HOUSING BOOMS, LABOR MARKET OUTCOMES, AND EDUCATIONAL ATTAINMENT*

Kerwin Kofi Charles† Erik Hurst‡ Matthew J. Notowidigdo§

April 2014

Abstract

We assess the extent to which the recent housing boom and bust affected employment, wages, and college enrollment and attainment during the 2000s. We exploit cross-city variation in local housing booms, and we identify plausibly exogenous variation in housing demand using sharp structural breaks in local housing prices. We find that positive housing demand shocks significantly increased wages and employment between 2000 and 2007, particularly for less-skilled workers. Consistent with the idea that the housing boom increased the opportunity cost of college for workers on the margin of college attendance, housing demand shocks during the boom reduced college enrollment and attainment for both young men and women, with the effects concentrated at community colleges. Over the longer time horizon spanning the housing boom and bust, we find that the positive wage and employment effects of the boom were generally undone during the bust. However, the negative effects of the housing boom on schooling persist, suggesting that reduced educational attainment may be an enduring effect of the large, temporary increase in housing demand. (JEL J24, I21, E24)

* This paper contains some results that appeared in a previously distributed working paper entitled "Manufacturing Decline, Housing Booms, and Non-Employment" (NBER Working Paper #18949), which has since been split into two separate projects. The remaining results from the working paper now appear in the paper "Masking and Aggregate Employment Changes: Housing Booms and Manufacturing Declines During the 2000s." We thank seminar participants at the UC Berkeley, Columbia, Duke, Harvard, Maryland, Northwestern, Princeton, Tulane, University of British Columbia, University of Chicago, UC Irvine, the University of Houston, University of Illinois at Chicago, the AEA, Einaudi Institute, the NBER Summer Institute (Macro Perspectives), Yale SOM, NYU, and the Atlanta, Chicago, Cleveland, and New York Federal Reserves for helpful comments. We are also grateful to Edward Glaeser, Tom Davidoff, Michael Lovenheim, and Matthew Gentzkow for useful suggestions and to Chris Mayer for sharing his data. We thank David Toniatti and Dan Zangri for excellent research assistance. We gratefully acknowledge the Initiative on Global Markets at the University of Chicago Booth School of Business for financial support. Hurst thanks the Peter Wall Institute for Advanced Studies at the University of British Columbia, and Notowidigo thanks the Einaudi Institute for both financial support and hospitality while working on this project.

† kerwin.charles@gmail.com, Harris School, University of Chicago and NBER
‡ erik.hurst@chicagobooth.edu, University of Chicago Booth School of Business and NBER
§ noto@chicagobooth.edu, University of Chicago Booth School of Business and NBER

I. INTRODUCTION

As Figure 1 shows, between 1997 and 2007, after decades of very little movement, national housing prices surged by about 50 percent in real terms, before abruptly and completely collapsing by 2010.[1] A growing literature studies the effect on this historically unprecedented boom and bust on various outcomes, including household consumption (Mian and Sufi 2011; Mian, Rao and Sufi 2013) and household defaults (Gerardi et al. 2008; Campbell and Cocco 2014), but there has been less work examining how the housing cycle affected labor market and related outcomes. Moreover, previous studies have mainly focused on the consequences of the collapse in prices after 2007, with little attention paid to the effect of the preceding boom.[2] This paper studies how the housing market boom and bust during the first decade of the 2000s affected workers' employment and wage outcomes, and their propensity to acquire schooling. We separately study both the boom and bust periods, as well as the longer time interval spanning the entirety of the housing cycle.

There is considerable suggestive evidence that changes in the housing market affected aggregate labor demand, and therefore workers' equilibrium labor market outcomes. Using data from the Current Population Survey (CPS), the first panel of Figure 2 shows that among persons aged 21-55 (henceforth, "prime-aged"), the fraction working in construction closely followed the time series in national housing prices, rising by 20 percent during the boom (from 5.5% to 6.6%), then falling by 28 percent (6.6% to 4.7%) from 2007 to 2011. The second panel shows the fraction of the population employed in construction in four sex×education groups: men and women, with and without at least a four-year college degree (henceforth, "college" and "non-college"). Virtually the entire employment shift into construction during the housing boom occurred among non-college men, for whom the share employed in construction increased massively from 12.8% to 15.6% during the boom, and collapsed from 15.6% to 11.3% over the bust. Since changes in household wealth caused by variation in housing prices may have also affected employment in other sectors, the patterns shown for construction employment may understate the overall employment effect from housing market changes. Figure 3 shows the trend in median wages for non-college men – a group that the aggregate construction patterns suggest experienced especially pronounced labor demand changes from the housing cycle. The figure shows that 25 years of consistent decline in

[1] The house price data used in this paper are from the Federal Housing Finance Agency (FHFA), which is a repeat-sales index of house prices across 363 MSAs. The patterns shown in Figure 1 are quite similar to those in other price series, like the Case-Shiller measure, which samples a much smaller set of MSAs.

[2] See, for example, Mian and Sufi (2012) for how the housing bust reduced employment in non-construction, non-tradeable sectors.

median real wages for these men abruptly stopped almost exactly at the start of the housing boom, and that median wages thereafter were either flat, or for a few years during the housing boom, even slightly *increasing*.

These aggregate trends point to a possible mechanism linking the housing market cycle to recent puzzling changes in schooling attainment. After increasing steadily for nearly two decades, the growth in the share of the young adult population having attended college abruptly slowed for women and halted entirely for men in the U.S in the late 1990s (Goldin and Katz 2008). Figure 4 depicts the pattern in college attainment for young men, showing the sharp break from trend, right around the start of the housing boom. Yet, as far as we are aware, no research has systematically explored the possibility that improvements in the market returns for non-college work from the housing boom may have contributed to the slowdown in attainment, by raising the opportunity cost of college-going for workers on the margin of attending college.

Moving beyond the suggestive relationships presented above, this paper uses a local labor market estimation strategy that exploits variation across metropolitan areas (MSAs) in the size of the housing boom and bust to estimate the effect of housing market demand shocks on labor market and schooling outcomes for different groups of workers. Whereas most previous work on the housing boom has proxied for the size of local housing demand shocks using the change in local housing prices, we use a simple model of housing to create a new proxy for housing demand shocks. This index is a function of price changes and housing demand and supply elasticities, and in principle captures both the effect of housing demand shocks on price changes as well as changes in the quantity and quality of housing supplied.

The key empirical challenge we confront is isolating exogenous variation in housing demand that is not the result of latent factors that also affect employment, wages, and schooling choices. Our approach builds on the emerging consensus that much of the variation in housing prices during the boom and bust derived from a speculative "bubble" and not from changes in standard determinants of housing values such as income, population, or construction costs (Mayer 2011; Sinai 2012). We introduce a new instrumental variable that exploits structural breaks in the evolution in housing prices in an MSA, arguing that these "sharp" breaks are plausibly exogenous to latent confounds, such as labor supply shocks or unobserved changes in labor demand, which are likely smoothly incorporated into price changes.[3] We show that these breaks are, in fact,

[3] Our approach is similar in spirit to recent work which uses structural breaks to identify economic effects of interest, such as the work by Card, Mas, and Rothstein (2008) on racial tipping.

systematically unrelated to a large set of observable local characteristics, and we provide several pieces of evidence consistent with them being the result of speculative activity. We further show that these breaks explain an important portion of the overall change in housing demand over the first decade of the 2000s. We present results of the estimated effect of changes in housing demand from both OLS and Two Stage Least Squares (TSLS) models, where we instrument for changes in housing demand using the estimated structural break measure.

Using information from the Census and from multiple years of the American Community Survey (ACS), we find that positive shocks to housing demand in an MSA during the 2000-2007 boom increased employment and wages.[4] Among all prime-age workers, a one standard deviation (or roughly 80 percent) increase in housing demand in an MSA raised the employment rate by 0.8 percentage points and increased wages by 1.9 percent. Separate results for different sex×education groups show that the estimated wage and employment effects were largest for non-college men and smallest for college women. Virtually all of the increase in employment among non-college men came from increased construction employment. By contrast, for non-college women, who also experienced large employment gains from the boom, these gains were mostly in local non-tradables such as real estate and insurance. Over the *entire* 2000-2011 period spanning the boom and bust, we find little change in long-term employment from a local housing boom, since gains during 2000-2007 were offset by declines of approximately equal size during the bust. There is, however, some evidence that housing demand shocks during the boom had persistent positive effects on wages over the longer term (2000-2011), as collapsing demand during the bust did not result in significant wage declines – perhaps because of downward wage rigidity. We estimate similar quantitative effects from OLS and Two Stage Least Squares (TSLS) models, and our labor market results are very similar across a number of robustness analyses and extensions.

We next present schooling results across several complementary data sets. First, we use data from Census/ACS self-reports about schooling attainment. Because of the rich information on age and location in this survey data, we are able to estimate the share of the young adult population in an MSA reporting particular levels of college attainment over time, by birth cohort. Using an estimation framework similar to that for the labor market outcomes discussed above, we find that among both men and women aged 18 to 21 at the start of the housing boom, the fraction holding an associate's degree grew less in markets with particularly large housing booms. Strikingly, we do

[4] Additionally, we assess how the housing boom affected mobility across MSAs. MSAs experiencing a housing boom had larger population increases than non-housing boom MSAs. We also show that the employment propensities of migrants into the housing boom MSAs were similar to the employment propensities of the natives within those MSAs.

not find similar effects during the boom years either for bachelor's degree holding among 18-21 year olds, or for associate's degree holding among relatively older workers (30-33 year olds). The strongly statistically significant and relatively large results suggest that improving labor market opportunities during the boom decreased advanced schooling attainment precisely for those workers who might have been expected to be on the margin between obtaining associate's degree training and not going to college at all when the boom began. Nor does it appear that these effects represented merely a temporary *delay* in schooling attainment. We show that, in fact, attainment rates remain significantly lower years after the collapse in housing, but only for a specific set of young people: the generation making college-going decisions at the time of the housing boom, and who lived in markets with large housing booms. Neither other generations of young people from boom markets nor young people in markets that did not experience a boom exhibit reduced attainment over the long-term.

While the Census/ACS data allow us to conduct detailed cohort-level analysis, the educational attainment reports from which the Census/ACS numbers are constructed do not perfectly measure all relevant educational investments. For example, for those attending some college but not receiving a college degree, the Census/ACS cannot distinguish between community college attendance and attendance at a four-year college or university. Additionally, self-reported schooling attainment may also be measured with error. We therefore present a complementary set of results that overcomes both of these shortcomings using rich administrative data on college enrollments from the Integrated Postsecondary Education Data System (IPEDS). Using this data set and the same local labor markets strategy, we find that increased housing demand significantly reduced overall college enrollment for both men and women between 2000 and 2007. These estimates are driven almost entirely by enrollment at community colleges, junior colleges, and technical colleges. We find no evidence that housing booms affect college enrollment at four-year colleges and universities. The declines in community college enrollments in response to housing booms thus did not arise from upgrading to four-year colleges and universities, but instead represent declines in enrollment in any postsecondary institution. Applying our local labor market estimates nationally, we find that the housing boom can explain approximately 40 percent of the national slowdown in college enrollment growth among both men and women. As with the Census/ACS survey evidence, the IPEDS results show strongly persistent effects of the housing boom on enrollment. While we find some evidence of "catch-up" during the bust, it is small and incomplete. Taken together, our schooling results suggest that reduced educational attainment may be an enduring effect of the large, temporary increase in housing demand.

Existing theoretical work posits a link between labor market conditions and educational attainment (Mincer 1958; Becker 1964), and previous empirical work has found that different types of labor demand shocks reduce college enrollment and educational attainment (Black, McKinnish and Sanders 2005; Atkin 2012).[5] Our focus on shocks originating in the housing sector extends this line of research, as does the fact that we separately identify effects for different types of colleges and universities. Finally, because we estimate the *total* effect of changes in local housing demand on college attainment, our estimates capture both any wealth effects on schooling among home-owning families from changing housing prices (Lovenheim 2011), and the effects of changing labor market conditions on attainment among all persons in a market. Our empirical results imply that the causal pathway operating through housing wealth is overwhelmed by that operating through labor market conditions, which is consistent with the empirical work on oil booms and export shocks cited above.

The remainder of the paper proceeds as follows. We present the overview of the theoretical and empirical framework that will guide the analysis in Section II. Section III discusses the data and presents summary statistics. Section IV discusses the instrumental variable used in much the analysis. Section V presents the labor market results, including various robustness results. Section VI presents the schooling results. Section VIII concludes.

II. EMPIRICAL FRAMEWORK

The empirical analysis focuses on comparisons across metropolitan statistical areas (MSAs), which we treat as different local labor markets. We assume that each labor market, k, is characterized by labor demand and labor supply functions given, respectively, by:

$$\begin{aligned} D_k^L &= \alpha_k^{D,L} - \eta_k^{D,L} \log W_k \\ S_k^L &= \alpha_k^{S,L} + \eta_k^{S,L} \log W_k \end{aligned} \tag{1}$$

[5] Black, McKinnish, and Sanders (2005) construct plausibly exogenous variation in local labor market conditions by interacting variation in coal prices (during coal boom and bust) with pre-existing differences in coal reserves. Atkin (2012) finds that sectoral shocks arising from trade reform affects the distribution of schooling attainment in Mexico.

where W_k is the wage in the labor market, and $\eta_k^{D,L}$ and $\eta_k^{S,L}$ are the (semi-)elasticities of labor demand and labor supply. The terms $\alpha_k^{D,L}$ and $\alpha_k^{S,L}$ represent shifters to labor demand and labor supply. We assume that the total shock to labor demand in a market between any two periods, $\Delta\alpha_k^{D,L}$, can be expressed as

$$\Delta\alpha_k^{D,L} = \delta^H \Delta\omega_k^H + \delta^O \Delta\omega_k^O \tag{2}$$

where $\Delta\omega_k^H$ and $\Delta\omega_k^O$ represent labor demand shocks originating in the housing sector and all "other" sectors, respectively, and the parameters δ measure how responsive total labor demand in a market is to shocks originating in the two sectors.

Substituting equation (2) into the two expressions in (1) and then solving for the equilibrium shows that the change in observed employment, E, and log wages in a local market are, respectively:

$$\begin{aligned} \Delta E_k^* &= \left(\frac{\delta^H \eta_k^{S,L}}{\eta_k^{D,L} + \eta_k^{S,L}}\right)\Delta\omega_k^H + \left(\frac{\delta^O \eta_k^{S,L}}{\eta_k^{D,L} + \eta_k^{S,L}}\right)\Delta\omega_k^O + \frac{\eta_k^{D,L}}{\eta_k^{D,L} + \eta_k^{S,L}}\Delta\alpha_k^{S,L} \\ \Delta \log W_k^* &= \left(\frac{\delta^H}{\eta_k^{D,L} + \eta_k^{S,L}}\right)\Delta\omega_k^H + \left(\frac{\delta^O}{\eta_k^{D,L} + \eta_k^{S,L}}\right)\Delta\omega_k^O - \frac{1}{\eta_k^{D,L} + \eta_k^{S,L}}\Delta\alpha_k^{S,L} \end{aligned} \tag{3}$$

Positive labor demand shocks, arising from housing or "other" sectors, increase employment and wages, while positive shocks to labor supply increase employment but reduce wages. We seek to empirically estimate the effect of housing demand shocks on equilibrium employment and wages, which equation (3) shows are given, respectively, by the reduced form parameters $\beta_1^E = \delta^H \eta_k^{S,L} \big/ \left(\eta_k^{D,L} + \eta_k^{S,L}\right)$ and $\beta_1^W = \delta^H \big/ \left(\eta_k^{D,L} + \eta_k^{S,L}\right)$ These parameters depend on the degree to which housing demand shocks in a market affect overall labor demand and the demand and supply elasticities of labor in the market.

