Local Scars of the US Housing Crisis[⋆]

Saroj Bhattarai[a], Felipe Schwartzman[b], Choongryul Yang[c]

[a] *University of Texas at Austin*
[b] *Federal Reserve Bank of Richmond*
[c] *Federal Reserve Board*

Abstract

The 2006–09 US housing crisis had scarring local effects. For a given county, a housing shock generating a 10% reduction in housing wealth from 2006 through 2009 led to a 4.4% decline in employment by 2018 and a commensurate decline in value added. This persistent local effect occurred despite the shock having no significant impact on labor productivity. The local labor market adjustment to the housing shock was particularly costly: local wages did not respond, and long-run convergence in the local labor market slack instead took place entirely through population losses in affected regions. Moreover, the 2002–06 housing boom does not generate significant employment gains, indicating that the employment losses relative to 2006 are also losses relative to the counterfactual case in which there was no housing cycle.

Keywords: US housing collapse, Scarring effects, Persistent regional effects, Local labor market slack, Downward wage rigidity
JEL classification: G01, R23, E24

1. Introduction

Can a temporary macroeconomic shock cast a long shadow even if it does not directly destroy capital or affect labor productivity? The housing crisis of 2006–09 suggests that this may be the case as, by many measures, the US economy appears to have taken very long to recover from it (Coibion et al., 2017).[1] As pointed out by Fernald et al. (2017), however, it can be hard to disentangle the effects of a one-time shock from underlying trends. Identifying persistent responses to the crisis, and shedding light on the mechanisms that may underlie them, can help inform targeted policies to mitigate the long-term impact of large shocks. For instance, as the world economy shuts down in response to a pandemic, policymakers need to worry about its aftermath. To the extent that much of the economic effect of the pandemic is through a severe but temporary reduction in demand for certain goods and services, some of its long-term impacts might resemble the ones observed after the 2006–09 housing crisis.

This paper provides causal evidence for very persistent local impact of the housing cycle in the US. In addition, we show that its local effect was highly asymmetric, with little local output or employment effect

⋆We thank Yuriy Gorodnichenko, the editor, an anonymous referee, Hassan Afrouzi, Mark Bils, Oli Coibion, Steve Davis, Rafael Dix-Carneiro, Stefano Eusepi, Greg Howard, Erik Hurst, Bob King, Nobu Kiyotaki, John Leahy, Andi Mueller, Vladimir Ponczek, Morten Ravn, Ricardo Reis, Esteban Rossi-Hansberg, Ayşegül Şahin, Matthew Shapiro, Mark Watson, Tao Zha, seminar participants at the UT Austin, Richmond Fed, UVA Darden brownbag, and IADB, and the audience at the Econometric Society Winter Meetings, SED, MMF, Dynare, CEF, and Midwest Macro conferences for valuable feedback. The views expressed here are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond, the Federal Reserve Board, or the Federal Reserve System.

Email addresses: saroj.bhattarai@austin.utexas.edu (Saroj Bhattarai), Felipe.Schwartzman@rich.frb.org (Felipe Schwartzman), choongryul.yang@frb.gov (Choongryul Yang)

[1] Such a slow recovery from a mostly transient demand shock is also consistent with cross-country evidence from Reinhardt and Rogoff (2009) and Jordà et al. (2020).

in the boom phase but persistent employment, output, and population losses during the bust. Its impact on the downturn appears to operate through the demand side since there are no significant changes in labor productivity and only temporary effects on measures of labor market slack. Moreover, the shock did not have such a durable impact on house prices and household leverage, lending credence to its temporary nature.

Regarding the labor market adjustment to these scarring effects on employment, we find no role for wage adjustment. In particular, although wages rose marginally with the housing boom, they did not react at all to the housing bust, implying a potential role for downward wage rigidity. Together, those findings imply that regional labor market adjustment took place entirely through population movements, for which we provide direct evidence. While the observation of permanent population movements leading to adjustment in slack is consistent with classic findings by Blanchard and Katz (1992) for unidentified local shocks, the lack of local wage reactions and asymmetries in labor market adjustment between boom and bust phases are novel findings that are specific to the identified housing shock.

Our analysis starts by documenting some general patterns: US counties with a more substantial housing decline during 2006–09 had a lower level of employment and output in 2018 relative to the pre-2002 trend. Critically, the divergence is a post-crisis phenomenon, with different locations behaving similarly in the boom years. The housing bust, therefore, plays a unique role in driving regional differences in employment and output. These permanent changes occur even though regional gaps in house prices and household leverage converge back to pre-boom baseline.

A formal econometric exercise at the county level follows to provide a causal interpretation of these patterns. We regress changes, over different horizons, in variables such as employment and wages on changes in housing net worth from 2006–09. A natural problem with such regressions is omitted variable bias: both housing net worth and other economic outcomes may have been caused by the same non-housing shock.

To deal with this issue, first, the specification is saturated with a rich set of controls to absorb location-specific effects of other shocks. Those include, among others, state fixed effects, local industrial composition, and local sensitivity to macro-shocks as measured by a factor model and identified aggregate shocks. Further controls are included to account for heterogeneous local ex-ante trends.

Given those controls, two instruments are then used for identification. The first instrumental variable (IV) is the Saiz (2010) housing supply elasticity, used to further eliminate the role of local shocks that may simultaneously affect local outcomes and housing wealth. While the Saiz (2010) instrument is by now an "industry standard,"[2] extra care is taken in precisely showing conditions for it to be valid in our application and various controls for determinants of local demand for land are added, which, as pointed out by Davidoff (2016), could conceivably invalidate the instrument. Our analysis thus addresses existing criticism of the instrument and shows that the results are robust to a wide range of stringent controls. A remaining issue is that with such stringent controls, standard diagnostics suggest that the Saiz (2010) instrument is potentially weak. Weak IV robust inference is therefore used throughout the paper.

As a second instrument, we use orthogonalized residuals to county-level house prices from 2002–05, obtained from a panel-VAR estimated using data from 1975–2006. In particular, by eliminating the variation in house prices that would be predicted by observable variables, such as employment (both total and in the construction sector), earnings, population, and wages, the goal is to isolate non-fundamental variation in house prices. One potential problem with this instrument is that such non-fundamental variation may be hard to disentangle from news that becomes capitalized in house prices. This problem is addressed, at least in part, by using construction employment and wages as conditioning variables, since those are also likely to react strongly to news that increases house prices. Moreover, fortunately, this source of bias is orthogonal to the Saiz (2010) instrument, which is based on local characteristics determined ex-ante. Since the sources of bias in the two instruments are unlikely to be correlated, it implies that their validity can be assessed through a test of overidentifying restrictions.

Impulse responses to the identified 2006–09 US county-level housing shock are estimated by adapting Jordà's (2005) local projection to a cross-sectional context. Results show that the initial 2006–09 housing

[2]Apart from Mian et al. (2013) and Mian and Sufi (2014), the instrument has been used recently to gauge the effects of the housing cycle by Stroebel and Vavra (2014) and Davis and Haltiwanger (2019).

shock has contractionary effects on employment and output as far out as 2018. In particular, at the county level, a housing shock that generates a 10% reduction in housing-wealth from 2006–09 leads to a 4.4% drop in employment in 2018 compared with 2006. There is also a commensurate drop in output. Moreover, there are no significant employment gains during the 2002–06 boom period, indicating that the employment losses relative to 2006 are also losses relative to the counterfactual case in which there was no housing cycle. This shows clearly the asymmetric nature of the housing shock. Those long-lasting local effects occur in spite of the fact that the shock is associated with a boom-bust cycle in house prices and household leverage that is finalized by 2014.

We next find that a regional slack measure, the employment-to-population ratio, returns to its pre-crisis (2002–04) average around 2014. Moreover, this convergence in slack occurs during a period in which the effects on employment continue to be high and significant. It follows that the convergence in regional slack happens because of slow population adjustment as workers move out of hard-hit areas. We indeed show direct evidence for such smooth population losses over time.

These findings on long-lasting effects on employment and output combined with more transient effects on regional slack raise the critical question of what happens to wages. Again, there is evidence for asymmetric effects. While the housing shock appears to lift wages marginally in the boom phase, there is no evidence of wage contraction in the bust. Identifying the housing shock is essential for this result, as OLS estimates would imply wage declines. The difference emerges because our IV procedure isolates the impact of the housing shock from that of productivity shocks, which are well-known to drive a positive co-movement between wages and employment or output.

We additionally show that with our identified shock, there are no significant short- or long-run effects on labor productivity, which complement our wage results. Moreover, like with wages, OLS estimates again show an effect on productivity, providing further evidence on the importance of separating out the housing shock from productivity shocks. Those results, in turn, imply that evidence on wage rigidity and, more generally, Phillips curve coefficients based on regional data, depend on the nature of the shock and should be interpreted with care even if they exploit a massive shock such as the 2006–09 housing crisis.[3]

Next, we investigate sectoral effects and show that the housing bust has a widespread effect across sectors that goes beyond those in construction. Revisiting Mian and Sufi's (2014) results regarding employment effects on non-tradables, it is shown that those are indeed significant in the short run as in their paper, and additionally, they continue to be significant in later years. This lends credence to the interpretation of the housing shock as a demand shock. In addition, there is some evidence for short- and long-run effects on the high-skilled services sector as well.[4]

Our results have implications for optimal currency areas as they highlight that local adjustment to asymmetric demand shocks in the US took place through labor mobility over several years rather than through wage movements. Therefore, even for the US economy, local adjustment to temporary asymmetric shocks can involve very long-lasting and costly changes.

Our paper connects to the literature on the local dynamic responses to shocks, building on seminal work by Blanchard and Katz (1992) and Davis et al. (1997). A recent application of their methodology to the Great Recession is in Yagan (2019). We add to that work by explicitly isolating the effects of the housing shock from other sources of local variation. In effect, in our second IV approach, our local house price shock is, by construction, uncorrelated with all shocks driving innovations to local employment. We find that, it is only when isolating the effect of the housing crisis from productivity shocks that the lack of local wage and productivity adjustment in response to the housing crisis can be uncovered. More broadly, the local scars of the housing crisis that are established echo findings that changes in trade tariffs have very persistent effects in local labor markets (Dix-Carneiro and Kovak, 2017), and that differences in local economic conditions are very persistent (Amior and Manning, 2018).

[3]For example, Beraja et al. (2019) also explore cross-sectional variation after the crisis and present results on wage adjustment but do not separate the housing shock from other local shocks.

[4]Generally, in our results, the employment responses are mirrored in sectoral output responses and that the employment results for the high-skilled services sector are noisier compared with output results. Moreover, other than in construction, the lack of downward adjustment in wages following the housing crash is a general phenomenon across sectors.

Recent empirical work in macroeconomics has frequently exploited regional variation to understand the labor market impact of the housing cycle. Crucially, Mian and Sufi (2014) in a seminal paper show the short-run effects of the housing crash on labor markets due to lower household demand. Papers following their study have focused, for the most part, on similar short-run dynamics. For instance, Gertler and Gilchrist (2018) examine the effect of housing shocks on local employment over two and a half years, Gilchrist et al. (2018) examine asymmetries in the two-year impact of house price fluctuations in boom and bust phases, and Guren et al. (2018) show how the one-year reaction of retail employment to house prices has changed over time. A similar focus on short-run variation also underlies estimates based on structural or quantitative models, such as Jones et al. (2018) and Beraja et al. (2019). In comparison, our paper directly estimates the dynamics of multiple local economic variables over the almost 20 years encompassing the housing boom-bust cycle and its aftermath.

The need for such a holistic view of the housing cycle, that is, a joint examination of both the housing boom and bust phases, is proposed by Charles et al. (2018). In particular, they find a symmetric movement of employment-to-population ratios between boom and bust, with labor market slack measured in that way converging back to its pre-housing boom levels by 2011. We add to their work by examining a wider range of variables over a longer time period, finding that effects on employment, output, population, and wages are, in fact, asymmetric over the housing cycle.[5] Those, in turn, lead to local scarring effects on employment and output, lasting for more than ten years after the pre-crisis peak. We then uncover a mechanism for the convergence in employment-to-population ratios: it occurs through population losses in the most-affected regions during the housing crisis.

Finally, from a methodological standpoint, our results highlight a key difference between local and aggregate elasticities and economies. Because of population movements, demand shocks can have persistent effects on aggregate slack even if that linkage is not apparent in regional data.[6] The findings in our paper should therefore help inform general equilibrium models of housing shocks by highlighting the relevance of labor mobility.

2. Data and Motivating Evidence

This Section describes in detail the data used in the paper as well as presents some stylized facts that serve as motivating evidence for the econometric analysis.

2.1. Data

Our primary dataset is the Quarterly Census of Employment and Wages (QCEW) from the Bureau of Labor and Statistics (BLS). It draws on employment and wages reported by establishments to unemployment insurance programs, and covers more than 95% of jobs in the US. It is the dataset of choice for the Bureau of Economic Analysis (BEA) for the production of national accounting estimates and for the BLS as a frame for the Current Employment Statistics.[7] The dataset includes total employment and wage bill by industry and county. In an extended analysis (in the Appendix), the American Community Survey (ACS) data is used to complement the wage-regression results by constructing an adjusted wage index.

For other important variables, additional data sources are used. We draw on the Local Area Unemployment Statistics (LAUS) dataset from BLS for the county-level unemployment rate and employment-to-population ratio. To examine the local responses of output to the housing shock, the Local Area Gross Domestic Product (LAGDP) dataset from BEA on county-level GDP that has been made available recently is used. Our analysis also draws on county-level personal income data from BEA to examine the local responses of income, and uses BEA state-level GDP deflator to construct a real measure of personal income. Moreover, in order to investigate migration patterns, population data from the County Resident Population

[5] In order to obtain this holistic view in terms of level variables, there is a need to control for heterogeneous local trends, which is done via controls for average growth rates in outcome variables between 1994–98 and 1998–2002.

[6] The results echo Dupor et al.'s (2018) point about spill-overs through trade.

[7] Compared to the County Business Patterns, it is more encompassing, since it includes government employees and a few other industries.

Estimates from the US Census Bureau after 2000, and the US Intercensal County Population data before that, is used. For some robustness checks and splits by worker demographics, the paper makes use of the Quarterly Workforce Indicators (QWI) from the US Census Bureau.

On the household finance side, debt-to-income (DTI) ratios for different counties is obtained using data on household debt from the Equifax/Federal Reserve Bank of New York Consumer Credit Panel (CCP) made available as part of the extended Financial Accounts of the United States on the Federal Reserve Board of Governors website.[8] For comparability with prior work, the change in housing net worth (defined below) made available in Mian and Sufi's (2014) replication files is used. For a robustness check, we use 2000 census data to construct a ratio of housing net wealth to income. Finally, county-level CoreLogic's HPI data serves as a measure of house prices. To construct HPI-to-income ratio, the county-level HPI data is divided by BEA personal income.

For more details on data sources and construction, see Appendix A.

2.2. Descriptive Facts

This Section shows suggestive evidence for large and persistent local effects of the housing crisis. In particular, it analyzes how changes to housing net worth around the housing crisis affected local outcomes, such as employment, output, house prices, and leverage over time. Moreover, it evaluates the extent to which these cross-county differences can be characterized as transitory or permanent.

We follow Mian and Sufi (2014) in defining the log change in housing net worth in a given region n from 2006 through 2009 ($\ln N_{n,2009} - \ln N_{n,2006}$) by

$$\begin{aligned}\ln N_{n,2009} - \ln N_{n,2006} &= (\ln p_{n,2009} - \ln p_{n,2006}) \\ &\quad \times \frac{\text{Housing Wealth}_{n,2006}}{\text{Housing Wealth}_{n,2006} + \text{Financial Wealth}_{n,2006} - \text{Debt}_{n,2006}},\end{aligned} \tag{1}$$

where $p_{n,t}$ is the house price in location n, year t. That is, the log change in household net worth due to housing is given by the log change in the house price index multiplied by a leverage term calculated using initial asset positions.

Focusing on the housing net-worth variation keeps our analysis consistent with a well established literature. It should not, however, be seen solely as a measure of changes in household wealth due to the housing crisis and thus as indicative only of a household demand channel. Instead it serves as a more general index of the size of the housing shock. That is, its main virtue is as a useful summary index that combines two important dimensions of affected counties: (i) house price declines (in the first term), and (ii) large housing leverage (in the second term).

To show basic stylized facts, counties are sorted by quantiles in terms of the size of the change in housing net worth from 2006 through 2009 and Figure 1 shows how various variables evolve over time in these groups. In Panels A and B, we show the evolution of employment. Panel A shows employment growth from 2002. While it shows convergence across counties in employment by 2014, it is also clear that the boom-bust cycle was most pronounced in counties that were growing fast ex-ante. Panel B corrects for these heterogeneous trends, by taking the 1994-2002 growth as baseline. What becomes clear in Panel B is that, relative to that baseline, there is no convergence across counties in employment. Panels C and D show the same facts for GDP. Here, the results are starker, since the long-run divergence between high and low housing net worth counties is also apparent without any detrending in Panel C.

Next, Panel E in Figure 1 shows the variation in house prices, relative to a 2002 baseline, for the different groups of counties. It reproduces a well known fact: the housing bust was largest in counties where the housing boom was also the most pronounced (Charles et al., 2018). It also shows that the housing bust completely and rapidly eliminated all relative gains generated by the boom: by 2009, relative house prices between counties with the largest and smallest house price booms were back to their 2002 baselines.

