February 2024

Dynamism Diminished: The Role of Housing Markets and Credit Conditions

Steven J. Davis* and John Haltiwanger#

* Thomas W. and Susan B. Ford Senior Fellow at the Hoover Institution, Senior Fellow at the Stanford Institute for Economic Policy Research, William H. Abbot Distinguished Service Professor Emeritus at the University of Chicago Booth School of Business, and Research Associate with the National Bureau of Economic Research. Webpage: https://stevenjdavis.com.

Dudley and Louisa Dillard Professor of Economics and Distinguished University Professor of Economics at the University of Maryland, College Park, and Research Associate with the National Bureau of Economic Research. Webpage: http://econweb.umd.edu/~haltiwan/.

Appendices Not Intended for Publication

Appendix A: Measurement of Young Firm Dynamics in the BDS

The Business Dynamic Statistics (BDS) reports tabulations from the Census Longitudinal Business Database (LBD). The LBD is a longitudinal establishment-level database with establishment and firm-level characteristics. Firms are defined based on operational control. As described in section II, firm age is based on the age of the oldest establishment when a new legal entity originates. Establishment-level net employment growth rates underlying the BDS tabulations use the Davis, Haltiwanger and Schuh (1996) (DHS) growth rate measure:

$$g_{et} = \left(\frac{E_{et} - E_{et-1}}{X_{et}}\right), \quad X_{et} = 0.5 * (E_{et} + E_{et-1}) \tag{A.1}$$

where *e* indexes establishments and *t* indexes years. The DHS growth rate measure is a 2nd order approximation of the log first difference, is bounded between -2 and 2, and accommodates zeros in *t* (exit) or *t-1* (entry). The employment at the establishment-level in the LBD in year *t* is the number of employees of workers on the payroll for the payroll period including March 12th. As such the net employment growth rates (and all change measures in the LBD and BDS) represent changes from March in *t-1* to March of *t*.

The net employment growth rate for establishments classified into a cell *S* in *t* (e.g., a firm age and state cell) is given by:

$$g_{st} = \sum_{e \in s} \frac{X_{est}}{X_{st}} g_{est} \tag{A.2}$$

where *S* is the characteristics of the establishment in year t. The BDS provides net employment growth rate statistics as well as the decomposition into job creation, job destruction (by continuing, entering and exiting establishments) by a wide range of cells *S* defined by industry, firm age, firm size, establishment age, establishment size, and geographic cells defined by state and MSA. The BDS also reports these changes in terms of levels as well as the levels of employment and number of firms in each of classification cells.

For any given classification into cells of type *S*, the aggregate net employment growth is defined as the employment-weighted average of the cell based growth rates:

$$g_t = \sum_{s} \frac{X_{st}}{X_t} g_{st} \tag{A.3}$$

Relating the above measurement concepts to the measures from the BDS used in the paper, Figures B3 and B5 exploit the BDS net employment growth rate statistics defined by firm age (specifically, we use broad firm age categories as described in section II). The measures used in these figures capture within firm age group net employment growth rates. While instructive, such within firm age group net growth rates don't permit a characterization of the changing composition of employment by firm age (and likewise the changing composition of firms by firm age). For the latter, we use the share of young firm employment and the share of young firms as described in the main text. These can be directly measured from the BDS since the number of employees and firms are reported for all classifications in the BDS. Section III includes discussion of how the changing employment by firm age is related to net change within firm age groups and the changing composition.

Firm age is censored in the BDS given that firm and establishment age cannot be determined for establishments that exist in 1976 (the first year of the LBD). This implies that in each year subsequent to 1976 more firm age categories can be defined. We use two broad age categories in most of the analysis: young (firm age 0-5) and mature (6+). Firm age 0 in a given year reflects a parent firm that has all new establishments in that reference year. Firm ages naturally for subsequent ages.

Much of our statistical analysis commences in 1992 or later where it is straightforward to compute directly directly the employment of young firms (less than six years old), log changes in the young-firm employment share and net differentials for young minus mature firms. For descriptive statistics (e.g., Figure 1) and some counterfactual exercises we commence the analysis starting in 1981. Given our definition of young (less than six years old), the young-firm share is readily computable in 1982 and the log change in 1983. Given the focus on cyclical episodes in our analysis, it is advantageous to include statistics starting in 1981. In 1980, the BDS yields the employment of firms less than four years old directly, but to measure the employment of firms less than six years old in 1980 we need an estimate of employment at firms age=4 and age=5 in 1980. We impute the latter in 1980 using the product of the share of employment of age=4 year old firms in 1982 and total employment in 1980. This imputation is feasible at the national, state and MSA levels of aggregation. Similar remarks apply to calculations for 1981. We note that these imputations are not critical for our results as the regression analysis commences in 1992 and later years. Moreover, these imputations have no impact on the depiction of the summary statistics from 1983 forward.