Natural empirical specifications based on equation (3) are

$$\Delta E_k^* = \beta_0^E + \beta_1^E \Delta\omega_k^H + \Gamma^E X_k + \underbrace{\left(\frac{\delta^O \eta_k^{S,L}}{\eta_k^{D,L}+\eta_k^{S,L}}\right)\Delta\omega_k^O + \frac{\eta_k^{D,L}}{\eta_k^{D,L}+\eta_k^{S,L}}\Delta\alpha_k^{S,L} + \upsilon_k^E}_{regression\ \ error}$$

$$\Delta \log W_k^* = \beta_0^W + \beta_1^W \Delta\omega_k^H + \Gamma^W X_k + \underbrace{\left(\frac{\delta^O}{\eta_k^{D,L}+\eta_k^{S,L}}\right)\Delta\omega_k^O - \frac{1}{\eta_k^{D,L}+\eta_k^{S,L}}\Delta\alpha_k^{S,K} + \upsilon_k^W}_{regression\ \ error} \tag{4}$$

where X_k is vector of observable controls for changes in local labor market activity, and the υ terms represent random sampling error. Importantly, (4) shows that the regression errors in these two estimating equations contain not only random error, but also two systematic latent components: demand shocks to "other" sectors in the local market and shocks to local labor supply.

Much of the previous literature has proxied for demand shocks originating in the housing sector, $\Delta\omega_k^H$, using simply the change in housing prices. One obvious limitation of this approach is that observed changes in prices do not accurately reflect changes in demand when the supply elasticity of housing in the local market is large. Indeed, if housing supply is perfectly elastic, large changes in housing demand will not be reflected in *any* change in the observed price of housing. We create an alternative proxy for housing demand shocks using a simple model of housing demand and supply. To derive the proxy, we assume that the log of housing demand and housing supply in a market are given by:

$$\begin{aligned}\log\left(H_k^D\right) &= \omega_k^H - \eta_k^{D,H}\log\left(P_k\right)\\ \log\left(H_k^S\right) &= \lambda_k^H + \eta_k^{S,H}\log\left(P_k\right)\end{aligned} \tag{5}$$

In (5) ω_k^H reflects, as before, factors that affect local housing demand and λ_k^H capture factors that affect local housing supply. P_k is the local housing price and $\eta_k^{D,H}$ and $\eta_k^{S,H}$ are the price elasticities of housing demand and supply, respectively. Log differentiating the equilibrium condition, $H_k^D\left(P_k\right) = H_k^S\left(P_k\right)$, the effect of a shock to housing demand is:

$$\Delta\omega_k^H = \eta_k^{D,H}\Delta P_k + \Delta H_k^S. \tag{6}$$

In general, a change in housing demand affects *both* the equilibrium price and the amount of housing units supplied in the market – each of which can affect local labor market outcomes. In particular, house price changes affect household wealth or liquidity and thus households' demand for goods and services produced in the local market (Mian and Sufi, 2012). In addition, changes in the amount (or quality) of housing necessarily entail changes in labor market outcomes through construction-related activity such as demolition, renovation, home improvements, or new construction. Our analysis does not disentangle the separate effects of household wealth and construction-related channels, but rather focuses on the combined effect of changes in housing demand.[6]

Recalling that $\eta_k^{S,H}$ in equation (5) is simply $\Delta H_k^S / \Delta P_k$, the effect of a *ceteris paribus* shock to housing demand may be written:

$$\Delta\omega_k^H = \eta_k^{D,H}\Delta P_k + \Delta H_k^S = \left(\eta_k^{D,H} + \eta_k^{S,H}\right)\Delta P_k \tag{7}$$

Equation (7) suggests that, so long as there are no shifts in the housing supply function, a natural proxy for the local change in housing demand is the elasticity-weighted change in prices. We denote this proxy $\widehat{\Delta\omega_k^H}$, and calculate it using changes in an MSA's housing prices, existing estimates of the local housing supply elasticity in each MSA, and estimates from the literature of the housing demand elasticity. In principle, $\widehat{\Delta\omega_k^H}$ captures *both* the price-related "wealth channel" (via $\eta_k^{D,H} / \Delta P_k$) and the resulting quantity-related "construction channel" (via ΔH_k^S). To see intuitively why our proxy is superior to just the change in the price, notice that if two MSAs experience the same change in house prices, our measure assigns a larger estimated housing demand change to the market with larger supply elasticity since a larger change in demand would have been necessary to generate the same observed change in price.

Having constructed this proxy, the simplest empirical strategy for estimating β_1^E and β_1^W would be to perform OLS regression of equation (4), with $\Delta\omega_k^H$ replaced by $\widehat{\Delta\omega_k^H}$. However, this approach may yield biased estimates of the parameters of interest, β_1^E and β_1^W, with the sign of the bias *a priori* ambiguous. One type of bias is positive endogeneity bias from the fact that the latent demand

[6] In the Online Appendix, we show that if the effects of housing prices and housing quantities on labor market outcomes are different in magnitude, then we will estimate a weighted average of the two effects, with the weights based on the price elasticities of housing demand and supply.

and supply shocks in the error term in (4) might be systematically correlated with the housing demand shocks experienced in a market. Additionally, there may be attenuation bias arising from the fact that existing estimates of housing supply and demand elasticity (from which $\widehat{\Delta\omega_k^H}$ is constructed) may be measured with error. To address both measurement error concerns and the potential endogeneity of $\widehat{\Delta\omega_k^H}$ in equation (4), we present Two Stage Least Squares (TSLS) estimates in addition to the baseline OLS results. Section IV describes the instrumental variable used in the TSLS analysis.

III. DATA AND SUMMARY STATISTICS

We study the years 2000-2011, a period that spans most of the years of the housing boom and all years of the bust. Although it is widely agreed that the boom began in the late 1990s, we study the 2000-2007 period because reliable information about the main outcome and control variables cannot be obtained before 2000 at a sufficiently fine level of geographic disaggregation for a representative sample of MSAs. We create a panel of MSAs using data from the 2000 Census, and from various years of the American Community Survey (ACS) individual-level and household-level extracts from the Integrated Public Use Microsamples (IPUMS) database (Ruggles et al., 2004). Restricting attention to persons living in metropolitan areas, we compute mean wages, overall employment shares, population shares employed in various occupations, and total population in each MSA. In 2000, these means are from the 2000 Census. For 2007, we pool ACS data from 2005 to 2007 to increase the precision of the MSA estimates. Similarly, we pool the 2009-2011 ACS for the MSA means for these variables in 2011, the end of the bust. Because of the large sample sizes, means can be reliably estimated for the separate sex $\times$ education groups we study (college/non-college and men/women). The primary sample from which outcome and control variables are drawn is the set of all non-institutionalized persons aged 21-55 in each MSA.

Data on local house prices are from the FHFA. We mapped the FHFA metro areas to the Census/ACS metro areas by hand.[7] To mirror the ACS data, we construct average house price growth between 2000 and the average of house price in the first quarter in 2005, 2006, and 2007. Similarly, when calculating house price changes between 2000 and 2011, we use the pooled FHFA data for 2009, 2010, and 2011. To compute $\widehat{\Delta\omega_k^H}$, our proxy for local housing demand change, we

[7] See the Online Appendix for details of this matching procedure.

combine information on price changes with estimates of housing supply and demand elasticities from the literature. MSA-specific housing supply elasticities come from Saiz (2010), who combines detailed topographic data and measures of land use regulation to construct the supply elasticity estimates. The measure of housing demand elasticity is the widely-used estimate derived by Polinsky and Ellwood (1979), whose calculation uses individual-level data on income, housing expenditures, and housing prices from thirty-one urban markets. We use Polinsky and Ellwood's preferred estimate of -0.7, which the main analysis treats as constant across MSAs.[8]

Table 1 reports summary statistics about housing market changes in the 237 MSAs with non-missing labor market and housing market data that constitute the main analysis sample. The first panel of the table shows that MSA house prices rose by 34.4 percent, on average, between 2000 and 2007. While prices increased across all markets during the boom, there was tremendous heterogeneity in this price growth: prices nearly doubled at the 90th percentile of the price change distribution across MSAs, but grew by only 3.5 percent at the 10th percentile. The second row shows the mean change during the boom in our housing demand proxy. On average, MSAs experienced an increase in housing demand of 85.1 percent during the boom. This is much larger than the growth in raw prices, but similar to what would be expected based on the expression in equation (7) and the fact that the average housing supply elasticity in the sample is roughly 2.5. Panel B shows that, on average, both estimated housing demand and prices fell dramatically during the 2007-2011 bust, with magnitudes roughly similar to the amount they had risen between 2000 and 2007: about 33 percent for average prices and approximately 84 percent for estimated housing demand.

Not surprisingly, given the approximate equivalence of the average 2000-2007 increases and average 2007-2011 declines, the bottom panel of the table shows that over the course of the entire boom and bust in housing, from 2000 to 2011, prices and estimated housing demand across all MSAs barely changed increasing only 1.5 and 1.3 percent on average, respectively. These relatively small long-term changes in price and estimated demand were true not only on average but also *within* MSAs. Figure 5, which plots an MSA's housing price reduction between 2007 and 2011

[8] Various robustness tests show that our results are not sensitive to other alternative specifications of the demand elasticity across markets. In Online Appendix Table OA.2, we present result from a series of robustness tests which show that our results are not sensitive to assuming reasonable alternative values for this measure. For example, the results are similar when we assume the housing demand elasticity is as low as -0.3 (among the lowest estimates in the literature) and as high as -1.9 (among the highest estimates in the literature). We also show that the results are similar when we assign to each MSA a demand elasticity drawn at random from a uniform distribution between -0.3 and -1.9.

against its price increase from 2000 to 2007 shows clearly that for the overwhelming majority of MSAs, price increases during the boom were nearly exactly offset by declines during the housing bust.[9]

IV. STRUCTURAL BREAK INSTRUMENTAL VARIABLE

As discussed above, some of the variation in the housing demand proxy in naïve OLS models is likely due to unobserved demand and supply shocks which may also affect labor market and education outcomes in a market. To account for the resulting endogeneity bias we estimate TSLS models that rely upon variation in housing demand that is arguably unrelated to these underlying latent fundamental factors.

The main motivation for the approach we employ comes from the active literature studying the large changes in housing prices during the recent boom and bust. Various explanations have been posited in the literature for dramatic housing price changes like those observed during the recent cycle might arise, including changes in national lending standards (Favilukis, Ludvigson, and Van Nieuwerburgh 2010), the growth of subprime mortgages (Mayer and Spence 2009), zoning rules (Glaeser, Gyourko, and Saiz 2008), and "fads" in beliefs about patterns of future prices (Burnside, Eichenbaum, and Rebelo 2011). In addition, an emerging body of work argues that, perhaps because of irrational exuberance (Shiller 2009) and the ability to use market products like interest-only mortgages (Barlevy and Fisher 2010), investors played a key role in driving price variation by virtue of speculative investment in over-valued housing assets until the bubble suddenly burst (Chinco and Mayer 2014). Many of these explanations, especially those related to investment behavior, suggest that part of recent price changes may reflect a housing "bubble": a sudden and extreme appreciation in housing prices than is larger than would be suggested by underlying fundamentals such as population, income, or productivity, followed by a sudden price decline that is also larger than would be suggested by fundamentals (Mayer 2011). Implicit in this definition is the presumption that underlying fundamental factors either do not change abruptly or else are smoothly incorporated into housing price series. This suggests that sharp breaks from the global trend in a market's price series represent variation that is thus arguably the result of the factors like speculative

[9] The percentage changes in house prices and housing demand for 2007-2011 are normalized by house prices in 2000, so that percentage increases and decreases of equal magnitude in the boom and bust imply zero net changes over the full 2000-2011 period. This similarly implies that cities located along the 45-degree line in Figure 5 experienced no net change in housing demand over the longer time horizon.

activity or the emergence of mortgage products and not the result of latent fundamental factors that are the major source of endogeneity concerns in OLS analysis of labor market and education outcomes.

Building on this insight, we implement a TSLS strategy, the intuition for which is illustrated in Figure 6. The figure plots quarterly housing price trends from the FHFA for six MSAs which experienced housing price increases of different amounts between the first quarter (Q1) of 2000 and the last quarter (Q4) of 2005. The potential endogeneity problem we confront is that in each case some of the price growth in these six cities was due to the influence of confounding latent demand or supply shocks. For the three cities on the left side of figure, prices evolved smoothly over time, consistent with the notion that the effect of fundamental factors is smoothly incorporated into prices. For each of the three cities on the right side of the figure, however, prices changed discontinuously, or "sharply", at some point in the 2000s, suggesting the influence of some factor different from standard fundamental determinants of housing price changes, such as drivers of the housing bubble discussed earlier. A TSLS strategy which uses as an instrumental variable the magnitude of the structural break in an MSA's quarterly price series could therefore isolate variation in housing demand that is plausibly exogenous in regressions of labor market or schooling outcomes.

We construct estimates of the structural break in the quarterly price series of each MSA between Q1, 2000 and Q4, 2005 by estimating MSA-specific OLS regressions with a single structural break, by searching for the location of the break which maximizes the R^2 of the following regression:[10]

$$\log\left(P_k^H(t)\right) = \alpha_k + \beta_k t + \gamma_k (t - t_k^*) 1\{t > t_k^*\} + \varepsilon_{k,t} \tag{8}$$

In equation (8) above, $P_k^H(t)$ represents the local housing price in MSA k at time t (where time indexes year-quarter observations), β_k represents the MSA-specific linear time trend before the structural break, t_k^* represents the MSA-specific location of structural break in the time series (which we restrict to be between Q1, 2002 and Q4, 2004), and γ_k represents the MSA-specific magnitude of the structural break, which can be interpreted as the change in (quarterly) house price

[10] See Ferreira and Gyourko (2011) for a detailed discussion of how many MSAs experienced discrete jumps in house prices during the mid-2000s relative to their historical trends. Given the discussion in that paper, in constructing the instrumental variable, we focus on time period between Q1, 2000 and Q4, 2005 to avoid identifying any structural breaks following the end of the housing boom.

growth rate at the location of the break. We extract γ_k and convert it to an annual growth rate, and we use it as an instrument for the change in housing demand between 2000 and 2007.

