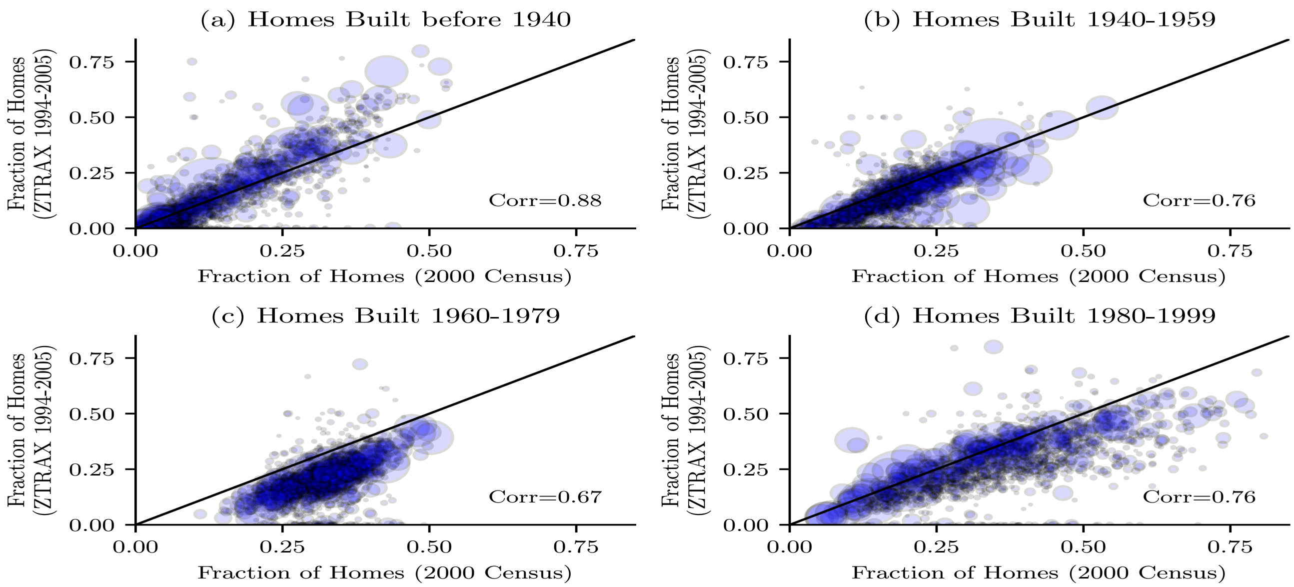

Notes: County shares of the housing stock by year built. Shares computed from the 2000 Census and ZTRAX data for unique houses sold between 1994 and 2005. Note that the Census reports data for occupied houses only, while the ZTRAX data is drawn from all houses sold. Observations are weighted by 2000 Census populations.
Sources: Census, ZTRAX

C Robustness of the Bartik-Like Instrument

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

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

Table B.8: Correlations of Local Characteristic Shares and Housing Supply

Built	Saiz	Wharton	Bedrooms	Saiz	Wharton	Bathrooms	Saiz	Wharton
pre-1940	−0.11	0.14	One	−0.17	0.20	Zero	−0.11	0.10
1940-1949	−0.10	0.01	Two	−0.17	0.13	One	0.14	−0.07
1950-1959	−0.15	0.04	Three	0.22	−0.23	Two	0.05	−0.03
1960-1969	−0.08	0.02	Four	−0.00	−0.07	Three	−0.06	0.05
1970-1979	−0.09	0.03	Five+	−0.16	0.11	Four+	−0.11	0.03
1980-1989	−0.13	0.10						
1990-1999	0.18	−0.10						
2000-2005	0.30	−0.23						

Notes: Correlation between local shares for housing characteristics and measures of local housing supply from Saiz (2010). County-population weighted correlations computed for the 233 cities (CBSAs) and 673 counties available for all three instruments.
Sources: Author's calculations using 2000 Census, Gyourko, Saiz and Summers (2008), Saiz (2010), ZTRAX.