[8] At the time of writing, the data was available at the source link: https://www.federalreserve.gov/releases/z1/dataviz/household_debt/county/map/#state:all;year:2018

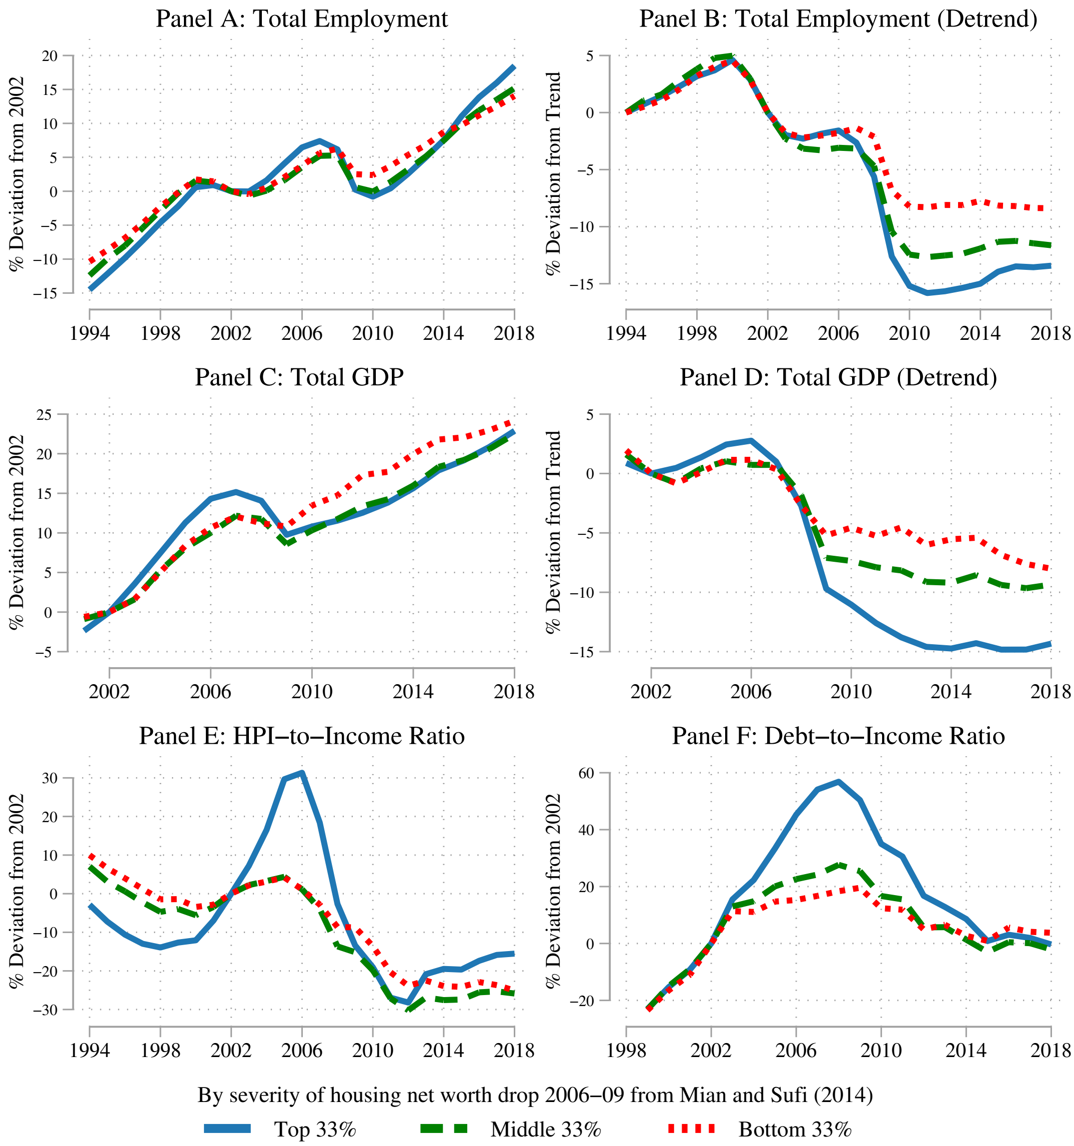

Notes: Panels A and C plot the percent deviation of employment and GDP from their 2002 levels by grouping counties in terms of the severity of housing-net-worth drop. Panels B and D plot the percent deviation of employment and GDP from their trends. Employment trend is calculated by taking average growth rates from 1994-2002 for each group and using those to project 2002 employment linearly into the future. The GDP trend is calculated by using average growth rates of BEA real personal income from 1998–2002 for each group. State-level GDP deflator is used to calculate the real personal income for each county. The lower panels plot the percent deviation of HPI-to-income ratio (Panel E) and debt-to-income ratio (Panel F) from their 2002 levels.

Figure 1: Changes in Variables by Housing Net Worth Quantiles

Finally, Panel F shows the evolution in debt-to-income ratio, which is the other important element in housing net-worth. Debt-to-income starts to increase in relative terms in the more affected counties around 2002, peaks in 2008, and then slowly declines back. While house prices are at similar levels by 2009, debt-to-income only converges back to baseline around 2015, as to be expected given the slow moving nature of the variable.

Taken together, the panels of Figure 1 imply that a transitory shock to house prices might generate a more persistent impact on debt and permanent reductions in local employment and output. We describe next how the effect of the housing shock is disentangled from other sources of local change to give this pattern a causal interpretation.

3. Disentangling the Effects of the Housing Shock

Figure 1 suggests that regions where the 2006–09 housing shock was more severe also exhibited relatively lower employment and output as late as 2018. This may not be a causal relationship, however. For example, a persistent increase in demand for products from a specific region would lead to local increases in both employment and house prices. How we disentangle the causal relationship from the housing shock through a combination of controls and instrumental variables is discussed now in detail.

3.1. The Basic Econometric Model

In order to estimate the impact of the housing shock on local outcomes, we assume that an outcome X in location n at time t follows the statistical relationships:

$$\ln X_{n,t} - \ln X_{n,2006} = g_n(t - 2006) + \gamma_t\left(\ln N_{n,2009} - \ln N_{n,2006}\right) + e^X_{n,t}, \tag{2}$$

$$\ln N_{n,2009} - \ln N_{n,2006} = \eta_n + e^N_{n,2009}, \tag{3}$$

where $\ln N_{n,t} - \ln N_{n,2006}$ is the log change in housing net-worth between 2006 and year t due to price changes, which, as equation (3) shows, is an index for the housing shock η_n. Furthermore, g_n is a region-specific trend-growth term. The parameter γ_t, our main object of interest, captures the time-varying effect of the housing shock on period t outcome variables.

The residuals $e^X_{n,t}$ and $e^N_{n,t}$ summarize all other shocks affecting the outcome variables X and housing variable N in location n at time t. More specifically,

$$e^X_{n,t} = \mu^X \sum_{r=1}^{R} \lambda^r_n z^r_t + \phi^X_t u_{n,t}, \tag{4}$$

where z^r_t is one out of R aggregate driving forces (such as nationwide increases in demand for certain products), λ^r_n is the local sensitivity to that aggregate shock (such as the share of the industry in the location), $u_{n,t}$ is a shock idiosyncratic to the location, (such as the opening of a new plant or a change in local regulations that were previously unexpected), and ϕ^X_t captures the effect of those idiosyncratic shocks on variable X at time t. Analogous structure as given in equation (4) for $e^X_{n,t}$ also holds for $e^N_{n,t}$.

Local trend-growth g_n is not observed either. In order to control for cross-sectional differences in growth rates, ex-ante growth rates are added as controls, with coefficients to be estimated.[9] The model is estimated for each year t separately, in a cross-sectional version of the Local Projection method proposed by Jordà (2005).[10] Since we measure the housing shock η_n with the housing net worth loss between 2006 and 2009, the more negative the change in housing net-worth, the larger is the housing shock. Therefore, if an outcome $X_{n,t}$ is house prices, for example, we would expect $\gamma_t < 0$ in the boom years and $\gamma_t > 0$ in years after the bust.

As in Section 2.2 above, the housing net worth loss between 2006 and 2009 is used as an index of the housing shock. As previously discussed, this variable is taken as a yardstick that is consistent with prior literature and with magnitudes that can be readily interpreted.

[9] In the baseline specification, both 1994–98 and 1998–2002 average growth rates are used, wherever possible. In a sensitivity analysis, 1990-94 average growth rates are also used, wherever possible.

[10] Apart from the extensive controls that was discussed in Section 3.2.1, we also include as controls residuals from the previous year (when available) to pick persistent shocks affecting the residuals.

Table 1: Instrumental Variables and Control Variables

Panel A. Instrumental Variables

- A dummy for upper tercile of housing supply elasticity (Saiz, 2010)
- A dummy for lower tercile of orthogonalized 2002–05 house price shocks from a panel-VAR

Panel B. Control Variables

- 1994–98 and 1998–2002 growth rates of outcome variables
  - 1998–2002 growth rates of real personal income (per worker) for GDP (per worker) regressions
- State-fixed effects
- 2002 QCEW 2-digit industry employment shares (20 industries)
- Aggregate shocks controls
  - Sensitivity of employment growth to monetary shocks and excess bond premium shocks
  - Three main factor loadings from a factor regression using 10-year employment growth rates
- 2002 Debt-to-Income ratio
- 2000 Housing wealth-to-Wage income ratio (Census and QCEW data)
- Davidoff (2016) controls and local land demand controls
  - Fraction of the population that had education greater than or equal to 4 years of college
  - Fraction of the population that were born outside the U.S.
  - "Bartik" measure of local demand pressure
  - Density measure which is housing units divided by land area
  - Geographical dummy variable for "Coastal" area
  - Quality of life index (Albouy, 2008)
  - Natural amenities scale (U.S. Department of Agriculture Economic Research Service)

Notes: This table shows our instrumental variables and a set of control variables in our baseline regressions. Data sources are available in Appendix A.2.

As equation (4) makes clear, the main problem with using housing net-worth as an index of the housing shock is that it is determined not only by the housing shock η_n, but also by the same aggregate and idiosyncratic shocks that determine other outcome variables X. How we handle those concerns is discussed next.

3.2. Handling Identification Concerns

When estimating γ_t in equation (2), the main identification concern is that a non-housing shock may simultaneously drive the housing net worth loss and appear in the residual term $e^X_{n,t}$. For example, a shock that increases local productivity, or demand for local products, might generate both an increase in housing net worth and in local output or employment.

The precise way in which these concerns are handled, with a mix of controls and instrumental variables, is described next. The various controls and instruments are summarized in Table 1.

3.2.1. Controls

The following controls are added to eliminate the effect of common shocks to housing net-worth and other local outcomes:

State effects:. In all specifications below, state fixed effects are used. This controls for any state-specific shocks, as well as any state-specific variation in the sensitivity to national shocks.

Aggregate shocks:. The following, more explicit, controls for the local effects of aggregate shocks, $\sum_r \lambda_n^r z_t^r$, are also included.

Shares of employment in 20 different 2-digit-level industries We control for the share of employment in 20 industries in 2002.[11] Industry shares are particularly well-suited to eliminate local differences in response to aggregate cost or demand shocks to particular industries. They also capture other systematic differences in local economies that could influence local response to aggregate shocks. For example, locations specializing in the production of durable manufacturing may be more susceptible to any national shock, since durables are more cyclically sensitive. In contrast, places that concentrate on financial services may be more responsive to monetary or financial shocks.

Local sensitivity to monetary and financial shocks Local (county-level) employment is regressed on identified aggregate monetary and financial shocks using pre-2002 data. The estimated coefficients are used as controls.

Local sensitivity to other aggregate shocks Note that $e_{n,t}$ has a factor structure, meaning that a large number of cross-sectional observations are in large part determined by a small number of aggregate factors. A rolling 10-year window of local employment changes is then used to estimate a principal component model with main three factors. The local factor loadings λ_n from this model are extracted and used as controls. See Appendix A.2 for details.

Initial conditions:. Lastly, we allow for the possibility that initial wealth conditions affect the dynamic response to the housing shock. Specifically, the debt-to-income ratio in 2002 and a measure of household wealth-to-income ratio in 2000 are used as controls.

3.2.2. Instrumental Variables

While the controls above can absorb a wide range of common sources of variation, an OLS estimate of equation (2) would still result in biased estimates if there are remaining sources of idiosyncratic shocks in the data. For example, the unexpected opening of a large plant can single-handedly affect local economies (Greenstone et al., 2010).

To deal with this problem, two instrumental variable strategies are combined. The first, which has been used before in the literature, is to use local measure of housing supply elasticities by Saiz (2010) as instruments, with enough additional controls to account for well-known criticism (Davidoff, 2016). The second is to use orthogonalized residuals of a house price index in a panel-VAR as a measure of non-fundamental variation in house prices. Each of these strategies is described in turn next:

Housing Supply Elasticities:. The Saiz (2010) instrument used by Mian and Sufi (2014) measures the local elasticity of housing supply given by geographical or regulatory constraints. Mian and Sufi (2014) propose it as an instrument for the housing shock because lower housing supply elasticity would allow house prices to increase more quickly in the run-up years from 2002–06, thus allowing households to raise more debt in comparison to their incomes.

A further motivation for the Saiz (2010) instrument comes again from the factor structure of the shocks $e_{n,t}$. Specifically, under an approximate factor structure (Chamberlain and Rothschild, 1982), which holds generally so long as the number of aggregate shocks driving local-level employment is not too large, the idiosyncratic components are such that $u_{n,t}$ cannot be predicted from fixed regional characteristics. That is, for any W_n that is fixed in time,

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} W_n u_{n,t} = 0. \tag{5}$$

[11]Those are also the primary set of controls used by Mian and Sufi (2014). A list of 20 industries is available in Appendix A.3.

Given equation (5), the local shock u_n is purely "random" in that it is not predictable based on fixed local characteristics.[12] Therefore, so long as our controls account for all aggregate sources of variation, any such characteristic that correlates with housing net worth changes around the crisis is a valid instrument. The Saiz (2010) instrument clearly satisfies that criterion. Following the findings of a nonlinear relationship between housing supply elasticity and local housing cycles (Gao et al., 2016), we use a discretized version of the instrument with a dummy for the highest house-price elasticity tercile.

Controlling for local land demand The use of the Saiz (2010) instrument has been criticized by Davidoff (2016), because the same geographical features that affect the supply of land may also affect the demand for land. In particular, Davidoff (2016) finds that the Saiz (2010) land supply elasticity correlates with various local characteristics that capture local demand for land. These local characteristics are therefore used as controls. They include the fraction of the population with more than 4 years of college, the fraction of the population born outside the US, a Bartik measure of local demand pressure, a measure of housing density, and a geographical dummy variable for "Coastal" area. The construction of these controls is described in further detail in Appendix A.2.

Further controls are added for land demand in the form of measures of local amenities and real wages. Specifically, we use (i) an index of local geographic amenities constructed by the US Department of Agriculture, combining six measures of climate, topography, and water area that reflect preferred environmental qualities (warm winter, winter sun, temperate summer, low summer humidity, topographic variation, and water area); and (ii) a measure of quality of life constructed by Albouy (2008), based on after-tax real wages in each location. In spatial equilibrium, differences in real wages between cities for a worker with the same attributes should reflect a compensating differential in local amenities. In other words, those real wages should capture any impact on the demand for living in those places from the geographical features captured by the Saiz (2010) instrument.

Orthogonalized Panel VAR House Price Shocks:. The second instrumental variable used is based on the notion that the housing shock appears saliently in increases in the house-price that are not easily traced back to observable indicators of local economic conditions.[13] Increases in that period are defined to be unusually large if they go beyond what would be normally predicted by current and past changes in employment (total and in construction), personal income per employee, 15–64 population, and wages per employee in construction.

More specifically, this strategy is implemented by: (i) running a panel-VAR at the county level from 1975–2006 with CoreLogic HPI index for house prices, employment (total and in construction), personal income per employee, 15–64 population, and wages per employee in construction; (ii) calculating the innovation for the house price index that is orthogonal to innovations to these other variables; and (iii) designating as an instrument for the housing shock a dummy variable for the orthogonalized house price residuals from 2002–05 that are in the bottom tercile of the distribution. In that period, the orthogonalized residuals in that tercile averaged to zero. By singling out the bottom tercile, the comparison is between counties where we can be confident there has not been a non-fundamental house price increase (since house prices were aligned with what fundamentals would predict) and those above it.

Let us discuss and justify the variable choice in the panel-VAR. In the panel-VAR, apart from the house price index, variables included are those that help summarize the fundamentals in a given locality. This naturally includes employment and population (which are also included in Blanchard and Katz (1992)) as well as total personal income per capita. The latter is especially important as, in combination with employment, it can capture productivity fluctuations. Construction employment and wages is additionally

[12] It holds without loss of generality so long as the number of aggregate shocks driving local-level employment is not too large, and enough aggregate factors are allowed for. If there is some W_n for which equation (5) does not hold, then we can define $z_t^{R+1} \equiv \frac{cov[u_{n,t}, W_n]}{var(W_n)}$ and $\lambda_n^{R+1} \equiv W_n$, and substitute $u_{n,t}$ for $\hat{u}_n, \equiv u_{n,t} - \frac{cov(u_{n,t}, W_n)}{var(W_n)} W_n$, in which case $\frac{1}{N} \sum_{n=1}^{N} \hat{u}_n W_n = 0$.

[13] A focus on unusually large house price increases underlies the instrumental variable approach in Charles et al. (2018). Fort et al. (2013) use orthogonalized panel VAR residuals as measures of regional house price shocks.

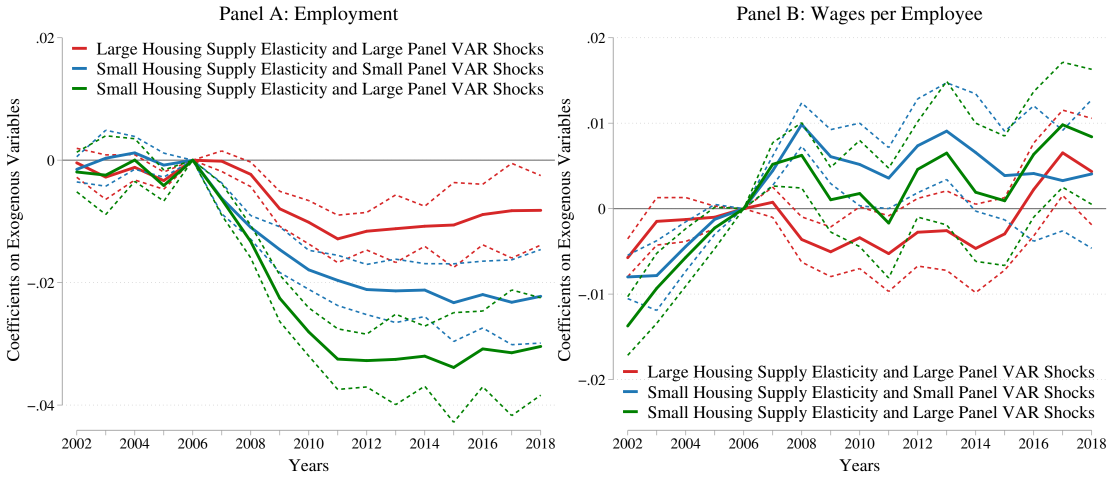

Notes: This figure shows the coefficients on exogenous variables in the reduced-form regressions. Dependent variables are employment (Panel A) and wages per employee (Panel B). Each line represents responses of outcome variables in each group of counties relative to those in a baseline group whose housing supply elasticity is above 33 percentile and orthogonalized panel VAR shocks are below 33 percentile. Red lines represent the relative responses of a group of counties with housing elasticity above 33 percentile and orthogonalized panel VAR shocks above 66 percentile. Blue lines represent the relative responses of a group of counties with housing elasticity below 66 percentile and orthogonalized panel VAR shocks below 33 percentile. Green lines represent the relative responses of a group of counties with housing elasticity below 66 percentile and orthogonalized panel VAR shocks above 66 percentile. Dashed lines are one standard deviation confidence intervals. All control variables listed in Table 1 are included. Prior trends for employment and wages per employee are the average growth rates in those variables from 1994-98 and from 1998-2002. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals.

Figure 2: Results from Reduced-Form Regressions

included to provide further fundamental information about local housing markets. In particular, to the extent that house prices follow news, one would expect those to be reflected in construction activity.[14]

A potential problem with this orthogonalized panel-VAR house price shocks based IV strategy is that an unusually large increase in house prices may also occur in response to news about future shocks. As mentioned above, we partially control for that possibility by including construction employment data as a conditioning variable, since that is also likely to respond to news. Importantly, moreover, this potential source of endogeneity is orthogonal to the potential sources of bias inherent in the Saiz (2010) instrument, which, instead, have to do with fixed local characteristics. This implies that the overidentifying restrictions test is likely to be appropriate to verify the validity of the two instruments. In what follows, results are reported using the two instruments simultaneously, and a J-test of overidentifying restrictions is used to verify that they are jointly valid.