Appendix B: Supplemental Figures and Tables

Figure B1. Firm Startup and Exit Rates, 1981 to 2018

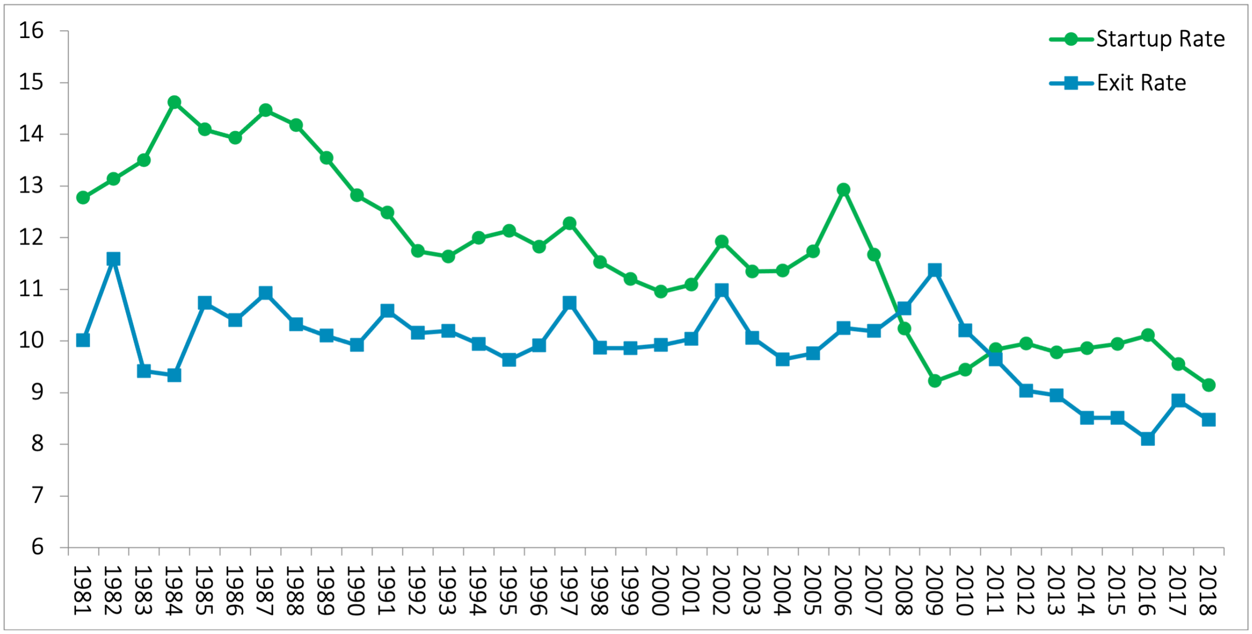

Figure B2. Share of Firms that are Young (<6 years old), 1981 to 2018

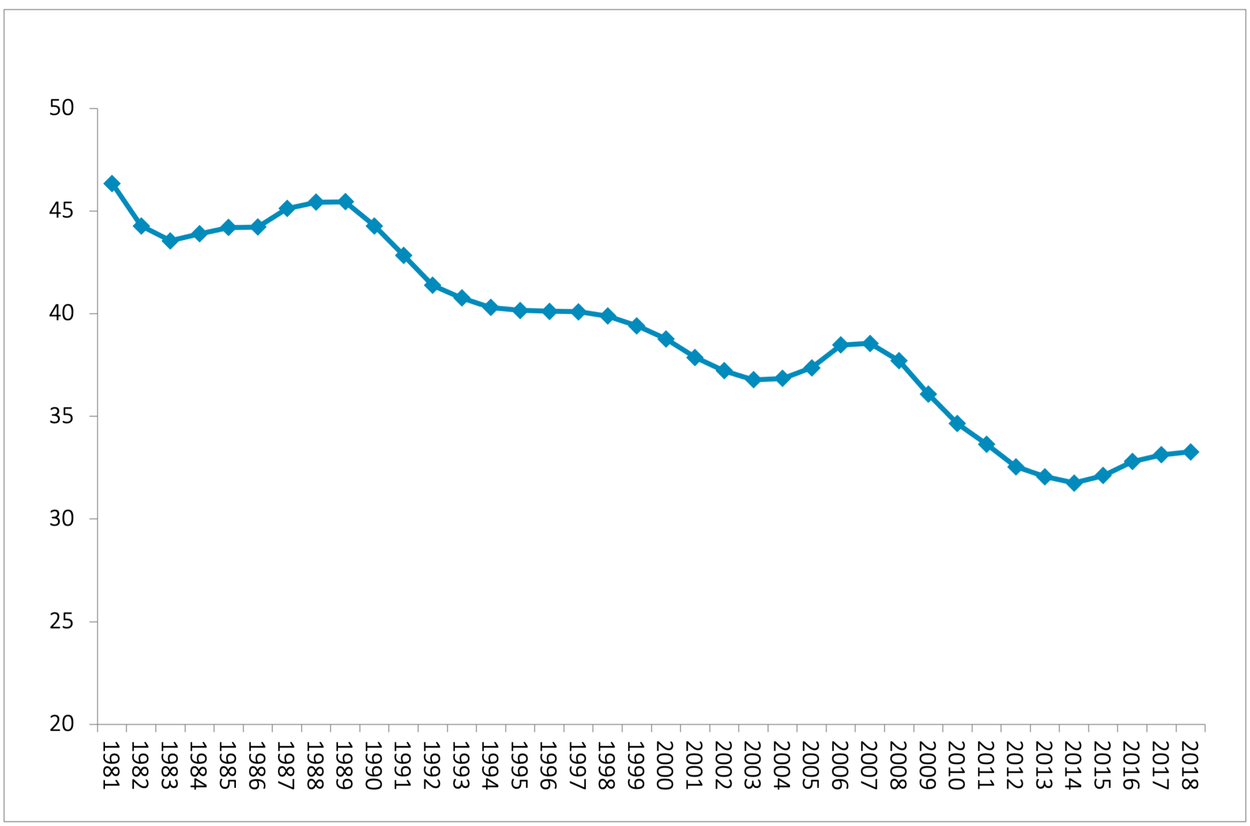

Figure B3. Annual Net Employment Growth Rates for Young and Mature Firms, 1981-2018

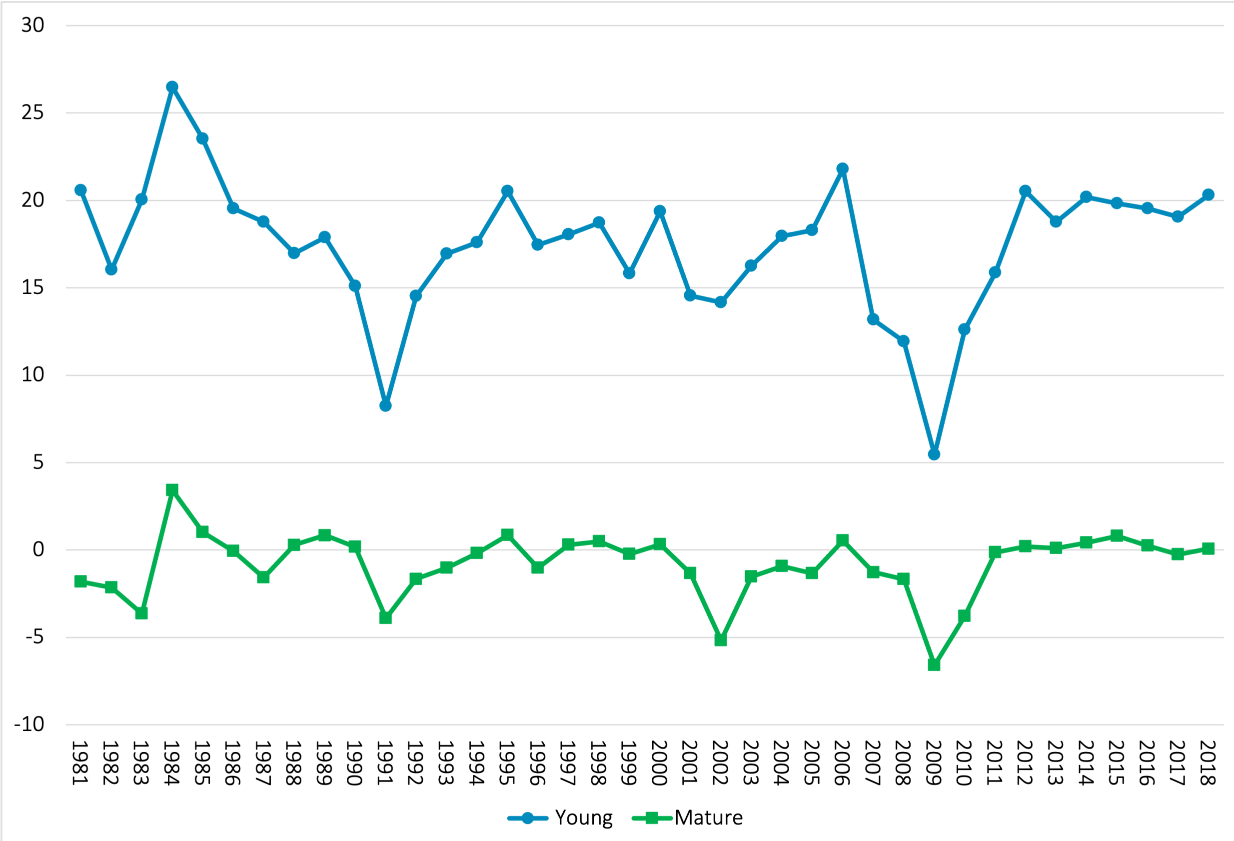

Notes: For each age group, the figure shows the employment-weighted DHS net growth rate from March of the previous year to March of the year reported on the horizontal scale. Net growth is inclusive of entry and exit of establishments.

Figure B4. Cyclicality of Log Changes in the Young-Firm Share of Firms

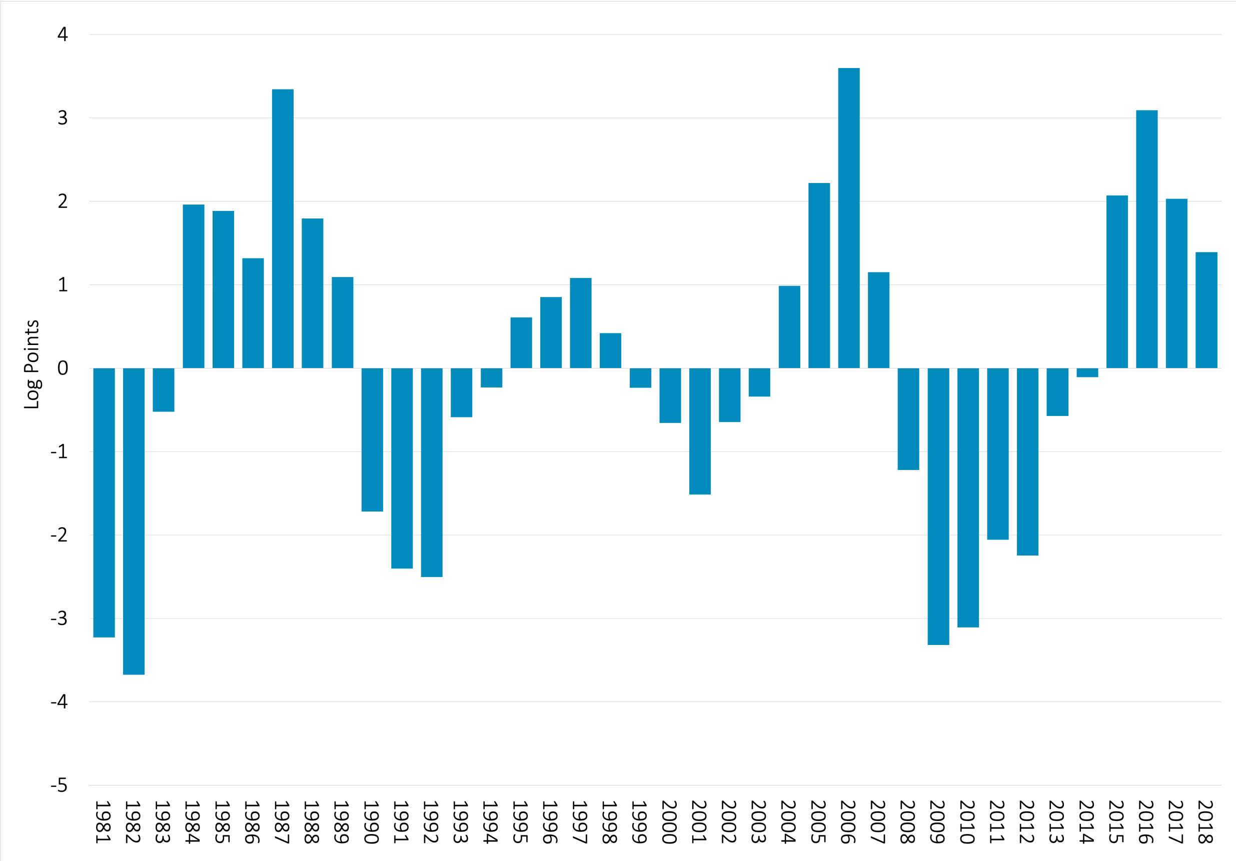

Notes: Each bar shows the annual average log change in the share of private sector firms that are young, deviated about the sample mean of minus 1.1 log change. See notes to Figures 1 and 2 for additional information.