Figure B.9: Distribution of Housing Age Across Zip Codes

(a) New York: Houses Build Pre-1960

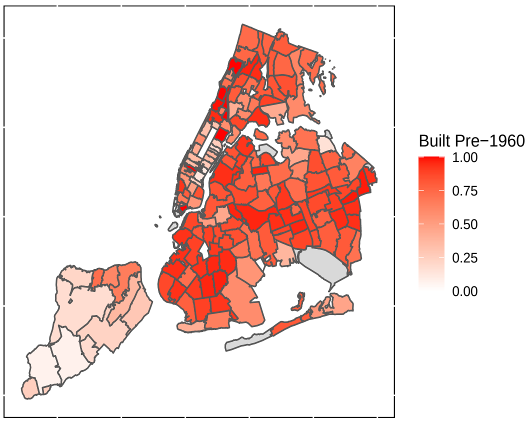

(b) New York: Houses Build Post-1990

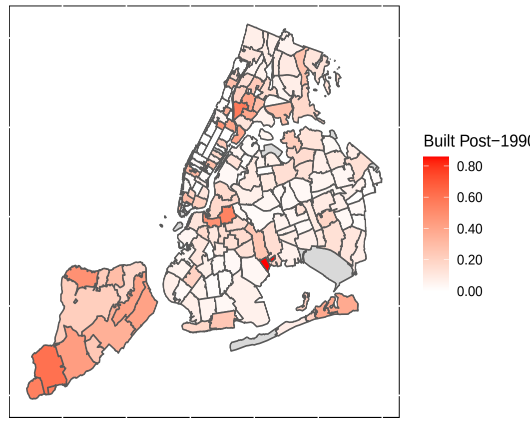

(c) Los Angeles: Houses Build Pre-1960

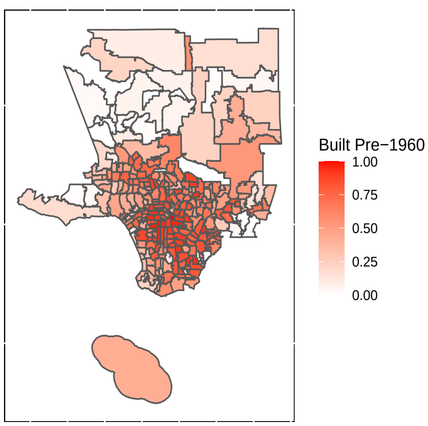

(d) Los Angeles: Houses Build Post-1990

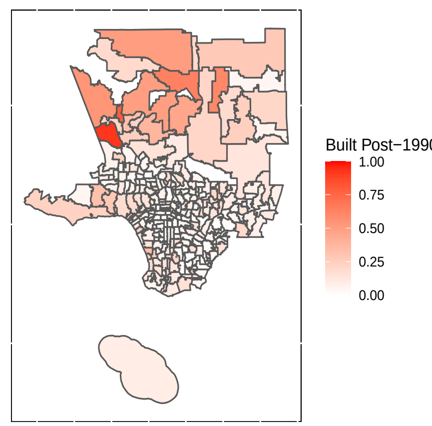

(e) Chicago: Houses Build Pre-1960

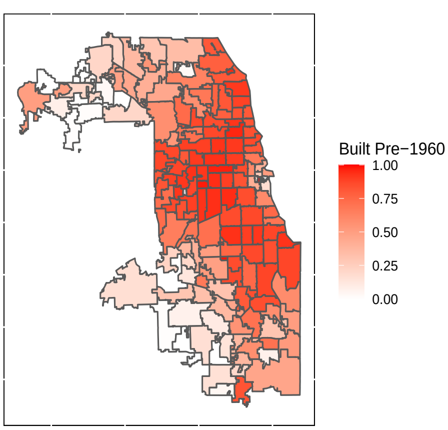

(f) Chicago: Houses Build Post-1990

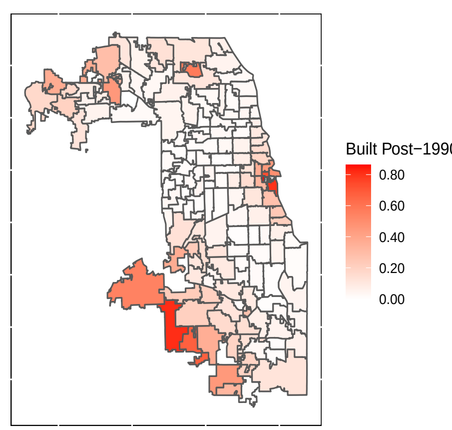

Notes: Maps show the proportion of houses of two vintages – pre-1960 and post-1990 – in each zip code within a city.
Sources: Author's calculations using ZTRAX.

