NBER WORKING PAPER SERIES

HOUSING BOOMS AND BUSTS, LABOR MARKET OPPORTUNITIES, AND COLLEGE ATTENDANCE

Kerwin Kofi Charles
Erik Hurst
Matthew J. Notowidigdo

Working Paper 21587
http://www.nber.org/papers/w21587

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2015

A prior version of this paper circulated as “Housing Booms, Labor Market Outcomes and Educational Attainment”. We thank seminar participants at the Columbia, Maryland, Northwestern, University of Chicago, the University of Houston, University of Illinois at Chicago, the AEA Annual Meetings, Yale - SOM, NYU, Princeton, University of British Columbia, UC - Berkeley, and New York Federal Reserve for helpful comments. We are also grateful to Tom Davidoff, Edward Glaeser, Matthew Gentzkow , Michael Lovenheim, Ofer Malamud, Atif Mian, Emily Oster, and Amir Sufi for useful suggestions and to Chris Mayer for sharing his data. We thank Mariel Schwartz, 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Housing Booms and Busts, Labor Market Opportunities, and College Attendance
Kerwin Kofi Charles, Erik Hurst, and Matthew J. Notowidigdo
NBER Working Paper No. 21587
September 2015
JEL No. E24,I21,J24

ABSTRACT

We study how the recent national housing boom and bust affected college enrollment and attainment during the 2000s. We exploit cross-city variation in local housing booms, and use a variety of data sources and empirical methods, including models that use plausibly exogenous variation in housing demand identified by sharp structural breaks in local housing prices. We show that the housing boom improved labor market opportunities for young men and women, thereby raising their opportunity cost of college-going. According to standard human capital theories, this effect should have reduced college-going overall, but especially for persons at the margin of attendance. We find that the boom substantially lowered college enrollment and attainment for both young men and women, with the effects concentrated at two-year colleges. We find that the positive employment and wage effects of the boom were generally undone during the bust. However, attainment for the particular cohorts of college-going age during the housing boom remain persistently low after the end of the bust, suggesting that reduced educational attainment may be an enduring effect of the housing cycle. We estimate that the housing boom explains roughly 30 percent of the recent slowdown in college attainment.

Kerwin Kofi Charles
Harris School of Public Policy
University of Chicago
1155 East 60th Street
Chicago, IL 60637
and NBER
kcharles@uchicago.edu

Erik Hurst
Booth School of Business
University of Chicago
Harper Center
Chicago, IL 60637
and NBER
erik.hurst@chicagobooth.edu

Matthew J. Notowidigdo
Northwestern University
Department of Economics
2001 Sheridan Road
Evanston, IL 60208-2600
and NBER
noto@northwestern.edu

1. INTRODUCTION

There is an active literature studying the consequences of the massive national boom and bust in housing that lasted from the late 1990s to the late 2000s, including an emerging body of work studying its effects on future economic growth. The creation of an "overhang" of debt that dampens future spending and investment is one possible mechanism by which the housing cycle may have affected future growth (Bhutta 2014; Jorda et al. 2014; Mian and Sufi 2014). Another possibility is that the cycle could have caused labor to be misallocated towards temporarily booming sectors with poor long-term growth prospects.[1] Curiously, how the boom and bust may have affected the distribution of schooling in the population has received little attention in this literature, despite education's key role in determining future individual economic wellbeing, and the fact that the level of schooling affects future productivity in the economy overall. This paper empirically assesses how housing demand shocks over the course of the housing cycle affected overall college attainment in the U.S, and adjudicates among alternative explanations for the patterns we document.

Suggestive evidence that the housing boom and bust changed college attainment comes from recent trends in overall college attainment that have received little attention. Using data from the Current Population Survey (CPS), the two panels of Figure 1 plot, separately for men and women, the share of persons aged 18-29 who reported ever having attended college training. The two panels show that while the share of young adults who have ever gone to college rose slowly and steadily since at least 1980, there was a noticeable slowdown relative to trend, for both men and women, beginning in the late 1990s – precisely when the national housing boom is generally agreed to have begun. The slowdown persisted through the peak of the national housing boom in 2006 and, despite some convergence during the bust period after 2006, attainment among young adult men and women had not fully reverted to trend as of 2013, years after end of the housing cycle.

In their seminal work on the much larger slowdown in attainment that occurred before the period we study, Goldin and Katz (2008) show educational attainment by birth year cohort up through the 1975 cohort. We follow their specification and use CPS data between 1994 and 2014 to examine college-attainment for year-of-birth cohorts from 1960 to 1990.[2] The year-of-birth effects from the Goldin-Katz-

[1] See 85th Annual Report of Bank of International Subsidies (BIS), www.bis.org/publ/arpdf/ar2015e.htm). In the popular press see "A New Explanation of America's Slow Productivity Growth," *Huffington Post*, August 8th 2015. David Adler. For similar arguments about misallocation in China see Chen and Wen (2014)

[2] These results can be interpreted as extending the Goldin and Katz (2010) results to birth cohorts after 1975, although their measure of college training is college degree completion rather than our measure of having attended college at all. The "second slowdown" in attainment that we focus on in this paper is much smaller than the very large slowdown from earlier in the century identified in earlier work. We pool the 1994 to 2014 waves of the CPS, restricting the sample to persons aged 25-54 in each year. We then estimate separate regressions for men and women separately on a dummy variable for whether the person has ever attended college on year-of-birth dummies, a quartic in age, and normalized year fixed effects where the first and last year effect are set to zero (as in Hall 1968).

style regression models are plotted in the two panels of Figure 2, which measure the predicted fraction of a birth cohort with any college training by age 25. The figure shows a clear slowdown in attainment somewhere between the 1970 and 1980 birth cohorts. For men, after steady cross-cohort growth of about 10 percentage points between the 1960 and 1970 cohorts, cohort-specific attainment rates were flat for the next ten birth cohorts, before starting to rise again, albeit at a much slower rate than before the slowdown. For women, the slowdown started around the 1975 birth cohort. Growth in cohort-specific attainment rates for the fifteen cohorts born after 1975 was about one-third the growth in cohort-specific rates for the fifteen cohorts born before 1975, with the cohorts between 1974 and 1980 essentially experiencing no growth in college propensity. Although the slowdown in cohort-specific attainment by age 25 roughly lines up with the start of the boom and bust, the figure also shows that at least some of the slowdown had nothing to do with the housing cycle since, especially for men, it began with cohorts that had already turned 25 *before* the boom began.