Figure 7 shows a strong positive relationship between the size of the estimated structural breaks in an MSA and change in housing demand in the city between 2000 and 2007. Table 2 assesses the relationship between the structural break measure and changes in housing demand proxy over the boom and bust. The first two columns show that, after accounting for a full set of standard controls, the structural breaks indeed strongly predict 2000-2007 housing demand growth. The second pair of columns shows the structural break in the price series during the boom also strongly predicts the 2007-2011 change in estimated housing demand, which is consistent with the fact shown earlier that the size of the boom in an MSA is close to perfectly negatively correlated with the size of the later bust in housing demand. Importantly, the *F*-statistic on the structural break measure is always far larger than 30 for housing demand changes during both the boom and bust. These first-stage results show that there are no "weak instrument" concerns with using the magnitude of the structural break as an instrumental variable for either the 2000-2007 and 2007-2011 change in housing demand in TSLS models.[11]

Whether the structural break instrumental variable isolates the effect of factors like speculative activity and *not* the effect of latent confounds depends on whether the effect of fundamental factors like demand or supply shocks are incorporated smoothly into housing prices. The various panels of Figure 8 plot the relationship between the size of the structural break for an MSA and various pre-existing features of the market as of 2000, including the initial non-employment rate, average wages, the share of women employed, the fraction of workers with a college degree, the total population in the market, and lagged housing price growth. Strikingly, the figure shows that the structural break measure does not systematically vary with *any* of these MSA-level variables. Though this evidence obviously does not rule out the possibility that the structural break is related to unobserved confounds, we find it reassuring that our instrumental variable exhibits no association with a large set of pre-existing observable variables that are likely closely related to latent factors that would raise obvious endogeneity concerns.

The results in the two panels of Figure 9 suggest that the structural breaks are the result of exogenous speculative activity, and are not related to changes in underlying factors that determine labor market or education outcomes. The first panel presents results concerning changes in the ratio

---

[11] We are grateful to Edward Glaeser for encouraging us to pursue this estimation strategy.

of the house price to rental price across MSAs, based on rental price information we have calculated for each market in our sample.[12] The logic for these results is as follows. Suppose that there is some latent amenity or productivity shock which raises the desirability of living in an MSA. This should increase the price of *all* housing in the MSA – including both prices and rents. If the structural break measures were the result of changes in these latent factors - the factors most relevant for endogeneity concerns regarding *current* employment, wages and schooling – then there should be no relationship between the price-to-rent ratio and the structural break in a city. If, by contrast, the structural break reflected price changes arising from speculative activity in which investors formed a (perhaps incorrect) judgment about the likely *future* desirability of the MSA and drove up prices through speculative purchases, the price-to-rent ratio should be positively related to the structural break in an MSA. This is precisely what the first part of Figure 9 shows, suggesting that the breaks do not reflect the changes in current amenities or productivity factors, at least to the extent these effects show up in rents.[13]

More evidence that the breaks represent price changes from speculative activity comes from the second panel in Figure 9. In recent work, Chinco and Mayer (2014) have carefully assembled data from transaction-level deeds records to identify purchases in several large housing markets made by "out-of-town buyers" – individuals with a primary residence in one market who nonetheless buy a house in another market. By examining differences between local and out-of-town buyers in exit timing and realized capital gains, they show clear evidence that out of town buyers across most housing markets during the 2000s were disproportionately misinformed speculators. Using the data they have assembled for the 21 markets they study, we show in the second panel of Figure 9 that for this sub-sample of markets with available data, our structural break variable is strongly correlated with growth in share of buyers who are (speculative) out-of-town buyers.

Taken together, we regard the patterns of evidence in Figure 8 and Figure 9 as strongly consistent with the idea that the estimated structural breaks identify plausible exogenous variation in housing demand and thus use these measures as instrumental variables in the TSLS portion of the

[12] See the Online Appendix for a detailed description of the construction of rental price data using the Census/ACS data. Online Appendix Table OA.3 reports analogous results to the 2SLS results in Table 3 using the change in the price-to-rent ratio as an instrumental variable (instead of the estimated structural break measure used in the main tables). The results are similar to the main results in Table 3.

[13] One concern with rental prices in the Census/ACS data is the fact that the quality of rental units may vary over time within an MSA, making quality-adjusted comparisons of rental price changes across MSAs difficult. Chico and Mayer (2014) construct quality-adjusted (residualized) rental price data using richer data, though the data only exist for 43 large MSAs. Using their data, we show in Online Appendix Figure OA.1 that there is a strong relationship between the structural break instrument and changes in the price-to-rent ratio, similar to what we show in Figure 9 for our full sample of 237 MSAs.

work that follows. The next section presents results for housing demand shocks on labor market outcomes during the boom, bust, and for the entire period spanning both the boom and bust.

V. LABOR MARKET RESULTS

Throughout the analysis to follow, standard errors in all regressions are clustered by state. The analysis is conducted in first differences and thus accounts for time-invariant differences across MSAs. In most specifications, the control vector includes controls for the share of employed workers with a college degree, the share of women in the labor force, and the MSA's population as measured in 2000.

Table 3 presents estimates of the effect of housing demand shocks in an MSA during the housing boom period of 2000-2007 on the change in the MSA employment rate (the share of the relevant population employed), based on OLS and TSLS estimation of equation (4). In these and all subsequent regressions, the change in housing demand in an MSA is proxied by $\widehat{\Delta\omega_k^H}$. For the TSLS results, the instrumental variable is the structural break in the quarterly house price series discussed in Section IV. To aid interpretation of the magnitudes implied by the point estimates, the rows beneath the estimated coefficients and standard errors show the implied effect of a one-standard deviation change in the housing demand measure.[14] The first four columns show results for each of four sex $\times$ education groups, and the fifth column presents results for all prime aged workers.

The OLS results shown in the top panel of the paper suggest that local housing booms during the 2000-2007 time period increased employment by statistically significant amounts for prime-aged men and women overall, and for each separate group of workers except college men. The results show that a standard deviation increase in housing demand increased the employment rate of non-college men by 1.2 percentage points. The standardized effects for college men shown in column 2 are only about one quarter the size, but are not statistically significant. The OLS results for women, shown in columns 3 and 4, are qualitatively similar to those for men, in that a standard deviation increase in housing demand during the boom raises employment rates for non-college women by about 1 percentage point, but had much smaller effect on college women. Overall, the OLS results

[14] The coefficients are always standardized by the cross-city standard deviation in magnitude of the housing demand change during the time period analyzed. One useful feature of this re-scaling is that the standard deviation happens to be similar in magnitude to the national change in house prices (for a city with average supply elasticity). Therefore, the standardized effects correspond to the estimated effect of the "national" housing boom.

imply that local housing demand shocks increased employment overall among all prime-aged workers by slightly less than 1 percentage point (0.8).

The preferred TSLS results in the second panel, based on the structural break instrumental variable, are meant to address both the endogeneity and measurement error concerns with the OLS results. Among men, the TSLS results imply standardized effects that are about 0.2 percentage points larger for both skill groups, while the standardized effects implied by the TSLS estimates are about 0.2 percentage points smaller than the corresponding OLS estimates for women. Qualitatively, these results are very similar to those from the OLS regressions. The TSLS analysis finds the largest effects for non-college men, and finds very small effects on employment rates among college women.

Because they are our preferred estimates, and to reduce clutter, we focus on TSLS results in all of the tables that follow.[15] It is thus useful at this point to discuss the robustness of the instrumental variables results. Table 4 presents results under several specifications of the instrumental variable and sample restrictions for non-college men, the group whose employment rate was most strongly affected by housing demand shocks. Columns 1 and 2 reproduce the main results from Table 3. Columns 3 and 4 present OLS and TSLS results when fixed effects for the nine Census divisions are added to the regressions. In column 5 we use the change in housing prices rather than $\widehat{\Delta\omega_k^H}$ to proxy for housing demand, using the same set of MSAs from the previous columns. The estimated coefficient is larger (since the price change is no longer scaled by sum of the supply elasticity and demand elasticity), but the standardized effect is extremely similar.[16] Once we use the change in prices as a proxy, we can also enlarge the sample to include all MSAs that have data on labor market outcomes and house prices; results for this larger sample are shown in column 6 and they are extremely similar to column 5, suggesting that restricting sample to set of MSAs with housing supply elasticity estimates does not affect the results. Next, we show that the TSLS results are similar in column 7, where the magnitude of the structural break is set to 0 if the estimated break is not statistically significantly different from 0 at the 5% level of significance. Finally, column 8 shows that the results are also similar when the instrumental variable is based on estimating the magnitude of the break using a cubic trend rather than a linear trend. The stability of estimates across these specifications gives us confidence that our preferred instrumental variable estimates are

[15] Analogous OLS results for each of our main tables are presented in the Online Appendix.
[16] The similarity in standardized effect in column 5 is readily explained by the fact that our structural break instrument is not strongly correlated with the estimated housing supply elasticity used to construct the housing demand proxy.

robust and leads us to conclude that housing demand shocks indeed had large effects on employment rates, especially for non-college men.[17]

What sectors account for these employment changes across different groups? Table 5 presents TSLS estimates for employment changes in two sectors that intuition suggests should be affected by housing market changes: the construction sector and the so-called "FIRE" sectors – finance, insurance, and real estate. In addition to the point estimates and standardized effects, the last row in each panel of the table shows the ratio of the standardized construction or FIRE employment estimate divided by the overall standardized employment effect from Table 3. These ratios measure how much of the total employment effect from housing shocks for a given type of worker occurs through employment in construction or in the FIRE sectors, respectively.[18] The results show that, for non-college men, a standard deviation increase in housing demand increased construction employment by 1.2 percentage points, which accounts for 91.5% of the overall change in employment from housing demand shocks for these men. For college men, 51.5% of the employment change attributable to housing demand shocks occurred in construction, although the total employment rate change for these men was small. Among women, only those without a bachelor's degree experienced any employment change at all from changes in housing demand, and 31.3% was in construction. Overall, 76.4% of the employment rate changes from the housing boom for all prime-aged workers came through construction employment.

Results for employment in FIRE sectors are shown in the bottom panel of Table 5. Overall, across all prime-aged workers, the 23.8% increases in employment in these sectors accounted for virtually all of employment increases from the boom that did not occur in construction. These estimates show that the bulk (56.6%) of the employment growth experienced by non-college women from the boom came in these sectors. Eighteen percent of the overall employment growth for non-college men occurred in these sectors. Interestingly, our estimates suggest that FIRE sectors account for very little of the small increases in employment experienced by college men and women from the boom.

One important consideration for the interpretation of our various employment results is the effect of migration. Previous work has shown that internal migrants move to booming markets and

[17] An additional concern about our employment results is that they are primarily driven by behavior of immigrants. We show in Online Appendix Table A.3 that this does not appear to be the case. The results in this table show that there is still a positive and significant increase on employment and wages, even when dropping all immigrants from the sample.
[18] Note that the ratios we present are based on the division of the actual point estimates, and thus they sometimes differ slightly from the ratio of the rounded point estimates in the tables.

leave markets that are declining.[19] To the extent that migration was selective, it may not be valid to infer from our estimates that housing booms change employment probability for people in affected markets. In particular, to the extent that in-migrants moving into MSAs experiencing housing booms had a higher propensity to work than the natives of those MSAs, our employment estimates would reflect this composition effect.

We examine directly both how net migration responded to housing booms and whether such migration as occurred was selective. In the top panel of Appendix Table A.1, we display the 2SLS results from a regression of an MSA's population growth between 2000 and 2007 on the housing boom. This regression is identical to the ones shown in Table 3 aside from the dependent variable. The results show that there was a relative increase in population in MSAs that experienced a relative increase in housing demand. However, the increase in population (or net in-migration) was larger for college-educated men and women. This pattern of migration results by sex and skill make it unlikely that migration is an important driver of the employment responses to booms that we estimate. This conjecture is confirmed in the bottom panel of Appendix Table A.1. Here we look at the 2007 employment propensities of migrants in the MSA relative to the 2007 employment propensities of natives in the MSA. The ACS data codes whether individuals moved into the MSA during the last year. The results show that the relative employment propensities between natives and migrants were unrelated to the size of the housing boom in the MSA. These findings suggest that composition considerations from selective migration are not an important concern for our estimates of labor market effects of the housing boom. These results are consistent with recent work that suggests that the bulk of incidence of labor demand shocks falls on the employment rate of pre-existing residents (Yagan 2013).

Table 6 presents estimates of the effect of local housing shocks on average wage growth between 2000 and 2007.[20] We find relatively large positive effects for the entire prime-aged population and for the separate subgroups. Except for college men, all of these estimated effects are strongly statistically significant. The standardized effects suggest that a standard deviation positive housing demand shock raised average local wages by 1.9 percent among all prime-aged

[19] See, for example, Blanchard and Katz (1992).

[20] When computing mean wages within an MSA during a given time period, we start with the same analysis samples described earlier. We then impose the following restrictions to the individual data: (1) the individual must be currently working at least 30 hours during a typical week at the time of the survey, (2) the individual's income in the year prior to the survey must exceed $5,000, and (3) the individual must have worked at least 48 weeks during the prior year. With these restrictions, we then complete mean wages at the MSA level in each of the time periods. Given these restrictions, our wage data should be considered for full-time workers with relatively few non-employment spells.

workers. Similar to the employment results, we find the largest estimated wage effects for non-college men, for whom our results imply that a standard deviation increase in housing demand increased average wages by 2.6 percent. Interestingly, the difference between the estimated effects for non-college men and other workers is smaller for wages than for employment – particularly for women. Thus, for example, average wages for non-college women rose by 2 percent during 2000-2007 in response to a standard deviation housing demand shock, while mean wages for college women increased by 1.8 percent.

The results in Tables 3-6 suggest that local housing demand shocks during the boom had significant positive effects on workers' employment and wages, with particularly large effects for non-college workers. The gains to these less-skilled workers seemed to come in different sectors, with those for non-college men coming through construction, while real estate and related non-tradable services appear to have been the source of non-college women's employment and wage gains.

Table 7 presents TSLS estimates of the effect of housing demand shocks during 2000-2007 on employment and wages during the 2007-2011 bust and during the longer 2000-2011 period spanning both the boom and bust. The first two columns of the table reproduce results from earlier tables showing the effect of demand shocks during the boom on labor market outcomes between 2000 and 2007 for ease of comparison. Consider first the results in Panel A. The results in columns 3 and 4 for the bust period imply that the size of housing demand increase during the boom (and thus, effectively, the size of the decrease in housing demand during the bust) sharply reduced employment rates among non-college men and among all prime-aged men and women. From the standardized effects, a standard deviation increase in the housing demand shock lowered employment rates among non-college men and all prime-aged workers by 1.4 and 1.0 percentage points, respectively. Notice, these estimates are opposite in sign and about equal magnitude to the gains during the boom period. The results in the final column should thus not be surprising: the size of the housing demand increase during the boom caused no statistically significant or economically meaningful change in employment rates in the longer-term, whether for non-college men or for all prime-aged workers. Essentially, employment gains during the boom were almost exactly wiped out by declines during the bust. The results in Panel B for average wage changes are noticeably different. In particular, we find that the negative housing demand shocks during the bust had no effect on average wages. The combination of the downward stickiness of wages during the bust and the positive wage effects during the boom mean that both non-college men and all prime-aged workers saw average increases in their wages of 2.6 and 2.1 percent, respectively, over the longer period spanning the boom and bust.