Table C.9: Further Heterogeneity in Consumption Response to House Prices

	Real annual non-durable household consumption growth					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta p_{county,t}$	0.100***	0.034	0.101***	0.053	0.100***	0.063
	(0.031)	(0.059)	(0.031)	(0.060)	(0.035)	(0.068)
$\Delta p_{county,t} \times \mathbb{1}(Homeownership > Median)$			−0.004	−0.060		
			(0.031)	(0.046)		
$\Delta p_{county,t} \times \mathbb{1}(RetailTrade > Median)$					−0.001	−0.052
					(0.028)	(0.047)
Method	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
Sample	Owners	Renters	Owners	Renters	Owners	Renters
Observations						
Total	289,665	66,064	289,665	66,064	289,665	66,064
Households	64,898	17,965	64,898	17,965	64,898	17,965
Counties	1,002	925	1,002	925	1,002	925
Adjusted R-squared	0.01	0.01	0.01	0.01	0.01	0.01

Notes: The table reports estimates of Equation (1) with household controls, county business cycle controls, county industry composition controls, and county and year fixed effects. Columns (1), (3), and (5) use the sample of inferred owner households, while Columns (2), (4) and (6) use the sample of inferred renter households. Columns (3) and (4) test for heterogeneity of consumption responses in counties with above-median homeownership rates. Columns (5) and (6) test for heterogeneity of consumption responses in counties with above-median retail trade employment sharaes. All columns are instrumented using the Bartik-like instrument discussed in the text. Standard errors are clustered at the county level.
Sources: BLS, CBP, FHFA, IRS, Nielsen, Zillow, ZTRAX.

Table C.10: Consumption Response to Zip Code House Prices Using the Bartik Instrument

	Real annual non-durable household consumption growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta p_{zipcode,t}$	0.062***	0.014*	0.073***	0.091***	0.095***	0.087***	0.112***	0.129***
	(0.006)	(0.008)	(0.018)	(0.015)	(0.017)	(0.015)	(0.016)	(0.020)
Method	OLS	OLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
Observations								
Total	201,846	201,846	201,846	201,846	201,846	201,846	201,846	201,846
Households	45,226	45,226	45,226	45,226	45,226	45,226	45,226	45,226
Zip Codes	6,654	6,654	6,654	6,654	6,654	6,654	6,654	6,654
Controls								
Household	N	Y	N	Y	N	N	N	Y
Local	N	N	N	N	Y	N	N	Y
Industry	N	N	N	N	N	Y	N	Y
Demographic	N	N	N	N	N	N	Y	Y
Zip Code FE	N	Y	N	Y	Y	Y	Y	Y
Year FE	N	Y	N	Y	Y	Y	Y	Y
Standard Errors								
Zipcode Clusters	Y	Y	N	N	N	N	N	N
Adão, Kolesár and Morales (2019)	N	N	Y	Y	Y	Y	Y	Y
F-statistic	–	–	366.93	304.20	312.55	304.95	413.52	362.21
Adjusted R-squared	0.00	0.01	0.00	0.01	0.00	0.00	0.00	0.01

Notes: The table reports estimates of Equation (1) using zip code-level house price growth. The regressions include household controls, county and zip code business cycle controls, county industry composition controls, zip code demographic controls, and zip code and year fixed effects. Household controls come from the Nielsen Consumer Panel, including: real household income growth, a quadratic in age, the change in household size, the presence of children, marital status, race, Hispanic or Latino origin, occupation, and education. Local business cycle controls include: county unemployment growth from the BLS and zip code-level real per capita income from the IRS. Local industry composition controls include: the employment share of construction, manufacturing, retail trade, and finance/real estate/insurance (FIRE) from the CBP. Local demographic controls include population shares of: black, Hispanic, foreign-born, those with at least some college education, homeowners, median age, household size, mean travel time to work, and employment shares in construction, manufacturing, retail trade, and FIRE. Each of these local demographic variables are interacted with year-dummy variables as suggested by Goldsmith-Pinkham, Sorkin and Swift (2020). Standard errors and F-statistics for 2SLS models are estimated following Adão, Kolesár and Morales (2019), also allowing for correlation in housing characteristics through time.
Sources: BLS, CBP, FHFA, IRS, Nielsen, Zillow, ZTRAX.

Table C.11: Consumption Response to House Prices Using Alternative Bartik Instruments

	Real annual non-durable household consumption growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta p_{county,t}$	0.102***	0.093**	0.106***	0.099**	0.157***	0.067	0.119***	0.332
	(0.014)	(0.046)	(0.014)	(0.044)	(0.014)	(0.069)	(0.015)	(0.388)
Instrument	Baseline	Baseline	Age Only	Age Only	Add Size	Add Size	National	National
Observations								
Total	289,665	289,665	289,665	289,665	289,665	289,665	289,665	289,665
Households	64,898	64,898	64,898	64,898	64,898	64,898	64,898	64,898
Counties	1,002	1,002	1,002	1,002	1,002	1,002	1,002	1,002
Controls								
Household	N	Y	N	Y	N	Y	N	Y
Local	N	Y	N	Y	N	Y	N	Y
Industry	N	Y	N	Y	N	Y	N	Y
Demographic	N	Y	N	Y	N	Y	N	Y
County FE	N	Y	N	Y	N	Y	N	Y
Year FE	N	Y	N	Y	N	Y	N	N
Standard Errors								
County Clusters	Y	Y	Y	Y	Y	Y	Y	Y
F-statistic	262.38	62.91	284.59	59.16	291.55	36.30	177.69	1.37
Adjusted R-squared	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.01