Figure 3 provides an initial assessment of whether the housing boom explains at least some of the slowdown. We combine the 1990 and 2000 Censuses with the 2005-2013 waves of the American Community Survey (ACS), and restrict attention to persons in this sample from the 1965-1987 birth cohorts living in their state of birth who were between the ages of 25 and 54.[3] We then compute the share of the birth cohort who had ever attended college, separately by whether or not the individual was living in a metropolitan area (MSA) that was in the top tercile of the increase in housing prices between 2000 and 2006.[4] The figure shows no cohort-specific differences in college attainment across the two groups of MSAs from the 1965 cohort through the first several cohorts of the slowdown. However, beginning with the 1979 birth cohort, who would have been eighteen when the national boom began and thus at the cusp of making college-going decisions, rates for persons in MSAs with especially big price increases fell behind rates for persons from the same birth cohorts in other markets. The difference in the propensity to attend college grew to a full two percentage points for the 1983 cohort.

The patterns in Figure 1-3 suggest that the housing boom reduced college attainment among both men and women. By what mechanism could this have occurred? To answer this question, we develop a simple conceptual model of college-going which shows that there are several different effects by which a housing boom might affect educational attainment. Since these effects are not of the same sign, and are differentially important for different population subgroups, the overall effect of a housing boom is theoretically ambiguous. We show, however, that a boom will tend to lower attainment if it improves current labor market opportunities for young adults so much that the labor market opportunity costs of college-going –

[3] The 2001-2004 waves of the ACS do not have MSA codes so could not be used for this exercise. As we discuss later, the restriction that individuals currently live in the state they were born mitigates concerns about endogenous migration.

[4] We use data from FHFA to measure local house price growth. As we formalize later, the standard deviation of housing price growth between 2000 and 2006 across MSAs was 0.36, relative to an average growth across MSAs of 0.40.

the earnings they must forego to acquire a college education instead of working – become large enough to override any other effect of the boom that might act in the other direction. One such counteracting effect is the loosening of liquidity constraints that persons from home-owning families might experience if their rising house values during a boom makes it easier for them to borrow to finance the direct costs of college.[5]

Our conceptual framework shows that, all else equal, if the boom did indeed improve labor market opportunities for young adults without a college education, then college attendance should have fallen during the boom, but mostly for the types of students and at the types of colleges where the gains from college-going are smallest. That is, increases in labor market opportunity costs should have reduced investment in two-year colleges that offer Associate's and similar degrees, but should have had little effect on investment in Bachelor's level training. Another insight from our framework is that the decision to not attend college in a given year due to a housing boom may be persistent because the time available to receive the gains from college-training shrinks while other opportunity costs, such as family obligations, rise mechanically and monotonically with age.

Figure 4 plots the massive changes over the housing cycle in three measures that affect labor market opportunities for persons without college training: housing prices; housing production, as measured by new residential construction permits; and total housing transactions.[6] Demand for workers providing local non-tradeable services, like waitresses, gardeners, hairdressers, nannies and retail clerks, has been shown to vary positively with changes in housing prices, perhaps because household wealth is driven substantially by changes in the value of housing and these services are normal goods (see Mian and Suf 2014). This suggests that the boom improved labor market opportunities for these kinds of workers. The increases in housing production depicted in Figure 4 would not have been possible without a substantial growth in labor market opportunities in blue-collar construction-related activities. Similarly, the massive surge in the number of houses bought and sold shown in the figure must have necessitated substantially greater activity in fields like real estate services, which are also conventionally open to non-college educated adults. These patterns suggest that the housing demand changes during the boom may have very substantially raised the opportunity costs of college-going for both young adult men and women, and could possibly partly explain the overall slowdown shown in Figures 1-3. As hinted at by Figure 3 above, the formal empirical work in this paper exploits differences across MSAs to assess how local housing demand shocks during the boom and bust affected educational attainment and labor market conditions. Most previous work on the housing boom has proxied for the size of housing demand change in a local market using only the size of the increase in housing prices in the area. We use instead a proxy for local demand that is the sum of changes in

[5] As we discuss later, Lovenheim (2011) studies this effect of the boom for home-owning families.
[6] We discuss these data sources in detail later in the paper.

both local prices and quantities, since large housing demand shocks might lead to very small price changes but large changes in supply in markets where housing is elastically supplied.

To account for potential measurement error and endogeneity in our measure of local housing demand, we isolate arguably exogenous variation in local housing demand. Our approach relies 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 (Shiller, 2008; Mayer 2011; Sinai 2012). Specifically, building on the work of Ferreira and Gyourko (2011) we estimate structural breaks in the evolution in housing prices in an MSA. We argue 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.[7] The estimated breaks are not, in fact, systematically related to any of 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 the size of an MSA's structural break explains an important portion of the overall change in housing demand in the MSA over the first decade of the 2000s. We use the size of structural break in an MSA to instrument for local changes in housing demand in Two Stage Least Squares (TSLS) first-difference models of the change of educational or labor market outcomes on the change in local housing demand.

Beginning with labor market outcomes, we find that increases in housing demand in an MSA during the 2000-2006 boom increased employment and wages for both young adult men and women without college training, raising their opportunity cost of college-going. Among young adult men, much of the improvements in labor market opportunity occurred in construction, whereas for young women the FIRE sectors of finance, real estate and insurance accounted for much of the gains. For both men and women, the remainder of the employment response was in local retail and service sectors. We also find that the boom either had no effect on or slightly lowered the expected future college/non-college earnings premium.