Figure B5. Cyclicality of the Net Growth Rate Differential Between Young and Mature Firms

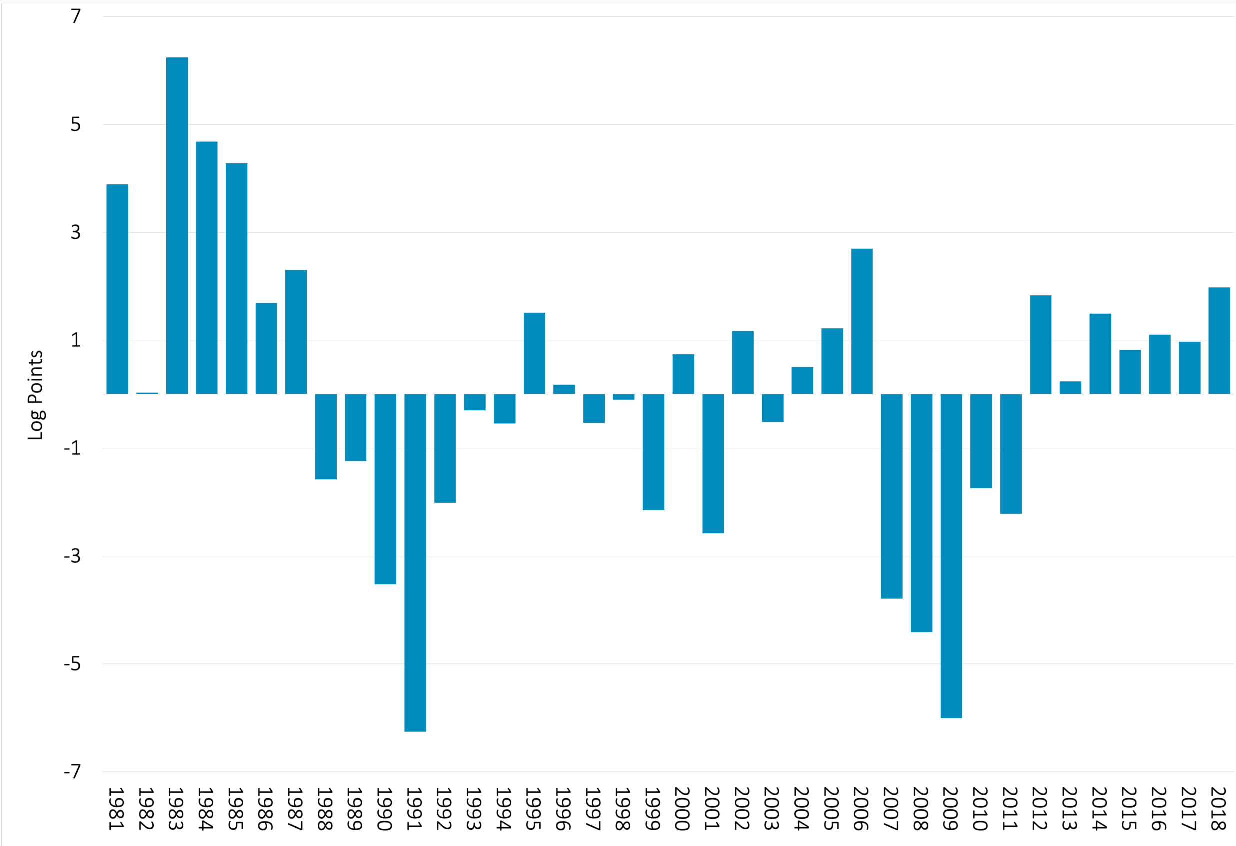

Notes: Each bar shows the annual average net employment growth differential between young and mature firms, deviated about the sample mean net differential of 18.2 log change. See notes to Figures 1 and 2 for additional information.

Figure B6A: Relationship Between Log Difference in Young-Firm Employment Share and the Change in the Unemployment Rate, State by Year Cells, 1981-2018

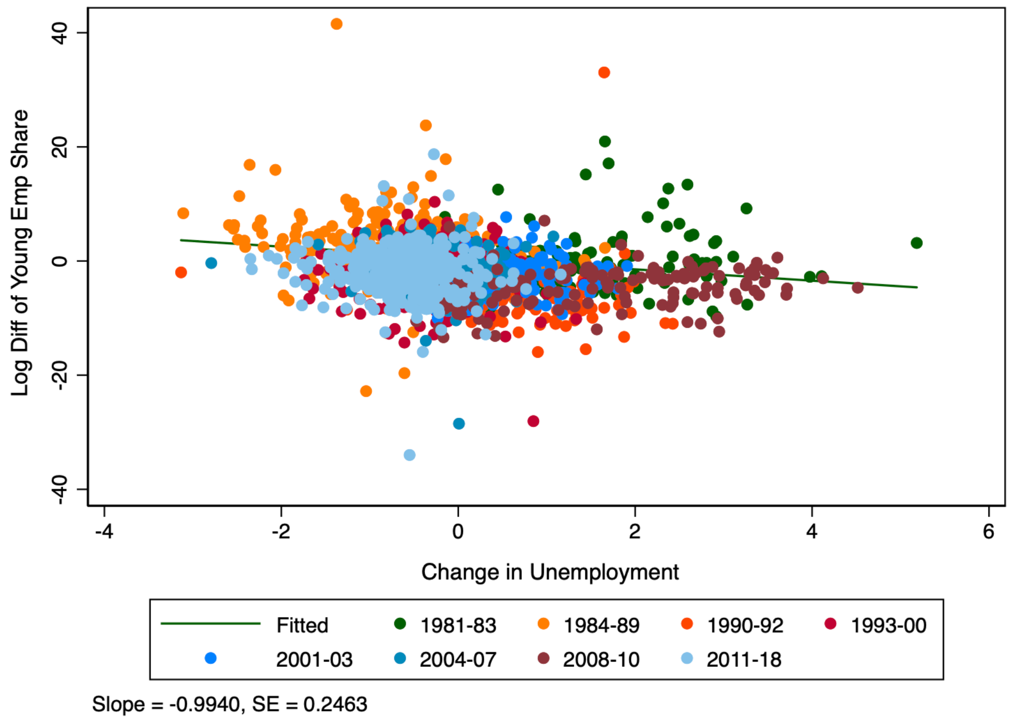

Figure B6B: Relationship Between Log Difference in Young Employment Share and Growth Rate of Real Housing Price, State by Year Cells, 1981-2018

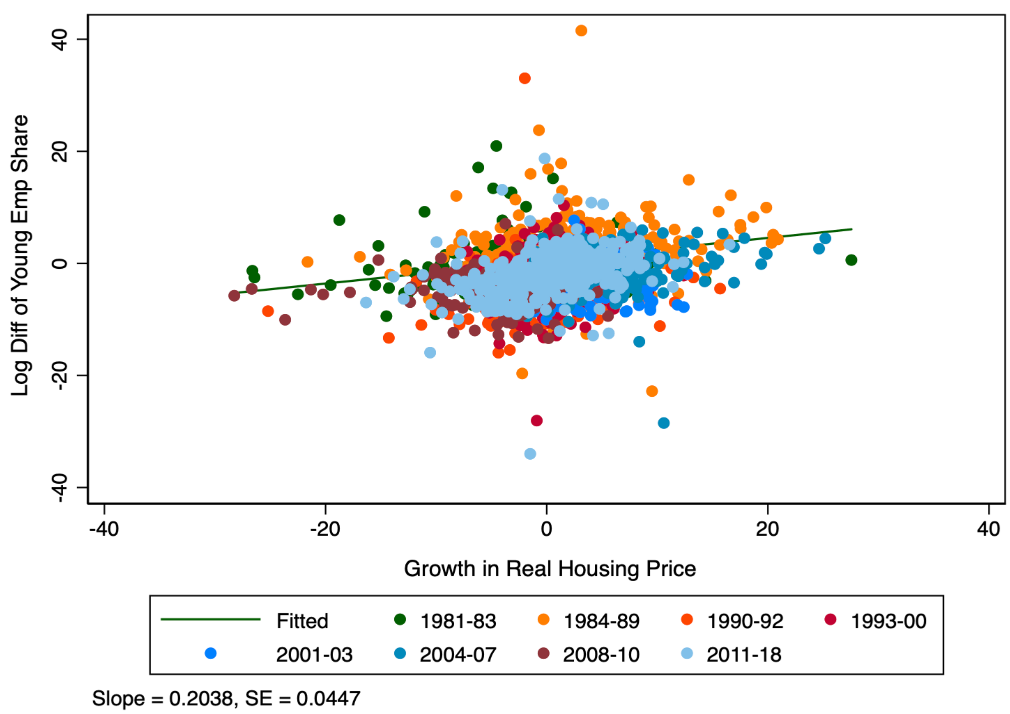

Notes: In Panel A, the scale is in log points on the vertical axis and percentage points on the horizontal axis. In Panel B, the scale is log points on both axis. See notes to Figure 2 for the timing convention of reported intervals.