We conclude the discussion of the effects of housing demand shocks on labor market outcomes with results showing average wage and employment effects for non-college men and all men and women, separately by age. For the results in Table 8, we re-estimate the TSLS models separately by two age-groups: ages 21-35 and ages 36-55.[21] We find that for both wages and employment, the estimated effects of housing demand shocks were slightly larger for younger persons. We turn below to an analysis of the effect of the housing demand shocks on educational outcomes – in particular the decision to obtain college training. The large labor market gains from the housing boom that we estimate for people without college training, especially for younger people in this group, suggests that the housing boom could have dissuaded some people in this group from foregoing those returns to invest in schooling. We examine this issue more formally in the next section.

VI. SCHOOLING RESULTS

The motivation for a possible relationship between housing demand shocks and schooling investment decisions is straightforward. Standard models of human capital investment argue that a forward-looking agent considering college weighs the expected future earnings gains from obtaining college training against the sum of the direct costs of college (tuition and other fees) plus the earnings she foregoes by spending time in college rather than working now – the opportunity costs of college (Becker 1964; Mincer 1958). For our purposes, the most straightforward implication of this simple yet powerful framework is that local labor market shocks that increase the employment and wages of non-college workers, like the housing demand shocks studied above, will decrease schooling investment for people at the margin of attending college. This basic prediction can, however, be refined in a number of ways to provide more powerful tests for whether local housing demand shocks during the boom and bust affect schooling decisions.

One refinement concerns the interaction of the timing of the housing demand shock with when people make schooling decisions. Because the acquisition of general and specific capital on the job

[21] Although these tables only highlight the results for non-college men and for all workers, analogous results for other subgroups (non-college women, college men, and college women) can be seen in the Online Appendix.

ensures that the earnings profile of the typical worker rises each year she works, a shock to positive wages or employment has a larger effect on the opportunity costs of college-going the younger the person. Thus, any negative effect of the 2000-2007 housing boom on educational investment should have been larger the younger the person was when the boom occurred. Another implication is that compared to other young people during periods when there was no local housing boom, young people deciding on college training during the boom should have been less likely to attend college than other young persons who did not experience a boom, either because they belong to a different birth cohort or because their local labor market experienced no housing boom.

Another refinement of the basic prediction has to do with the *type* of college training relevant for persons whose decision to attend college at all could be affected, at the margin, by the labor market effects associated with housing demand shocks. In particular, we expect that the types of college training for which improvements in wages and employment opportunities associated with the housing boom would make a marginal person less likely to go to college would be post-secondary training with market returns not appreciably larger than those in sectors like construction or less-skilled services. This suggests that housing demand shocks should have had disproportionately larger negative effects on community college or other associate's degree-level training, as compared to bachelor's degree-level training, which arguably may not be expected to be affected much at all.

To investigate the extent to which the housing boom and bust affected schooling, we use three complementary data sources. The first data source is the March Current Population Survey (CPS), which contains information on highest level of schooling attained. This survey allows us to construct annual educational attainment estimates (which was used to construct Figure 4 above). The small sample sizes of CPS, however, preclude a detailed local labor market analysis. For that we turn to the Census/ACS, which was also used for the labor market outcomes analysis above. This data set contains detailed information about the highest level of education completed, and it also distinguishes between associate's degrees and bachelor's degrees. The primary advantage of this data source is that we can classify respondents by birth cohort, age, and geographic location, and we have sufficient sample size to construct reasonably accurate estimates for each MSA in our main sample. We use the Census/ACS data to estimate changes in educational attainment (within each MSA over time) for persons who were of college-going age during the housing boom. Lastly, because the educational attainment information in the CPS and the Census/ACS is self-reported, and because the data are not well-suited to measuring educational investments that do not result in degree completion (such as individuals who enroll in community college but do not graduate), we turn to the rich administrative data from the IPEDS. This data set provides detailed information on enrollments for a very large share of all colleges and universities in the U.S., including both community colleges and four-year colleges and universities. Additionally, the data carefully tracks first-time, full-year enrollments, which allow us to measure persons who are enrolling in college for the first time during the boom and bust periods that we study. We match colleges and universities to MSAs by hand and we compute MSA-specific estimates of first-time, full-year enrollments for different types of colleges and universities.

Turning to our empirical results, we begin with the CPS data, which we use to investigate in more detail the trends in Figure 4 which show a slowdown (nationally) in educational attainment around the same time as the start of housing boom. In Figure 10, we group together the four states that had the most substantial housing booms (Arizona, California, Florida, and Nevada), and we show that the trends in educational attainment in this group of states appears to be strongly related to timing of the housing boom and bust. In these states, the share of men aged 18-29 who report having attended at least one year of college does not just slow down during the housing boom, but instead abruptly declines after more than a decade of steadily increasing. In the other states, there is an apparent slowdown in educational attainment, but the slowdown is substantially weaker than the overall pattern in Figure 4. We therefore conclude that the overall slowdown in educational attainment observed nationally appears to be driven in large part by the sharp trend reversal in the states that experienced the strongest housing booms, and that the timing of this trend reversal is very much in line with the timing of the housing boom itself. While the limited sample size of the CPS precludes strong conclusions, we interpret these trends in the CPS data as highly suggestive of a connection between the housing boom and the slowdown in college-going.

To quantify the contribution of local housing booms to schooling, we first turn to the Census/ACS data. Specifically, we estimate TSLS versions of equation (4) above except that instead of using changes in employment or wages as the key dependent variables, we construct several alternative measures of changes in educational attainment for different age groups and birth cohorts. Focusing on a main sample of 18-21 year olds in 2000, we conduct two types of tests, both of which are based on comparisons over time within the same MSA. First, we examine how attainment evolves over time for people aged 18-21 in 2000, as a function of the local housing demand shock. This is a "within cohort" analysis. Second, we compare how attainment for people aged 18-21 in 2000 compares to attainment compares to attainment for persons aged 18-21 in later years (2007 or 2011), again as a function of the housing demand shock in the market. These are "cross cohort" results. Table 9 presents the results, with Panel A showing results for men and women combined and Panels B and C showing results for men and women separately.

Column 1 shows the 2000-2007 within-cohort results for associate's degree holding for 18-21 year olds. We estimate a statistically significant negative effect of the housing demand shock, meaning that within the generation of persons aged 18-21 in 2000 (around the time the housing boom began), those living in markets with larger positive changes in housing demand during the boom experienced smaller growth over 2000-2007 in the share of them receiving associate's degrees. The standardized effects imply that a one standard deviation housing demand change reduced the share completing associate's degrees by 1.7 percentage points (relative to a mean of 26.0 percent). The estimates in Panel B and C of the table show similar effects for men and women. Column 2 presents within-cohort results for this same generation of 18-21 year olds, but over the longer 2000-2011 period. These longer run estimates assess whether, by the end of the bust, persons from this generation in markets with larger housing demand shocks during the boom had "caught up" in terms of their schooling by the end of the full boom and bust cycle or whether the lowered attainment for this generation during the boom years was persistent. The standardized effects of these strongly significant point estimates show that there was only slight catch-up, with the little that there was occurring only for men. A full decade after the start of the housing boom, a standard deviation increase in growth in local housing demand continued to reduce the share with associate's degrees by 1.3 percentage points for men and 1.8 percentage points for women.

Column 3 presents 2000-2007 cross-cohort results for different generations of 18-21 year olds. The results show that the larger the growth in housing demand during the boom, the smaller the rates of associate's degree attainment among person aged 18-21 at the peak of the boom in 2007, compared to attainment rates among 18-21 year olds at the start of the boom in 2000. The standardized effect for men and women together, which is very similar to the effects for them separately, show that a standard deviation larger increase in housing demand resulted in 0.6 percentage point lower fraction holding associate's degrees. Column 4 shows results for the cross-cohort analysis for 18-21 year olds over the longer 2000-2011 period, and the results are small in magnitude and they are not significantly different from zero. In other words, 18-21 year olds at the end of boom/bust cycle obtained associate's degree training at exactly the same rate as had 18-21 year olds when the boom began, relative to those in non-housing boom areas.

Taken together, the results in the first four columns strongly suggest that the housing boom significantly negatively affected associate's degree completion, but seemed to only do so for the particular generation who were making college-going decision during the period of the run-up in housing demand. Although later generations in markets that *had* boomed returned during the bust to associate's degree attainment rates that were similar to rates at the start of the boom, persons who were young at the start of the boom have persistently lower attainment. For this particular group there appears to have only a small amount of catch-up, even several years after the bust in housing demand began.

Although the precision of some of our results are limited, we believe the estimates themselves are credible because of the results in the last two columns in Table 9, which present results for two counterfactual exercises. First, in column 5 we present 2000-2007 within-cohort results for people who were aged 30-33 at the start of the housing boom, and thus beyond the usual college-going age (for reasons noted earlier). Reassuringly, these results show no evidence that the size of the housing demand shock in a market had an effect on the share with associate's degrees. This addresses concerns that other unobserved factors (correlated with our instrumental variable) were increasing educational attainment across the entire local population, perhaps because of skill-specific changes in migration in response to changes in housing demand. Second, the counterfactual estimates in column 6 look at changes in share with bachelor's degrees. This investigation is motivated by the fact that the changes in labor market opportunities were primarily in construction and less-skilled services, and therefore would not likely have affected college-going decisions for young people who were deciding about attending four-year colleges and universities. Reassuringly, the results for the 2000-2007 within-cohort analysis for persons aged 18-21 at the start of the boom show no evidence that changes in housing demand affected the percentage holding bachelor's degrees. The absence of any evidence of skill-upgrading for 30-year-olds or a change in share with bachelor's degrees suggests that the housing boom directly affected college-going decisions for persons who were of typical college-going age and were on the margin of going to college at all.

Lastly, while the Census/ACS data allow us to conduct both within- and cross-cohort analyses, an important limitation of these data is that they are based on self-reports about completed schooling. Additionally, the Census/ACS data measure associate's degree attainment but does not measure college attendance at two-year colleges that does not lead to degree completion. Given that associate's degree graduation rates are around 30%, the decision to go to community college at all by persons on the margin may not typically translate into receiving an associate's degree. Therefore, the Census/ACS will not measure some of the overall educational investment responses to changes in housing demand. To address this issue, we turn to the IPEDS data, which measures first-time, full-year enrollments for a very large share of colleges and universities in the U.S. We present results of TSLS estimates of equation (4) in Table 10, focusing on college enrollments across different types of colleges. As with Table 9, Panel A shows results for total enrollments, while Panels B and C present results for men and women separately.[22] As before, we show only the preferred TSLS results to conserve space but the corresponding OLS results can be seen in the Online Appendix. The outcome variable in these regressions is the log of the ratio of average annual first-time enrollments divided by the total enrollment in 2000. This variable is calculated separately by time period (2000-2007 and 2000-2010) and type of school (community colleges and four-year colleges and universities). The results for change in first time community college enrollment for 2000-2007 are shown in column 1. We find the markets experiencing large positive changes in housing demand show sharp declines in enrollments. Specifically, a standard deviation increase in the housing demand shock in a market reduced first-time enrollment in community colleges each year by approximately 22.5 and 21.8 percent for men and women.

These seemingly large negative effects of housing demand shocks on enrollment are not necessarily inconsistent with the attainment numbers in Table 9. The most immediate explanation is the low probability of completing an associate's degree conditional on enrollment. Therefore, if we multiply the standardized effects in column 1 of Table 10 by $A \times B \times 0.30$, where A is the average annual community college enrollment / total student body in 2000 (=0.50), B is the initial 2000 enrollment / total 18-21 population in the MSA in 2000 (=0.27), and 0.3 is the degree completion probability, then we would get an estimate of the standardized effect of the housing demand shocks on the share of 18-21 year olds in the market getting an associate's degrees. Doing this exercise yields an estimate of -0.8 percentage points per year, which is much closer to the standardized effects in Table 9 that is estimated with Census/ACS data. Adjusting this estimate further to

[22] For both Table 9 and Table 10, analogous OLS results are shown in the Online Appendix. Unlike all of the preceding wage and employment results, the general pattern of results is similar across the OLS and TSLS results, but the magnitudes are sometimes different by as much as a factor of 2-3. We do not have a simple explanation for the source of these differences, in part because the magnitudes of the OLS results are generally higher than TSLS in Table 9 (when studying degree attainment) while the magnitude of the OLS results are generally lower than the TSLS results in Table 10 (when studying college enrollment).

account for age differences and the number of years of exposure to the housing boom brings these two estimates even closer together.[23]

Column 2 of Table 10 shows that the size of the housing boom in a market had *no* significant effect on the change in first-time enrollment in four-year colleges and universities between 2000 and 2007. These results are again strongly consistent with the attainment patterns shown in Table 9 (based on different data set) showing no effect of the size of the boom on share with bachelor's degrees. Finally, column 3 examines the change in first-time community college enrollment over the longer 2000-2010 period. The standardized effect estimates show that the a standard deviation increase in the size of the boom reduced first-time community college enrollments by about 15 percent for both men and women. These results show that while there is some suggestive evidence of "catch up" in community college enrollments after the end of the boom, the catch-up is far from complete. Taken together, we interpret these results as pointing to an enduring negative effect of the housing boom on college attainment, with the effects coming exclusively at community colleges and affecting individuals deciding between working towards an associate's degree or not attending any college at all. This empirical result is consistent with standard human capital models (Becker 1964; Mincer 1958, Ben-Porath 1967), where transitory shocks to wages can have permanent effects on human capital investments because the returns to education are now discounted over a shorter time period. Additionally, the results are also consistent with age-dependent costs of college attendance, since even though the wage returns to college return to normal following the boom and bust in housing, the true economic return for individuals on the margin of college attendance will have fallen because the costs of attending college have risen with age.

VII. Conclusion

This paper studies the economic effects of the recent boom and bust in housing on labor market and schooling outcomes. We compare changes in outcomes across cities using plausibly exogenous variation in housing demand and find that housing demand shocks raise employment and wages in

[23] Specifically, the IPEDS data on first-time full-year enrollments does not distinguish by age, while the results in Table 9 focus on particular age groups. Approximately 60 percent of community college enrollees are younger than 24, so this is rough estimate of the appropriate age adjustment to the standardized effect. Additionally, the enrollment estimates refer to annual enrollment, so to compare to Table 9 we must multiply by the number of years of exposure (3 years of first-time enrollment exposure, in order to give enough time to complete degree by 2007). Doing both of these rough adjustments produces an adjusted estimate of $-0.008\times0.6\times3 = -1.4$ percentage point decline (as compared to 1.7 percentage points in column 1 of Table 9).

the local labor market, particularly for less-skilled men and women. For men, these employment effects are almost entirely concentrated in construction employment, while for women the employment effects are somewhat concentrated in the FIRE sectors (i.e., Finance, Insurance, and Real Estate).

Our results indicate that both the housing bust *and* the housing boom affected aggregate employment, meaning that the specific sectoral shifts occurring during that housing boom may have stimulated aggregate employment. Understanding why these particular sectoral shifts had such an effect on aggregate employment represents an important area of future work.