We present results for college attainment that are based on a variety of complementary data sources and different estimation methods. First, using data from Census/ACS, we relate the 2000-2006 change in an MSA in attainment among young adults aged 18-25. We restrict our analysis to samples where the individuals currently reside in the state they were born to mitigate concerns about endogenous migration. Both OLS and TSLS estimates show that the growth in the fraction of young adults with any college training was lower the larger the MSA's housing booms. Strikingly, we find that the change in an MSA's housing demand during the boom had no effect on the change in the fraction of young adults with a Bachelor's degree. The results suggest that improving labor market opportunities during the boom

7 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.

decreased advanced schooling attainment precisely for those persons who our conceptual model suggests should have been on the margin between obtaining "Associates-level" training and not going to college at all when the boom began.

We next turn to the rich, annual administrative enrollment data in the Integrated Postsecondary Education Data System (IPEDS). After hand-coding each institution in the data set to its MSA, we estimate TSLS first difference models of change in average annual enrollment from the years immediately before the start of the boom, to the years immediately before the peak of the boom. We find large and strongly significantly lower growth in enrollment in two-year colleges and universities the larger the housing demand growth in the MSA, but no difference in the change in enrolment in four-year, Bachelor-degree-granting institutions over the same time.

When estimating the size of the structural break in an MSA (which we use for the TSLS first-difference results), we also identify the precise date when the break is estimated to have occurred in an MSA. Using these two pieces of information (the size and timing of estimated structural break) and exploiting the annual frequency of the IPEDS data, we next estimate difference-in-difference-style regressions that assess whether enrollment in an MSA changed after the particular year in which the MSA experienced a structural break compared to enrollment in the years before MSA-specific break, and whether this effect differed by the size of the structural break. These models control for both MSA and year fixed effects, so they relate within-city, over-time variation in college enrollment to both the size and timing of the estimated structural break in housing prices. We estimate that 2-year college enrollment was lower after the specific year when an MSA had its structural break, and the effect for enrollment in 4 year colleges and universities exhibits no statistically significant change relative to enrollment during the pre-break period.

For our third set of education results we obtained permission for the restricted-use version of the NLSY97, a panel data set which follows a nationally representative sample of young adults who reached late adolescence and early adulthood during the early 2000s, right around the beginning of the national boom. With observations on only a few thousand individuals, this data set is very small compared to the other data sets we study. However, the individual panel feature of the data set allows us to track specific individuals as they age, with exact information about their MSA at each point of the housing cycle that is simply not available with other types of data. This data set allows us to better account for the potential bias from endogenous migration, and the data set also contains a rich set of individual- and family-level controls not present in the other data sources.

We find that NLSY97 respondents who, at the start of the housing boom were in MSAs that experienced especially large housing demand shocks, were much more likely to be employed at age 20 compared to their counterparts from other markets. This direct individual evidence about the opportunity cost channel is quite reassuring and consistent with the aggregate employment estimates from the Census/ACS. With respect to college training, respondents from high boom markets were substantially less likely to have obtained any college training at all by 2006, but did not differ from people in other markets in their propensity to have obtained a Bachelor's degree. Taken together, the results for educational attainment over the course of the boom are strikingly similar across datasets, specifications, and empirical methods. We consistently find that the boom lowered attainment at 2-year or Associates' level training but not at 4-year or Bachelor's level. Applying our local labor market estimates nationally, we find that the housing boom can explain approximately 30 percent of the national slowdown in college enrollment growth among both men and women.

We next investigate educational attainment over the housing bust and over the full course of the boom and bust cycle. Consistent with our conceptual model, we find that the bigger the growth in an MSA's housing demand during the boom (and thus the larger its decline during the bust), the larger the *increase* in attainment between the generation of young adults who made their schooling decisions at the peak of the boom to the generation of young people whose college decisions were made after the bust, when market opportunities for young people without college training had essentially returned to levels from before the cycle began. Again, we find no statistically significant difference in Bachelor's degree attainment across these two generations of young adults. Our final cross-generational difference estimates show that once the housing cycle had ended, new generations of young adults from boom markets appear to be investing in all types of college training no differently to young adults elsewhere or to generations of young adults in their markets from before the boom and bust cycle began. This, too, is consistent with our conceptual framework.

Our empirical work concludes with an assessment of persistence: whether the *particular generation* of young adults who obtained less schooling during the boom reversed this pattern by obtaining more schooling during the bust as labor market opportunities collapsed. Was the reduction in attainment we find for these particular people during the boom merely a delay, or does their schooling reduction seem to be permanent? The evidence from both Census/ACS and from the individual panel NLYS97 data is that young adults who invested less in college during the nearly ten years of the housing boom did not make up their lost college going propensity during the bust. Our evidence suggests that these cohorts have experienced a sort of "educational scarring", whereby their rates of attainment are permanently lower than would have been true had there been no boom. Their reduced educational attainment appears to be an enduring effect of the boom and bust cycle.

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 2015).[8] 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. Indeed, whereas none of the results we estimate for the effect of the boom (or bust) on Bachelor's degree attainment is statistically different from zero, some of our point estimates are positive, consistent with the idea that boom may have eased credit constraints at these more expensive colleges. Our empirical results imply that the causal pathway operating through housing wealth is overwhelmed by that operating through labor market conditions.

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 2. Section 2 discusses how we measure and isolate exogenous variation in local housing demand shocks. In Section 4 assess the effect of shocks on labor market opportunity costs and the expected future earnings return from college-going. Section 5 present the educations results over the course of the boom. Results for persistence, attainment over the bust and over the full housing cycle are presented in Section 6. Section 7 concludes.