3.2.3. Reduced-Form Results

Before proceeding to our main results, the reduced-form is examined, that is, the relationship between the instruments and outcome variables.[15] Figure 2 shows the estimated paths for employment and wages,

[14] In the Appendix, several additional results related to the panel-VAR are presented, using impulse response and forecast error variance decomposition analysis. First, Appendix Figures B.1 and B.2 present impulse responses to the identified house price residuals and the forecast error variance accounted by the identified house price residual. This is done both for the sample period 1975–2006 as well as 1975–1999, to check for sub-sample stability in propagation of the house price shock, especially when the boom years in the 2000s are excluded. The impulse responses and variance decomposition results are very similar in the two sample periods. Second, Appendix Figures B.3 and B.4 shows how the impulse response and variance decomposition results change as various variables are included in the panel-VAR. To make this clear, we first start with just employment and HPI index, and then progressively add one variable at a time, thereby providing a sense of how various variables affect the propagation of the house price shock. These results show the importance of including personal income and construction data.

[15] That is, equation (3) is estimated with the instrumental variables on the right hand side, instead of $\ln N_{n,2009} - \ln N_{n,2006}$.

the most important outcome variables, conditional on different values for the instrumental variables. The baseline case is the one in which the least amount of variation in housing net-worth is expected, including the counties with top house price elasticities and low non-fundamental house price variation between 2002 and 2005. The expected values refer to differences between this baseline and other combinations. For instance, the green line refers to the case in which the most non-fundamental variation is expected.

The reduced-form results in Figure 2 show that there is no pre-tend in employment, but a progressive increase in wages before 2006 in the most affected areas. Conversely, after 2006, it shows a a clear ranking of employment across counties according to this classification, but no such difference for wages.[16]

4. Results

The impulse responses of various outcomes to the housing shock is now presented. These are computed by estimating equation (2) separately for each year, including all controls, as described above. The impulse response functions are then just the estimated coefficients on the housing net worth loss. All Figures in this Section thus show the estimated values of γ_t in equation (2), together with 95% confidence intervals.[17]

For all variables, OLS and IV results are shown.[18] As discussed before in the introduction and Section 3.2, OLS results mix the effects of shocks to housing wealth on local outcomes with the simultaneous effect of productivity shocks (and, more generally, other shocks on all observables). That is, one of the main concerns for us is of omitted variable bias. As will be seen, results for both estimators are qualitatively similar in many, but not all, instances.

In what follows, IV results are presented using both instruments simultaneously. As previously discussed, this allows us to use J-tests to evaluate the validity of the instruments, since their potential sources of bias occur over a-priori orthogonal dimensions.[19] The role of each IV individually is explored in detail in the Appendix, together with the standard diagnostics.[20] Broadly, the same main results are obtained with both instruments individually. It should be noted however that, given the state fixed effects and other stringent controls, the Saiz (2010) housing supply elasticity is a potentially weak instrument, and standard errors for estimates using only that instrument are large.[21]

4.1. Scarring Effects on Economic Activity

We now show that the housing shock had very persistent effects on employment and GDP. In particular, Panels A and B of Figure 3 confirms the basic descriptive findings of long-run effects from Section 2.2: While up to 2006, the housing cycle did not appear to generate a discernible difference in employment levels between counties, after the bust, the most affected counties experienced significantly larger employment losses, which persisted in the long-run. A similar behavior is obtained in county-level GDP, as shown in Panels C and D of Figure 3, and to a slightly lesser extent, also in county-level personal income, as shown in Panels E and F of Figure 3.

Interestingly, the IV results imply larger employment effects over the long-run as compared to OLS estimates. This may happen if local productivity shocks are relatively short-lived, so that they have a larger

[16]The reduced-form results for the two instruments separately is in Appendix Figure B.5.

[17]In all impulse response figures, we include 95% weak IV robust confidence intervals with coverage distortion bounded by 10%. The `twostepweakiv` package in STATA written by Sun (2018) is used to implement the two-step identification-robust confidence intervals proposed by Andrews (2018), based on the Wald tests and the linear combination tests in Andrews (2016).

[18]The same baseline sample is restricted in both OLS and IV regressions.

[19]While in the results in this section, the p-value for the J-statistics is presented only for 2018 to keep them uncluttered, the year-by-year p-values are in Appendix Table B.1.

[20]In particular, the F-statistics, separately by instruments, and year-by-year are in Appendix Table B.2. The first-stage coefficients, separately by instruments, are in Appendix Table B.3 for 2018 and in Appendix Figure B.6 for all years. In terms of results, Appendix Figure B.7 shows some of our key findings using the two instruments separately. Appendix Figure B.7 reports the F-statistics and the p-value for the J-statistics for year 2018 only, to keep it uncluttered, but more details is provided elsewhere as mentioned above.

[21]Those standard errors are still interpretable, however, since weak IV robust inference is used throughout.

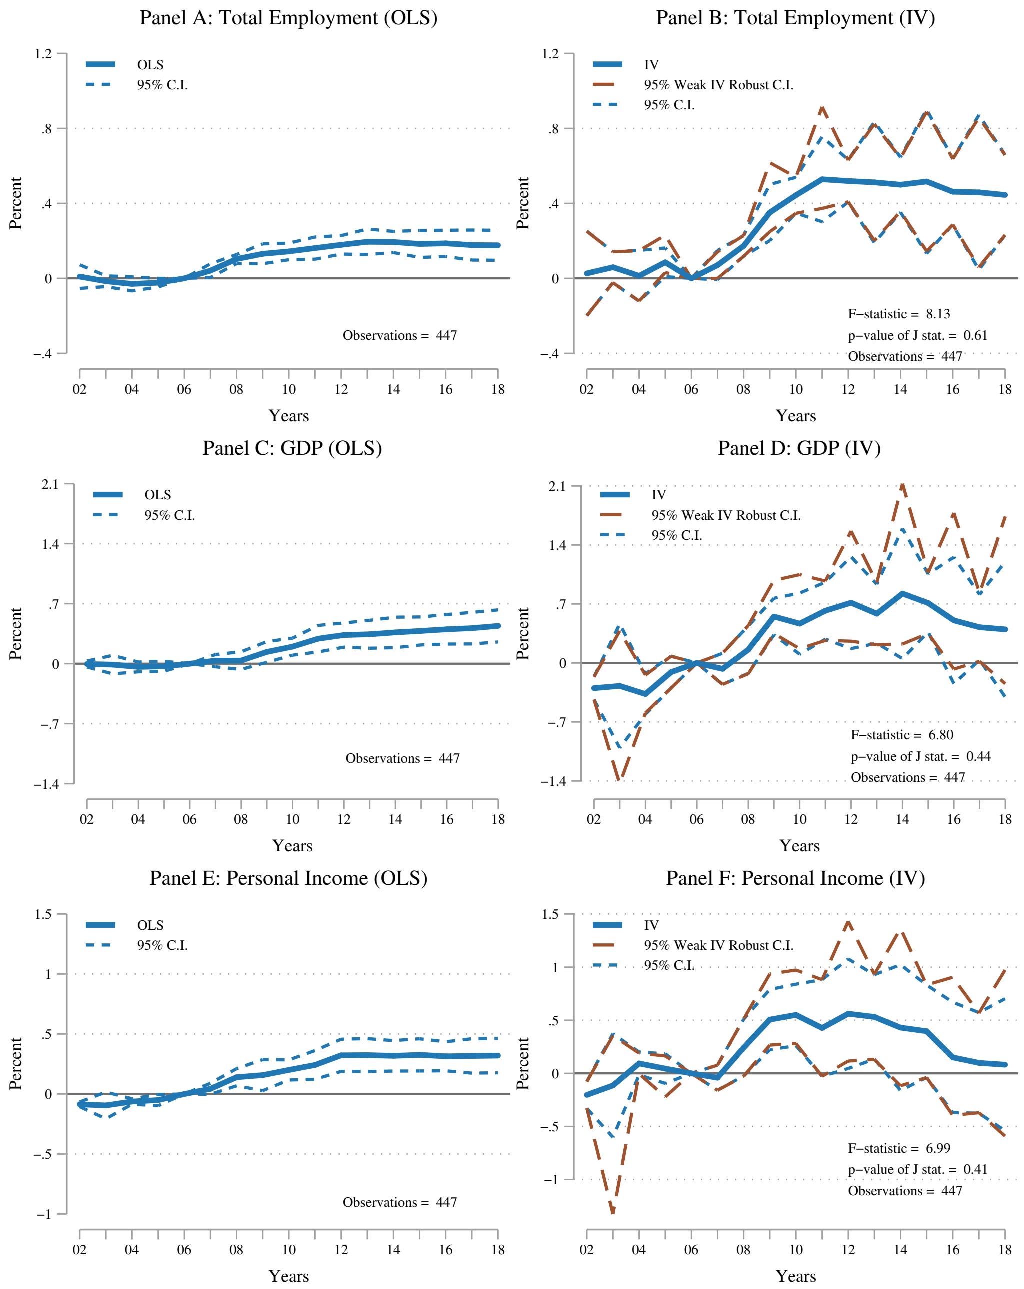

Notes: The figure plots the impulse responses of total employment (Panels A and B), total GDP (Panels C and D) and real personal income (Panels E and F) to the 2006–09 housing shocks. The left columns are results from OLS estimations, and the right columns are results from IV estimations. All control variables listed in Table 1 are included. Prior trends for employment are the average growth rates in employment from 1994–98 and from 1998–2002. Prior trends for GDP and real personal income are average growth rate in real personal income from 1998–2002. We divide BEA county-level personal income by state-level GDP deflator to calculate the real person income. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Figure 3: Changes in Employment, GDP, and Income

effect on housing net worth losses over a three-year period than on employment over 12 years. To see that, consider the simplified model:

$$\ln X_{n,t} - \ln X_{2006} = \gamma_t \left(\ln N_{n,2009} - \ln N_{n,2006}\right) + \phi_t^X u_{n,t},$$
$$\ln N_{n,2009} - \ln N_{n,2006} = \eta_n + \phi_t^N u_{n,2009},$$

where X is a local outcome, N is the housing net worth, η is the housing shock, u is a local productivity shock, and γ_t, ϕ_t^X, and ϕ_t^N are strictly positive. Assuming that u and η are orthogonal, if β is estimated by running an OLS regression of change in $\ln X$ on $\ln N$, then $\gamma_t^{OLS} = \gamma_t + \left(\frac{\phi_t^X}{\phi_t^N} - \gamma_t\right) \frac{\left(\phi_t^N\right)^2 var(u_{n,t})}{var(\eta_n) + \left(\phi_t^N\right)^2 var(u_{n,t})}$. The bias is downward if $\frac{\phi_t^X}{\phi_t^N} < \gamma_t$ and is upward otherwise. For example, a downward bias will occur if productivity shocks have an impact on housing net worth changes from 2006 through 2009 ($\phi_{2009}^N > 0$) coupled with no effect on local employment in 2018 ($\phi_{2018}^X = 0$).

In terms of magnitudes, our IV results imply that a housing shock that generates a 10% reduction in housing wealth in 2006–09 leads to a 4.4% drop in employment, and a 4.0% drop in output, in 2018 compared to 2006. For a sense of economic importance, the estimates imply that going from the 90th to the 10th percentile of change in housing net worth distribution reduces employment by 7.7%, and GDP by 6.9%, in 2018 compared to 2006. For comparison, going from the 90th to the 10th percentile of the 2006–18 employment-growth distribution reduces employment growth rate by 31.7 percentage points and GDP growth rate by 33.3 percentage points.[22]

Overall, the dynamic reaction of employment mirrors classic findings by Blanchard and Katz (1992). The IV results show that this is true also when we separately identify the housing shock. We further find the same persistent impact on local GDP using newly available data constructed by the BEA, as well as to a slightly less extent, persistent effects also on personal income.

4.2. Mean Reversion in Labor Market Slack

Having established long-run effects on employment and GDP of the housing shock, we now turn to the effects on local labor market slack. This is an important question that was also examined by Blanchard and Katz (1992). They find that while local shocks have permanent effects on employment levels, they have only a temporary impact on measures of local labor market slack, such as the employment-to-population ratio and the unemployment rate. They interpret those results with population changes across regions in response to the shock, which leads to mean reversion in local slack.

Such mean-reverting dynamics for local slack in response to the housing shock appear clearly in Figure 4, both for the employment-to-population ratio (Panels A and B) and the unemployment rate (Panels C and D).[23] If employment changes permanently while the employment-to-population ratio does not, then the adjustment must take place through population movements. Panels E and F in Figure 4 verify that to be true. Population reacts smoothly, but persistently, to the shock in both OLS and IV specifications.

4.3. No Effects on Wages and Productivity

Our results above on population changes playing a key role in regional slack adjustment raise a natural question on the behavior of wages. We, therefore, investigate the role that wages play in helping equilibrate local labor markets as house prices fluctuate. Responses of local aggregate wage per worker (from QCEW) are depicted in Figure 5 (Panels A and B).

[22] In terms of short-run effects, at the county level, a housing shock that generates a 10% reduction in housing wealth in 2006–09 leads to a 3.5% drop in employment, and a 5.5% drop in output, in 2009 compared to 2006. This short-run employment elasticity is very similar to the estimate in Mian and Sufi (2014). Focusing ten years out, until 2016, at the county level, a housing shock that generates a 10% reduction in housing wealth in 2006–09 leads to a 4.6% drop in employment, and a 5.1% drop in output, in 2016 compared to 2006. These ten-year estimates imply that going from 90th to 10th percentile of change in housing net worth distribution reduces employment by 8.0%, and GDP by 8.8%, in 2016 compared to 2006.

[23] These results are in line with Charles et al. (2018), who show labor market participation converging back to pre-boom baselines in localities most affected by the housing bubble.

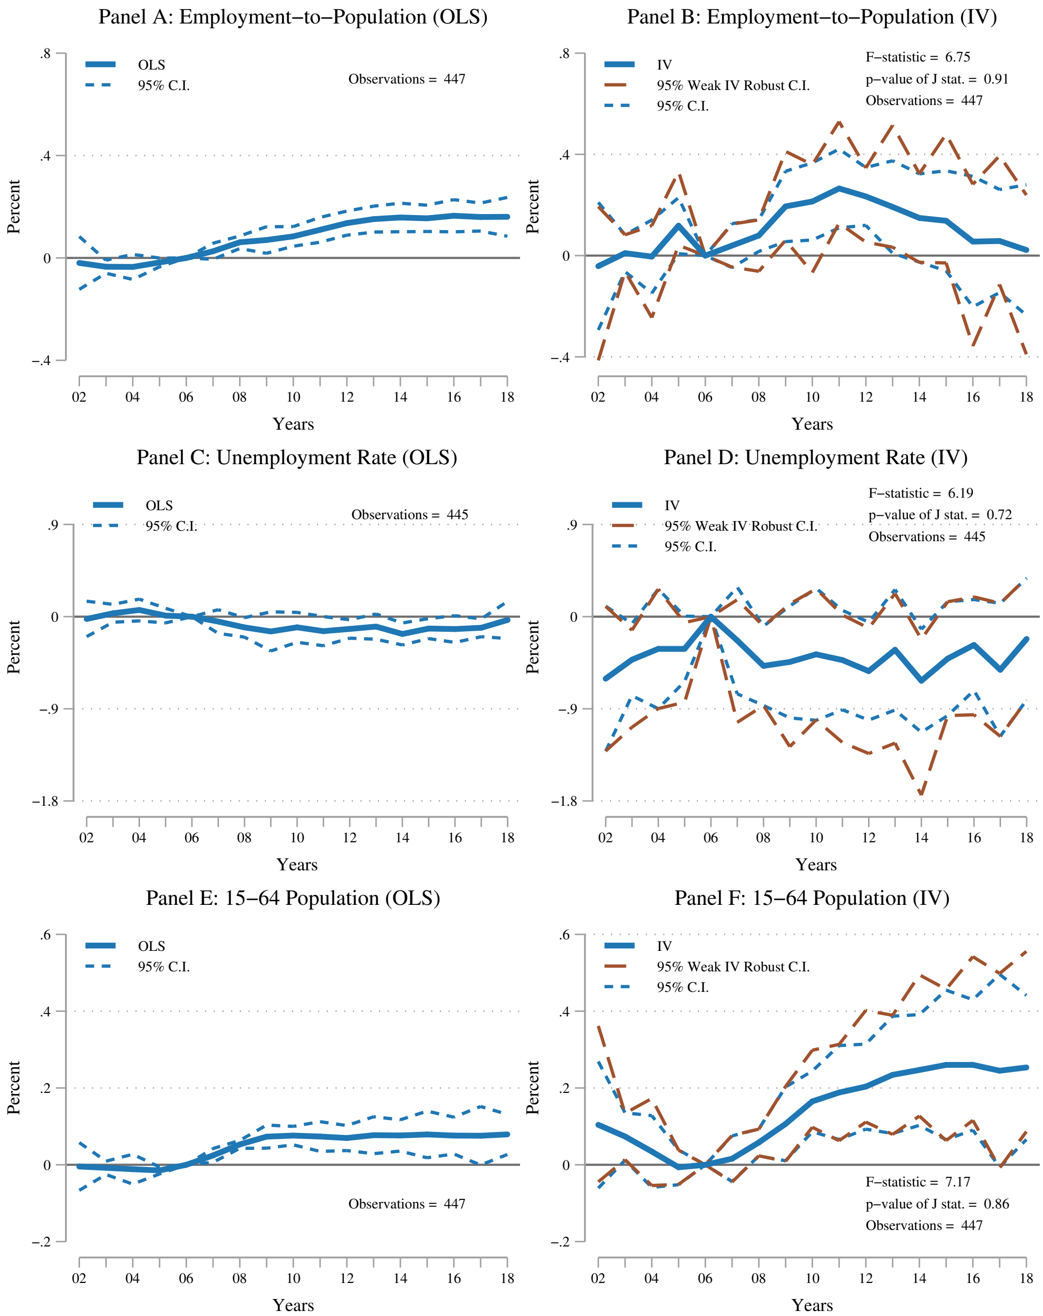

Notes: The figure plots the impulse responses of employment-to-population ratio (Panels A and B), unemployment rate (Panels C and D), and 15–64 population (Panels E and F) to the 2006–09 housing shocks. The left columns are results from OLS estimations, and the right columns are results from IV estimations. All control variables listed in Table 1 are included. Prior trends are average growth rates of outcome variables from 1994–98 and from 1998–2002. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Figure 4: Changes in Employment-to-Population Ratio, Unemployment Rate and Population

These results contain the most meaningful differences between OLS and IV estimates. With OLS, there is no difference in wages before the housing peak, but afterward, wages decrease persistently in more-affected locations. In contrast, the IV results have the opposite pattern: wages at first increase faster in places that are more affected by the housing boom, but then they do not adjust downward as the boom turns into a bust.

These results suggest an asymmetric adjustment of wages consistent with the literature emphasizing downward wage rigidity. In particular, downward wage rigidity has recently been documented in microeconomic data by Grigsby et al. (2019) within this same context. Moreover, it can play a very important role in hindering the adjustment of regions within a currency union to asymmetric shocks in the presence of limited labor mobility, as shown in Schmitt-Grohé and Uribe (2016). The contrast between OLS and IV highlights that while wages may react to some shocks, they do not seem to react to the exogenous negative housing shock suffered by many localities in the recession.[24]

We now look at effects of the housing shock on productivity. First, this serves as a complementary evidence for the results on wages. Second, it helps assess whether the productivity based channel emphasized in Anzoategui et al. (2019) through which transitory shocks can have persistent effects is relevant for the housing shock. Panels C and D of Figure 5 show the effects on one measure of labor productivity (GDP per worker), while Panels E and F show the effects on another measure (Personal Income per worker).