Figure B7. Relationship between Real Housing Price Growth in Boom (1998-2006) and Bust (2007-10) periods, MSA-Level Data

A. **Boom**

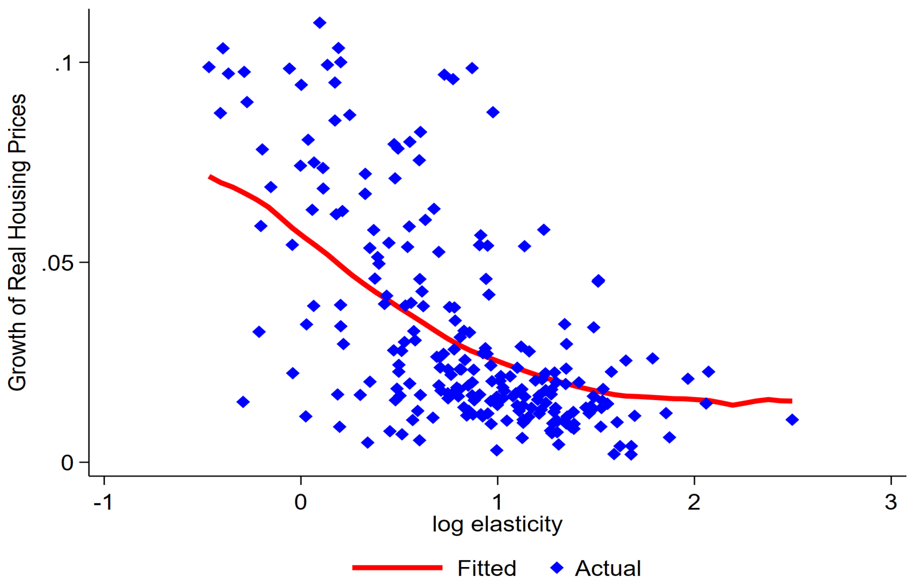

B. **Bust**

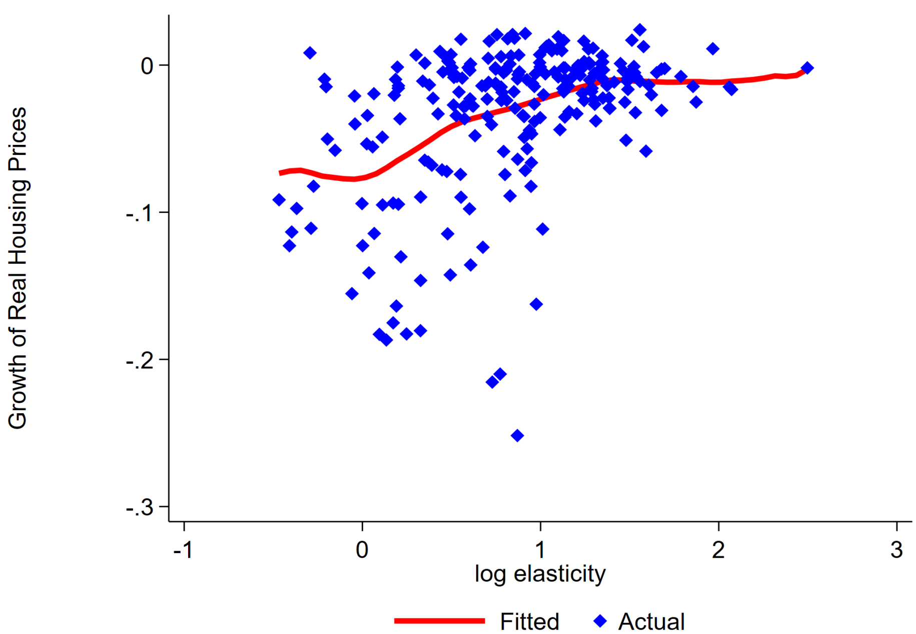

Notes: The vertical scale is the average annual MSA-level log change during the indicated period. The fitted line is a locally smoothed polynomial.

Figure B8: Relationship Between Log Difference in Young-Firm Employment Share and Growth Rate of Real Housing Prices, MSA-level by Year (Bin Scatter)

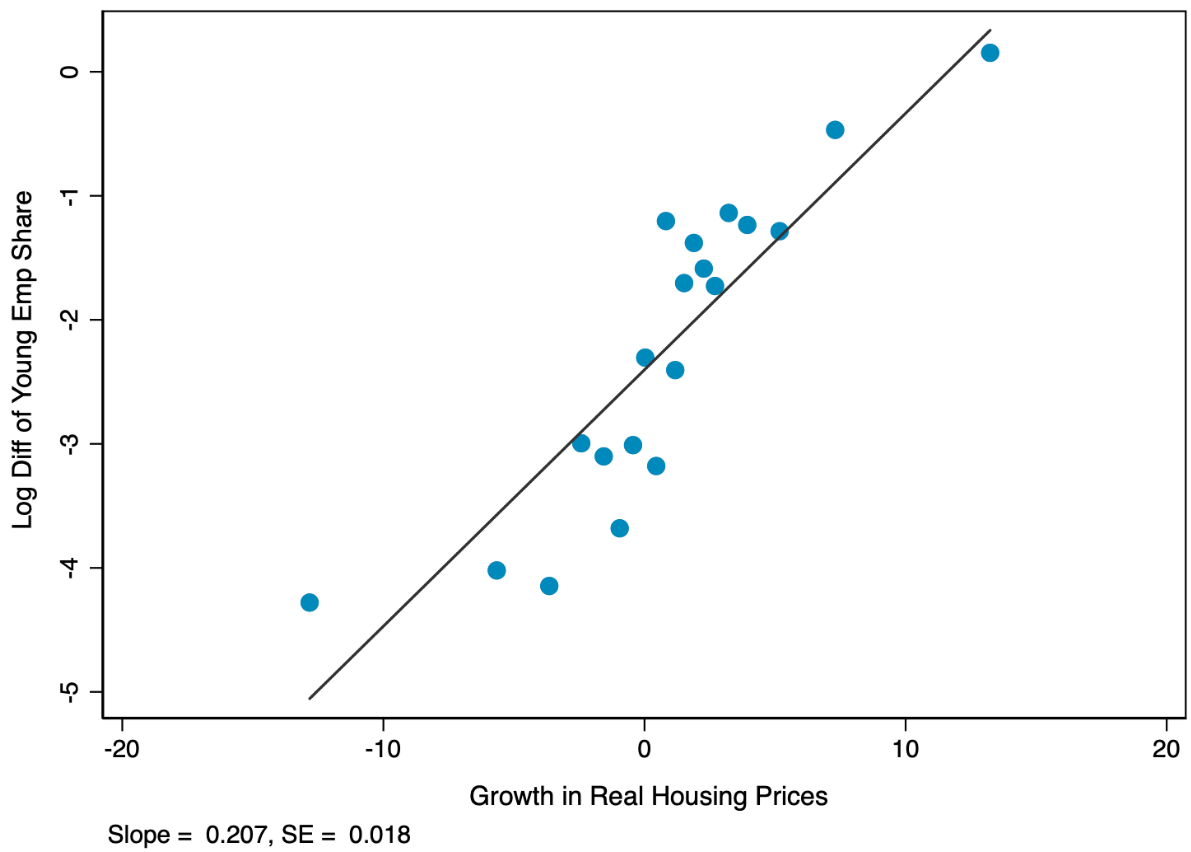

Notes: This chart shows a bin scatter of the raw data displayed in Figure 4.

Figure B9 Relationship between Saiz and Price Sensitivity Housing Supply Elasticities