2. THEORETICAL OVERVIEW

To motivate our empirical work, we present a simple conceptual model that illustrates how housing demand shocks affect college-going by exploring the key considerations emphasized in existing models of human capital investment. We highlight, in particular, the effect of shocks on potential students' opportunity cost of college attendance (Willis and Rosen 1979; Cameron and Taber 2004).

The potential students in our framework are young adults, who are aged α_t in year t and live until age L. They have completed the minimum required amount of schooling and now can either participate in the labor market or attend one of the two type of colleges, c, in the economy: "Associates" colleges $(c = A)$, or "Bachelor's" colleges $(c = B)$. Young adults differ in academic ability θ_i, which is distributed according to some distribution Φ_θ, over the interval [0,1]. As a college student, a young person incurs psychic costs of learning each year given by $\kappa_c(1-\theta_i)$. Training in a type- B college is inherently more difficult, and

[8] 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 (2015) finds that sectoral shocks arising from trade reform affects the distribution of schooling attainment in Mexico.

especially so for less able students, so $\kappa_B > \kappa_A$. It costs F_c in annual fees and tuition for c colleges, and students can borrow at an interest rate of b. In any year t, labor market participants with and without college training receive labor market incomes of Y_t^c and Y_t^0, respectively, which vary from one year to the next because of macro-economic and other shocks. The college premium in a given year for persons educated at a given type of college c is thus $\Pi_t^c = Y_t^c - Y_t^0 \geq 0$.

Given the model setup above, the lifetime payoff that a person of ability θ_i gets from attending a type-c college in year t, $R_{it}^c\left(\theta_i\right)$, is given by:

$$R_{it}^c(\theta_i) = \sum_{k=1}^{L-\alpha_t} E_t\left[\Pi_{t+k}^c \mid \Lambda_t\right] - \left(1+b\right)F - \kappa_c(1-\theta_i) - Y_t^0. \qquad (1)$$

The first term in (1) is the person's expected future lifetime income premium from college at date t, or the sum of their expectation of the college premium for every year of their future working life, given their current information Λ_t.[9] The last three terms in (1) represent, in order, the direct, psychic and opportunity costs of college-going, with the latter being the labor income the student foregoes in year t by enrolling in college and not working. Which, if any, college someone of ability θ_i attends in year t is determined by $\max[0, R_i^A(\theta_i), R_i^B(\theta_i)]$: he does not attend college if his expected lifetime payoffs from the both types of colleges are negative, and enrolls in the college where his expected lifetime payoffs are larger otherwise.

In all that follows, we focus on equilibria where some young adults enroll in each of the two types of college while others do not attend college at all. Since foregoing college is the preferred choice for some, the payoff functions for the two types of colleges must both simultaneously be negative for some ability levels. Similarly, since students attend both types of college there must be a range of ability levels where a college's payoff function is both positive and larger than the payoff function for the other type of college. That is, the payoff functions for the two types of college must cross at some point when both are positive. The two conditions

$$0 > R_{it}^A\left(\theta_i = 0\right) > R_{it}^B\left(\theta_i = 0\right) \quad \text{and} \quad 0 < R_{it}^A\left(\theta_i = 1\right) < R_{it}^B\left(\theta_i = 1\right)$$

guarantee that the payoffs from both type of college are negative at very low ability levels, and are both positive for very high levels of ability. These conditions also guarantee that the functions cross, since $R_{it}^B(\theta_i)$

[9] We make several assumptions for expositional simplicity which have no bearing on the main insights of this framework. In particular: we ignore discounting and assume agents are risk neutral; we assume that enrollees in a given type of college ultimately receive a degree, thereby ignoring the fact that there are differences across colleges in completion rates; and we do not allow students to work and attend school simultaneously. We also assume the opportunity cost for going to "Associates" college and "Bachelors" college is the same. In practice, the opportunity cost of "Bachelors" is likely higher because it takes longer to complete. Again, this abstraction does not alter the main takeaways of the simple model.

is smaller than $R^A_{it}(\theta_i)$ at ability 0, the inequality is reversed at ability 1, and both functions are strictly increasing in θ_i. If the payoff level $\overline{R}$ is the value at which the two functions cross and $\hat{\theta}^{AB}$ the associated level of ability, then the final assumption necessary for the baseline equilibrium is that

$$\overline{R} = R^A_{it}\left(\hat{\theta}^{AB}\right) = R^B_{it}\left(\hat{\theta}^{AB}\right) > 0.$$

Figure 5 illustrates an equilibrium satisfying the above conditions. The payoff functions for the two types of colleges are negative at the lowest levels of academic ability, with the lower intercept for the "Bachelor's" colleges indicating the greater inherent difficulty of that type of college. The functions strictly increase with ability, with the steeper slope for the "Bachelor's" function indicating the larger marginal benefit for this type of college for individuals with higher academic ability. Both functions eventually become positive, with the flatter "Associates" function becoming positive at a lower level of ability, θ^{A^*}, than the corresponding ability level for the "Bachelor's" function. The two functions eventually intersect at the value $\overline{R} > 0$ at ability level $\hat{\theta}^{AB}$. A person of ability θ^{A^*} is just indifferent between attending "Associates" college and not going to college at all; this person is the "marginal college-goer". Someone with ability $\hat{\theta}^{AB}$ is just indifferent between going to "Bachelor's" and "Associates" college; this is the "marginal Bachelor's student". These two thresholds completely characterize college-going in the population in our set-up: low ability individuals in the population – those with ability less than θ^{A^*} – do not attend college; persons with ability between θ^{A^*} and $\hat{\theta}^{AB}$ attend "Associate" college; and the most able persons, whose ability is at least $\hat{\theta}^{AB}$, enroll in "Bachelor's" college.