As with wage results in Panels A and B, it is clear that while the OLS results show a relationship between housing net worth losses from 2006 through 2009 and labor productivity changes over time, that relationship is absent in the IV estimates. This finding is important for two reasons in order to interpret both previous, as well as, the rest of the results. First, they show that the long-term effects of the housing crisis that we document below do not arise from a reduction in productivity but instead, operate through other channels. Second, the difference between OLS and IV again indicates that OLS results are likely to be contaminated by other shocks, especially those that have effects on labor productivity.

4.4. Short-Lived Effects on House Prices and Leverage

This Section assesses the results on variables that are likely to mediate the response of employment and output to the housing shock. First, almost by definition, the housing shock should have an impact on local house prices. Second, theories of protracted propagation such as Guerrieri and Lorenzoni (2017) emphasize that financial or wealth shocks can have protracted demand-side effects as households are forced to de-lever.[25] Thus, our focus is on house prices and leverage, and in particular, on investigating whether the effects of the housing shock on house prices and leverage were as long-lived as those on employment.

Our analysis starts by checking that the housing net worth losses indeed capture the boom-bust cycle in house prices. Here, differences of house prices from 2002 is shown to capture the full cycle. Panels A and B of Figure 6 confirm this to be the case. Counties which experienced the largest reduction in housing wealth from 2006 through 2009 were also subject to the strongest boom-bust cycle in house prices. IV responses are more pronounced, indicating that those are more effective at singling out the boom-bust cycle. Conversely, the OLS estimates are likely to be contaminated by the simultaneous response of household net worth and house prices to productivity shocks. Also, they drop below the 2002 baseline, indicating that OLS captures more than a reversal of the housing boom.

Looking at dynamic implications, the losses in house prices captured by the IV bottom out around 2010. Then, by 2011, the differences in house prices across counties stabilize at close to 2002 levels, after which the difference is no longer statistically significant.

Much of the post-crisis literature has emphasized the role of household deleveraging in delaying the recovery from the recession. For comparison with house price results, we show difference in leverage from 2002 to capture the full cycle. Panels C and D of Figure 6 show that during the boom years, household leverage rises relatively more in the more affected regions, peaking in 2009, three years after the peak in

[24]Our OLS results are in line with those found by Beraja et al. (2019), who find a positive correlation between wages and employment outcomes at the state level during the recession, using ACS data.

[25]Berger et al. (2017), Jones et al. (2018), and Justiniano et al. (2015) exploit the interaction between debt and housing values in quantitative models.

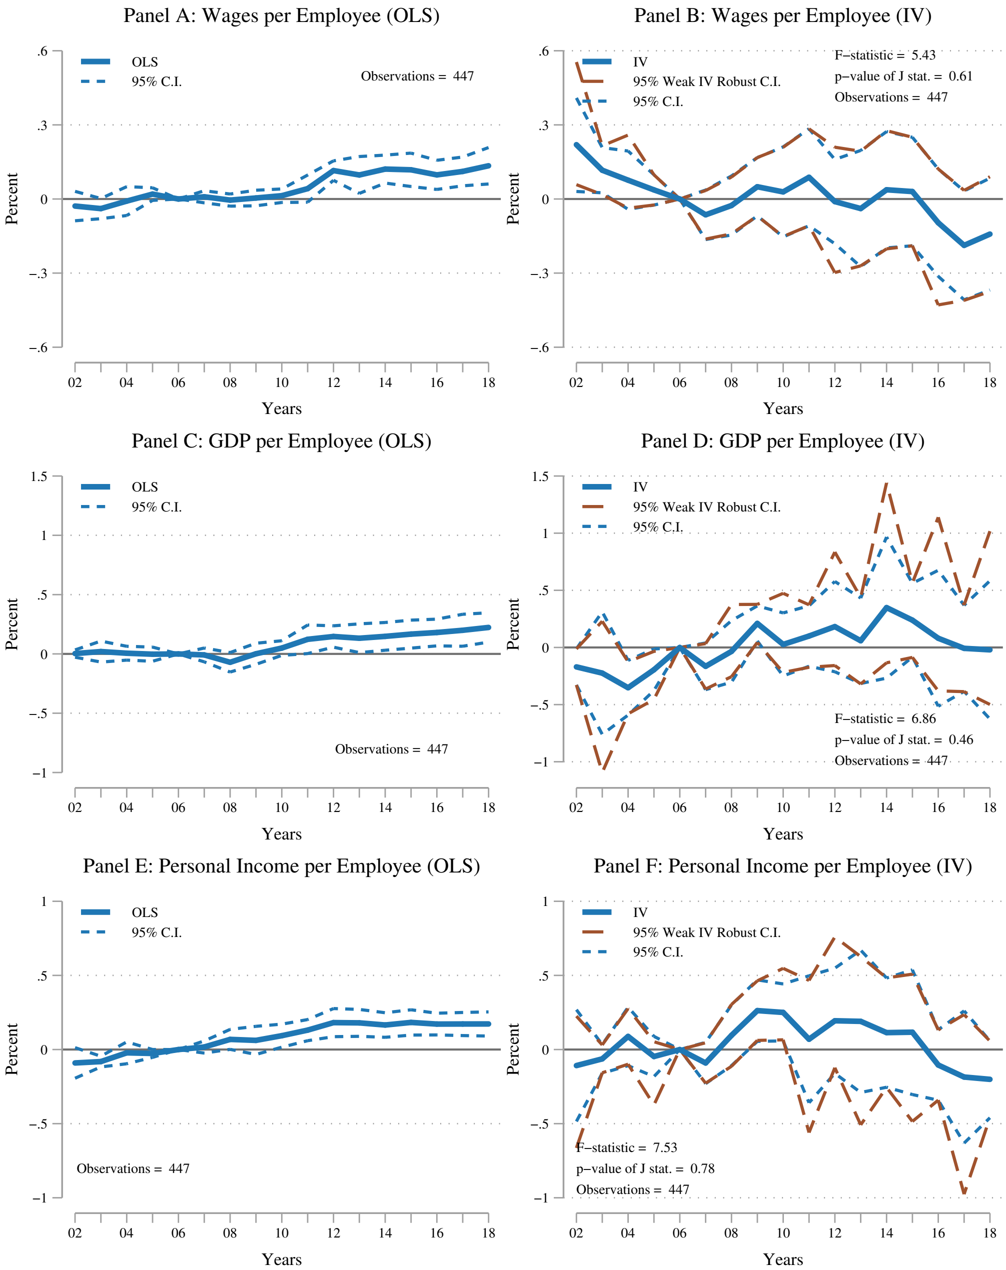

Notes: The figure plots the impulse responses of QCEW wages per employee (Panels A and B), GDP per employee (Panels C and D), and BEA real personal income per employee (Panels E and F) to the 2006–09 housing shocks. The left columns are results from OLS estimations, and the right columns are results from IV estimations. All control variables listed in Table 1 are included. Prior trends for wages per employee are the average growth rates from 1994–98 and from 1998–2002. Prior trends for real personal income per employee and GDP per employee are average growth rate in real personal income per employee from 1998–2002. We divide BEA county-level personal income by state-level GDP deflator to calculate the real person income. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Figure 5: Changes in Wages per Employee and GDP per Employee

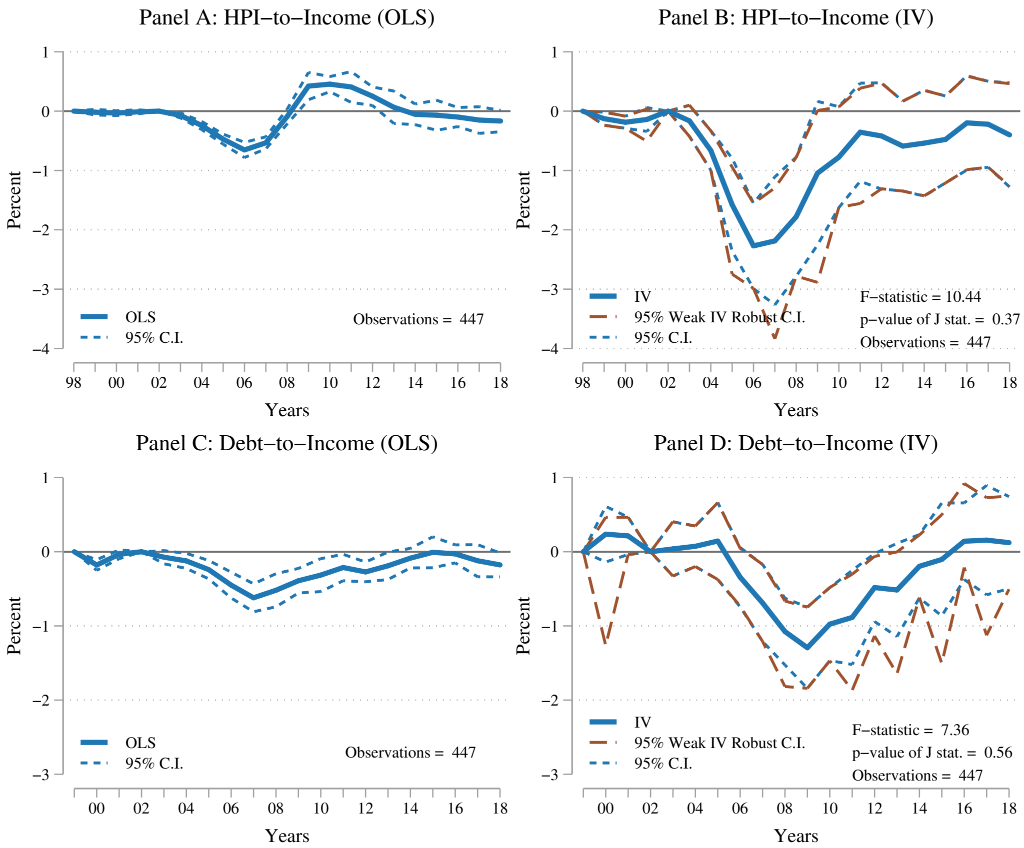

Notes: The figure plots the impulse responses of HPI-to-income ratio (Panels A and B) and debt-to-income ratio (Panels C and D) to the 2006–09 housing shocks. Outcome variables are expressed as deviations from 2002 levels. The left columns are results from OLS estimations, and the right columns are results from IV estimations. All control variables listed in Table 1 are included. Prior trends for HPI-to-income ratio are captured by the average growth rates from 1994–98 and from 1998–2002, while prior trends for debt-to-income ratio are the average growth rate from 1999-2002. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Figure 6: Changes in Housing Prices and Debt-to-Income

house prices. Deleveraging takes over after that, but leverage is mostly back to 2002 levels by 2014–15 and remains so after that. Therefore, even if deleveraging helped propagate the impact of the housing shock, it could not explain the continuing short-fall in employment as of 2018.

4.5. Broad-based Sectoral Effects

Finally, we investigate the impact of the housing shock on employment within sectors. Those can be useful to evaluate if our results are broad-based or particular to specific sectors. For example, Mian and Sufi (2014) show that the short-term impact of the housing shock was particularly relevant among non-tradables, reinforcing the interpretation of the shock as having its main impact through household demand.

The sample is split into five sectoral groupings: tradable (mainly manufacturing), non-tradable (retail and restaurants), construction, high-skilled services (professional and business services, educational services, and health services), and others (including, among others, wholesalers and transportation services). In these sectoral splits, Mian and Sufi (2014) is followed directly, except that the "others" sector is further split from

Notes: The figure plots the impulse responses of employment to the 2006–09 housing shocks by sectors. All the results are from IV estimations. All control variables listed in Table 1 are included. Prior trends for sectoral employment are the growth rates in employment in each sector from 1994–98 and from 1998–2002. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions. See Appendix A.3 for the details of sectoral splits.

Figure 7: Changes in Employment by Sector

their decomposition into two: a high-skilled and the rest. The details of these splits is described in Appendix A.3.

These sectoral-level employment results are presented in Figure 7 (the same exercise is repeated for wages in Appendix Figure B.8 and for output in Appendix Figure B.9). First, as is clear from Panel D, the housing crash had both short- and long-run effects on construction employment. These effects however, were not restricted to the construction sector only, and in fact spilled over to other sectors.

Thus, as in Mian and Sufi (2014), in Panel C, there are sizable estimated effects on non-tradable employment over the first few years of the recession. Moreover, these effects on non-tradable employment persist over the long-run, lending credence to the housing shock as a demand shock. Intriguingly, as Panels E and F makes clear, there are also large and sustained estimated effects on the high-skilled services and others sectors.[26] Lastly, like in Mian and Sufi (2014), Panel B shows that there is no statistically significant effect on tradable sectors. These findings for employment effects are mirrored in output responses (Appendix Figure B.9). Likewise, the lack of downward adjustment in wages following the housing crash is also similarly

[26] The employment results on the high-skilled services sector are noisier compared to output results shown in Appendix Figure B.9.

broad-based (Appendix Figure B.8).

To summarize our sectoral results, we find that the housing bust had effects that spilled over to other sectors beyond construction, such as non-tradables, the high skilled sector, and others.

4.6. Sensitivity Analysis

Appendix B reports results from several robustness and sensitivity exercises. In the baseline IV results, the two instrumental variables are jointly used and tests of over-identifying restrictions are reported. For completeness, Appendix Figure B.7 presents results for employment and wages while using the two instruments separately. The results are similar to our baseline results. This is evidence for the validity of the two instruments. For example, if news was an important driver of the panel VAR residual, the results would diverge from the ones obtained from using the Saiz (2010) elasticity, since the latter are not influenced by news. Conversely, since land demand factors that are correlated with the Saiz (2010) instrument are fixed local categories, they are unlikely to be correlated with a one-time panel VAR residual. The statistical similarities between the two specifications is verified formally by the J-statistics reported previously. The one caveat with the separate instrument results is that, as mentioned before, the Saiz (2010) elasticity based IV estimates lead to wider standard errors as the instrument is potentially weak given standard diagnostics.

Next, additional sectoral results are presented. Appendix Figure B.8 shows the responses of wages per employee to the housing shock by sectors. While wages do not decline following the housing crash either in the aggregate or in other sectors, there is a substantial decline in the construction sector. Appendix Figure B.9 shows the responses of value added to the housing shock by sectors. It is found that GDP responds persistently in the non-tradable and high-skilled sector, similar to our baseline sectoral employment results.

For our baseline results on employment and wages, Appendix Figure B.10 presents results while including an additional pre-trend control using growth rates from 1990–94 (our benchmark results use as controls, growth rates from 1994–98 and 1998–2002, as there is data on a wider range of variables for later time periods). The results are indistinguishable from our baseline results.

Some sensitivity analysis regarding our weighting procedure is presented next, where note that in our baseline specification, we weighted our regressions with number of households, following Mian and Sufi (2014), for clear comparability. Some additional econometric justification is now explored for using weights. In Appendix Figure B.11, we compare our main results, those of employment and wages per employee, with and without weighting. The results show that precision improves with weights and thus they are consistent with efficiency gains coming from appropriate handling of heteroskedasticity through weighting.[27] That is, while the point estimates for employment and wages are robust to weighting, the standard errors are tighter with weights than without. For completeness, in Appendix Figure B.12, we report house prices-to-income and debt-to-income results with and without weighting. Overall, point estimates are still similar, but here, the efficiency gains through weighting are not visible.

Next, using ACS micro-data, a wage series is computed that allows for shifts in labor force composition following Katz and Murphy (1992). The adjustment method is described in more detail in Appendix A.1. Appendix Figure B.13 presents our results on these adjusted ACS hourly wages, where for comparison, the baseline QCEW wage results are also shown. For this new, composition adjusted measure for wages, the same results that they did not respond to the housing crash are obtained. Furthermore, Appendix Figure B.14 examines responses of ACS employment at the regional level split by education and age, while Appendix Figure B.15 examines whether changes in ACS wages at the regional level differ by education and age. They suggest that the employment results are quite broad based while the wage results are the same as our baseline results of no response.

Finally, additional results using the QWI are presented, which not only gives us an alternate series of employment and earnings, but also further allows us to split the analysis by worker characteristics to get another view on compositional issues. First, Appendix Figure B.16 shows regression results for employment and earnings per employee using QWI data, which are very similar to our baseline results. Appendix Figures

[27] In fact, regressing squared residuals for employment on the inverse weights shows a relationship which is positive and highly significant with a t-stat of 6.63.

B.17 and B.18 next show the impacts of employment and earnings per employee to the housing shock by workers' education, age, and gender groups. They suggest that employment losses are mostly broad-based, while earnings do not respond generally.

5. Conclusion

The housing collapse of 2006–09 had scarring effects across US counties. To show this, this paper used an instrumental variable strategy to establish causality for the dynamic and long-run effect of the initial (2006–09) housing shock on future regional outcomes. Counties that had a larger loss in housing net worth in that period had more depressed employment and output as late as 2018. In addition, the local housing boom-bust cycle had asymmetric effects with little local output or employment effect in the boom phase but very persistent employment, GDP, and population losses during the bust. The effect of the housing crisis was well-characterized as mostly operating through the demand side since there is no significant change in labor productivity and a persistent impact on non-tradable employment.

Interestingly, there is only a temporary impact on measures of labor market slack, such as the employment-to-population ratio. Moreover, the negative housing shock had a comparatively short-lived impact on house prices and household leverage, lending credence to its temporary nature. On the labor market adjustment to these scarring effects on employment, our analysis finds no role for wage adjustment. In fact, we find indications that downward wage rigidity may have played a role since wages did increase marginally with the housing boom but did not react at all to the housing bust. Together, those findings imply that local labor market adjustment took place entirely through population movements, for which we provide direct evidence.

Our results suggest that future work leveraging regional US data to understand macroeconomic responses to temporary shocks might consider modeling labor movements explicitly since those constitute an adjustment mechanism that is at work at the local level but is not available at the national level. It also calls attention to asymmetric local effects of aggregate shocks, possibly due to downward wage rigidity. Importantly, it shows that those shocks can have have very persistent effects and as such, their distributive and allocative implications might be of interest for further analysis. Relatedly, as the world economy faces another large scale shock in the form of a pandemic with strong consumption demand effects, our results suggest that the most affected places could change in a permanent way.

References

Albouy, D., 2008. Are big cities bad places to live? Estimating quality of life across metropolitan areas. Technical Report. National Bureau of Economic Research.

Amior, M., Manning, A., 2018. The persistence of local joblessness. American Economic Review 108, 1942–70.

Andrews, I., 2016. Conditional linear combination tests for weakly identified models. Econometrica 84, 2155–2182.

Andrews, I., 2018. Valid two-step identification-robust confidence sets for gmm. The Review of Economics and Statistics 100, 337–348.

Anzoategui, D., Comin, D., Gertler, M., Martinez, J., 2019. Endogenous technology adoption and r&d as sources of business cycle persistence. American Economic Journal: Macroeconomics 11, 67–110.