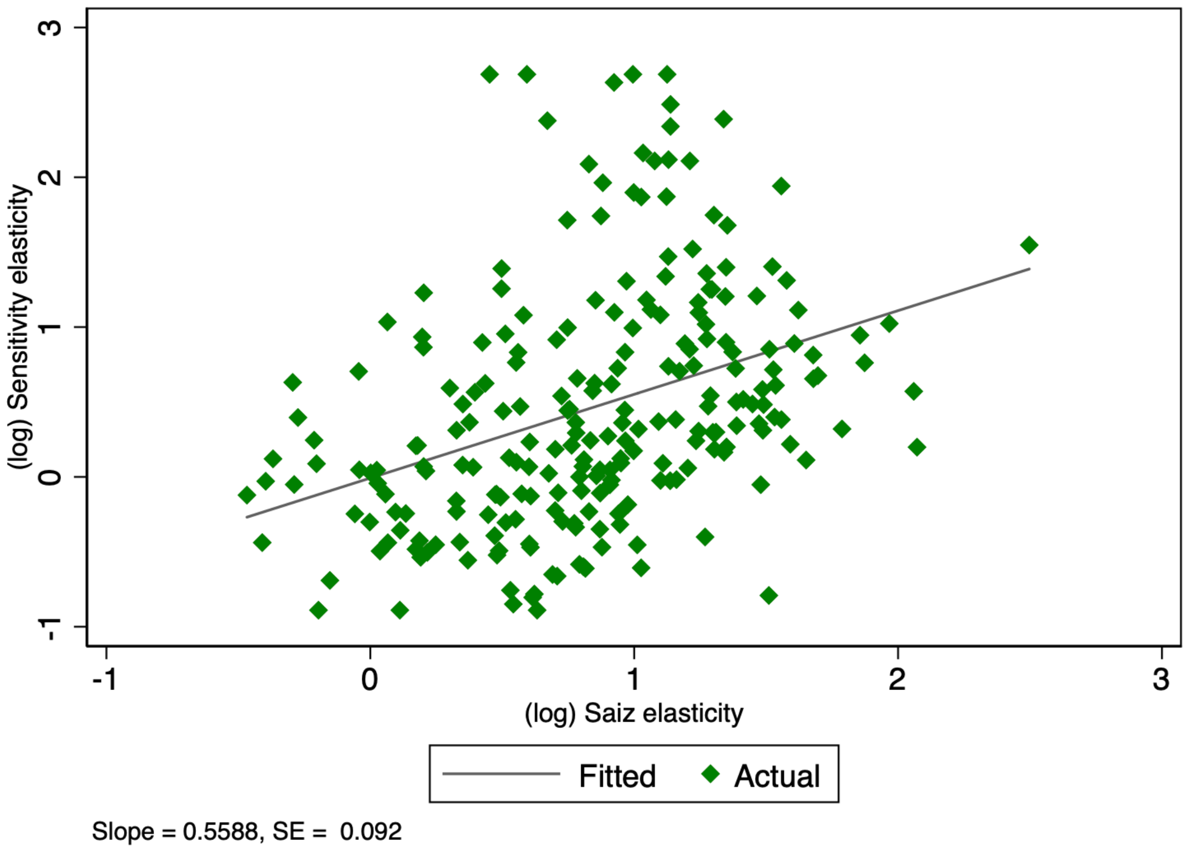

Figure B10: Relationship Between Log Difference in Young-Firm Employment Share and Growth Rate of Real Housing Prices, MSA-level data in Boom and Bust Periods

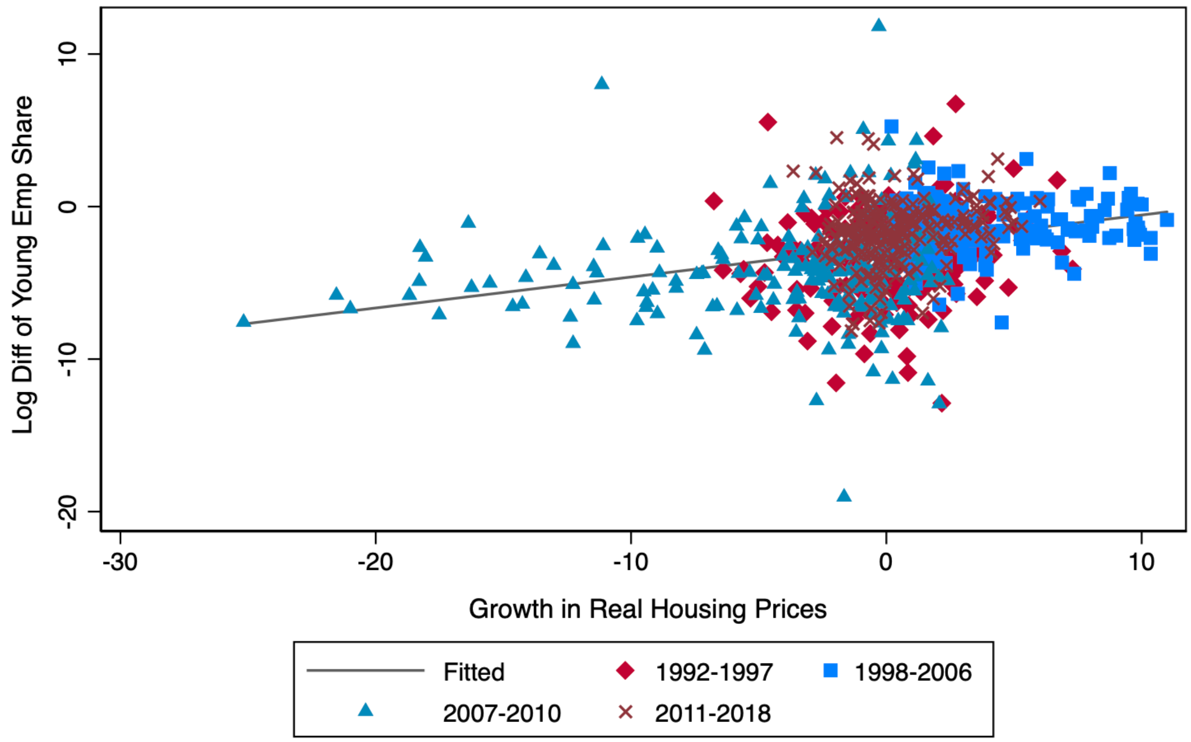

Slope = 0.203, SE = 0.021

Mean Growth Rate of Housing Prices:
Full Sample = 0.192, 1992-1997 = 0.495 1998-2006 = 3.281, 2007-2010 = -3.349, 2011-2018 = 0.344

Figure B11: Second-Stage Relationship Between the Log Change in the Local Young-Firm Employment Share and the IV-Predicted Real Growth Rate of Local Housing Prices

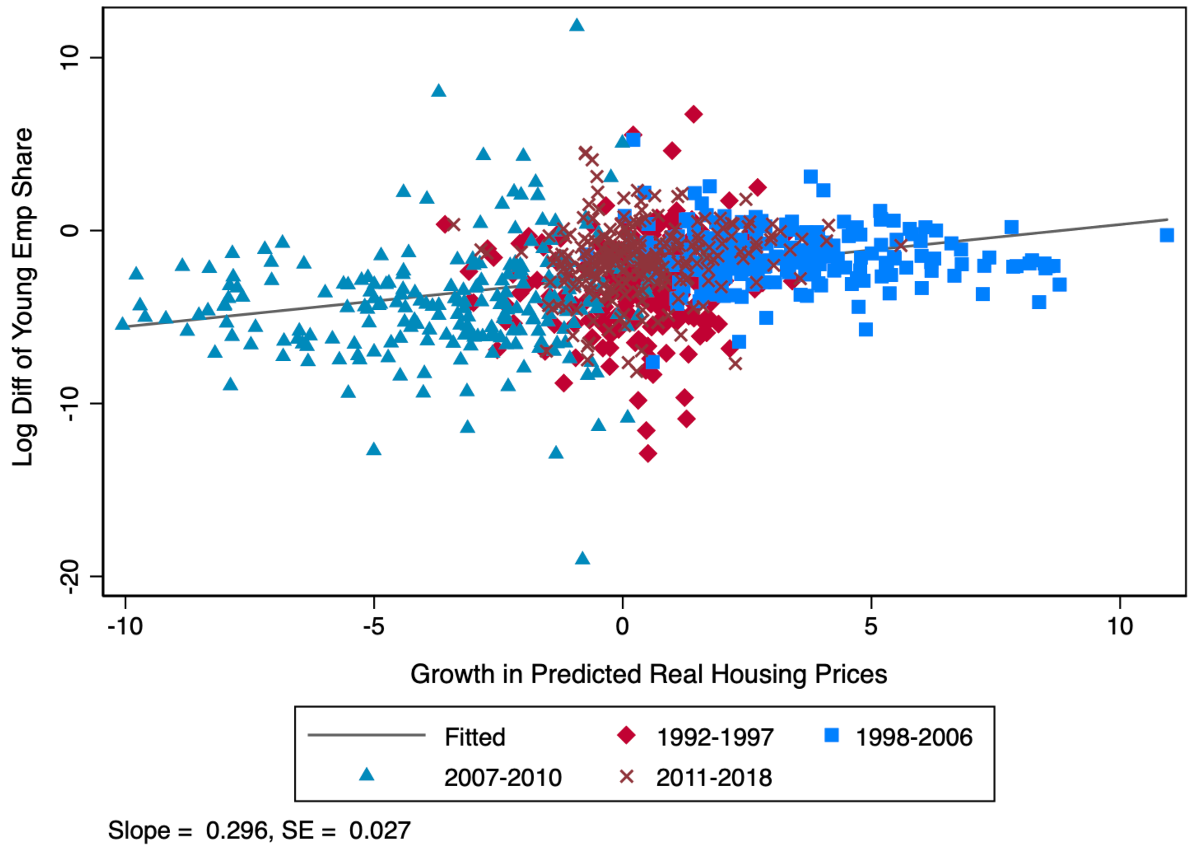

Notes: This figure reflects the estimated specification reported in Column (2) of Table 1. See Table 1 for the appropriately adjusted standard error. The analogous figure for Column (4) is very similar.