While we find that employment rates return roughly to pre-boom levels following the boom and bust in housing, we document persistent declines in both college enrollment and educational attainment. These persistent effects are concentrated at community colleges and for birth cohorts who were of college-going age during the boom. Our schooling results are also consistent with marginal individuals having fairly low completion rates, since the estimated enrollment effects are larger than the estimated associate's degree completion effects. However, both the estimated college enrollment effects and educational attainment effects are large. Applying our local labor market effects nationally, we find that the national housing boom can account for as much as 40 percent of the observed slowdown in college-going for young men.[24]

An important question left unanswered by this paper is what the consequences are of the persistent negative schooling effects on individual and social welfare. The answer would appear to partly depend on the magnitude of the marginal returns to additional schooling for individuals on the margin of college attendance. Recent work suggests that this marginal return is very high for academically marginal students who would seem to be fairly representative of the marginal individuals whose college-going decisions are affected by local housing demand shocks (Zimmerman, forthcoming).[25] If true, then our results suggest that the housing boom created a "lost generation" of individuals who had the bad luck of being college-going age during the historically unprecedented boom and bust in housing.

[24] This estimate comes from scaling the TSLS estimate in column 1 of Table 9 by the national change in housing prices (multiplied by the average demand and supply elasticity) and dividing by the "gap" relative to trend in Figure 4.

[25] Of course, this logic implicitly assumes that observably similar individuals have similar marginal returns. Given recent work documenting large heterogeneity in marginal returns to college across individuals (see, e.g., Carneiro, Heckman, and Vytlacil 2011), such a claim is necessarily speculative.

REFERENCES

Atkin, David (2012). "Endogenous Skill Acquisition and Export Manufacturing in Mexico." Working Paper.

Barlevy, Gadi and Jonas D. M. Fisher (2010) "Mortgage Choices and Housing Speculation," Federal Reserve Bank of Chicago Working Paper.

Becker, Gary (1964). *Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education*. New York: Columbia University Press.

Ben-Porath, Yoram. (1967). "The Production of Human Capital and the Life Cycle of Earnings," *Journal of Political Economy*, 75(4): 352-365.

Black, Dan A., Terra G. McKinnish, and Seth G. Sanders (2005). "Tight Labor Markets and the Demand for Education: Evidence from the Coal Boom and Bust," *Industrial and Labor Relations Review*, 59(1): 3-16.

Blanchard, Olivier and Lawrence F. Katz (1992). "Regional Evolutions," *Brookings Papers on Economic Activity*, 23(1): 1-76.

Burnside, Craig, Martin Eichenbaum and Sergio Rebelo (2013). "Understanding Boom and Busts in Housing Markets," NBER Working Paper #19817.

Campbell, John and Joao Cocco (2014). "A Model of Mortgage Default", NBER Working Paper #17516.

Card, David, Alex Mas, and Jesse Rothstein (2008). "Tipping and the Dynamics of Segregation," *Quarterly Journal of Economics* 123(1), 177-218.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo (2013). "Manufacturing Decline, Housing Booms, and Non-Employment," NBER Working Paper #18949.

Chico, Alex and Chris Mayer (2014). "Misinformed Speculators and Mispricing in the Housing Market," NBER Working Paper #19817.

Favilukis Jack, Sydney C. Ludvigson., Stijn Van Nieuwerburgh (2010). "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium," NBER Working Paper #15988.

Ferreira, Ferando and Joe Gyourko (2011). "Anatomy of the Beginning of the Housing Boom: U.S. Neighborhoods and Metropolitan Areas, 1993-2009," *NBER Working Paper* 17374.

Gerardi, Kristopher, Andreas Lehnert, Shane Sherlund, and Paul Willen (2008). "Making Sense of the Subprime Crisis", *Brookings Papers on Economic Activity*, Fall, 69-145.

Glaeser, Edward, Joseph Gyourko, and Albert Saiz (2008). "Housing Supply and Housing Bubbles," *Journal of Urban Economics* 64(2): 198-217.

Goldin, Claudia and Lawrence Katz (2008). *The Race Between Education and Technology*, Bekknap Press.

Lovenheim, Michael F. (2011). "The Effect of Liquid Housing Wealth on College Enrollment." *Journal of Labor Economics*, 29(4), 741-771.

Mayer, Chris (2011). "Housing Bubbles: A Survey," *Annual Review of Economics*, 3, 559-577.

Mayer, Chris and Karen Pence (2009). 'Subprime Mortgages: What, Where, and to Whom?" In *Housing Markets and the Economy: Risk, Regulation, and Policy. Essays in Honor of Karl E. Case*, ed. E. Glaeser, J. Quigley, pp. 149-96. Cambridge, MA: Lincoln Land Ins.

Mian, Atif and Amir Sufi (2012). "What Explains High Unemployment? The Aggregate Demand Channel," NBER Working Paper #17830.

Mian, Atif and Amir Sufi (2011). "House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crises," *American Economic Review*, 101, 2132-56.

Mian, Atif, Kamalesh Rao, and Amir Sufi (2013). "Household Balance Sheets, Consumption, and the Economic Slump", *Quarterly Journal of Economics*, 128(4): 1687-726.

Mincer, Jacob (1958). "Investment in Human Capital and Personal Income Distribution." *Journal of Political Economy*, 66(4): 281-302.

Polinsky A. Mitchell and Ellwood David T. (1979). "An Empirical Reconciliation of Micro and Group Estimates of the Demand for Housing," *Review of Economics and Statistics*, vol. 61, pp. 199-205.

Ruggles, Steven, Matthew Sobek, Trent Alexander, Catherine Fitch, Ronald Goeken, Patricia Kelly Hall, Miriam King, and Chad Ronnander (2004). "Integrated Public Use Microdata Series: Version 3.0". Minneapolis, MN: Minnesota Population Center.

Saiz, Albert (2010). "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 125(3), 1253-96.

Sinai, Todd M. (2012). " House Price Moments in Boom-Bust Cycles," NBER Working Paper #18059.Shiller, Robert J. (2005). *Irrational Exuberance*. Princeton, NJ: Princeton Univ. Press

Shiller, Robert J. (2009). "Unlearned Lessons From the Housing Bubble," *Economist's Voice* 6(7):6.

Yagan, Danny (2014). "Moving to Opportunity? Migratory Insurance over the Great Recession", Working Paper.

Zimmerman, Seth D. "The Returns to College Admission for Academically Marginal Students." *Journal of Labor Economics*, forthcoming.

Table 1
Descriptive Statistics of Changes in Housing Market, 2000-2011

	N	Mean	Std. Dev.	Percentiles 10th	50th	90th
Panel A: Change Between 2000-2007						
Observed Change in Housing Prices	237	0.344	0.349	0.035	0.180	0.918
Estimated Housing Demand Change, $\widehat{\Delta\omega_k^H}$	237	0.851	0.763	0.147	0.584	1.852
Panel B: Change Between 2007-2011						
Observed Change in Housing Prices	237	-0.329	0.337	-0.884	-0.190	-0.052
Estimated Housing Demand Change, $\widehat{\Delta\omega_k^H}$	237	-0.838	0.743	-1.867	-0.622	-0.186
Panel C: Change Between 2000-2011						
Observed Change in Housing Prices	237	0.015	0.142	-0.157	0.015	0.182
Estimated Housing Demand Change, $\widehat{\Delta\omega_k^H}$	237	0.013	0.444	-0.578	0.047	0.481

Notes: This table reports the summary statistics for the baseline sample of 237 metropolitan areas (MSAs) across the time periods studied in the regressions that follow. The Estimated Housing Demand Change is constructed by multiplying the change in housing prices (from FHFA house price index) by the sum of the price elasticity of housing demand (assumed to be 0.7 based on Polinsky and Ellwood 1976) and a MSA-specific housing supply elasticity estimate (from Saiz 2010). This procedure creates a proxy for the change in housing demand in an MSA. All of the reported sample statistics are computed using the 2000 population of prime-aged men and women in the MSA (from Census) as weights, since these weights are used in all of the regressions that follow.

Table 2
First Stage for Housing Demand Change Using Magnitude of Structural Break in House Prices as Instrumental Variable

	Dependent Variable is Estimated Housing Demand Change Across…			
	… 2000-2007		… 2007-2011	
	(1)	(2)	(3)	(4)
Magnitude of Structural Break in House Prices	7.436	7.000	-5.953	-5.705
[Housing Boom Instrument]	(0.898)	(0.862)	(0.789)	(0.794)
	[0.000]	[0.000]	[0.000]	[0.000]
Standardized (1σ) effect:	0.494	0.465	-0.396	-0.379
First-stage F-statistic	68.55	66.00	56.86	51.66
N	237	237	237	237
R^2	0.59	0.61	0.47	0.50
Include baseline controls		y		y

Notes: This table reports OLS estimates of first stage. The baseline control variables included are initial (year 2000) values of the share of employed workers with a college degree, the share of women in labor force, and log population. The Magnitude of Structural Break in House Prices corresponds to the estimated MSA-specific magnitude of structural break in house price as estimated using quarterly house price data (from FHFA) between Q1, 2000 and Q4, 2005, where the structural break is constrained to be between Q1, 2002 and Q4, 2004 (inclusive). The structural break procedure is carried out separately for each MSA by regressing log house prices on a linear time trend and a structural break term, where the timing of the structural break is selected to maximize the R^2 of the time-series regression. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

Table 3
Employment Response to Housing Demand Change: OLS Estimates and TSLS Estimates using Magnitude of Structral Break in House Prices as Instrument

Dependent Variable is Change in Employment Rate, 2000-2007

Sample:	Non-College Men	College Men	Non-College Women	College Women	All Men and Women
	(1)	(2)	(3)	(4)	(5)
Panel A: OLS Estimates					
Estimated Housing Demand Change	0.015	0.004	0.011	0.004	0.011
$\widehat{\Delta\omega_k^H}$	(0.004)	(0.003)	(0.002)	(0.002)	(0.002)
	[0.001]	[0.135]	[0.000]	[0.071]	[0.000]
Standardized (1σ) effect:	0.012	0.003	0.008	0.003	0.008
R^2	0.71	0.17	0.67	0.11	0.76
Panel B: TSLS Estimates					
Estimated Housing Demand Change	0.018	0.007	0.008	0.002	0.011
$\widehat{\Delta\omega_k^H}$	(0.006)	(0.002)	(0.004)	(0.003)	(0.003)
	[0.002]	[0.002]	[0.038]	[0.637]	[0.000]
Standardized (1σ) effect:	0.013	0.005	0.006	0.001	0.008
First stage F-statistic	66.00	66.00	66.00	66.00	66.00
Mean employment rate in 2000	0.822	0.925	0.695	0.814	0.758
N	237	237	237	237	237
Include baseline controls	y	y	y	y	y

Notes: This table reports OLS and TSLS estimates for alternative demographic groups using the same set of baseline controls as in Table 2. See Table 2 for more details. In Panel B, the Magnitude of Structural Break in House Prices is used as an instrument for the Estimated Housing Demand Change. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

Table 4
Robustness to Alternative Specifications: Employment of Non-College Men

	Dependent Variable is Change in Employment Rate of Non-College Men, 2000-2007							
Estimation:	OLS	TSLS	OLS	TSLS	TSLS	TSLS	TSLS	TSLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Estimated Housing Demand Change	0.015	0.018	0.014	0.017	0.032	0.032	0.018	0.021
$\widehat{\Delta\omega_k^H}$	(0.004)	(0.006)	(0.004)	(0.006)	(0.011)	(0.011)	(0.006)	(0.007)
	[0.001]	[0.002]	[0.001]	[0.004]	[0.004]	[0.003]	[0.002]	[0.004]
Standardized (1σ) effect:	0.012	0.013	0.011	0.013	0.011	0.011	0.014	0.016
First stage F-statistic		66.00		38.24	32.79	34.10	66.92	21.14
N	237	237	237	237	237	283	237	237
R^2	0.708	0.706	0.750	0.748	0.735	0.719	0.706	0.698
Specification:								
Include baseline controls	y	y	y	y	y	y	y	y
IV with Magnitude of Structural Break in House Prices		y		y	y	y		
Include census region fixed effects (9 regions)			y	y				
Use change in house prices instead of housing demand change					y	y		
Add in MSAs with missing supply elasticity						y		
IV setting structural break to 0 if not significant							y	
Structural break estimated using cubic instead of linear trend								y

Notes: This table reports alternative specifications of the main results in column (1) of Table 3 (which are reproduced in columns (1) and (2), respectively). Columns (3) and (4) extend columns (1) and (2) by adding region fixed effects for the 9 census divisions. In column (5), the Estimated Housing Demand Change is replaced with the Change in House Prices (but the same instrumental variable is used). In column (6), the sample is enlarged to include all MSAs with labor market outcome data and house price data (the restriction that housing supply elasticity estimate is available for MSA is no longer imposed). In column (7), the estimated structural break values are set to 0 if they are not statistically significant at the 5% level in the MSA-specific regressions. In column (8), the structural break variable is estimated using a cubic trend instead of a linear trend. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

Table 5

TSLS Estimates of Effect of Housing Demand Change on Construction Employment and FIRE Employment

Sample:	Non-College Men	College Men	Non-College Women	College Women	All Men and Women
	(1)	(2)	(3)	(4)	(5)
Panel A: Change in Share Employed in Construction, 2000-2007					
Estimated Housing Demand Change	0.016	0.003	0.002	0.001	0.008
$\widehat{\Delta\omega_k^H}$	(0.004)	(0.002)	(0.001)	(0.001)	(0.002)
	[0.000]	[0.160]	[0.000]	[0.145]	[0.001]
Standardized (1σ) effect:	0.012	0.003	0.002	0.001	0.006
Construction Empl. Change / Empl. Rate Change [Table 3, Panel B]	89.9%	49.8%	29.3%	N/A	73.4%
Panel B: Change in Share Employed in Finance, Insurance, and Real Estate (FIRE), 2000-2007					
Estimated Housing Demand Change	0.003	0.000	0.004	0.003	0.002
$\widehat{\Delta\omega_k^H}$	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
	[0.003]	[0.861]	[0.001]	[0.173]	[0.001]
Standardized (1σ) effect:	0.002	0.000	0.003	0.002	0.002
FIRE Employment Change / Empl. Rate Change [Table 3, Panel B]	17.5%	3.9%	50.8%	N/A	22.3%
First stage F-statistic	66.00	66.00	66.00	66.00	66.00
N	237	237	237	237	237
Include baseline controls	y	y	y	y	y

Notes: This table reports TSLS estimates analogous to columns (1) through (5) in Panel B of Table 3 for alternative dependent variables. The Magnitude of Structural Break in House Prices is used as an instrument for Estimated Housing Demand Change; see Table 2 and Table 3 for more details. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

Table 6
Effect of Housing Demand Changes on Average Wages:
TSLS Estimates using Magnitude of Structral Break in House Prices as Instrument