Beraja, M., Hurst, E., Ospina, J., 2019. The aggregate implications of regional business cycles. Econometrica 87, 1789–1833.

Berger, D., Guerrieri, V., Lorenzoni, G., Vavra, J., 2017. House prices and consumer spending. The Review of Economic Studies 85, 1502–42.

Blanchard, O.J., Katz, L.F., 1992. Regional evolutions. Brookings Papers on Economic Activity 1, 1–75.

Chamberlain, G., Rothschild, M., 1982. Arbitrage, factor structure, and mean-variance analysis on large asset markets. Technical Report. National Bureau of Economic Research.

Charles, K.K., Hurst, E., Notowidigdo, M.J., 2018. Housing booms, manufacturing decline and labour market outcomes. The Economic Journal 129, 209–48.

Coibion, O., Gorodnichenko, Y., Ulate, M., 2017. The cyclical sensitivity in estimates of potential output. Technical Report. National Bureau of Economic Research.

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

Davis, S.J., Haltiwanger, J.C., 2019. Dynamism diminished: The role of housing markets and credit conditions. Technical Report. National Bureau of Economic Research.

Davis, S.J., Loungani, P., Mahidhara, R., 1997. Regional labor fluctuations: oil shocks, military spending, and other driving forces. International Finance Discussion Papers 578. Board of Governors of the Federal Reserve System.
Dix-Carneiro, R., Kovak, B.K., 2017. Trade liberalization and regional dynamics. American Economic Review 107, 2908–46.
Dupor, B., Karabarbounis, M., Kudlyak, M., Mehkari, M.S., 2018. Regional consumption responses and the aggregate fiscal multiplier. Working Paper Series 2018-4. Federal Reserve Bank of San Francisco.
Fernald, J.G., Hall, R.E., Stock, J.H., Watson, M.W., 2017. The disappointing recovery of output after 2009. Technical Report. National Bureau of Economic Research.
Fort, T.C., Haltiwanger, J., Jarmin, R.S., Miranda, J., 2013. How firms respond to business cycles: The role of firm age and firm size. IMF Economic Review 61, 520–559.
Gao, Z., Sockin, M., Xiong, W., 2016. Economic consequences of housing speculation. Technical Report. Working paper.
Gertler, M., Gilchrist, S., 2018. What happened: Financial factors in the great recession. Journal of Economic Perspectives 32, 3–30.
Gilchrist, S., Siemer, M., Zakrajšek, E., 2018. The real effects of credit booms and busts. Technical Report. Working Paper.
Greenstone, M., Hornbeck, R., Moretti, E., 2010. Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. Journal of Political Economy 118, 536–598.
Grigsby, J., Hurst, E., Yildirmaz, A., 2019. Aggregate nominal wage adjustments: New evidence from administrative payroll data. Technical Report. National Bureau of Economic Research.
Guerrieri, V., Lorenzoni, G., 2017. Credit crises, precautionary savings, and the liquidity trap. The Quarterly Journal of Economics 132, 1427–67.
Guren, A.M., McKay, A., Nakamura, E., Steinsson, J., 2018. Housing wealth effects: The long view. Technical Report. National Bureau of Economic Research.
Jones, C., Midrigan, V., Philippon, T., 2018. Household leverage and the recession. IMF Working Papers 18/194. International Monetary Fund.
Jordà, Ò., 2005. Estimation and inference of impulse responses by local projections. American Economic Review 95, 161–182.
Jordà, Ò., Singh, S.R., Taylor, A.M., 2020. The long-run effects of monetary policy. Working Paper 26666. National Bureau of Economic Research.
Justiniano, A., Primiceri, G.E., Tambalotti, A., 2015. Household leveraging and deleveraging. Review of Economic Dynamics 18, 3–20.
Katz, L.F., Murphy, K.M., 1992. Changes in relative wages, 1963-1987: Supply and demand factors. The Quarterly Journal of Economics 107, 35–78.
Mian, A., Rao, K., Sufi, A., 2013. Household balance sheets, consumption, and the economic slump. The Quarterly Journal of Economics 128, 1687–1726.
Mian, A., Sufi, A., 2014. What explains the 2007–2009 drop in employment? Econometrica 82, 2197–2223.
Reinhardt, C., Rogoff, K., 2009. This time is different: six hundred years of financial folly. Cambridge MA: Harvard University Press.
Saiz, A., 2010. The geographic determinants of housing supply. The Quarterly Journal of Economics 125, 1253–96.
Schmitt-Grohé, S., Uribe, M., 2016. Downward nominal wage rigidity, currency pegs, and involuntary unemployment. Journal of Political Economy 124, 1466–1514.
Stroebel, J., Vavra, J., 2014. House prices, local demand, and retail prices. Technical Report. National Bureau of Economic Research.
Sun, L., 2018. Implementing valid two-step identification-robust confidence sets for linear instrumental-variables models. The Stata Journal 18, 803–825.
Yagan, D., 2019. Employment hysteresis from the great recession. Journal of Political Economy 127, 2505–58.

Online Appendix for Local Scars of the US Housing Crisis*

Appendix A. Data Construction

Appendix A.1. Outcome Variables and Housing Net Worth

1. **Employment, Unemployment, Wages, and Population**
   (a) **QCEW county-level employment**
   - QCEW monthly employment data represent the number of covered workers who worked during, or received pay for, the pay period that included the 12th day of the month. We use annual averages of county-level employment data.
   - Sample period 1990–2018
     - Main analysis: 2006–09(18) changes in employment
     - Control for pre-trends: 1994–98 and 1998–2002 changes in employment
   - 5 sectoral employment from NAICS 2-digit industry classification
     - Tradable / Nontradable / Construction / High-skilled service sectors / Others
     - NAICS 2-digit QCEW codes are in Appendix A.3.
   - Industry controls (employment share controls)
     - NAICS 2-digit QCEW sectoral employment shares of private employment (23 industries)

   (b) **QCEW wages data**
   - QCEW wages data represent the total compensation paid during the calendar quarter regardless of when the services were performed.
   - We use annual average wages in each county.

   (c) **BLS Local Area Unemployment Statistics**

*The views expressed here are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond, the Federal Reserve Board, or the Federal Reserve System. First version: Dec 2018. This version: Feb 2021.

- The Local Area Unemployment Statistics (LAUS) produces monthly and annual employment, unemployment, and labor force data for counties.
- We use annual average unemployment rate and employment-to population ratio in each county.

(d) **Quarterly Workforce Indicators**

- The Quarterly Workforce Indicators (QWI) provide local labor market statistics by industry, worker demographics, employer age and size.
- We use annual average of beginning of quarter employment and annual average of monthly earnings of employees who worked at the beginning of the reference quarter in each county.
- We use QWI data from 1998 through 2018 because many states had participated in QWI program after 1998.

(e) **ACS Employment and Adjusted Hourly Wages Data**

- To construct adjusted wage data, we use data from the 2000 census and the 2001-14 American Community Surveys (ACS). Following Beraja et al. (2019), we calculate hourly wages for prime-age males by restricting our sample to only males ages 25-54, who live outside of group quarters, have no self-employment income, and who are not in the military. We calculate the hours worked by multiplying weeks worked last year and usual hours worked per week. We divide wage and salary income by the hours worked to calculate the hourly wages for each individual. We exclude any individual with a zero wage and truncate the measured wage distribution at the top and bottom one percent.

  We adjust the hourly wages by creating a composition-adjusted wage measure following Katz and Murphy (1992). We divide our sample into six age bins (25-29, 30-34, 35-39, 40-44, 45-49, 50-54) and four education bins (completed years of schooling < 12, $= 12$, between 12 and 16, and 16 and more). We then adjust the wage index by averaging over those wages for 24 groups with fixed weights to calculate the wage for different educational and age groups within each geographic unit and estimate an adjusted wage index by averaging over those wages with fixed weights. We use the share of each demographic group in each geographic level during 2005 as the fixed weights.

- To construct an ACS employment measure, we restrict our sample to people (both male and female) who live outside of group quarters.

(f) **Population**

- US Census Bureau Annual County Resident Population Estimates (from 2000-2016)
- For pre-2000, use Census US Intercensal County Population Data, 1970-2014 from NBER (http://www.nber.org/data/census-intercensal-county-population.html)
- Use 15-64 population by each county

(g) We exclude Orleans Parish county from our sample since employment and population in the county decreased by more than 50% in 2006 due to Hurricane Katrina.

2. **GDP and Income**

(a) **BEA Local Gross Domestic Product**

- GDP by county is the value of goods and services produced by the county's economy less the value of goods and services used up in their production. It is the substate counterpart of the nation's GDP. GDP by county statistics are also the foundation for metropolitan and micropolitan GDP statistics.
- Sample period 2001–18
  - Main analysis: 2006–09(18) changes in GDP
  - Control for prior trends: We use 1998–2002 growth rates in BEA real personal income as prior trends controls for GDP regressions. Also, we use 1998–2002 growth rates of BEA personal income per employee as prior trends controls for GDP per employee regressions.
- Five sectoral GDP from NAICS 2-digit industry classification
  - Tradable / Nontradable / Construction / High-skilled service sectors / Others

- For sectoral GDP regressions, we use 2002–2006 growth rates of sector's GDP as prior trend controls.

(b) **BEA Personal Income by County, Metro, and Other Areas**

- Personal income for an area is the income received by, or on behalf of all persons resident in the area, regardless of the duration of residence, except for foreign nationals employed by their home governments in the United States. Personal income can be defined as the sum of wages and salaries, supplements to wages and salaries, proprietors' income, dividends, interest, and rent, and personal current transfer receipts, less contributions for government social insurance.
- Sample period 1990–2018
  - Real personal income data are defined as personal income divided by state-level BEA GDP deflator.
  - Main analysis: 2006–09(18) changes in personal income
  - Control for pre-trends: 1998-2002 changes in real personal income

3. **House Price Data**

- We use county-level CoreLogic's HPI data as a measure of house prices. We divide the CoreLogic's HPI data by BEA personal income to construct HPI-to-income ratio.

4. **Housing Net Worth**

(a) We use the measure of housing net worth shocks constructed by Mian and Sufi (2014). Below is the brief description of how they construct the housing net worth shocks in Mian and Sufi (2014).

(b) "One of our key right-hand-side variables is the change in household net worth between the end of 2006 and 2009. We define net worth for households living in county i at time t as $NW_{it} = S_{it} + B_{it} + H_{it} - D_{it}$, where the four terms on the right hand side represent market values of stocks, bonds, housing, and debt owed, respectively. We compute the market value of stock and bond holdings (including deposits) in a given county using IRS Statistics of Income (SOI) data.

We estimate the value of housing stock owned by households in a county using the 2000 Decennial Census data as the product of the number of homeowners and the median home value. We then project the housing value into later years using the CoreLogic zip code level house price index and an estimate of the change in home ownership and population growth. Finally, we measure debt using data from Equifax Predictive Services that tells us the total borrowing by households in each county in a given year." (Mian and Sufi (2014) p. 2200.)

Appendix A.2. Control Variables

1. **Industry Employment Shares**
   - Using 2002 QCEW 2-digit level industry data, we define each industry's employment share as the ratio of employment in each industry to total number of private employment in 2002
   - A list of 20 industries is in Appendix A.3.
2. **Debt-to-Income**
   - Compute DTI at different geographical levels using data on household debt from the Equifax/Federal Reserve Bank of New York Consumer Credit Panel (CCP) made available as part of the extended Financial Accounts of the United States on the Federal Reserve Board of Governors website and the data on household income from the Bureau of Labor Statistics (BLS). At the time of writing, the Equifax/FRB NY CCP data was available at the source link: `https://www.federalreserve.gov/releases/z1/dataviz/household_debt/county/map/#state:all;year:2018`
   - Calculate DTI as the ratio of aggregate household debt from Equifax (excluding student loans) to aggregate income (from BLS).
     - Calculate aggregate household debt by summing individual household debt in the CCP within each geographical area and multiplying by the sampling ratio.

  - Use data from the BLS, which reports income earned by workers covered by unemployment insurance programs overseen by the Department of Labor. Income is reported quarterly and aggregated to annual amounts for each geographic region, including counties, CBSAs, and states.

3. **Quality of life data by Albouy (2008)**
   - Table A.1. in http://davidalbouy.net/PDF/improvingqol.pdf

4. **Amenities index (Natural amenities scale)**
   - https://www.ers.usda.gov/data-products/natural-amenities-scale/

5. **2000 housing wealth to total wages**
   - We calculate the housing wealth for each county by multiplying each county's median home value and total number of home owners from Census 2000. Then, we divide it by 2000 QCEW total wages to calculate the housing wealth to wages ratio.

6. **Davidoff (2016) controls**
   - Fraction of the population that had education greater than or equal to 4 years of college
   - Fraction of the population that were born outside the U.S.
   - "Bartik" measure that approximates local demand pressure based on national industrial employment growth
   - Density measure which is housing units divided by land area
   - A geographical dummy variable, "Coastal" (metropolitan areas with at least one county adjacent to the Pacific Ocean in California, Oregon, or Washington; or stops on the Acela line)
   - Replication files are available in the author's webpage (https://sites.google.com/site/tomdavidoff/)

7. **Sensitivity to Aggregate Shocks**
   (a) Local sensitivity to monetary and financial shocks

i. To calculate local sensitivity to monetary and financial shocks, we use quarterly QCEW employment data from 1990 through 2002. We separately regress each county's quarterly employment growth rate on monetary and excess bond premium shocks. Then, we define the coefficients on the both shocks from each county regression as the county's sensitivity to monetary and financial shocks.

ii. We use an identified monetary shock series constructed by Romer and Romer (2004). excess bond premium shocks constructed by Gilchrist and Zakrajšek (2012).

(b) Local sensitivity to other macroeconomic shocks

i. We construct county-level 10-year growth employment rates ($g_{i,t}$) using annual QCEW employment data from 1988 through 2002. Then, we define a normalized employment growth rate ($g_{i,t}^N$) as the deviation of $g_{i,t}$ from its average over time ($\bar{g}_i$), that is, $g_{i,t}^N = \frac{g_{i,t} - \bar{g}_i}{sd(g_i)}$, where $sd(g_i)$ is the standard deviation of county i's growth rate from its time average.

ii. We do a factor analysis using these county-level normalized employment growth rates and use loadings of the three main factors for each county as controls.

Appendix A.3. Industry Categorization

- Tradable sector:
  - NAICS 11 Agriculture, forestry, fishing and hunting
  - NAICS 21 Mining, quarrying, and oil and gas extraction
  - NAICS 31-33 Manufacturing
- Nontradable sector:
  - NAICS 44-45 Retail trade
  - NAICS 72 Accommodation and food services
- Construction sector:
  - NAICS 23 Construction
  - NAICS 53 Real estate and rental and leasing

- High-skilled services sector:
  - NAICS 51 Information
  - NAICS 52 Finance and insurance
  - NAICS 54 Professional and technical services
  - NAICS 55 Management of companies and enterprises
  - NAICS 56 Administrative and waste services
  - NAICS 61 Educational services
  - NAICS 62 Health care and social assistance
- Others:
  - NAICS 22 Utilities
  - NAICS 42 Wholesale trade
  - NAICS 48-49 Transportation and warehousing
  - NAICS 71 Arts, entertainment, and recreation
  - NAICS 81 Other services, except public administration
  - NAICS 92 Public administration

Appendix A.4. Panel VAR

The instrument we construct identifies the housing bubble as a large increase in house prices during the 2002–05 period that cannot be attributed to fundamentals. Our approach is to use the panel VAR with the Cholesky decomposition to identify a housing price shock that is orthogonal to general business conditions in each county.

We first run a panel-VAR at the annual county level from 1975 through 2006 with CoreLogic's county-level house prices, QCEW employment (total and in construction), BEA personal income per employee, the number of 15-64 population, and QCEW wages per employee in construction. We use three-year changes of those six variables in the panel-VAR analysis. Notice that QCEW industry-level data are available in SIC from 1975 through 2000 and in NAICS from 1990 onward. We use construction employment and wages data from 1975 through 1990 in SIC and from 2001 through 2006. Then we take an average of employment and wages between SIC data and NAICS data from 1991 through 2000 to construct historical data.

We use a STATA package `pvar2` used in Fort et al. (2013). They modify a package `pvar` developed by Abrigo and Love (2016). We use three lags for the panel-VAR estimation.

We calculate the innovation for the house price index that is orthogonal to innovations to these other variables. Finally, we designate as instruments for the 2006–09 housing crash, a dummy variable for the orthogonalized house price residuals from 2002–05 that are in the bottom tercile of the distribution.

Appendix B. Appendix Tables and Figures

Appendix Table B.1: P-Value for J-Statistic

Dependent variable:	2002	2003	2004	2005	2007	2008	2009	2010
- Employment	0.59	0.90	0.66	0.72	0.03	0.56	0.19	0.54
- Wages per Employee	0.96	0.15	0.96	0.62	0.11	0.03	0.66	0.22
Dependent variable:	2011	2012	2013	2014	2015	2016	2017	2018
- Employment	0.28	0.34	0.24	0.30	0.24	0.41	0.21	0.61
- Wages per Employee	0.65	0.31	0.44	0.43	0.95	0.76	0.55	0.61

Notes: This table shows p-values for J-statistics for employment and wages per employee for each year from our baseline specification. All control variables listed in Table 1 in the paper are included. Prior trends for employment and wages per employee are the average growth rates in those variables from 1994–98 and from 1998–2002. Sample weights (by the number of households) are applied to all specifications.

Appendix Table B.2: First-Stage F-Statistic (Kleibergen and Paap F-statistic)

	Second-Stage Dependent Variable: Employment			Second-Stage Dependent Variable: Wages Per Employee		
Year	(1) Both Instruments	(2) Housing Elasticity IV	(3) Panel VAR Shock IV	(4) Both Instruments	(5) Housing Elasticity IV	(6) Panel VAR Shock IV
2002	6.95	7.32	11.41	5.50	5.41	10.9
2003	6.55	3.77	12.39	7.59	3.68	14.99
2004	6.60	3.81	12.29	5.49	6.17	12.94
2005	6.61	4.42	12.38	5.85	3.93	11.47
2007	11.63	4.87	23.61	5.53	2.98	10.9
2008	6.60	5.03	12.45	6.29	7.36	11.41
2009	6.19	5.58	12.19	5.01	4.53	9.97
2010	7.97	6.50	13.69	5.53	3.64	10.98
2011	6.39	6.68	11.58	5.35	2.53	10.12
2012	8.73	6.77	15.02	5.44	3.39	10.86
2013	6.72	5.89	12.42	6.00	3.54	11.46
2014	8.64	6.32	14.65	5.44	3.38	10.76
2015	6.85	4.74	12.43	5.64	3.26	11.02
2016	8.47	6.24	14.44	5.56	3.00	10.99
2017	6.59	4.12	13.04	5.79	3.65	11.19
2018	8.13	6.34	13.99	5.43	2.93	10.51

Notes: This table shows first-stage F-statistics for employment and wages per employee regressions for each year. First-stage regressions depend on pre-trend controls, which are different for different second-stage dependent variables. Columns (1) and (4) are F-statistics from the baseline two instruments. Columns (2) and (5) are from housing elasticity instrument and Columns (3) and (6) are from panel VAR shocks instrument. All control variables listed in Table 1 in the paper are included. Prior trends for employment and wages per employee are the average growth rates in those variables from 1994–98 and from 1998–2002. Sample weights (by the number of households) are applied to all specifications.