Not for Publication

Figure B12. Share of Employment at Small Firms, 1981-2018, U.S. Nonfarm Private Economy

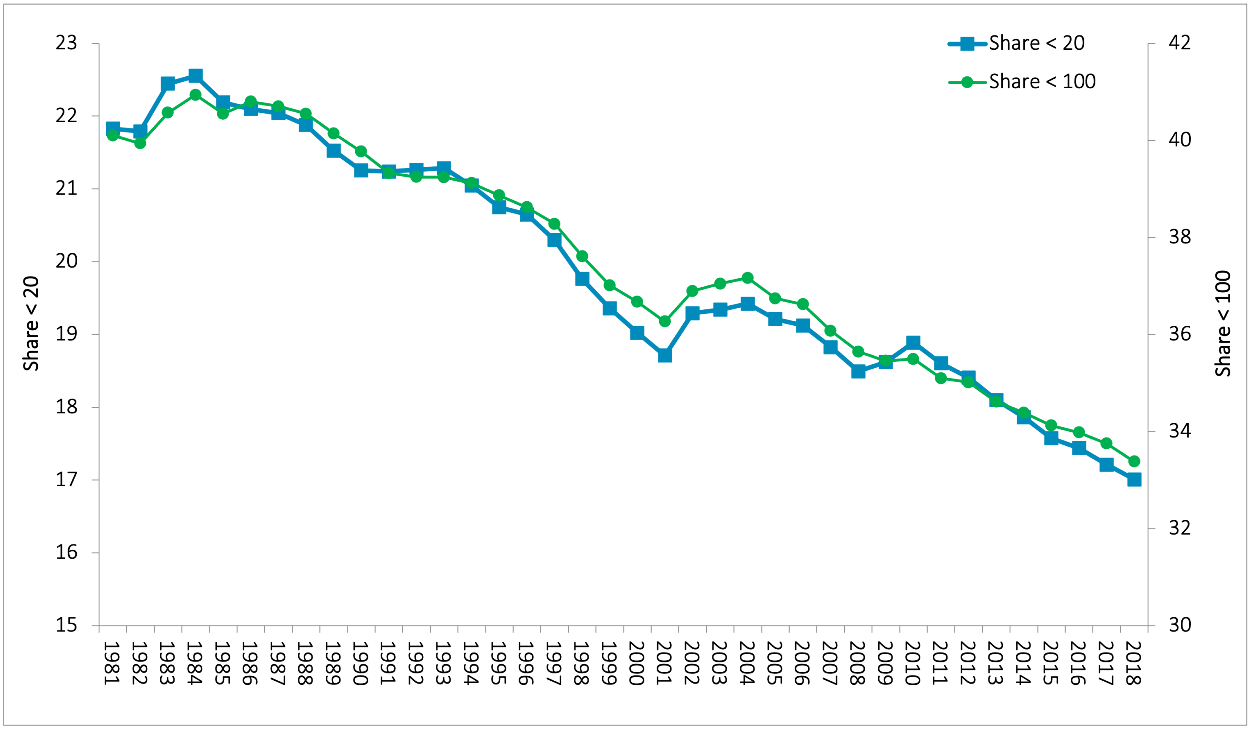

Figure B13. Log Differences in Employment Share of Small Firms by Cycle Episode

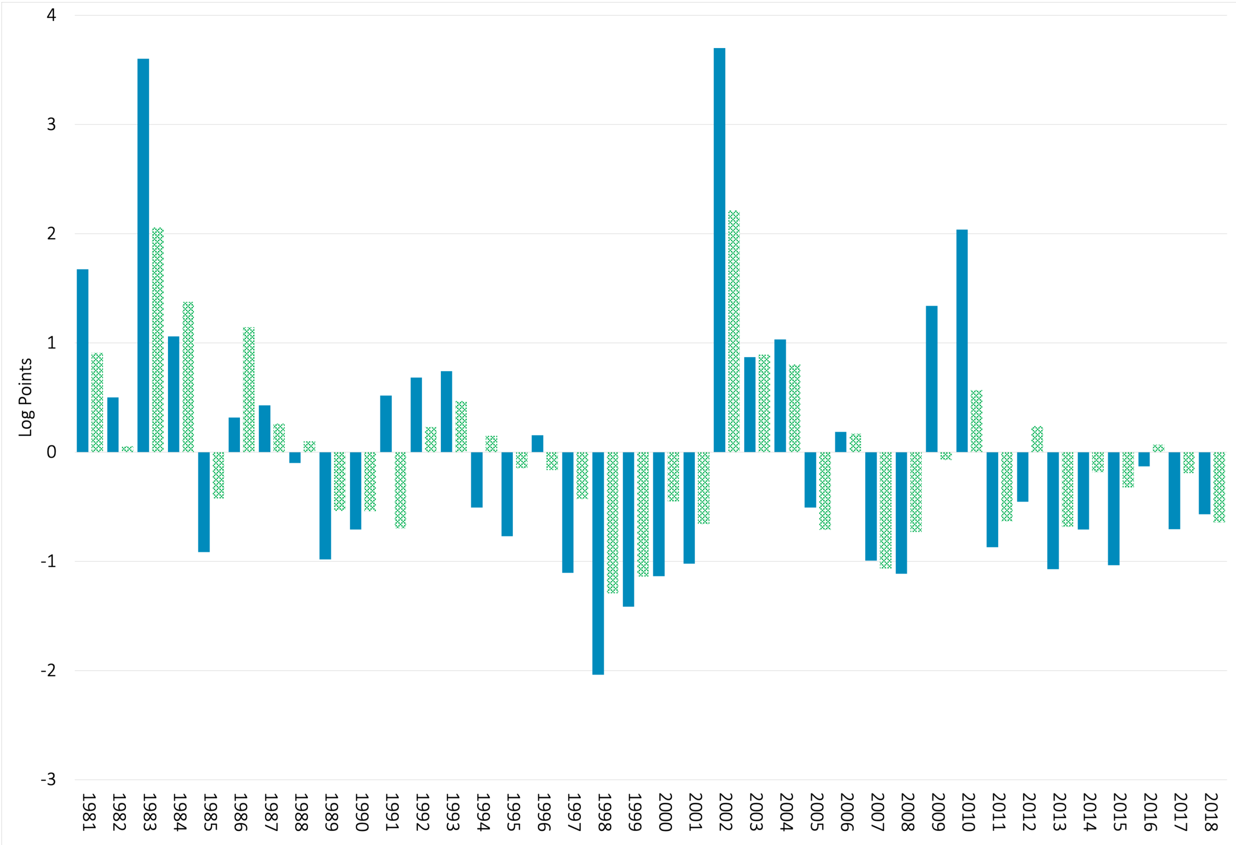

Notes: Sold (blue) Bar is Actual for small firms less than 20, Dotted (green) Bar is Actual for small firms less than 100. Annualized deviations from overall mean of minus 0.6 (less than 20) and 0.5 (less than 100) log change depicted.

Figure B14. Log Differences in Employment Share of Small Firms (Less than 20), Actual and Predicted