Notes: The table reports estimates of Equation (1) with household controls, county business cycle controls, county industry composition controls, county demographic controls, and county and year fixed effects. Columns (1) and (2) use the baseline Bartik instrument described in the text. Columns (3) and (4) use a restricted version of the Bartik instrument that only exploits the housing age characteristic. Columns (5) and (6) use an extended version of the Bartik instrument that adds continuous measures of floor size and lot size as housing characteristics. Column (7) uses a version of the Bartik instrument that makes use of national, rather than regional, variation in housing characteristic prices.
Sources: BLS, CBP, FHFA, IRS, Nielsen, Zillow, ZTRAX.

Table C.12: Consumption Response to House Prices Using Housing Characteristic Share-Year Dummy Interaction Instrument

	Real annual non-durable household consumption growth							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta p_{county,t}$	0.119***	0.032***	0.165***	0.067***	0.055***	0.065***	0.090**	0.088**
	(0.008)	(0.009)	(0.011)	(0.020)	(0.021)	(0.020)	(0.037)	(0.037)
Method	OLS	OLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
Observations								
Total	289,665	289,665	289,665	289,665	289,665	289,665	289,665	289,665
Households	64,898	64,898	64,898	64,898	64,898	64,898	64,898	64,898
Counties	1,002	1,002	1,002	1,002	1,002	1,002	1,002	1,002
Controls								
Household	N	Y	N	Y	N	N	N	Y
Local	N	N	N	N	Y	N	N	Y
Industry	N	N	N	N	N	Y	N	Y
Demographic	N	N	N	N	N	N	Y	Y
County FE	N	Y	N	Y	Y	Y	Y	Y
Year FE	N	Y	N	Y	Y	Y	Y	Y
Standard Errors								
County Clusters	Y	Y	Y	Y	Y	Y	Y	Y
F-statistic	–	–	201.59	26.41	34.97	31.63	8.68	9.68
Adjusted R-squared	0.00	0.01	0.00	0.01	0.01	0.01	0.01	0.01

Notes: The table reports estimates of Equation (1) using as an instrument the interaction between the housing characteristic shares (see Section 4.1) with year dummies. The specifications include: household controls, county business cycle controls, county industry composition controls, county demographic controls, and county and year fixed effects. Household controls come from the Nielsen Consumer Panel, including: real household income growth, a quadratic in age, the change in household size, the presence of children, marital status, race, Hispanic or Latino origin, occupation, and education. Local business cycle controls include: county unemployment growth from the BLS and real per capita income from the IRS. Local industry composition controls include: the employment share of construction, manufacturing, retail trade, and finance/real estate/insurance (FIRE) from the CBP. Local demographic controls include population shares of: black, Hispanic, foreign-born, those with at least some college education, homeowners, median age, household size, mean travel time to work, and employment shares in construction, manufacturing, retail trade, and FIRE. Each of these local demographic variables are interacted with year-dummy variables as suggested by Goldsmith-Pinkham, Sorkin and Swift (2020). Standard errors and F-statistics for 2SLS models are clustered at the county level.
Sources: BLS, CBP, FHFA, IRS, Nielsen, Zillow, ZTRAX.

D Estimation of Rotemberg Weights

We follow the suggestion from Goldsmith-Pinkham, Sorkin and Swift (2020) that shift-share instruments can be recast over-identified GMM estimators 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). Together with their Rotemberg weights, the local shares denote their contribution to the overall Bartik-like estimates. To see this, recall that a simplification of our two-stage least squares estimator is summarized by:

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

where $B_{l,t}$ denotes our Bartik-like instrument and $\widehat{p_{l,t}}$ denotes the predicted values obtained from the instrument on housing price growth. Suppose only one household is observed in each location, that there is only one time period, and that the exclusion restriction holds.