The effect of any shock on average college-going in the population is determined by how the shock shifts payoff functions and thus the two threshold ability values. The sign and magnitude of the shift in the payoff function for type-c colleges from a housing demand shock, dH, is the sum of three separate effects:

$$\frac{dR^c_i\left(\theta_i\right)}{dH} = \frac{d\sum_{k=1}^{L-\alpha_t} E_t\left[\Pi^c_{t+k} \mid \Lambda_t\right]}{dH} - \frac{db}{dH}F_c - \frac{dY^0_t}{dH}. \tag{2}$$

The first term in (2) measures how a housing demand shock affects the expected future premium from having attended college c. The second term is the change in borrowing cost from the rising housing wealth associated with the housing boom. The third term, which is the main focus of our paper, is the effect of housing shocks on potential students' opportunity costs.

Leaving aside for the moment any other effects it might cause, if housing demand shocks increase the labor market income that a young adult foregoes by enrolling in college in year t, college attendance becomes less appealing. This is reflected in the downward shift in the payoff function for both types of colleges shown in Panel B of Figure 5 from increases in Y_t^0. The ability threshold defining the marginal college-goer, θ^{A*}, rises, and some students who would have gone to an "Associates" college now forego college altogether.[10] By contrast, the marginal "Bachelor's" student is not changed by a rise in Y_t^0. This result stems from the fact that a young person in our framework forgoes the same amount in labor market returns whether he attends a type-A or a type-B school in a given year, which means that the relative attractiveness of attending one type of college versus another does not depend on the opportunity cost, Y_t^0 .[11] The two payoff functions shift downwards by the same amount and the threshold for "Bachelor's" college-going, $\hat{\theta}^{AB}$, is unchanged.[12] The model therefore captures the intuitively appealing idea that an increase in the opportunity cost of college attendance, all else equal, should have a greater effect on the propensity of individuals to pursue an "Associates" degree as opposed to a "Bachelor's" degree. In other words, if a housing demand shock improves labor market opportunities for persons without any college education, our conceptual model predicts that the share of the population attending college should fall, with most of the effect coming from reduced enrollment in "Associates" colleges, and little to no effect in "Bachelor's" colleges.[13] We would expect to see these enrollment responses unless there were offsetting influences from the other two effects in (2). What are the expected signs and magnitudes of these two other effects?

How a housing demand shock affects the expected future lifetime earnings gain from college depends in part on how young adults form expectations. If they believe that the shock will be temporary and produce no persistent effects on earnings in the future, then their expectations about the future college/non-college income premium will not change meaningfully. If young adults instead believe that a boom today will have persistent effects, affecting the labor income of college educated and non-college educated persons in the *future*, when today's young adults are older, the sign of the first term in (2) will depend on people's beliefs

---

[10] The formal statement of this claim is that $d\theta^{A*} \,/\, dH < 0$. This is straightforward to prove because if $dR^A(\theta) \,/\, dH = -dY_t^0 \,/\, dH < 0$, then the statement is true because $dR^A(\theta)$ is strictly increasing in θ .

[11] In practice, the opportunity cost channel may be larger for "Bachelor's" colleges since it takes longer to complete such a degree. Therefore, we will explore empirically the extent to which a housing boom affects the propensity to complete a bachelor's degree, as well.

[12] The formal statement of this claim is that $d\hat{\theta}^{AB} \,/\, dH = 0$. This is true because $\hat{\theta}^{AB}$ is implicitly defined by $R^A(\hat{\theta}^{AB}) = R^B(\hat{\theta}^{AB})$ and equation (2) implies that $dR^A(\theta) \,/\, dH = dR^B(\theta) \,/\, dH = -dY_t^0 \,/\, dH$.

[13] As we show in Online Appendix Figure OA.6, a large enough change in Y_t^0 could shift the functions down so much that their new intersection point is negative $\left(\overline{R} < 0\right)$. In this case, "Associates" enrollment falls to zero and, if the shock is big enough, even some persons who previously attended "Bachelor's" colleges could forgo college enrollment. As noted above, we ignore equilibria with corner solutions such as this case.

about the relative size of the effect of the boom on future skilled versus future less-skilled labor income. Only if it is expected that the boom will increase the future labor income gap between college and non-college educated people could the first term possibly over-ride the opportunity cost effect. Otherwise, an expected decline in the future college earnings premium will complement and reinforce the opportunity cost mechanism.

The other effect that could, in principle, offset the negative opportunity cost mechanism is if the positive shock to housing values and family wealth reduces borrowing costs or relaxes liquidity constraints. The existing evidence of the importance of liquidity constraints is mixed. Work by Cameron and Taber (2004) suggests that in the U.S. most persons wishing to attend college are not liquidity constrained, which is consistent with more recent work by Hilger (2014) that estimates very small effects of parental income on the probability of college enrollment. By contrast, Manoli and Turner (2015) find evidence that tax refunds have meaningful effects on college enrollment and Lovenheim (2011) finds some evidence of increased college attendance among person from families experiencing increases in housing wealth during the boom, with the effects concentrated among low-income families. Lovenheim's analysis, which uses data from the Panel Study of Income Dynamics, compares outcomes among homeowners to those of renters *within* a given market to capture the liquidity effect, netting any effect of the boom common to both groups in a market. Since changes in average labor market conditions affect opportunity costs for *all* potential college students in an MSA, the effect that is the main focus of our paper could not have been addressed in Lovenheim's study.

Our analysis compares the change in average outcome *across* different MSA. Our estimates therefore capture the opportunity cost difference between one MSA and another *plus* the offsetting borrowing/liquidity cost effects, if any, among homeowners compared to renters in the different MSAs. Our estimates thus can provide a measure of relative quantitative importance of two effects on college going in the aggregate. In particular, if increases in housing demand cause aggregate reductions in college enrollment, then this would imply that the opportunity cost mechanism was overwhelmingly large compared to any liquidity relaxation effect the boom might have caused.