Dependent Variable is Change in Average Wages, 2000-2007

Sample:	Non-College Men	College Men	Non-College Women	College Women	All Men and Women
	(1)	(2)	(3)	(4)	(5)
Estimated Housing Demand Change	0.034	0.015	0.026	0.023	0.026
$\widehat{\Delta\omega_k^H}$	(0.007)	(0.009)	(0.003)	(0.006)	(0.005)
	[0.000]	[0.092]	[0.000]	[0.000]	[0.000]
Standardized (1σ) effect:	0.026	0.012	0.020	0.018	0.020
First stage F-statistic	66.00	66.00	66.00	66.00	66.00
N	237	237	237	237	237
Include baseline controls	y	y	y	y	y

Notes: This table reports TSLS estimates analogous to columns (1) through (5) in Panel B of Table 3 for an alternative dependent variable. The Magnitude of Structural Break in House Prices is used as an instrument for Estimated Housing Demand Change; see Table 2 and Table 3 for more details. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

Table 7

TSLS Estimates of Wage and Employment Effects: Longer Run Results

Change in Dependent Variable Between …	… 2000-2007		… 2007-2011		… 2000-2011	
Sample:	Non-College Men	All Men and Women	Non-College Men	All Men and Women	Non-College Men	All Men and Women
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Dependent Variable is Change in Employment Rate						
Estimated Housing Demand Change, 2000-2007	0.018	0.011	-0.018	-0.012	0.000	0.002
$\widehat{\Delta\omega_k^H}$	(0.006)	(0.003)	(0.007)	(0.005)	(0.011)	(0.007)
	[0.002]	[0.000]	[0.012]	[0.013]	[0.986]	[0.804]
Standardized (1σ) effect:	0.013	0.008	-0.014	-0.009	0.000	0.001
Panel B: Dependent Variable is Change in Average Wage						
Estimated Housing Demand Change, 2000-2007	0.034	0.026	0.000	0.002	0.034	0.028
$\widehat{\Delta\omega_k^H}$	(0.007)	(0.005)	(0.011)	(0.007)	(0.012)	(0.009)
	[0.000]	[0.000]	[0.994]	[0.765]	[0.006]	[0.002]
Standardized (1σ) effect:	0.026	0.020	0.000	0.002	0.026	0.021
First stage F-statistic	66.00	66.00	66.00	66.00	66.00	66.00
Include baseline controls	y	y	y	y	y	y

Notes: N = 237 in all columns. This table reports TSLS estimates analogous to columns (1) and (5) in Panel B of Table 3 for alternative time periods. See Table 2 and Table 3 for more details. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

Table 8
TSLS Estimates of Wage and Employment Effects by Age Group

Restriction:	Age 21-35		Age 36-55	
Sample:	Non-College Men	All Men and Women	Non-College Men	All Men and Women
	(1)	(2)	(3)	(4)
Panel A: Dependent Variable is Change in Employment Rate, 2000-2007				
Estimated Housing Demand Change	0.022	0.015	0.016	0.009
$\widehat{\Delta\omega_k^H}$	(0.008)	(0.004)	(0.006)	(0.004)
	[0.005]	[0.000]	[0.005]	[0.018]
Standardized (1σ) effect:	0.017	0.011	0.012	0.007
Panel B: Dependent Variable is Change in Average Wages, 2000-2007				
Estimated Housing Demand Change	0.041	0.034	0.034	0.026
$\widehat{\Delta\omega_k^H}$	(0.008)	(0.007)	(0.006)	(0.004)
	[0.000]	[0.000]	[0.000]	[0.000]
Standardized (1σ) effect:	0.031	0.026	0.026	0.020
Include baseline controls	y	y	y	y

Notes: N = 237 in all columns. This table reports TSLS estimates analogous to columns (1) and (5) in Panel B of Table 3 for alternative samples. See Table 2 and Table 3 for more details. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

Table 9
Housing Booms and Educational Attainment Within and Across Cohorts

Dependent variable: Change in Share of Men and Women with …

	… Associate's Degree					… Bachelor's Degree
Sample Period	2000-2007	2000-2011	2000-2007	2000-2011	2000-2007	2000-2007
Age Group for Cohort Change	18 to 21	18 to 21	18 to 21	18 to 21	30 to 33	18 to 21
Within Cohort Change	y	y			y	y
Across Cohort Change			y	y		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: TSLS Estimates for Men and Women						
Estimated Housing Demand Change	-0.023	-0.020	-0.008	0.002	0.000	0.004
$\widehat{\Delta\omega_k^H}$	(0.008)	(0.009)	(0.004)	(0.006)	(0.005)	(0.006)
	[0.005]	[0.033]	[0.023]	[0.694]	[0.946]	[0.496]
Standardized (1σ) effect:	-0.017	-0.015	-0.006	0.002	0.000	0.003
Panel B: TSLS Estimates for Men Only						
Estimated Housing Demand Change	-0.023	-0.017	-0.007	0.005	0.006	0.003
$\widehat{\Delta\omega_k^H}$	(0.009)	(0.009)	(0.004)	(0.008)	(0.006)	(0.006)
	[0.006]	[0.056]	[0.072]	[0.522]	[0.327]	[0.603]
Standardized (1σ) effect:	-0.018	-0.013	-0.005	0.004	0.004	0.002
Panel C: TSLS Estimates for Women Only						
Estimated Housing Demand Change	-0.023	-0.023	-0.010	0.000	-0.005	0.005
$\widehat{\Delta\omega_k^H}$	(0.009)	(0.010)	(0.006)	(0.007)	(0.006)	(0.006)
	[0.011]	[0.025]	[0.084]	[0.954]	[0.440]	[0.441]
Standardized (1σ) effect:	-0.018	-0.018	-0.007	0.000	-0.004	0.004
Include baseline controls	y	y	y	y	y	y

Notes: N = 237 in all columns. This table reports TSLS estimates analogous to the columns in Table 3 with alternative dependent variables. The baseline controls include the same controls in Table 2 and Table 3 as well as the share (in 1990) of 18-25 year olds with an associate's degree and the share in 1990 of 22-25 year olds with a bachelor's degree (as well as an indicator for missing data in 1990 and interactions between this indicator and all of the other controls). See Table 2 and Table 3 for more details. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

Table 10
Housing Booms and College Enrollment at Community Colleges and Four-Year Colleges and Universities

Dependent Variable: Log(Average Annual First-Time Full-Year Enrollment / Total Enrollment in 2000)	Community Colleges, 2000-2007	Four-Year Colleges and Universities, 2000-2007	Community Colleges, 2000-2010
	(1)	(2)	(3)
Panel A: TSLS Estimates for Men and Women			
Estimated Housing Demand Change	-0.284	-0.002	-0.199
$\widehat{\Delta\omega_k^H}$	(0.082)	(0.053)	(0.077)
	[0.001]	[0.976]	[0.009]
Standardized (1σ) effect:	-0.218	-0.001	-0.153
Panel B: TSLS Estimates for Men Only			
Estimated Housing Demand Change	-0.293	-0.058	-0.203
$\widehat{\Delta\omega_k^H}$	(0.086)	(0.049)	(0.080)
	[0.001]	[0.236]	[0.011]
Standardized (1σ) effect:	-0.225	-0.040	-0.156
Panel C: TSLS Estimates for Women Only			
Estimated Housing Demand Change	-0.283	0.046	-0.199
$\widehat{\Delta\omega_k^H}$	(0.080)	(0.060)	(0.078)
	[0.000]	[0.442]	[0.011]
Standardized (1σ) effect:	-0.218	0.032	-0.153
First stage F-statistic	63.42	68.72	63.42
N	219	216	219
Include baseline controls	y	y	y

Notes: This table reports TSLS results. The baseline controls include the initial (year 2000) values of the share of employed workers with a college degree, the share of women in the labor force, and the log population in the MSA. The standardized effects rescale the coefficient by a one standard deviation change using the cross-state standard deviation. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

Appendix Table A.1
Changes in Population and the Relative Employment Rate of In-Migrants: TSLS Estimates using Magnitude of Structral Break in House Prices as Instrument

Sample:	Non-College Men	College Men	Non-College Women	College Women	All Men and Women
	(1)	(2)	(3)	(4)	(5)
Panel A: Dependent Variable is Change in Population, 2000-2007					
Estimated Housing Demand Change	0.041	0.048	0.039	0.051	0.043
$\widehat{\Delta\omega_k^H}$	(0.035)	(0.018)	(0.034)	(0.017)	(0.029)
	[0.237]	[0.010]	[0.244]	[0.002]	[0.134]
Standardized (1σ) effect:	0.032	0.036	0.030	0.039	0.033
Panel B: Dependent Variable is Emplyoment Rate of In-Migrants Relative to Non-Migrants, 2007					
Estimated Housing Demand Change	-0.008	0.009	-0.007	0.002	-0.004
$\widehat{\Delta\omega_k^H}$	(0.008)	(0.009)	(0.008)	(0.008)	(0.005)
	[0.334]	[0.314]	[0.374]	[0.829]	[0.471]
Standardized (1σ) effect:	-0.006	0.007	-0.005	0.001	-0.003
First stage F-statistic	66.00	66.00	66.00	66.00	66.00
N	237	237	237	237	237
Include baseline controls	y	y	y	y	y

Notes: This table reports TSLS estimates analogous to columns (1) through (5) in Panel B of Table 3 for alternative dependent variables. The Magnitude of Structural Break in House Prices is used as an instrument for Estimated Housing Demand Change; see Table 2 and Table 3 for more details. The change in population is calculated as the log difference between 2000 and 2007. The relative employment rate of in-migrants is computed based on the subsample of individuals who reported living in a different MSA in the previous year. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

Appendix Table A.2
List of Metropolitan Areas with Largest Structural Breaks

Rank	Metropolitan Area	Magnitude of Structural Break [Housing Demand Instrument]	Estimaed Housing Demand Change, 2000-2007
1	Phoenix, AZ	0.272	1.869
2	Yuma, AZ	0.271	2.388
3	Fort Walton Beach, FL	0.228	2.206
4	Visalia-Tulare-Porterville, CA	0.214	2.662
5	Lakeland-Winterhaven, FL	0.213	1.663
6	Naples, FL	0.213	2.395
7	Fort Myers-Cape Coral, FL	0.191	2.296
8	Orlando, FL	0.190	1.587
9	Bakersfield, CA	0.189	2.848
10	Pensacola, FL	0.186	1.185
11	Las Vegas, NV	0.183	1.794
12	Reno, NV	0.181	1.691
13	Ocala, FL	0.178	1.969
14	Tucson, AZ	0.177	1.410
15	Flagstaff, AZ	0.177	3.053
16	Melbourne, FL	0.172	1.852
17	Boise City, ID	0.163	1.162
18	Daytona Beach, FL	0.158	1.795
19	Norfolk, VA	0.140	1.157
20	Medford, OR	0.137	1.644
21	Merced, CA	0.136	3.812
22	Eugene-Springfield, OR	0.135	0.828
23	Riverside-San Bernardino,CA	0.129	2.054
24	Fresno, CA	0.129	3.079
25	West Palm Beach, FL	0.127	1.781

Notes: This table reports the top 25 MSAs in the main sample by the magnitude of structural break instrumental variable. See notes to Table 2 for more details on construction of the structural break variable. The units of the structural break variable represent the discontinuous change in (annualized) house price growth rates at the location of the break. The average estimated magnitude of structural break is 0.041 and the average estimated housing demand change is 0.851 (which corresponds to an 85% increase in housing demand between 2000 and 2007).

Figure 1: Home Prices in the U.S., 1980-2011

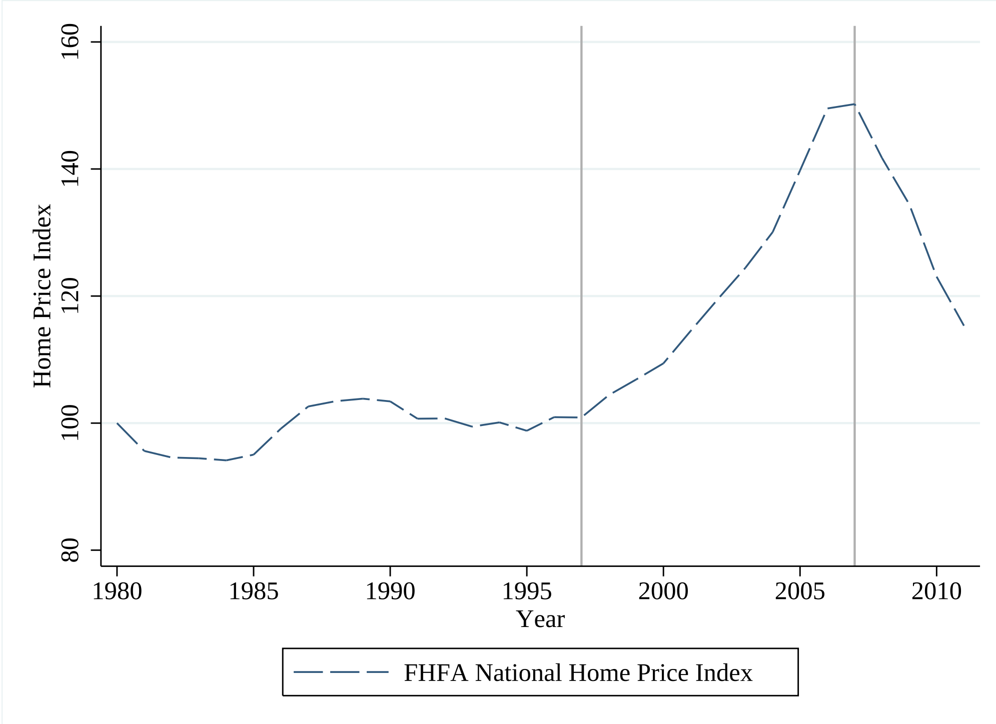

Notes: This figure reports trends in the FHFA National Home Price Index (1980 = 100). This series is a weighted, repeat-sales index that measures average changes in house prices across 363 metropolitan areas.