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-like instruments is $B = \Lambda Q$. The IV estimator of the consumption elasticity using the Bartik-like instrument has the familiar form:

$$\beta^{bartik} = \frac{B'C}{B'P} = \frac{Q'\Lambda' C}{Q'\Lambda' P} \tag{1}$$

Following Goldsmith-Pinkham, Sorkin and Swift (2020), the Bartik-like estimate can then be decomposed into the just-identified estimates β_k^{bartik} and the associated Rotemberg weights α_k. Then the IV 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, Sorkin and Swift (2020) notes that individual Rotemberg weights α_k may be negative, which means that the over-identified IV estimate using the full Bartik-like 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-like 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-like instrument for prediction.

Table D.13: 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	0.105	0.06	0.032	-0.321	0.284	0.338
α_k	–	0.001	0	0	0.001	0.274

Panel B: Correlations

	α_k	Δq_k	β_k	var(λ_k)
α_k	1			
Δq_k	-0.394	1		
β_k	-0.004	-0.001	1	
var(λ_k)	-0.048	0.024	0.027	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
Decade to 2005	2008	West	0.1	0.04	-1.74	0.19
Decade to 1989	2008	West	0.1	0.04	-1.74	0.1
Decade to 1999	2008	West	0.09	0.04	-1.68	0.18
Decade to 1979	2008	West	0.06	0.03	-1.68	0.13
Decade to 1999	2009	West	0.06	0.03	-0.8	0.01
Decade to 2005	2009	West	0.05	0.02	-0.82	-0.01
Decade to 1999	2013	West	0.04	0.02	0.95	0
Decade to 2005	2008	South	-0.04	0.02	-0.77	0.03
Decade to 2005	2013	West	0.04	0.02	0.97	-0.02
Decade to 1989	2009	West	0.04	0.02	-0.84	-0.02
Decade to 1959	2008	West	0.04	0.02	-1.73	0.14
Decade to 1989	2013	West	0.04	0.02	0.97	0.09
Decade to 1969	2008	West	0.03	0.02	-1.68	0.13
Decade to 1999	2008	South	-0.03	0.01	-0.74	0.11
Decade to 1979	2009	West	0.03	0.01	-0.79	0.03
Decade to 1979	2013	West	0.03	0.01	0.95	0.07
Decade to 1999	2014	West	0.02	0.01	0.6	-0.03
Decade to 2005	2014	West	0.02	0.01	0.61	-0.1
Bathrooms: 2	2013	South	-0.02	0.01	0.26	0.07
Decade to 1989	2014	West	0.02	0.01	0.61	-0.1

Notes: Panel A reports summary statistics for the just-identified estimates β_k and Rotemberg weights α_k. Panel B reports correlations between these variables, housing quality price growth rates Δq_k and the cross-sectional variance of the local housing shares var(λ_k). Panel C reports the top 20 housing characteristics sorted by share of absolute Rotemberg weight associated with the just-identified estimates.

Table D.13 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 characteristic shares. The Rotemberg weights and housing quality price growth are negatively correlated, which suggests that more weight is placed observations in which housing quality prices are declining, as they are during the housing bust. The Rotemberg weights and the variance of the local housing shares are only weakly negatively correlated. Additionally, the variance of the local housing shares is only weakly negatively correlated with quality price movements. This is important, as it shows that the identifying variation in the instrument contained in the housing shares is not tied to the potentially endogenous time-series variation produced by the housing quality prices.

Panel C reports the components of the instrument with the largest share of absolute Rotemberg weights, decomposed into variation due to the house characteristics shares, region, and year. Strikingly, virtually all of the Rotemberg weight is associated with the Western region, with housing age characteristics, and is largely concentrated in the bust years of 2008 and 2009, but also the recovery years of 2013 and 2014. We graphically illustrate these results in Figure **??** by overlaying the evolution of national house prices over this period. We see 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). Moreover, much of the variation is associated with price fluctuations occurring in the West of the US, which is perhaps unsurprising given that states such as Arizona, California, and Nevada had some of the largest house price fluctuations in the entire country during this period.

References

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales. 2019. "Shift-share designs: Theory and inference." *The Quarterly Journal of Economics*, 134(4): 1949–2010.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik instruments: What, when, why, and how." *American Economic Review*, 110(8): 2586–2624.

Gyourko, Joseph, Albert Saiz, and Anita Summers. 2008. "A new measure of the local regulatory environment for housing markets: The Wharton Residential Land Use Regulatory Index." *Urban Studies*, 45(3): 693–729.

Rotemberg, Julio. 1983. "Instrument variable estimation of misspecified models."

Saiz, Albert. 2010. "The geographic determinants of housing supply." *The Quarterly Journal of Economics*, 125(3): 1253–1296.

Stroebel, Johannes, and Joseph Vavra. 2019. "House prices, local demand and retail prices." *Journal of Political Economy*, 127(3): 1391–1436.