Lastly, our conceptual framework suggests that any reductions in educational attainment from positive housing demand shocks at a point in time could, for some persons, represent permanent reductions rather than temporary delays. As in every life-cycle human capital model, young adults in our conceptual model are less likely to invest in schooling the older they get because their horizon to receive the expected lifetime earnings premium from college training, $L-a_t$, shrinks at each higher age.[14] One implication of this

[14] This is only one reason for the age effect in human capital models. In addition life events like marriage, the birth of children, infirmity of parents, expenditure commitments (as for durable goods), and any number of similar events, are all more likely to have occurred at higher ages, reducing the likelihood of college-going or indeed of any type of human capital investment at every

mechanical effect of aging is that if a share *ds* of the population decides not to enroll in Associate's colleges in a given year as a result of a housing boom, then it is unlikely that the entire mass *ds* will decide to enroll in a subsequent year if there is a negative housing demand shock that is equal in size to the preceding boom. Graphically, this mechanical aging effect causes the payoff functions to shift vertically downward each year as the person ages. The upward vertical shift in the two payoff functions caused by a housing bust would move them to a lower intersection point than where they intersected before a preceding boom of equal size.

In the empirical work below we study how local housing demand shocks during the national boom and bust cycle in housing affected college attainment. Before turning to the main education results, we offer direct evidence about how opportunity cost of (and expected future college earnings premium from) attending college were affected by the shocks. In terms of the education estimates, we assess both effects during the housing boom and the extent to which those effects persist after the end of the boom.

3. LOCAL HOUSING DEMAND SHOCKS

Our empirical work exploits variation across metropolitan statistical areas (MSAs), k, in the size of the housing demand shock that the MSA experienced during the national housing boom and bust. Many papers in the recent literature have concluded that the housing boom during the 2000s was caused primarily by larger changes in housing demand (see, for example, Shiller 2008). Furthermore, there is an emerging consensus that different MSAs in the U.S. experienced different house price appreciations during the boom primarily because of a combination of differences in the magnitude of changes in local housing demand (Davidoff 2015; Ferreira and Gyourko 2011) and differences across MSAs in the local housing supply elasticity (Mian and Sufi 2011).

To create a measure of local housing demand shocks, consider a log linear model of housing demand and housing supply. Holding local housing supply shocks constant, a local housing demand shock, ΔH_k^D, produces both a price and quantity change given by:

$$\Delta H_k^D = \eta_k^D \Delta P_k + \Delta Q_k, \tag{3}$$

where ΔP_k is the change in the log of local housing prices in MSA *k*, η_k^D is the price elasticity of housing demand and ΔQ_k is the change in log of new housing produced.[15] Using the fact that existing estimates of

higher age. This would be easy to capture in our framework by allowing the psychic costs of college attendance to vary with age as well as ability.

[15] The equation follows from a log linear housing demand function of the form $Q_k^D = H_k^D - \eta_k^D P_k$. In equilibrium, housing demand is equated to housing supply at the equilibrium level of housing Q_k. This equilibrium model is also used to motivate alternate proxies for changes in housing demand including $\Delta H_k^D = (\eta_k^D + \eta_k^S)\Delta P_k$, where η_k^S is the local housing supply elasticity.

the elasticity of housing demand in the literature suggest that $\eta_k^D \approx 1$, we create a proxy for local housing demand in any period, $\widehat{H_k^D}$, which is simply the sum of the log of price and the log of housing produced in the MSA.[16] The change in this proxy over any two periods, $\widehat{\Delta H_k^D}$, measures the change in local housing demand over that interval.

Our proxy for changes in housing demand is a function of both changes in local housing prices (ΔP_k) and changes in local housing supply (ΔQ_k). Theory says both changes in housing prices and changes in housing supply should affect local labor markets. Increases in housing supply can directly stimulate the local construction industry. Increases in housing prices can stimulate local employment either a housing wealth effect on consumer spending or through a relaxation of liquidity constraints (Mian and Sufi 2014). Additionally, both the housing price and housing supply channels can increase the volume of housing transactions which stimulates sectors associated with the selling and financing of housing (e.g., mortgage brokers, real estate agents, etc.). This discussion makes clear that it is theoretically ambiguous whether the housing supply effect on local labor markets is weaker or stronger than the housing price effect on local labor markets. In our baseline specification, we combine the two effects together in one metric, which assumes that the labor market effects are similar. We describe evidence below that is consistent with this assumption being approximately true in our setting.

We get local housing price information from the Federal Housing Finance Agency (FHFA) annual series on prices in FHFA metro areas, which we match by hand to Census/ACS MSAs.[17] We measure local housing supplied by the number of new privately owned housing units authorized via permits within the market. We match information on building permits from the Census Building Permits Survey to Census/ACS metro areas by hand using the MSA codes in the permits data. Merging the Census/ACS data with the FHFA and Building Permits Survey data produces 275 MSAs, which constitute our analysis sample of local labor markets.[18]

Figure 6 plots trends over time at the median, 10th percentile and 90th percentile for our local housing demand measure (the sum of log permits and log prices in an MSA). The figure shows variation at all three percentiles points, with particularly dramatic changes at the 90th percentile of MSA over the course of the boom and bust compared to changes at the median and 10th percentile. It is this large variation across MSAs

---

Our results are very similar under both specifications. We prefer equation (3) because it does not rely on assuming that the housing market is always in supply-demand equilibrium, which may be a poor assumption during the housing boom.

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

[16] The assumption of a unitary housing demand elasticity is justified by taking the average of the two most widely-cited estimates of the housing demand elasticity in the literature: 0.7 from Polinsky and Ellowood (1979) and 1.2 from Houthakker and Taylor (1970).

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

[18] Of the 283 MSAs with labor market data, 8 of them have missing house price data and/or missing housing permits data during the 2000-2006 time period.

that our analysis exploits. A 1 log point change in the housing demand measure between 2000 and 2006 corresponds to approximately the 90/10 percentile difference in the distribution of the 2000-2006 log changes across MSAs. The standard deviation across MSAs in the 2000-2006 changes in 0.57.