Figure 2: Trends in Construction Employment

Panel A - Overall Construction Employment Share, 1980-2011

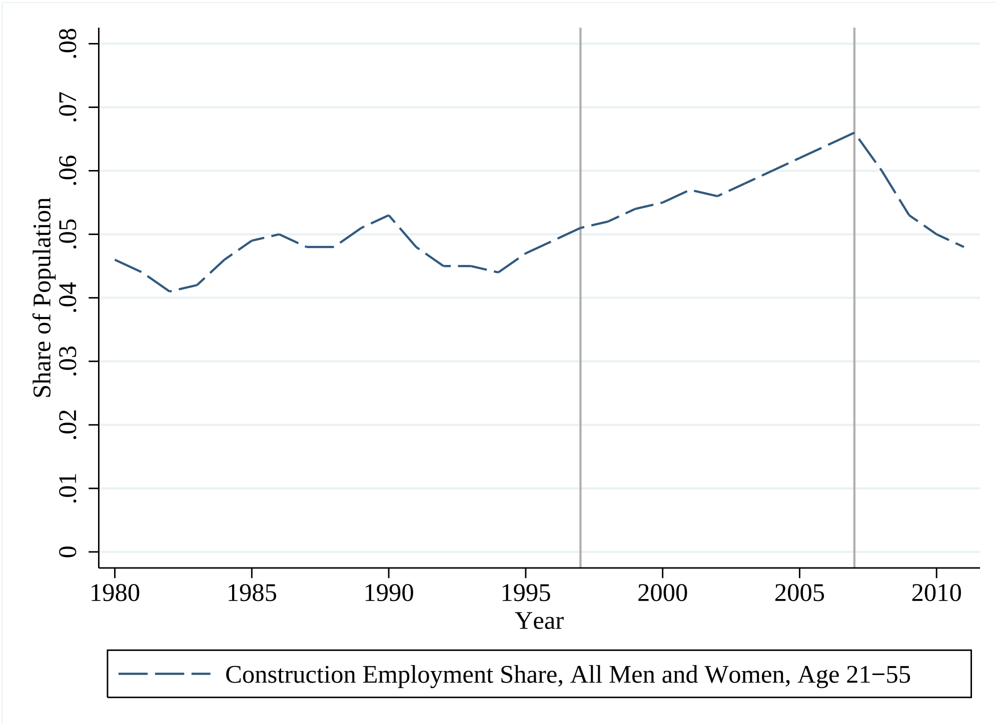

Panel B - Construction Employment Share by Gender and Education, 1980-2011

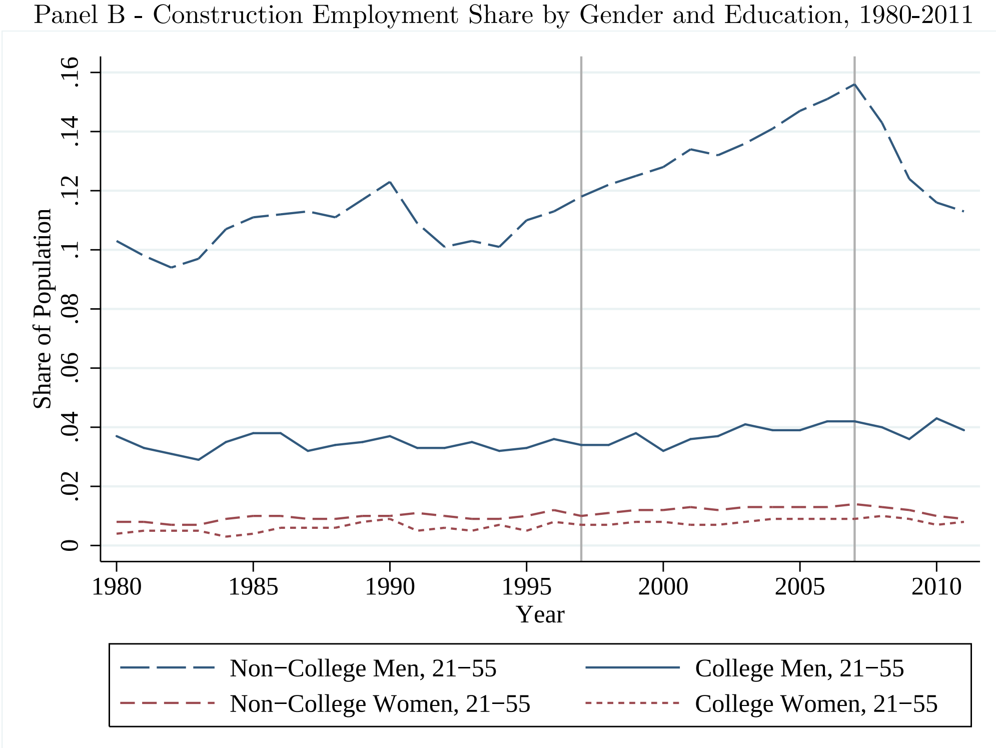

Notes: This figure reports trends in the share of total population employed in construction using data from the Current Population Survey.

Figure 3: Trends in Median Wages for Non-College Men, 1980-2011

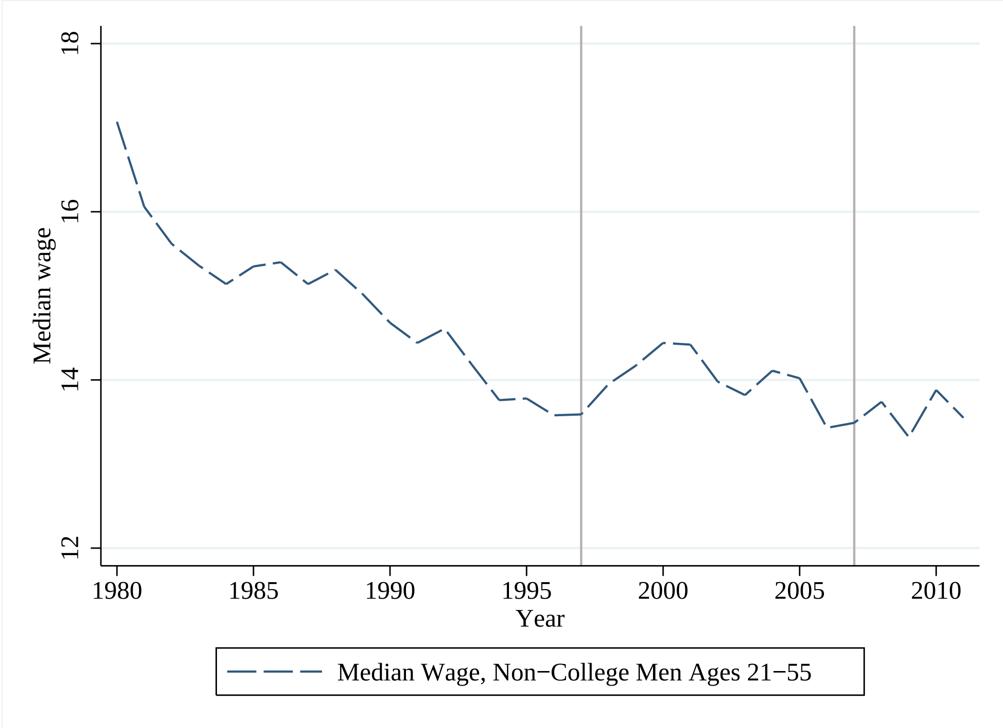

Notes: This figure reports trends in the median wage of men aged 21-55 without a college degree. The series is constructed from the Current Population Survey.

Figure 4: Trends in Educational Attainment for Non-College Men, 1980-2011

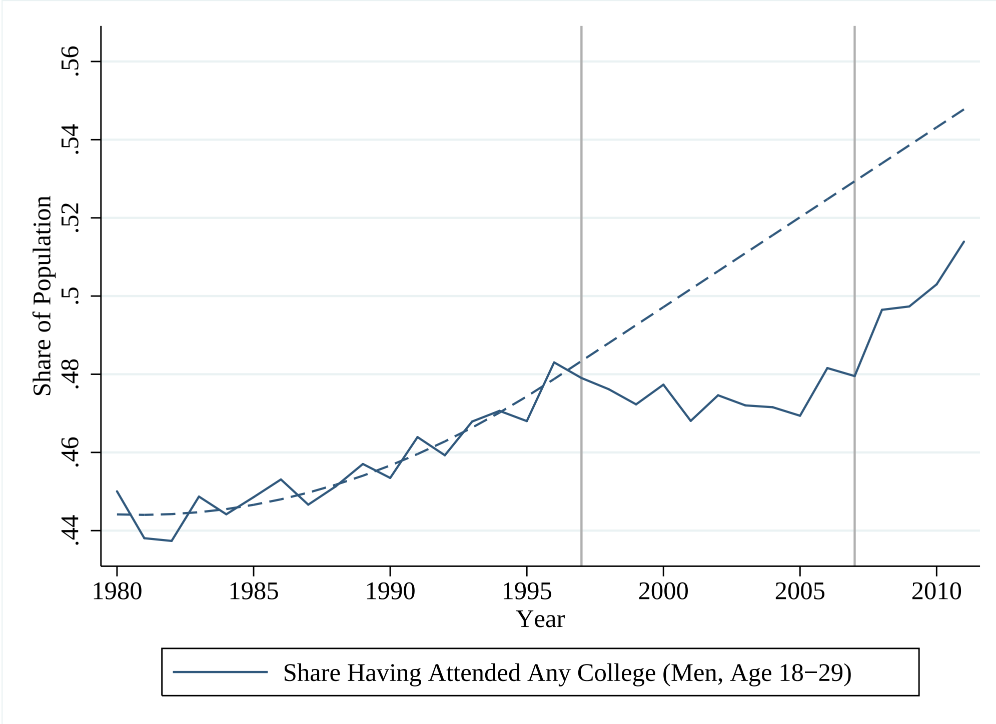

Notes: This figure reports trends in the share of men aged 18-29 who have attended at least one year of college. This series is constructed from the Current Population Survey. The dashed line is the predicted college attendance rates based on a quadratic fit for 1980-1995 period.

Figure 5: House Price Growth, 2007-2011 versus 2000-2007

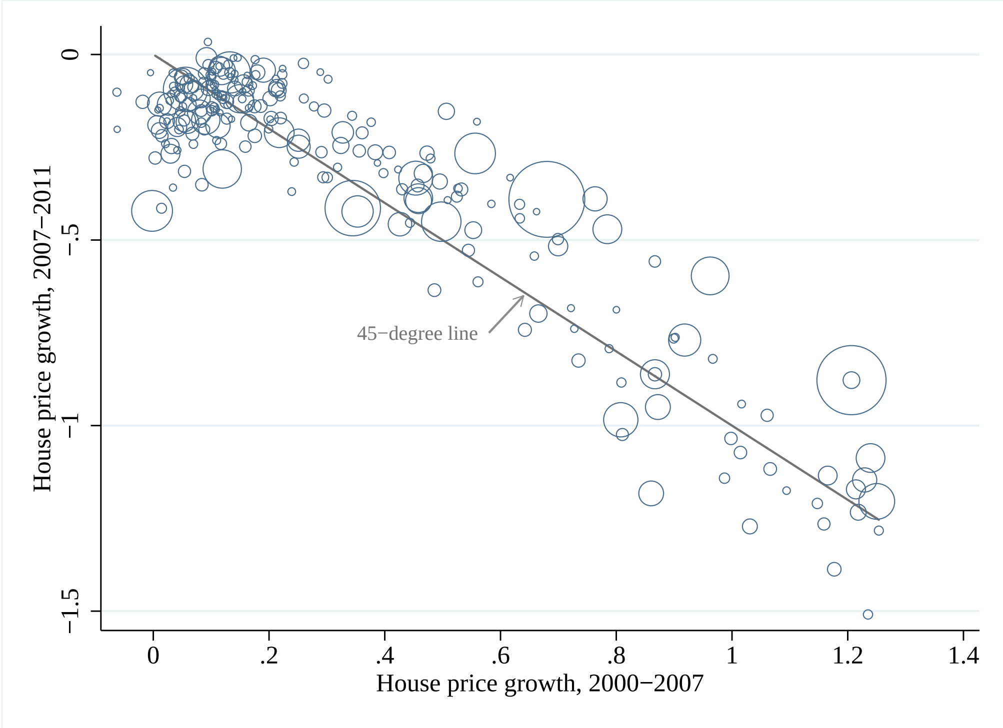

Notes: These figure shows the correlation between the change in house prices in 2000-2007 and the change in house prices in 2007-2011 for the 237 MSAs in the baseline sample. The solid line is a 45-degree line (i.e., slope of -1). The house price growth for each time period is calculated as log differences between the two endpoints, so that any city located on the 45-degree line experienced no change in house prices over the entire 2000-2011 period.