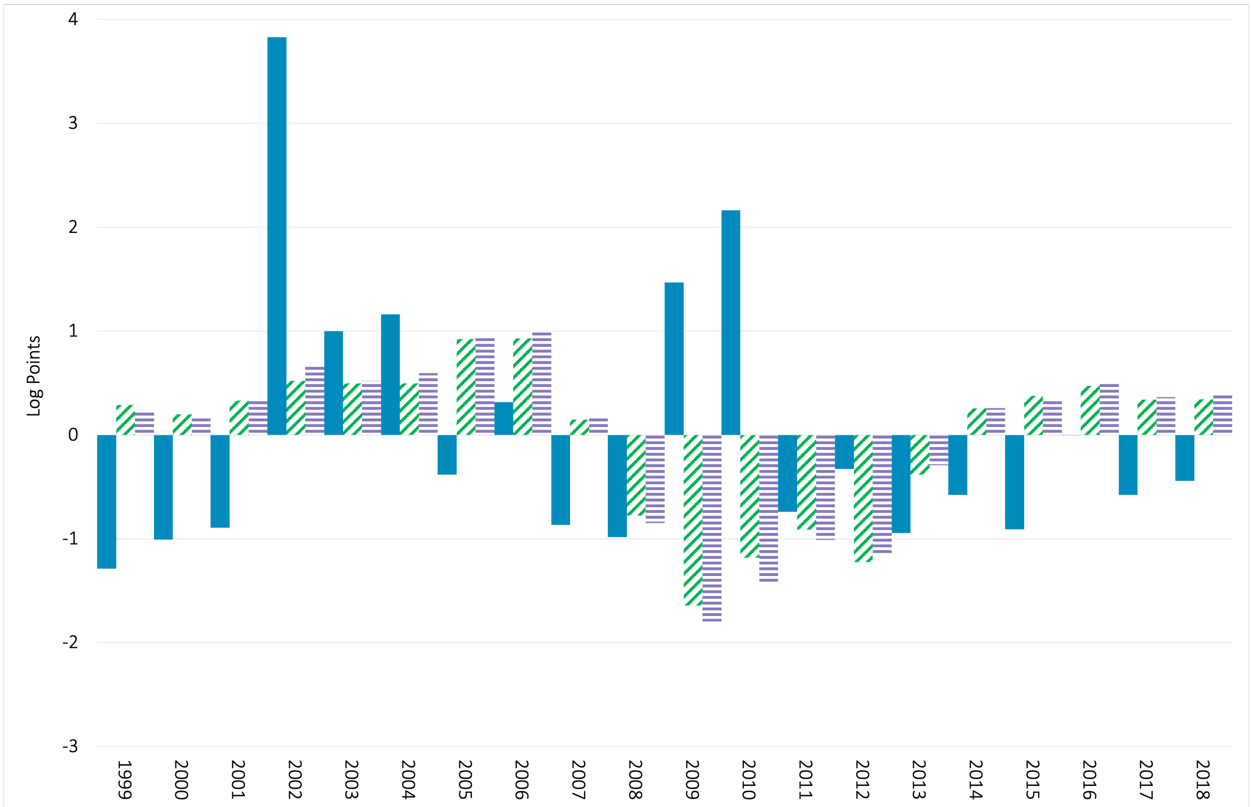

Notes: Sold (blue) Bar is Actual, Diagonal Striped (green) Bar is Counterfactual (Housing Prices only), Horizontal Striped (orange) Bar is (Housing Prices + Loan Supply). Using IV estimates from last column of Table B3. Annualized deviations from overall mean of minus 0.8 log change depicted.

Table B1. Coefficient Estimates on Cubic Polynomial of the Saiz Measure of Housing Supply Elasticities in the First Stage Specifications of the IV Regression Models

	Specification			
Coefficient on:	(2)	(2)	(6)	(6)
Log(Elasticity)*Chg_MSA_UR	1.541	1.245	1.581	1.323
	(0.570)	(0.650)	(0.708)	(0.798)
$(\text{Log(Elasticity)})^2$*Chg_MSA_UR	1.515	1.517	2.108	2.081
	(0.809)	(0.856)	(0.995)	(1.063)
$(\text{Log(Elasticity)})^3$*Chg_MSA_UR	-0.752	-0.691	-0.989	-0.927
	(0.317)	(0.313)	(0.388)	(0.390)
MSA Bartik Demand Control	No	Yes	No	Yes
MSA Population Growth Rate	No	Yes	No	Yes

Notes: Table entries are estimated coefficients on $CYC_{mt}Z'_m$ in the first-stage regressions (3) and (7), as indicated. In addition to the controls indicated in the bottom two rows of the table, the first-stage specifications include MSA effects, year effects, and local cycle controls (the MSA-level change in the unemployment rate). We cluster errors at the MSA level in computing standard errors.

Table B2. Young-Firm Employment Share Response to Local Housing Price Growth, Adapting the IV Approach of Mian and Sufi by Collapsing the MSA-level Data Stacked within Four Periods and Stacking It Over the Periods

Dependent Variable: Average annual log change in MSA young-firm employment share during the period

	(1) OLS	(2) IV	(3) OLS	(4) IV
Coefficient on log real housing price change (β)	0.100 (0.025)	0.138 (0.052)	0.098 (0.027)	0.141 (0.048)
F-Test for Excluded Instruments		23.1		20.4
Period Effects	Yes	Yes	Yes	Yes
MSA Effects	No	No	Yes	Yes
R^2	0.172	0.170	0.271	0.186
Observations	932	932	932	932

Notes: We estimate specifications (D.2) and (D.3) from Appendix D using MSA-level data for four periods: 1992-97, 1998-06, 2007-10, and 2011-18. To construct instruments for MSA-level changes in housing prices, we interact period effects with a cubic polynomial in the log of the Saiz measure of the MSA housing supply elasticity. Robust standard errors in parentheses.

Table B3. Young-Firm Employment Share Response to Local Housing Price Growth, Adapting the IV Approach of Mian and Sufi and Adding More Controls

Dependent Variable: Average annual log change in MSA young-firm employment share in the period

	(1) IV	(2) IV	(3) IV	(4) IV
Coefficient on log real housing price change (β)	0.141 (0.048)	0.147 (0.059)	0.147 (0.058)	0.131 (0.061)
F-Test for Excluded Instruments	20.4	17.4	17.4	15.3
Period & MSA Effects	Yes	Yes	Yes	Yes
MSA Unemployment Rate Change	No	Yes	Yes	Yes
MSA Bartik Demand Control	No	No	Yes	Yes
MSA Population Growth Rate	No	No	No	Yes
R^2	0.186	0.186	0.186	0.187

Notes: See notes to Table B2. There are 932 observations in each column.

Table B4. Testing the Age-Invariance Proposition: Industry-Level Log Employment Growth Response to Local House Price Growth, Annual Data from 1992 to 2018 (Omitting the Construction Sector)

Dependent Variable: Annual Log Employment Change at the MSA-Industry Level (GR_{jmt})

	(1) OLS	(2) OLS	(3) OLS	(4) IV
Change in Local Unemployment Rate (CYC_{mt})	-0.72 (0.100)			-0.540 (0.109)
Log Change in Local House Prices (HP_{mt})	0.092 (0.011)			0.215 (0.041)
Lagged Young-Firm Employment Share ($YoungSh_{jm,t-1}$)	0.039 (0.005)	0.041 (0.006)	0.023 (0.006)	0.038 (0.005)
Interaction Term ($HP_{mt} * YoungSh_{jm,t-1}$)	0.354 (0.058)	0.323 (0.072)	0.162 (0.083)	0.548 (0.099)
R-squared Value	0.069	0.141	0.173	0.067
MSA Fixed Effects	Yes	No	No	Yes
Year Fixed Effects	Yes	No	No	Yes
Industry Fixed Effects	Yes	Yes	No	Yes
MSA-by-Year Fixed Effects	No	Yes	Yes	No
Industry-by-Year Fixed Effects	No	No	Yes	No

Notes: The sample, which covers 233 MSAs and 13 industries, contains 81,319 observations at the industry-MSA level from 1992 to 2018. It replicates Table 6 excluding the construction industry, 23. We drop cells with no employment. See text for list of industries. Column (4) instruments for HP_{mt} and the interaction term using the IV approach used in Tables 1 and 2. Standard errors in parentheses, clustered at MSA level.

Table B5. Quantifying the Departures from Age Invariance (Omitting the Construction Sector)

A. Dispersion in Young-Firm Employment Shares Across Local Industries

Industry-MSA Young-Firm Share	1992-2018	Boom Period	Bust Period
90th Percentile	1.442	1.441	1.378
10th Percentile	1.061	1.077	1.059
Standard Deviation	0.155	0.149	0.130
90-10 Difference	0.381	0.364	0.319

B. Dispersion in Annual MSA-Level Log House Price Changes

Log MSA House Price Change	1992-2018	Boom Period	Bust Period
90th Percentile	0.061	0.088	0.036
10th Percentile	-0.044	-0.002	-0.123
Standard Deviation	0.054	0.043	0.077
90-10 Difference	0.105	0.090	0.159

C. Calculating the Departures from Age Invariance, Using $\hat{c} = 0.548$

	Boom Period		Bust Period	
$HP(p)$ from Panel B	P90	P10	P90	P10
Average Annual Log Changes	0.088	-0.002	0.036	-0.123
$YoungSh^{90-10}$ from Panel A	0.364		0.319	
Response Differential Per Equation (9)				
Annual, Percentage Points	1.7	-0.0	0.6	-2.2
Cumulative, Percentage Points	7.2	0.3	1.9	-6.3

Notes: Panel C implements the calculation expressed in equation (9) in the main text. $\hat{c}$ is the coefficient for the interaction term from column 4 of Table B4. Panels A and B report inputs to the calculation and related summary statistics. The Boom Period runs from 1998 to 2006, and the Bust Period runs from 2007 to 2010.

Table B6. Testing the Age-Invariance Proposition with a Dynamic Specification: Industry-Level Log Employment Growth Responses to Local House Price Growth, Annual Data, 1992 to 2018

Reporting only the coefficients on the current and lagged interaction terms.

Dependent Variable: Annual Log Employment Change at the MSA-Industry Level (GR_{jmt})

	(1) OLS	(2) OLS	(3) OLS	(4) IV
Interaction Term ($HP_{mt} * YoungSh_{jm,t-1}$)	0.892 (0.075)	0.872 (0.100)	0.603 (0.112)	1.379 (0.161)
Interaction Term ($HP_{mt-1} * YoungSh_{jm,t-1}$)	-0.319 (0.084)	-0.280 (0.101)	-0.293 (0.115)	-0.560 (0.159)
MSA Fixed Effects	Yes	No	No	Yes
Year Fixed Effects	Yes	No	No	Yes
Industry Fixed Effects	Yes	Yes	No	Yes
MSA-by-Year Fixed Effects	No	Yes	Yes	No
Industry-by-Year Fixed Effects	No	No	Yes	No

Notes: See notes to Table 6.