Our empirical work examines how different measures of education attainment and labor market outcomes are affected by local housing demand shocks. A key problem we face is measurement error in our housing demand shock which could lead our estimates to be attenuated. There is some unavoidable error in the dating of the start and end of the boom in an MSA, and the information on prices and permits that we use to create the measure of housing demand are only noisy proxies of underlying housing demand. A second challenge is that changes in housing demand in an MSA might be correlated with latent factors, such as latent amenity shocks, labor demand shocks, or labor supply shocks that could independently affect education or labor market outcomes. This would cause bias of indeterminate sign in the OLS estimates. To account for both measurement error and endogeneity problems, we supplement our OLS analyses with Two Stage Least Squares (TSLS) models that exploit plausibly exogenous variation in local housing demand arising from speculative activity.

Our strategy for isolating this exogenous variation draws upon the emerging consensus that much of the variation in housing prices, production and transactions during the national boom and bust was not the result of changes in traditional fundamentals like latent productivity, income or population, but rather was the result of factors specific to the housing market. These explanations include irrational exuberance and "bubbles" (Shiller 2009, Mayer 2011, Chinco and Mayer 2014, Glaeser and Nathanson 2014), the introduction of market products like interest-only mortgages (Barlevy and Fisher 2010), and changes in national lending standards (Favilukis, Ludvigson, and Van Nieuwerburgh 2010).[19] The combination of these forces caused widespread speculative investment in housing assets, with dramatic increases in housing prices, production, and sales until the bubble eventually burst.

To create our instrument, we build on the work of Ferreira and Gyourko (2011) by looking at rapid changes in housing prices that occurred in the local area between the 2000 and 2006 period. Since it is usually assumed that underlying fundamentals do not change abruptly, and are smoothly incorporated into prices when they do change, sharp breaks from the trend in a market's quarterly housing price arguably reflects variation that is the result of exogenous speculative activity or other housing specific forces, rather than unobserved changes in fundamental factors that are the major source of endogeneity concerns in OLS analysis of labor market and education outcomes. As Ferreira and Gyourko (2011) show, these sharp

[19] Burnside, Eichenbaum, and Rebelo (2015) also discuss "fads" in beliefs about patterns of future prices as an important force during the recent housing boom.

changes occurred at different times in different locations, which casts doubt on the extent to which national shocks can explain the sharp change in prices within local areas. [20]

Figure 7 illustrates how we use this insight to create an instrumental variable for local housing demand changes.[21] The figure plots quarterly housing prices for six MSAs between the first quarter (Q1) of 2000 and the last quarter (Q4) of 2005. For the three cities on the left side of figure, the smooth evolution of prices over time suggests that all or most of this change could have been the result of latent unmeasured fundamental factors, which "smoothly" affect demand. By contrast, for each of the three cities on the right side of the figure, the price series changed discontinuously ("sharply") at some point in the 2000s, suggesting the influence of some factor different from smooth changes in fundamentals, such as the effect of a speculative bubble.

Using the quarterly price series of each MSA between 2000Q1 and 2005Q4, we estimate MSA-specific OLS regressions with a single structural break, and search for the location of the break which maximizes the R^2 of the following regression:

$$P_k^H(t) = \omega_k + \tau_k t + \lambda_k (t - t_k^*) 1\{t > t_k^*\} + \zeta_{k,t} \tag{4}$$

In equation (4), $P_k^H(t)$ represents the log local housing price in MSA k in period t, which is a given year-quarter; t_k^* is the date of the structural break in the MSA's time series, restricted to be between 2001Q1 and 2005Q4; τ_k is an MSA-specific linear time trend before the structural break; and λ_k is the size of the MSA-specific structural break - the extent to which the growth rate of MSA's quarterly house price series changed at the break. Like Ferreira and Gyourko (2011), we find that the structural breaks occurred at different times for different MSAs. For MSAs whose house prices evolved fairly linearly over the time period, our estimates of λ_k will be close to zero.

Figure 8 shows the very strong positive relationship between the size of an MSA's estimated structural break and the 2000-2006 growth in housing demand in the MSA. We conduct of variety of formal econometric investigations of the "first-stage" relationship shown in the figure, all of which confirm the pattern evident in the figure. In particular, the structural break strongly predicts the 2000-2006 MSA change in housing demand after accounting for full set of standard controls. The *F*-statistic on the structural break measure in these analyses is always far larger than 20 for housing demand changes, removing any "weak

[20] Ferreira and Gyourko (2011) conclude that many of the MSAs that experienced a structural break in housing prices during the 2002-2005 period experienced no contemporaneous jump in income or any other traditional demand shifters when the structural break in house prices occurred. These "late booming" MSAs – which are the focus of our instrument – have price movements consistent with a speculative bubble.

[21] We are grateful to Edward Glaeser for discussions that encouraged us to formulate this empirical strategy.

instrument" concerns.[22] The structural break also strongly predicts the housing demand change during the 2006-2012 bust period, a result that follows from the fact that the size of the boom an MSA experienced was nearly perfectly correlated with the size of its later housing bust.[23]

Our interpretation is that these structural breaks that we identify in the 2000-2005 period are the result of speculative activity in the local areas. It is natural to wonder if these structural breaks are actually capturing exogenous shifts in speculative activity, as we argue, or are they instead reflect changes in some latent confound in the MSA. Formally, our goal is to use these structural breaks as instruments for changes in housing demand that are orthogonal to other factors that would drive local labor markets and/or educational choices. The six graphs in Figure 9 plot the relationship between the size of an MSAs structural break, λ_k, and pre-existing features of the MSA: average housing prices in the MSA in 1990; lagged housing price growth in the MSA between 1990 and 1995; average employment and wages in the MSA in 1990; and the growth in per capita enrollment in the MSA in two- and four-year colleges from 1990 to 1995. Strikingly, the figure shows that the structural break does not systematically vary with *any* of these pre-existing MSA-level variables, which is consistent with the findings in Ferreira and Gyourko (2011). Of course, these patterns obviously do not rule out the possibility that the structural break is related to *some* latent confound, but it is reassuring that λ_k exhibits no association with key pre-existing observable variables that one would think are likely closely related to latent factors that would raise obvious endogeneity concerns.