Figure 6: Variation in Structural Break Across Cities

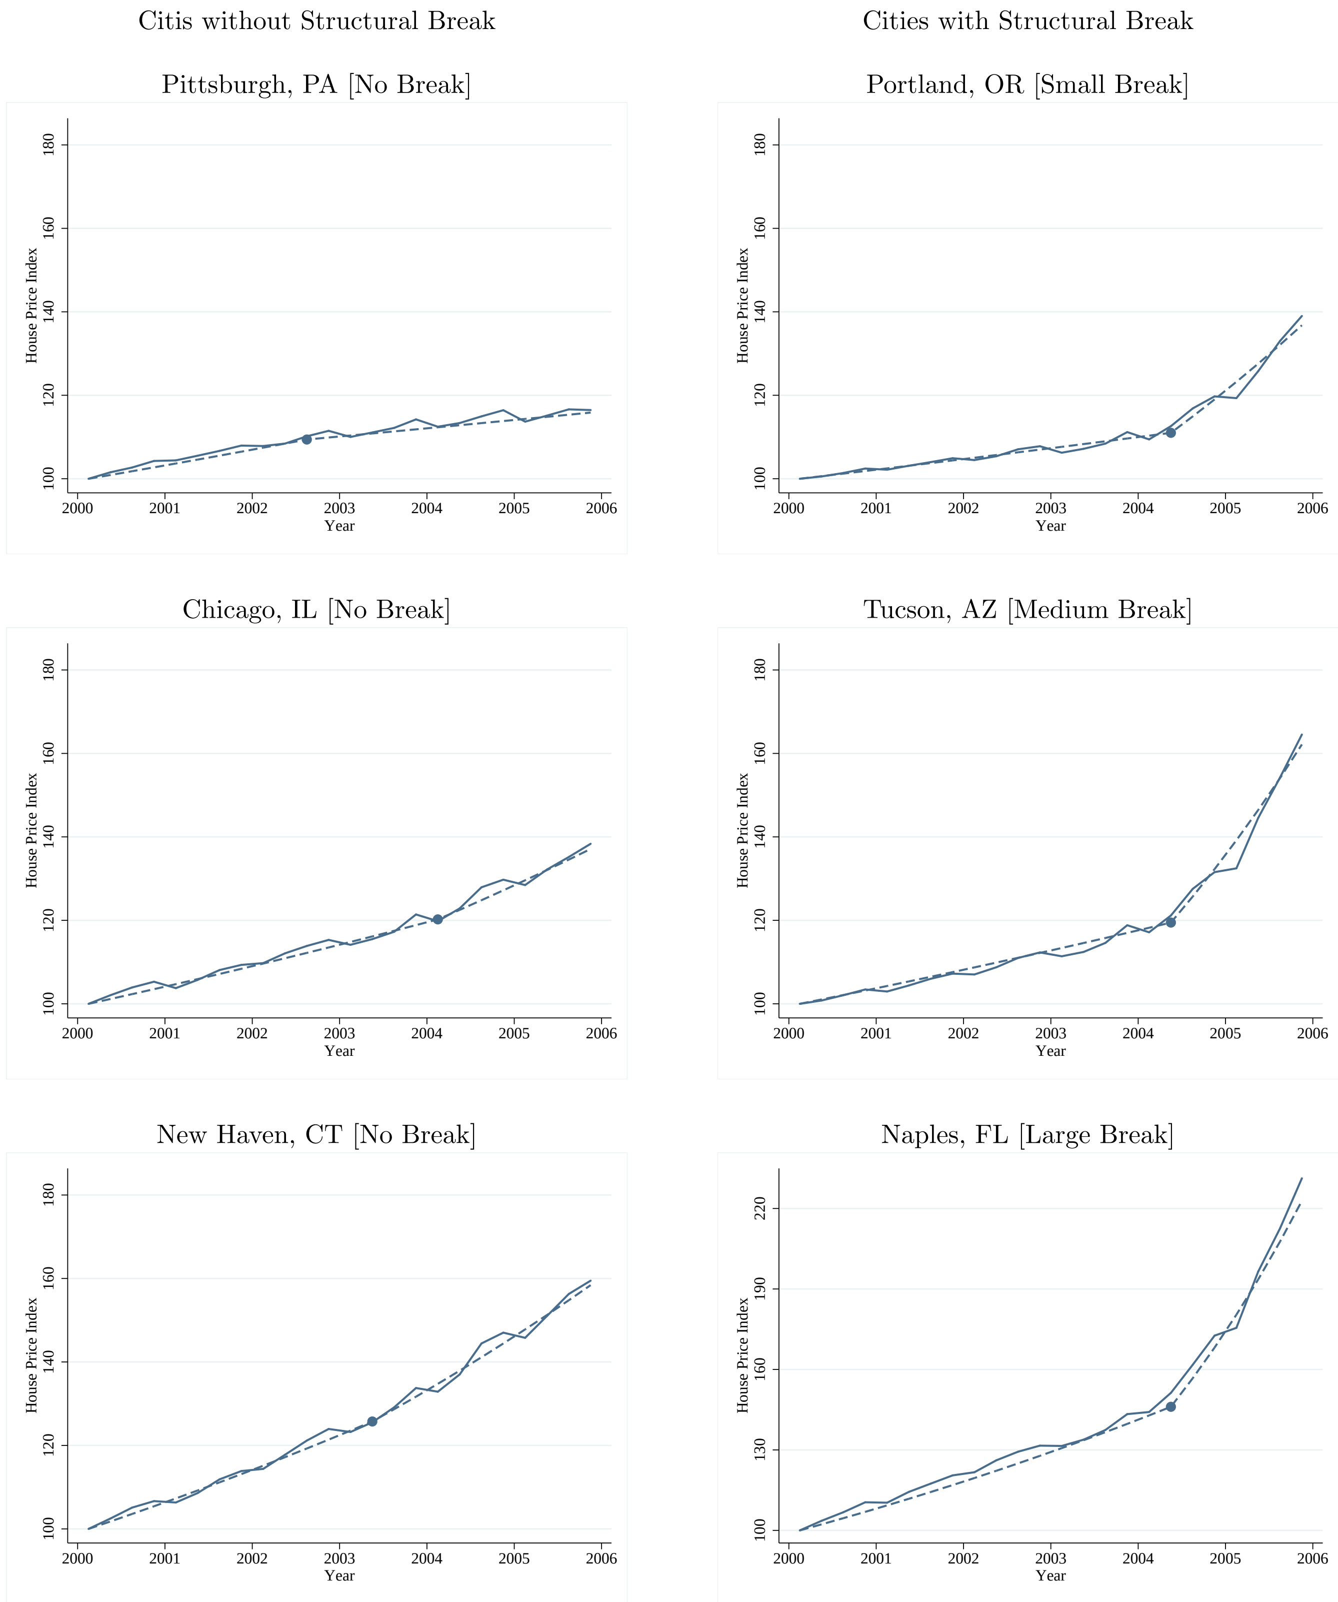

Notes: This figure shows graphs of quarterly house price data for six MSAs. The house price index for each city is normalized so that Q1, 2000 = 100. The solid lines report the house price series, while the dashed lines reports the structural break estimates, with a solid dot indicating the estimated quarter of the structural break. The six MSAs are grouped by overall house price growth, but differ in magnitude of estimated structural break.

Figure 7: First Stage Relationship Between Structural Break Instrument and Change in Housing Demand

Notes: This figure shows the correlation across cities between the magnitude of structural break (based on quarterly data between Q1, 2000 and Q4, 2005) and the estimaed housing demand change across 2000-2007. The Magnitude of Structural Break variable corresponds to the (annualized) coefficient from the city-specific structural break regression. The higher the value of the instrument, the larger the estimated structural break.

Figure 8: Lack of Correlation Between Structural Break Instrument and ...

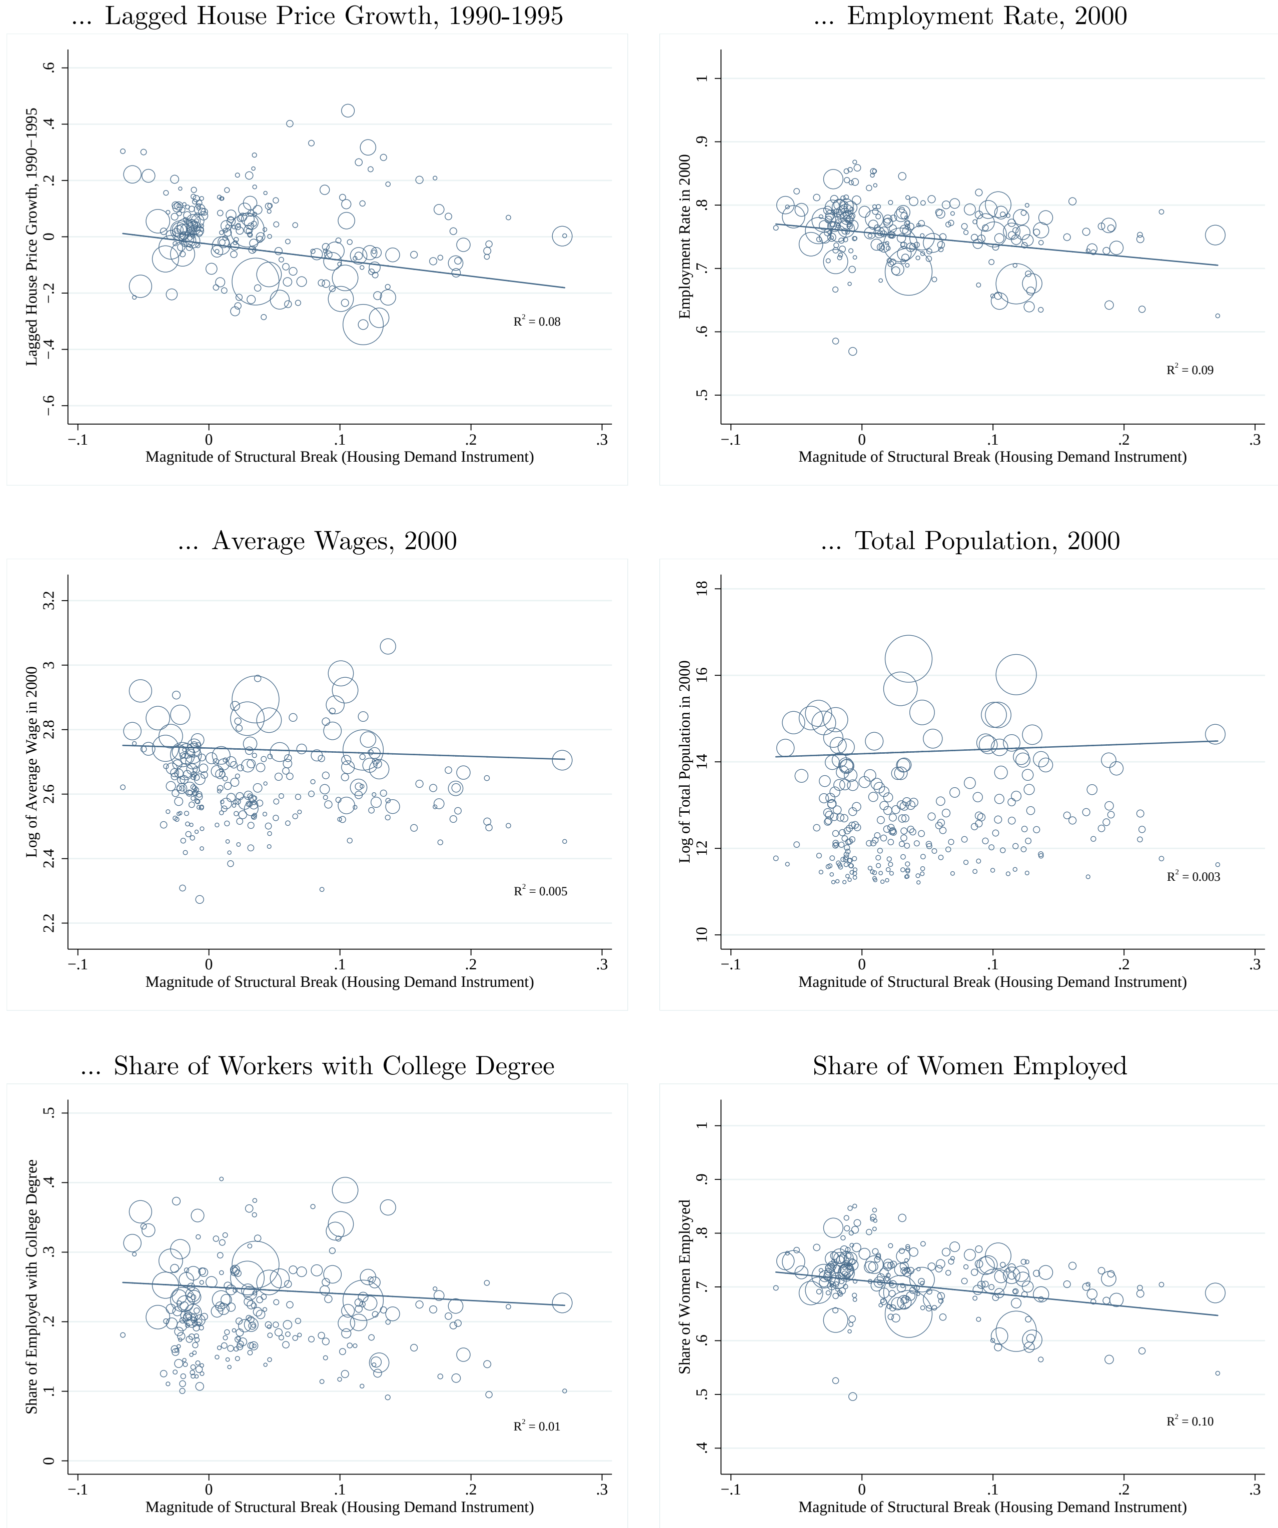

Notes: This figure reports the correlation between the structural break instrument used in the TSLS specifications and the following: lagged (1990-1995) house price growth and year 2000 values of the employment rate for all prime age men and women, log of average wages, log of total population, share of workers with college degree, and the share of women employed in the labor force.

Figure 9: Significant Correlation Between Structural Break Instrument and ...

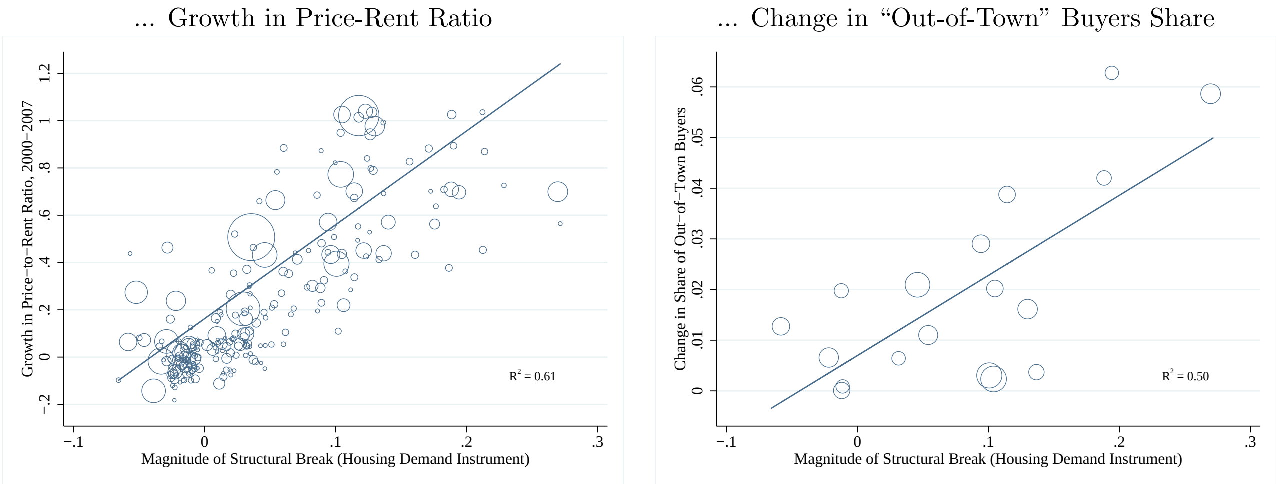

Notes: This figure reports the correlation between the structural break instrument used in the TSLS specifications and the growth in the price-rent ratio and the change in the share of "out-of-town" buyers.

Figure 10: Trends in Educational Attainment in States with Housing Booms, 1985-2011

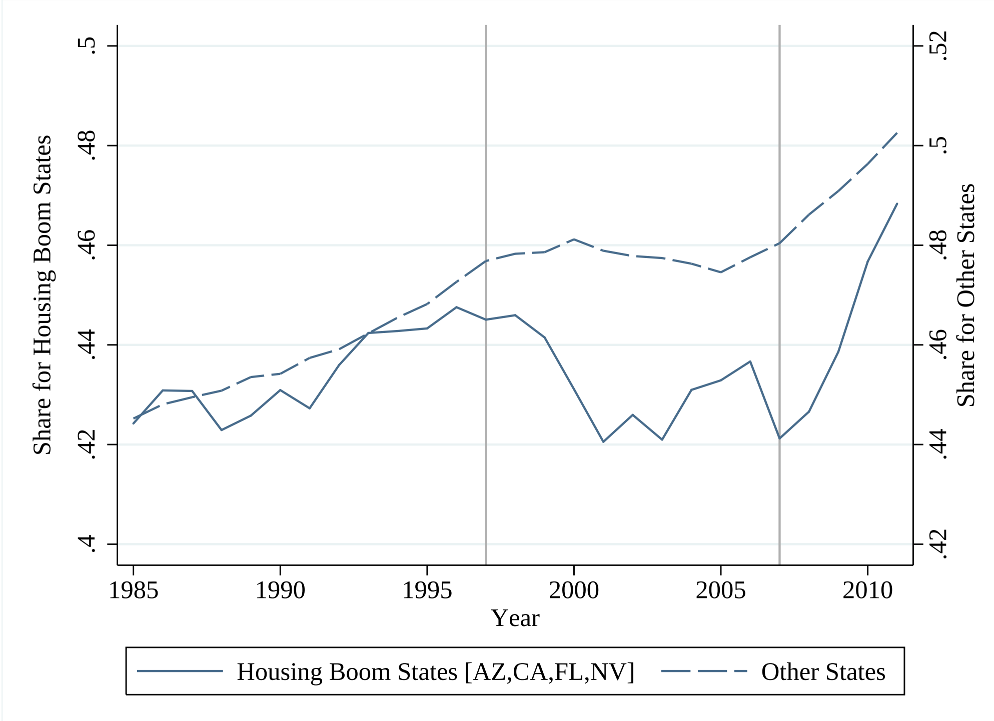

Notes: This figure reports trends in the share of men aged 18-29 who have attended at least one year of college for states that had largest housing booms compared to all other states. This series is constructed from the Current Population Survey, and both series are smoothed by calculating five-year (backward-looking) moving averages.