Appendix Table B.3: First-Stage Regressions

Second-Stage Dependent Variable	Both Instruments: (1) Coefficient on Housing Elasticity	Both Instruments: (2) Coefficient on Panel VAR Shocks	(3) Housing Elasticity IV	(4) Panel VAR Shock IV
Employment	0.024	0.032	0.024	0.034
	(0.011)	(0.008)	(0.010)	(0.009)
Wages Per Employee	0.020	0.033	0.022	0.034
	(0.012)	(0.010)	(0.013)	(0.010)

Notes: This table shows the first-stage regressions results for employment and wages per employee for 2018. First-stage regressions depend on pre-trend controls, which are different for different second-stage dependent variables. Columns (1) and (2) shows the results using the baseline two instruments. Column (3) shows the results using housing elasticity as an instrument and Column (4) shows the results using panel VAR shocks as an instrument. All control variables listed in Table 1 in the paper are included. Prior trends for employment and wages per employee are the average growth rates in those variables from 1994–98 and from 1998–2002. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are reported in parentheses.

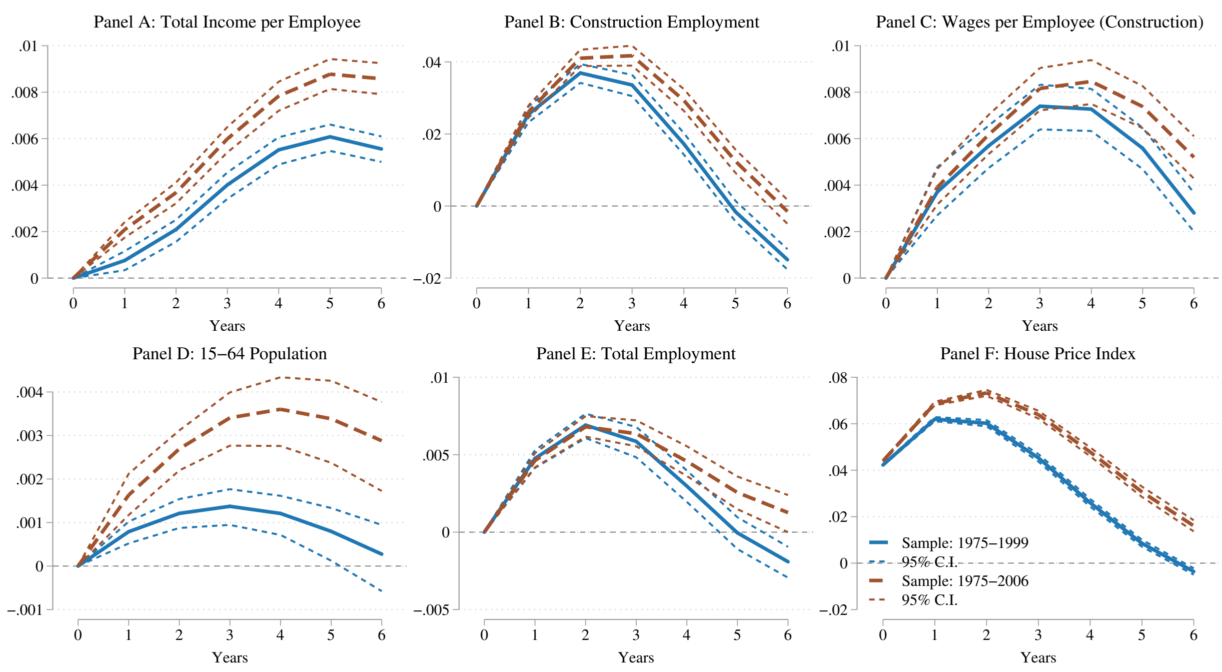

Notes: The figure plots the impulse responses of variables to orthogonalized HPI shocks in the panel VAR regressions. Blue lines are the results from sample period 1975–1999 and red lines are the results from sample period 1975–2006. Dashed lines are 95% confidence intervals.

Appendix Figure B.1: Impulse Response Functions to HPI Shocks in the Panel-VAR

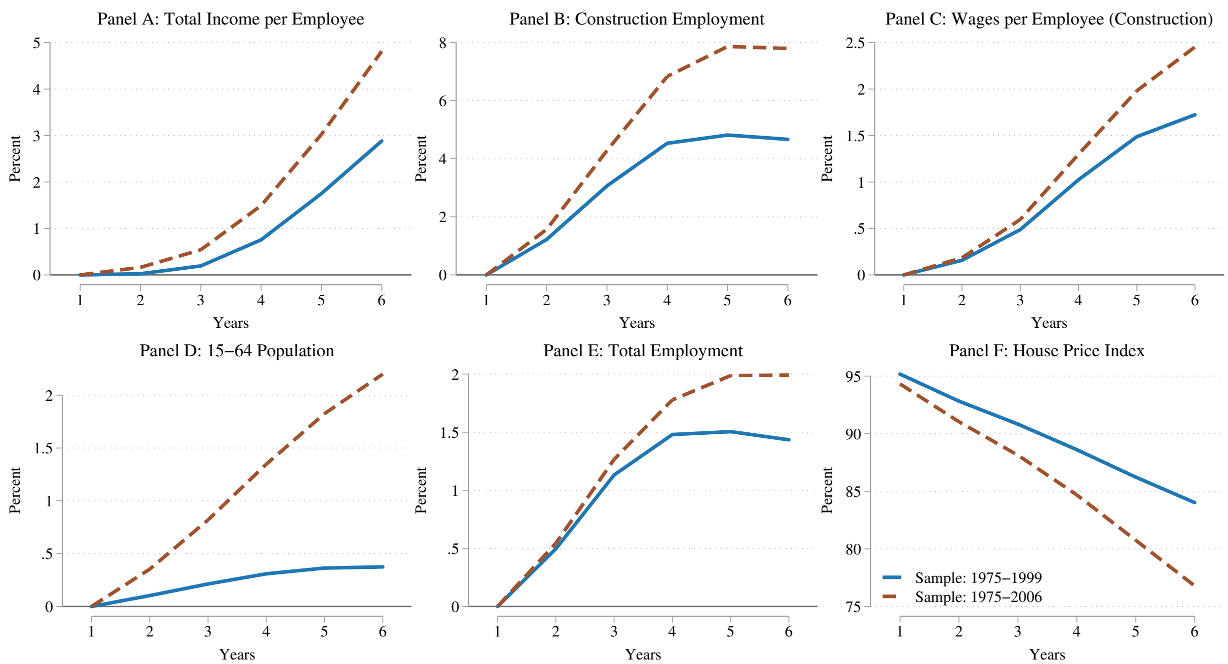

Notes: The figure plots forecast error variance decomposition of variables for HPI shocks. Blue lines are the results from sample period 1975–1999 and red lines are the results from sample period 1975–2006.

Appendix Figure B.2: Forecast Error Variance Decomposition for HPI Shocks in the Panel-VAR

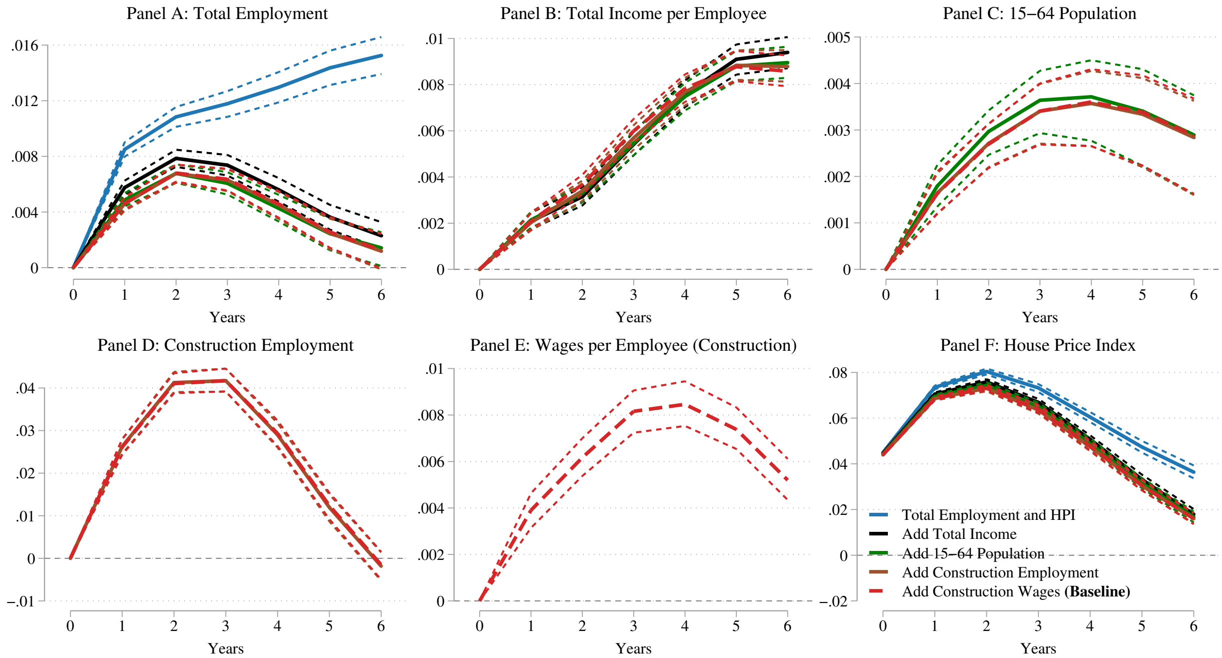

Notes: The figure plots the impulse responses of variables to orthogonalized HPI shocks in the panel VAR regressions. Blue lines are the results from the panel VAR using total employment and HPI variables. From the two variables, we add total income (black lines), 15-64 population (green lines), construction employment (brown lines), and construction wages (red lines). Dashed lines are 95% confidence intervals.

Appendix Figure B.3: Impulse Response Functions to HPI Shocks with Different Variables in the Panel-VAR

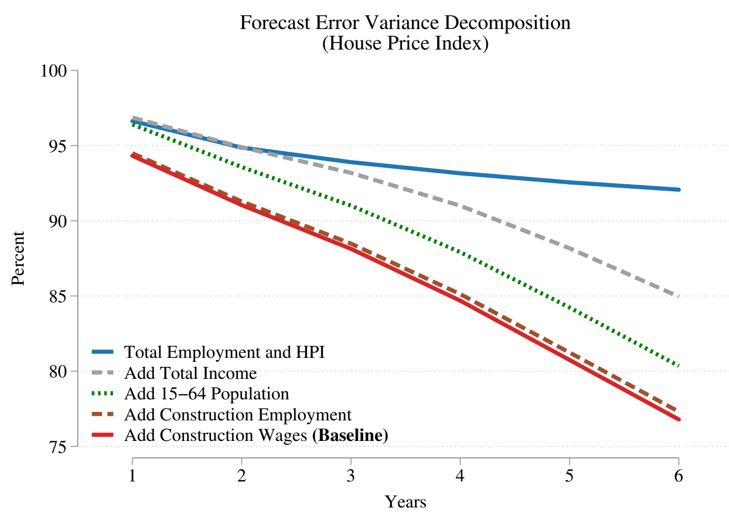

Notes: This figure shows forecast error variance decomposition of house prices for HPI shocks. Blue line are the results from the panel VAR using total employment and HPI variables. From the two variables, we add total income (gray line), 15-64 population (green line), construction employment (brown line), and construction wages (red line). Dashed lines are 95% confidence intervals.

Appendix Figure B.4: Forecast Error Variance Decomposition of House Prices accounted by HPI shocks with Different Variables in the Panel-VAR

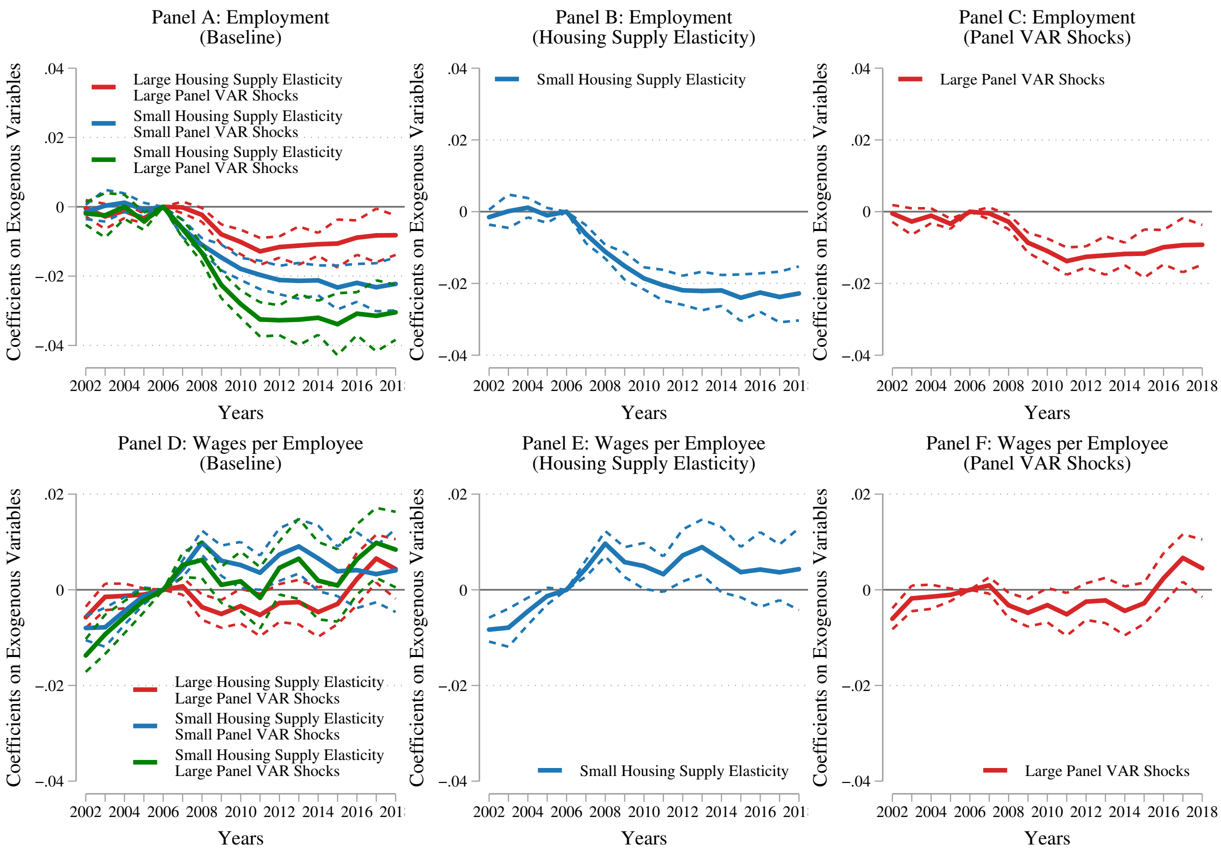

Notes: The figure plots the reduced-form regression results for employment and wages per employee. Panels A and D reproduces Figure 2 in the paper, which use the baseline two instrument. Panels B and E are the reduced-form results using only housing elasticity as an instrument. Panels C and F are the reduced-form results using only panel VAR shocks as an instrument. Dashed lines are one standard deviation confidence intervals. All control variables listed in Table 1 in the paper are included. Prior trends for employment and wages per employee are the average growth rates in those variables from 1994–98 and from 1998–2002. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals.

Appendix Figure B.5: Reduced-Form Regressions

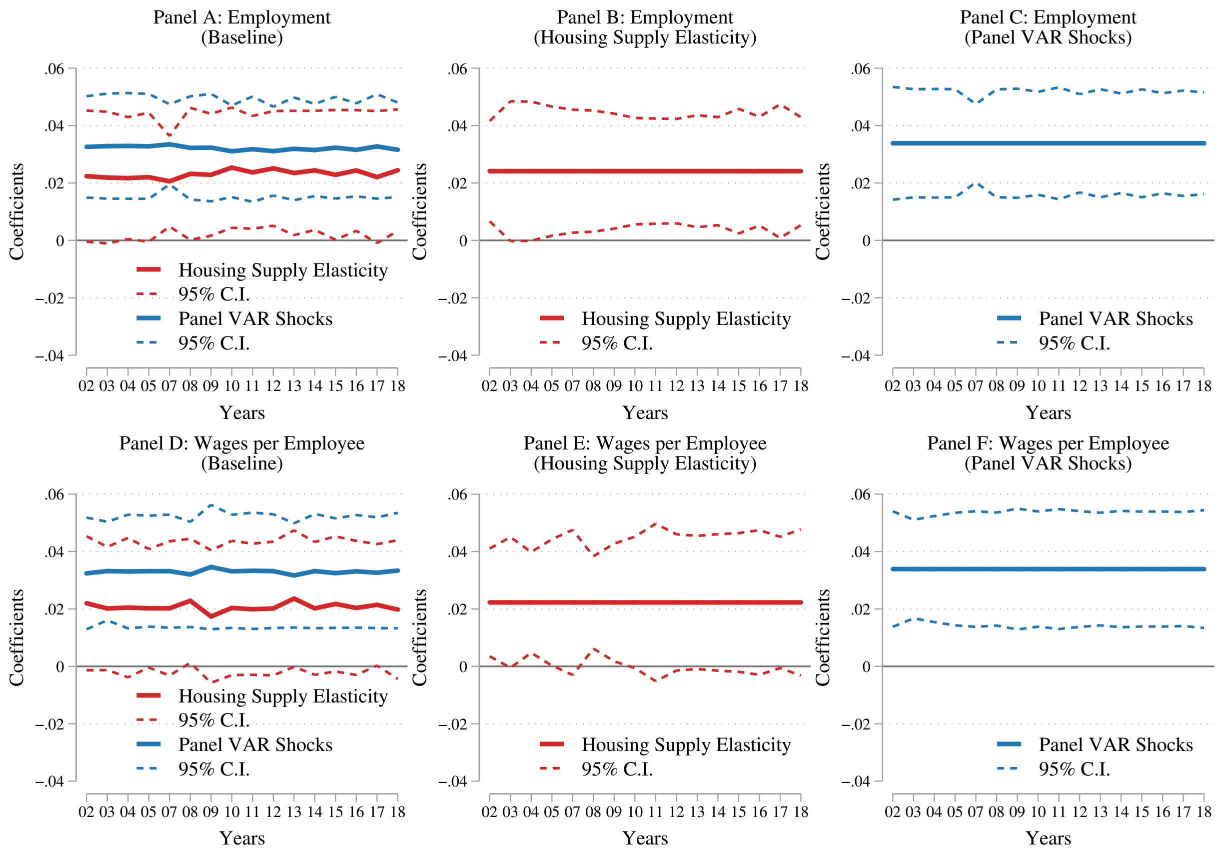

Notes: This figure shows the first-stage regressions results for employment and wages per employee for each year. Panels A and D reproduces Figure 2 in the paper, which uses the baseline two instrument. Panels B and E are the reduced-form results using only housing elasticity as an instrument. Panels C and F are the reduced-form results using only panel-VAR shocks as an instrument. All control variables listed in Table 1 in the paper are included. Prior trends for employment and wages per employee are the average growth rates in those variables from 1994–98 and from 1998–2002. Sample weights (by the number of households) are applied to all specifications. Dashed lines are one standard deviation confidence intervals. Robust standard errors (clustered by state) are used to calculate the confidence intervals.