Appendix C: Details Related to the CRA Data and How We Use Them

CRA data provide bank-level information by local area and year on the volume of business loan originations to firms with less than $1 million in revenue. We deflate nominal loan volumes by the same-year GDP implicit price deflator from the Bureau of Economic Analysis to obtain real small business loan volumes. As in GMN (2015), we roll up the bank level data to the bank holding company, using data sources from the FDIC and Federal Reserve call reports. We also use data from the Federal Reserve Bank of Chicago that tracks mergers and acquisitions, so that for any pair of years *t-1* and *t* we assign a bank to its owner in year *t*.

We measure the growth rate of small business loan volume for a given bank holding company in a particular MSA – what we call g_{mjt} in equation (4) – using the symmetric growth rate measure in Davis, Haltiwanger and Schuh (1996). The DHS measure is equivalent to the log first difference up to a second-order Taylor Series approximation, but the DHS measure down weights outliers relative to log changes. As it turns out, DHS growth rates and log changes produce similar econometric results in Section IV.d.

One additional detail: We re-time the calendar-year SBL_{it} measure in equation (6) to align it with the March-to-March employment changes in the BDS. Specifically, in our regression analysis, the loan supply shock for MSA *m* and year *t* is $0.75 *$ Raw SBL_{it-1} + $0.25 *$ Raw SBL_{it}. The correlation between this re-timed measure for year *t* and the unadjusted measure for year *t-1* is 0.95 for the MSA-level data used to produce Table 3. Replacing the re-timed SBL measure for year *t* with the corresponding lagged Raw SBL measure yields very similar results.

Appendix D: Alternative Approaches to IV Estimation of House Price Effects

D.1 Using Housing Price-Sensitivity Measures to Construct Instruments

Recall that our main approach uses MSA-level housing supply elasticities from Saiz (2010) in constructing our instruments for MSA-level house price growth rates. As a robustness check, we follow Guren et al. (2021) and obtain housing supply elasticities as the inverse of MSA-level price-sensitivity estimates. Specifically, we fit the following regression specification by least squares to annual data from 1992 to 2018, pooled over MSAs and years:

$$HP_{mt} = \sum_m \delta_m I_m + \sum_s \delta_{tr} I_{tr} + \psi_m HP_{dt}^{(m)} + \theta_m CYC_{dt} + \chi_m CYC_{mt} + X'_{mt}\phi_m + \xi_{mt}, \qquad \text{(D.1)}$$

where *r* indexes Census regions, *d* indexes Census divisions, HP_{mt} is the housing price growth rate from year *t-1* to year *t* in MSA m, $HP_{dt}^{(m)}$ is the contemporaneous housing price growth rate in the Census division *d* that contains MSA *m,* and the other variables and subscripts are defined as in the main text. We compute $HP_{dt}^{(m)}$ using the leave-out method, i.e., excluding MSA *m.*

The key coefficients of interest in (D.1) are the ψ_m – the estimated sensitivity of housing prices in MSA *m* to contemporaneous changes in housing prices in the rest of the same Census division. The idea is that a similar-size housing demand shift in the relevant Census division produces a larger price response in MSA *m* when *m* has a low housing supply elasticity. Thus, we interpret ψ_m as inversely related to the local housing supply elasticity of MSA *m.* As Guren et al. stress, including controls in (D.1) helps address concerns about reverse causality and omitted variables. Thus, X_{mt} includes the Bartik-type and population growth controls described in the main text. (D.1) also includes MSA fixed effects, time-varying cyclical controls for the MSA *m* and Census division *d,* plus Census region-time fixed effects. There are four Census regions, indexed by *r,* each of which contains 2 or 3 Census divisions.

As Guren et al. (2021) discuss, ψ_m is a proxy for the inverse of the housing supply elasticity. Figure B9 shows how the inverse of our estimated ψ_m values correlate with the (log) Saiz elasticity measures. The two measures correlate strongly, but they also differ a good deal. The R-squared value for the regression line in Figure B9 is 0.13, the same as in Guren et al. We also find a similar pattern of discrepancies. For example, the inverse price-sensitivity estimates are lower than the Saiz elasticities in Detroit and Las Vegas.

D.2 An Alternative Approach that Is Closer to Mian and Sufi (2011)

Our main IV approach uses annual observations at the MSA level. We also implemented an alternative approach that more closely follows that of Mian and Sufi (2011). The alternative approach uses 932 observations on the *average* annual log changes in MSA-level data during four subperiods: a pre-boom period from 1992 to 1997, the housing market boom from 1998 to 2006, the bust from 2007 to 2010, and the post-bust period from 2011 to 2018. As seen in Figure B10, most MSAs experienced modest log price changes in the pre-boom and post-bust periods, strong price appreciation during the boom period, and sharp price declines during the bust.

We use these collapsed data to estimate the following statistical model:

$$Y_{ms} = \sum_s \lambda_s I_s + \sum_m \lambda_m I_m + \beta HP_{ms} + \varepsilon_{ms} \qquad \text{(D.2) (Second stage)}$$

$$HP_{ms} = \sum_m \delta_m I_m + \sum_s \delta_s I_s + \sum_s I_s Z_m' \gamma_s + \eta_{ms} \qquad \text{(D.3) (First stage)}$$

where Y_{ms} is the annual average log change in the young-firm employment share for MSA *m* during period *s*, HP_{ms} is the contemporaneous annual average log change in the MSA's house price index, I_s is a dummy for period *s,* I_m is dummy for MSA *m,* Z_m is a cubic polynomial in the Saiz housing supply elasticity, and λ_i and δ_i are coefficients on dummy variables. The chief parameter of interest is β, the response of the change in the local young-firm employment share to the local house price change.

To identify β we rely on the exclusion restrictions, $\mathrm{E}(I_s Z_m', \varepsilon_{ms}) = 0$, which says that $I_s Z_m'$ influences young-firm employment shares only through house price growth, conditional on period and MSA effects. In their papers, Mian and Sufi typically consider cross-sectional regressions in either the boom period or the bust period. Stacking the periods lets us control for MSA-specific trends, addressing concerns that these trends reflect other factors that happen to correlate with local housing supply elasticities, as argued by Davidoff (2016).

Table B2 reports regression results for specification (D.2) and (D.3) fit to the MSA-level data. We find a positive, statistically significant effect of local housing price growth on local young-firm activity shares. According to the IV estimates in column (3), which controls for common period effects and MSA fixed effects, an increase in local real housing prices of 10 log points per year yields a gain of 1.4 log points per year in the local young-firm employment share. IV estimates for β are somewhat larger than the corresponding OLS estimates, in line with the view that measurement error in the local housing price indices produces some attenuation under OLS. F-tests show a very strong first stage, with test statistics well above 10. As seen in Figure

B11, the IV-estimated relationship in Table B2 is similar in the boom and bust periods, and it is not driven by a few outliers.

Following the same logic as in the main text, we also consider specifications with additional local controls. See Table B3. We again find a strong positive impact of local housing price growth on young-firm activity shares for the IV results. Conditioning on all three additional controls in Column (4), an increase in local real housing prices of 10 log points per year raises the local young-firm employment share by 1.31 log points per year.[1] This effect is statistically significant at the 5 percent level, and an F-test again provides strong evidence against the hypothesis of weak instruments.

D.3 A Dynamic Extension to Specification (8)

We now extend (8) to include the lagged main effect for local house price changes and its interaction with the young-firm employment share in the local industry. This dynamic specification allows the effects of interest to unfold over time, possibly attenuating or amplifying the contemporaneous effect captured by coefficient, *c,* in (8). Accordingly, Appendix Table B6 reports the coefficients on the contemporaneous and lagged interaction terms in the dynamic extension to (8). The other coefficients are similar to the ones reported in Table 6.

The dynamic extension yields two additional results: First, the contemporaneous local-industry response to an increase in local house prices now rises more steeply with the local industry's young-firm share. For example, the coefficient on the contemporaneous interaction term is 1.38 in column (4) of Table B6, as compared to 0.95 in Table 6. Second, the dynamic extension implies that this impact effect is amplified in the following year. To see this point, note first that local house price changes are highly persistent, with an AR1 coefficient of 0.73, conditional on MSA fixed effects. Combining this AR coefficient with the results in Column (4) of Table B6, the effect associated with the interaction terms one year after a local house price increase is $(1.38 * 0.73) - 0.56 = 0.45$. That is, the effect of local house price gains in period t on local industry employment growth in t with the local industry's young-firm share, *and* it rises even further in period $t+1$. The average cumulative effect equals 1.38 + 0.45 = 1.83, which is nearly twice as large as the $\hat{\boldsymbol{c}} = 0.95$ value yielded by the static specification (8). Thus, the dynamic extension to (8) implies that local industry employment responses to local house price

[1] Because of the dotcom bust, San Francisco stands out as an MSA with a large drop in the young-firm employment share from 2002 to 2006, even as local housing prices appreciated. In unreported results that exclude data for San Francisco, the estimate of β corresponding to Column (4) in Table B2 is .154, somewhat larger than the full-sample estimate.

changes vary with the firm-age structure of employment in the local industry by nearly twice as much as indicated by the calculations in Panel C of Table 7.