Appendix Figure B.6: First-Stage Regressions for Each Year

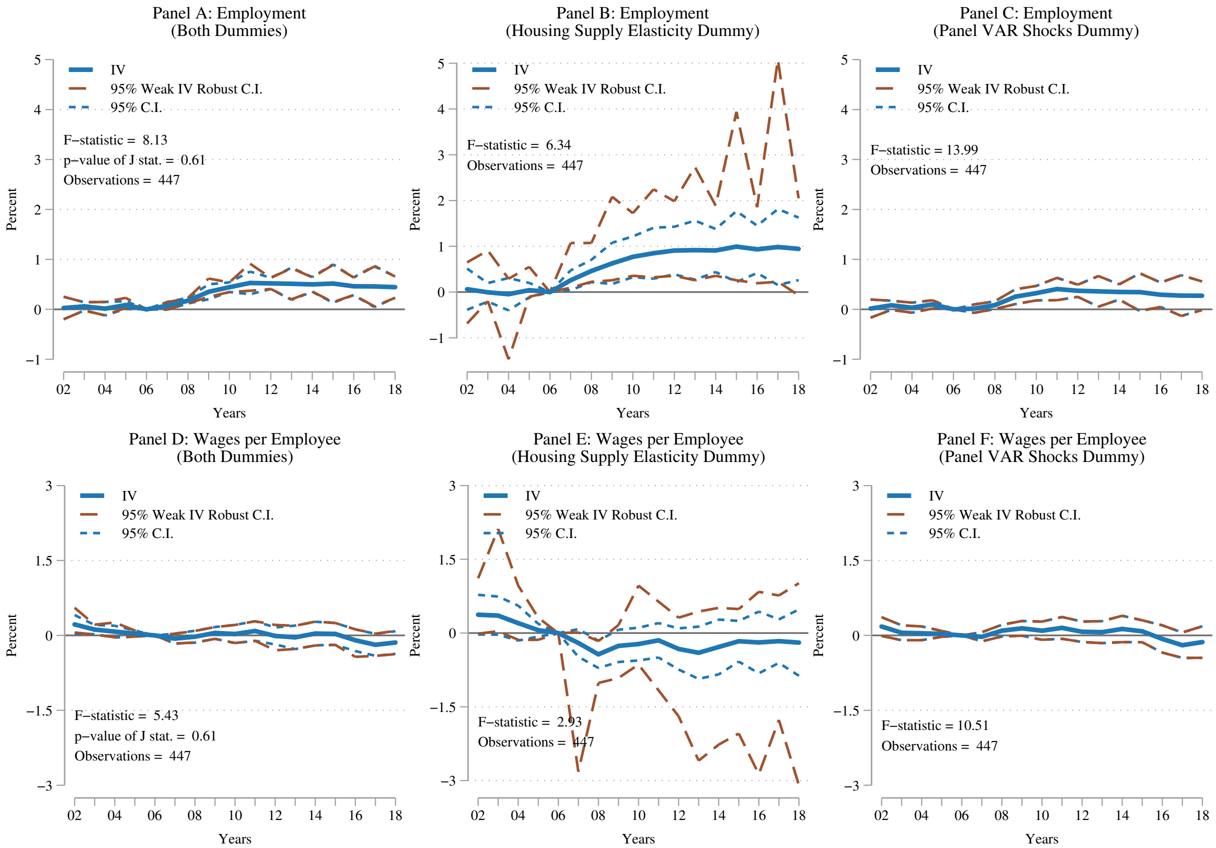

Notes: The figure plots the impulse responses of employment and wages per employee to the 2006–09 housing shocks with different instruments. Panels A and D use the baseline instrumental variables which are a dummy for upper tercile of housing elasticity and a dummy for lower tercile of panel VAR orthogonalized shocks. Panels B and E use only the dummy for upper tercile of housing elasticity and Panels C and F use only the dummy for lower tercile of panel VAR orthogonalized shocks. All control variables listed in Table 1 in the paper are included. Prior trends for employment and wages per employee are the average growth rates in those variables from 1994–98 and from 1998–2002. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Appendix Figure B.7: Changes in Employment and Wages per Employee with Different Instruments

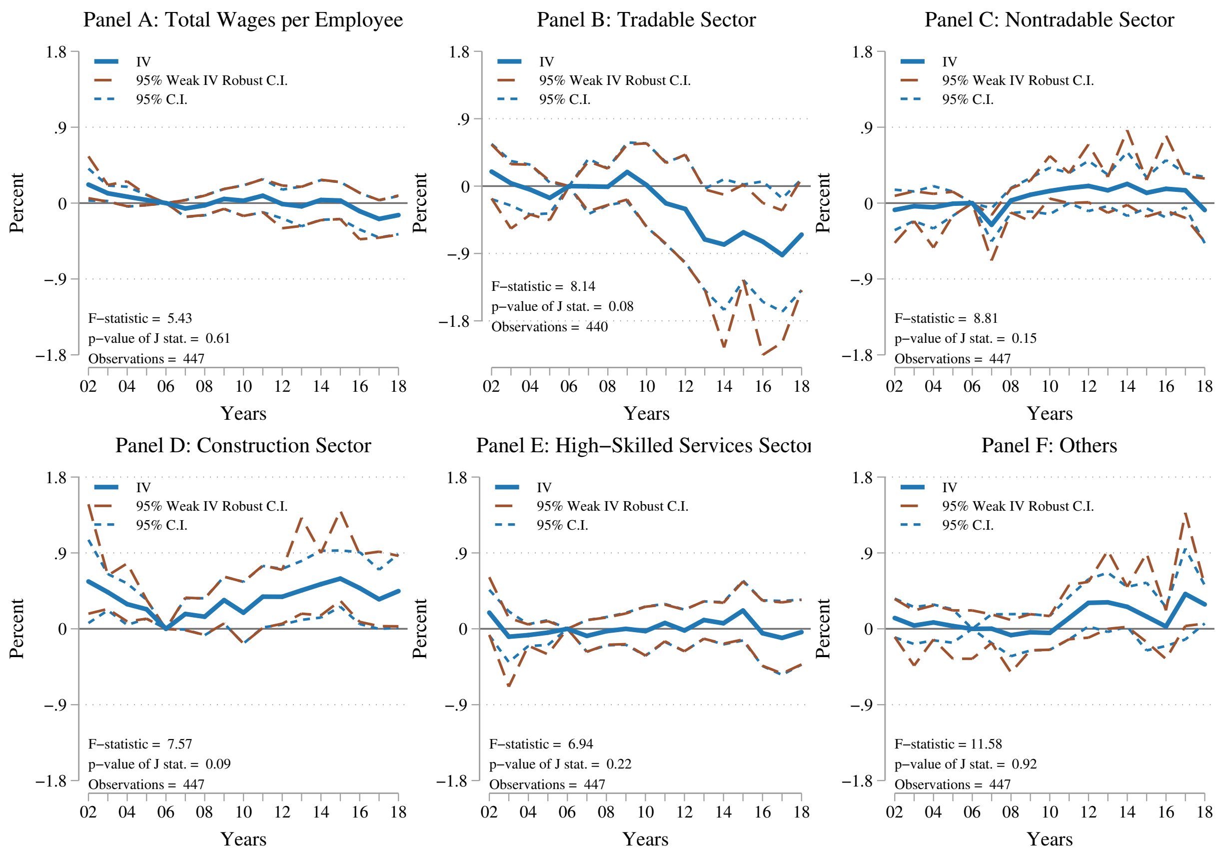

Notes: The figure plots the impulse responses of wages per employee to the 2006–09 housing shocks by sectors. All the results are from IV estimations. All control variables listed in Table 1 in the paper are included. Prior trends for sectoral wages per employee are the average growth rates in wages per employee in each sector from 1994–98 and from 1998–2002. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions. See Appendix A.3 for the details of sectoral splits.

Appendix Figure B.8: Changes in Wages per Employee by Sector

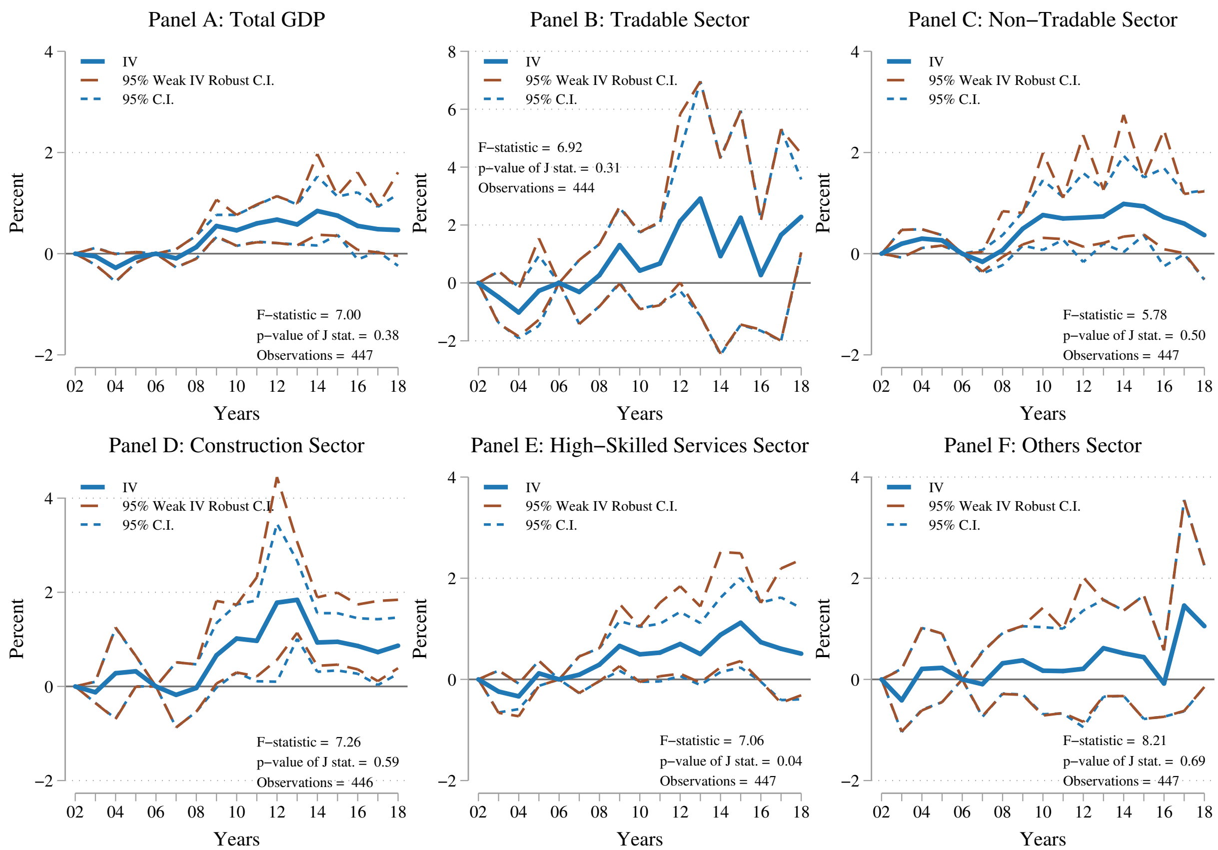

Notes: The figure plots the impulse responses of GDP to the 2006–09 housing shocks by sectors. All the results are from IV estimations. All control variables listed in Table 1 in the paper are included. Prior trends for sectoral GDP are the average growth rates in GDP in each sector from 2002-2006. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions. See Appendix A.3 for the details of sectoral splits.

Appendix Figure B.9: Changes in GDP by Sector

Notes: The figure plots the impulse responses of employment (Panels A and B) and wages per employee (Panels C and D) to the 2006–09 housing shocks. All control variables listed in Table 1 in the paper are included. Three prior trends are included: the average growth rates in outcome variables from 1990–94, 1994–98, and 1998–2002. Sample weights (by the number of households) are applied to all specifications. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Appendix Figure B.10: Changes in Employment and Wages per Employee with 1990–1994 Prior Trends

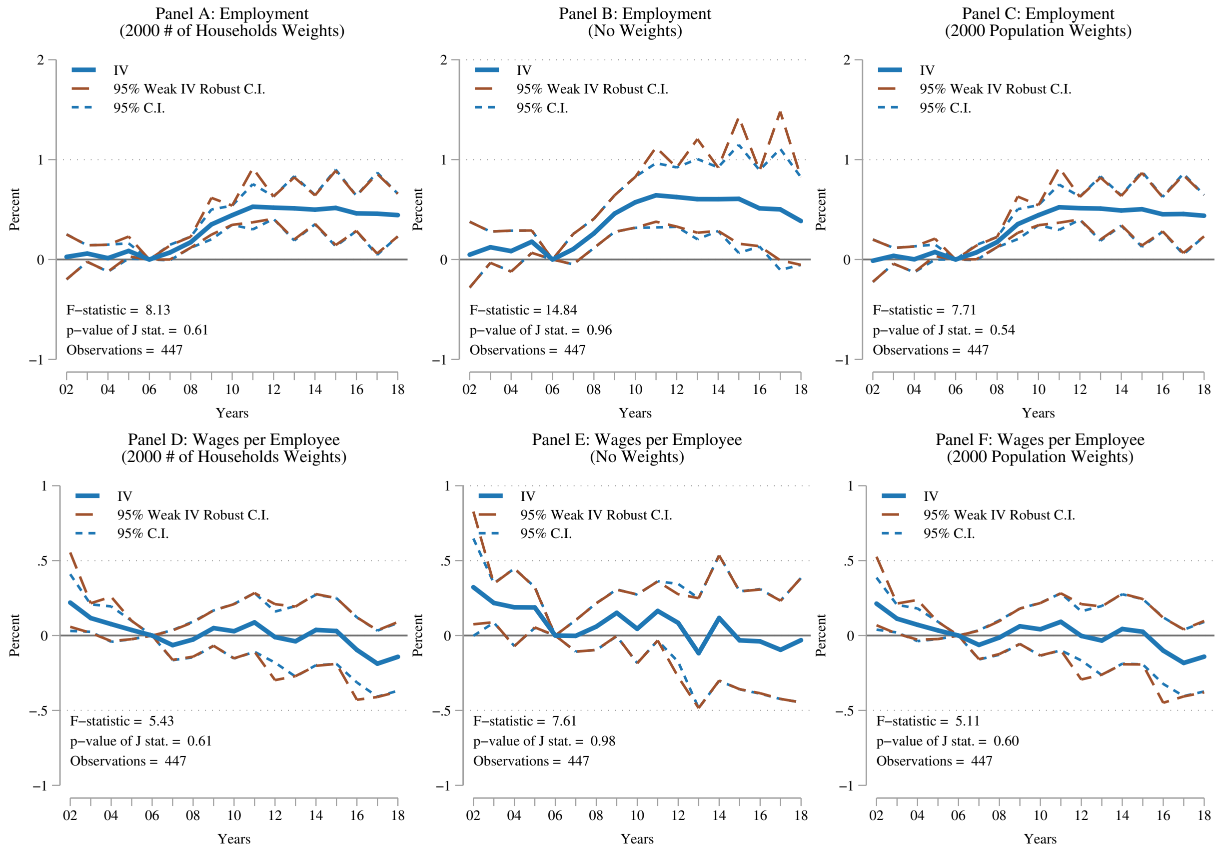

Notes: The figure plots the impulse responses of total employment and wages per employee to the 2006–09 housing shocks with and without weightings instruments. Panels A and D apply sample weights by 2000 number of households. Panels B and E are results without weighting. Panels C and F apply sample weights by 2000 15-64 population. All control variables listed in Table 1 in the paper are included. Prior trends for employment and wages per employee are the average growth rates in those variables from 1994–98 and from 1998–2002. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Appendix Figure B.11: Changes in Employment and Wages per Employee with and without Weighting

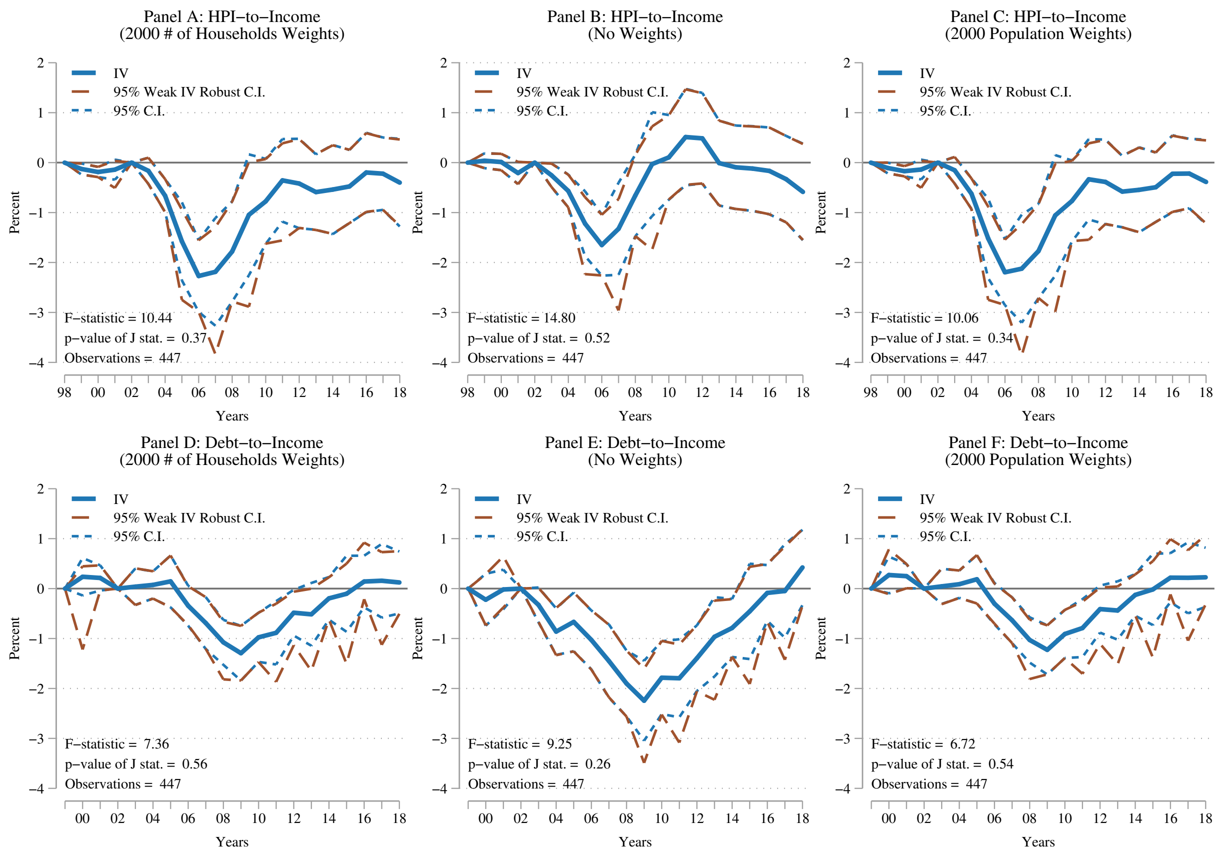

Notes: The figure plots the impulse responses of HPI-to-income ratio and debt-to-income ratio to the 2006–09 housing shocks with and without weightings instruments. Panels A and D apply sample weights by 2000 number of households. Panels B and E are results without weighting. Panels C and F apply sample weights by 2000 15-64 population. All control variables listed in Table 1 in the paper are included. Prior trends for HPI-to-income ratio are the average growth rates from 1994–98 and from 1998–2002, while prior trends for debt-to-income ratio are the average growth rate from 1999-2002. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions

Appendix Figure B.12: Changes in HPI-to-Income and Debt-to-Income with and without Weighting

Notes: The figure plots the impulse responses of QCEW wages per employee at the county level (Panels A and B) and hourly wages at PUMA level using adjusted ACS data (Panels C and D) to the 2006–09 housing shocks. The adjustment procedure for ACS data follows Beraja, Hurst and Ospina (2019) and is described in Appendix A. The left columns are results from OLS estimations, and the right columns are results from IV estimations. All control variables listed in Table 1 in the paper are included for QCEW wages per employee regressions while we exclude a set of controls (prior trends, quality of life index, natural amenities scale, and Davidoff (2016) controls) for the ACS wages regressions due to data limitation. Prior trends for wages per employee are the average growth rates from 1994–98 and from 1998–2002. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Appendix Figure B.13: Changes in QCEW Wages per Employee and ACS Adjusted Hourly Wages

Notes: The figure plots the impulse responses of employment to the 2006–09 housing shocks by education and age groups using ACS data at PUMA level. Panel A shows results from the group with less than a college degree, while Panel B shows results from the group with a bachelor's degree or more. Panel C shows results from the group with ages 25-40, while Panel D shows results from the group with ages from 41-55. All the results are from IV estimations. All control variables listed in Table 1 in the paper are included, except for a set of controls (prior trends, quality of life index, natural amenities scale, and Davidoff (2016) controls) due to data limitation. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Appendix Figure B.14: Changes in ACS Employment by Education and Age Groups

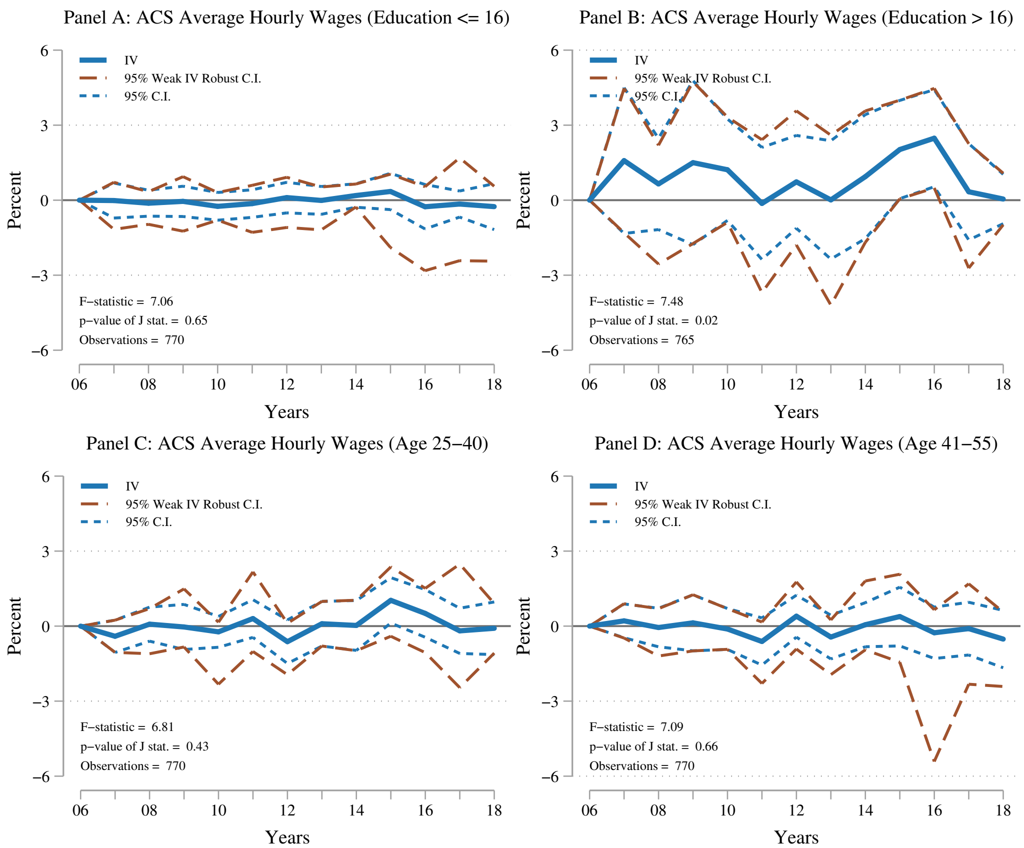

Notes: The figure plots the impulse responses of ACS adjusted hourly wages to the 2006–09 housing shocks by education and age groups. Panel A shows results from the group with less than a college degree, while Panel B shows results from the group a with bachelor's degree or more. Panel C shows results from the group with ages 25-40, while Panel D shows results from the group with ages from 41-55. All the results are from IV estimations. All control variables listed in Table 1 in the paper are included, except for a set of controls (prior trends, quality of life index, natural amenities scale, and Davidoff (2016) controls) due to data limitation. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Appendix Figure B.15: Changes in ACS Hourly Wages by Education and Age Groups

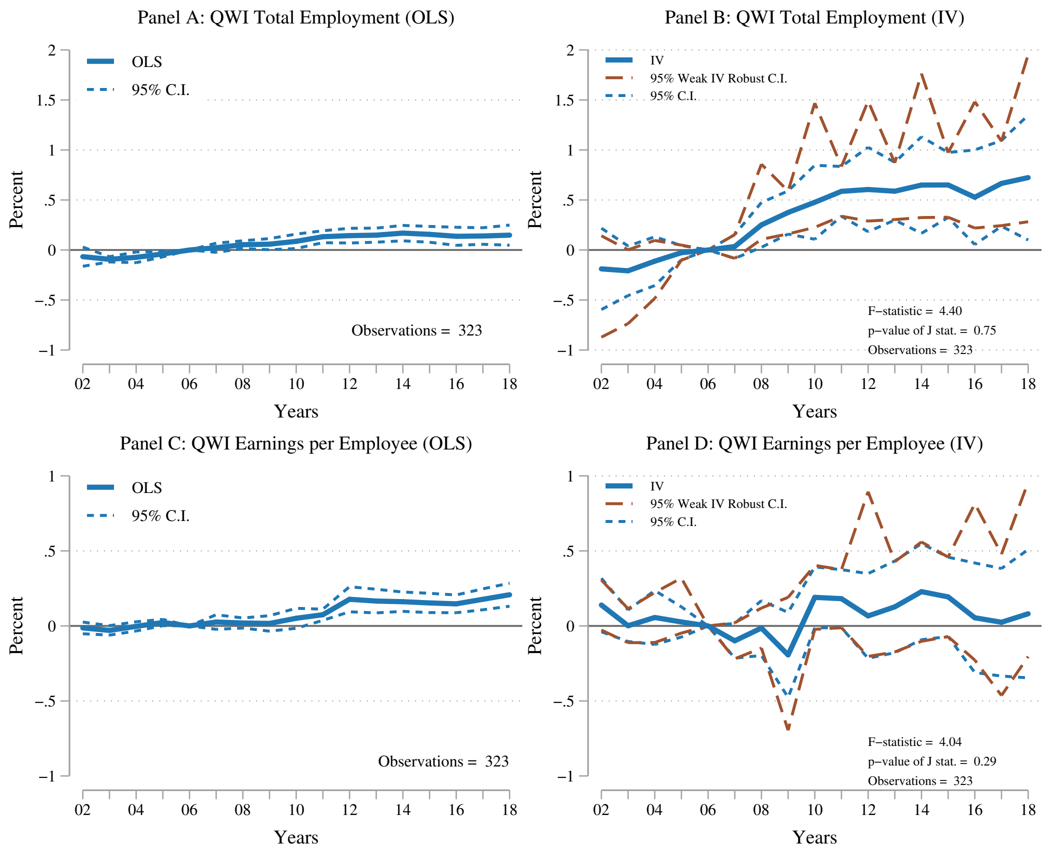

Notes: The figure plots the impulse responses of employment (Panels A and B) and earnings per employee (Panels C and D) to the 2006–09 housing shocks using QWI data. The left columns are results from OLS estimations and the right columns are results from IV estimations. All the results are from IV estimations. All control variables listed in Table 1 in the paper are included. Prior trends are the average growth rates in outcome variables from 1998–2002. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Appendix Figure B.16: Changes in QWI Employment and Earnings per Employee by Year

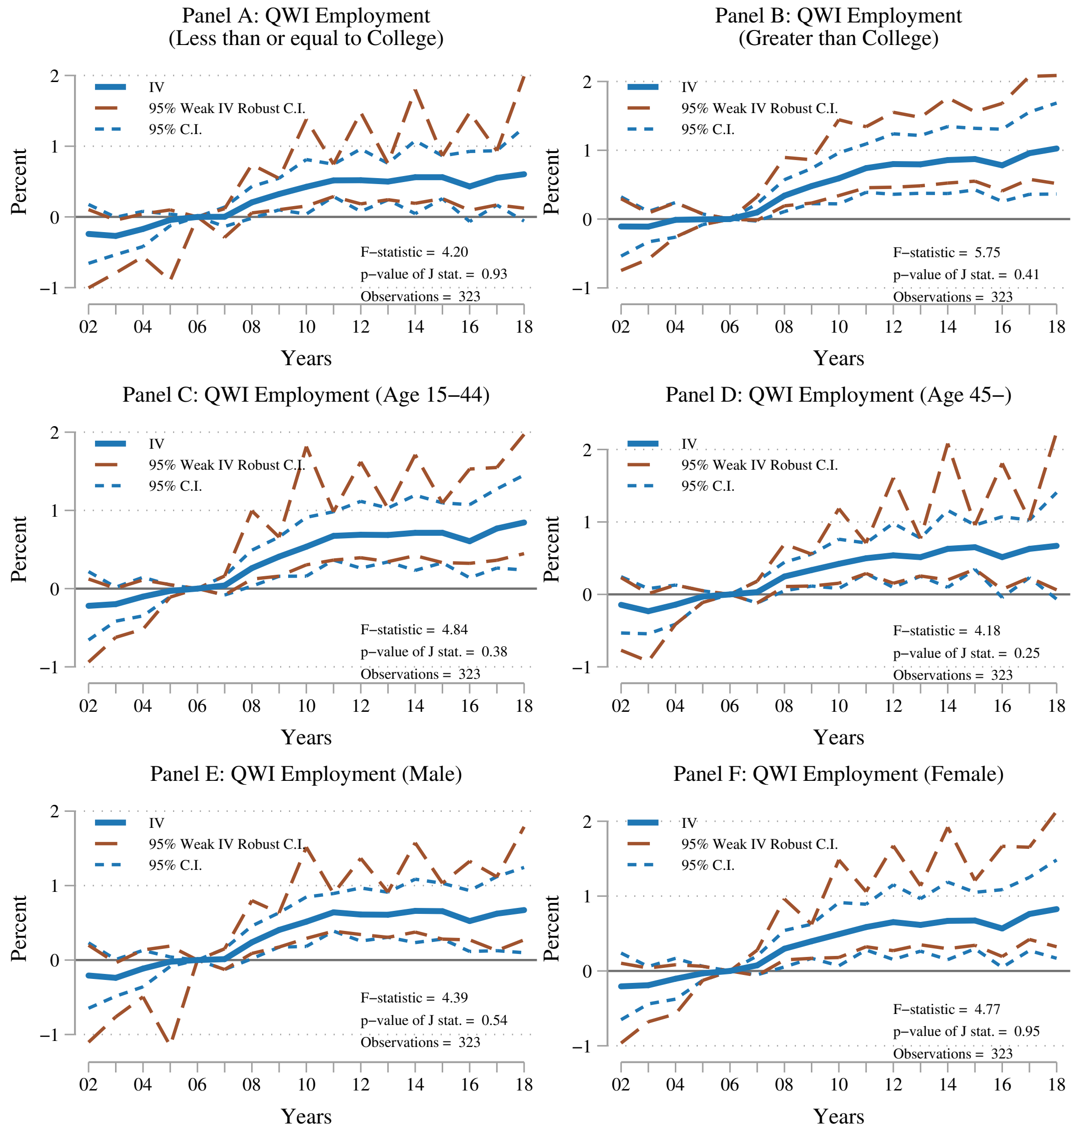

Notes: The figure plots the impulse responses of employment to the 2006–09 housing shocks by workers' age and education groups using QWI data. Panel A shows results from the group with a college degree or less, while Panel B shows results from the group with more than a bachelor's degree. Panel C shows results from the group with ages 15-44 while Panel D shows results from the group of ages 45-plus. Panel E shows results from the group of males, and Panel F shows results from the group of females. All the results are from IV estimations. All control variables listed in Table 1 in the paper are included. Prior trends are the average growth rates in outcome variables from 1998–2002. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Appendix Figure B.17: Changes in QWI Employment by Education and Age Groups

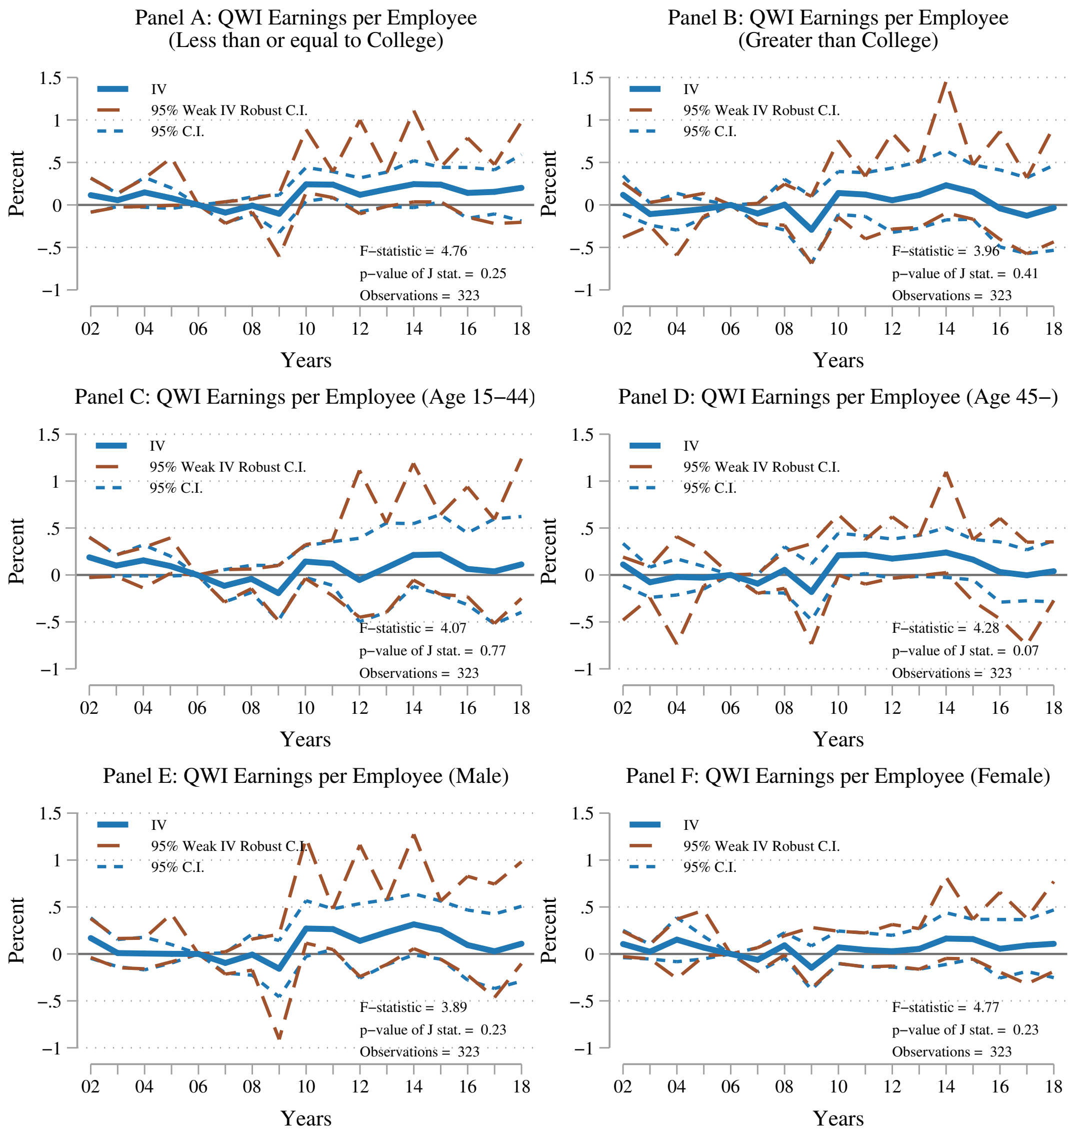

Notes: The figure plots the impulse responses of earnings per employee to the 2006–09 housing shocks by workers' age and education groups using QWI data. Panel A shows results from the group with a college degree or less, while Panel B shows results from the group with more than a bachelor's degree. Panel C shows results from the group with ages 15-44 while Panel D shows results from the group of ages 45-plus. Panel E shows results from the group of males, and Panel F shows results from the group of females. All control variables listed in Table 1 in the paper are included. Prior trends are the average growth rates in outcome variables from 1998–2002. Sample weights (by the number of households) are applied to all specifications. Robust standard errors (clustered by state) are used to calculate the confidence intervals. Red lines are weak IV robust confidence intervals. F-statistics and p-values for J-statistics are from 2018 regressions.

Appendix Figure B.18: Changes in QWI Earnings per Employee by Education and Age Groups

References

Abrigo, M.R.M., Love, I., 2016. Estimation of panel vector autoregression in stata. The Stata Journal 16, 778–804.

Albouy, D., 2008. Are big cities bad places to live? Estimating quality of life across metropolitan areas. Technical Report. National Bureau of Economic Research.

Beraja, M., Hurst, E., Ospina, J., 2019. The aggregate implications of regional business cycles. Econometrica 87, 1789–1833.

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

Fort, T.C., Haltiwanger, J., Jarmin, R.S., Miranda, J., 2013. How firms respond to business cycles: The role of firm age and firm size. IMF Economic Review 61, 520–559.

Gilchrist, S., Zakrajšek, E., 2012. Credit spreads and business cycle fluctuations. American Economic Review 102, 1692–1720.

Katz, L.F., Murphy, K.M., 1992. Changes in relative wages, 1963-1987: Supply and demand factors. The Quarterly Journal of Economics 107, 35–78.

Mian, A., Sufi, A., 2014. What explains the 2007–2009 drop in employment? Econometrica 82, 2197–2223.

Romer, C.D., Romer, D.H., 2004. A new measure of monetary shocks: Derivation and implications. American Economic Review 94, 1055–84.