The two graphs in Figure 10 offer some evidence that the structural breaks indeed capture exogenous speculative activity rather than sharp changes in the underlying factors that determine labor market or education outcomes. The first graph relates the size of the structural break to the change in the price/rent ratio in an MSA, using data on rental price information that we have calculated for each MSA.[24] To understand what this graph tests for, assume that there is a sudden change in amenities, productivity or similar latent "fundamental", which discontinuously raises the desirability of living in an MSA. The current price of *all* housing in the MSA, whether to own or rent, should rise discontinuously in this case. In other words, there should be no relationship between λ_k and the price-rent ratio in an MSA if the break identified sudden changes in the latent fundamentals that give rise to endogeneity concerns regarding *current* employment, wages and schooling. By contrast, if the structural break reflects price changes from speculative investment purchase, based on investors' (perhaps incorrect) judgments about the likely *future*

[22] We present first stage regression results in Appendix Table A1.

[23] Online Appendix Figure OA.1 illustrates this pattern.

[24] See the Online Appendix for a detailed description of the construction of rental price data using the Census/ACS data. In the Online Appendix we report 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 very similar to the main results in Table 3.

desirability of the MSA, the price of owning should rise relative to that paid by renters, and an MSA's structural break should be positively related to growth in its price to rent ratio. This is precisely what the graph 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.[25]

More evidence that the structural breaks represent changes from speculative activity comes from the second graph in Figure 10. In recent work, Chinco and Mayer (2014) have carefully assembled data from transaction-level deed 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 present 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 twenty-one markets, the second graph in the last row of Figure 10 shows that, at least for this sub-sample of MSAs with available data, our structural break variable is strongly correlated with growth in the share of buyers who are speculative out-of-town buyers.

Taken together, the evidence in Figures 9 and 10 suggests that the estimated structural breaks identify plausibly exogenous variation in housing demand. Using TSLS, we estimate first-difference models of the effect of the 2000-2006 change in housing demand on the change over the same time period in education and labor market outcomes, using the structural break we estimate (converted to an annualized growth rate) as an instrumental variable for the change in housing demand. In addition, we exploit the information provided by the size and precise timing of structural breaks about exogenous changes in housing demand to estimate other regression models described later in the paper.

4. CHANGES IN OPPORTUNITY COSTS AND EXPECTED LIFETIME PREMIUM FROM HOUSING BOOMS

In this section, we assess how young adults' opportunity cost of (and their expected future lifetime earnings gain from) college attendance were affected by the boom, before turning to our main analysis studying different aspects of educational attainment.

We assume that a young adult who attends college in a given year foregoes the equivalent of the average labor market income received that year by persons in his MSA of roughly the same age who have no college

[25] 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 that allows them to follow rents for specific properties, though the data only exist for 43 large MSAs. Using their data, we find that there is a strong relationship between the structural break instrument and changes in the price-to-rent ratio, similar to what we show for the full sample of 275 MSAs.

training. His best estimate of the future lifetime premium from having gone to college is taken to be the current mean difference in labor market outcomes between older adults in his MSA with and without a college education.

We estimate mean labor market outcomes in an MSA from the 2000 Census and from several years of data from the American Community Survey (ACS).[26] We restrict the Census/ACS sample to non-institutionalized persons living in an MSA in their state of birth. This "same state" sample restriction partially accounts for any potential confounding from endogenous migration of the type shown to accompany other types of local demand shocks (Blanchard and Katz 1992; Bound and Holzer 2000; Notowidigdo 2013). Likewise, this restriction excludes all foreign-born individuals, mitigating the concern that our results are being driven by compositional changes in the local area due to both international migration and the intrastate migration of immigrants (Cadena and Kovak 2015). We also always exclude from Census/ACS samples anyone living in group quarters. Using the Census/ACS samples, we explore three separate time periods: 2000, 2006, and 2012. Averages for the year 2000 are estimated using the 2000 Census. To compute the labor market and education averages in the years 2006 and 2012, respectively, we pool ACS data from 2005 to 2007 (and refer to it as 2006) and from 2011 to 2013 (and refer to it as 2012). We pool the data in the ACS to increase precision given that our analysis is always conducted at the level of MSA observations.[27] For the Census/ACS analysis, we cannot explore the years between 2000 and 2005 because the ACS does not provide information on the individual's MSA of residence.

Using the Census/ACS sample, we estimate first-difference regressions of the form:

$$\Delta\overline{Y}_{kt} = \gamma_0 + \gamma_1 \widehat{\Delta H_{kt}}^D + X_{kt}\Gamma + \nu_{kt} \qquad (5)$$

where $\widehat{\Delta H_{kt}}^D$ and $\Delta\overline{Y}_{tk}$ are, respectively, the change in housing demand (defined above) and the change in the average labor market conditions that proxy for opportunity costs and expected future lifetime college premium in MSA k between periods t and $t+s$. The first difference specification in (5) accounts for the effect of latent fixed MSA-specific factors. The control vector X_{kt} in (5) is designed to control for any factors that could cause differential trends in labor market conditions across MSAs. This vector includes controls for the share of employed workers with a college degree, the share of women in the labor force, the fraction of the MSA that is foreign-born, and the log of the MSA's total population as measured in 2000.[28]

---

[26] We use the Census and ACS individual-level and household-level extracts from the Integrated Public Use Microsamples (IPUMS) database (Ruggles et al., 2004).

[27] When computing house price growth over the boom, we examine the change between 2000 and the relevant measure in the first quarter in 2005, 2006 and 2007. For the housing supply proxy, we calculate the change between annual average between 2001-2006 and 1998-2000.

[28] When calculating these MSA-level control variables, we use the "same state" restriction on our Census/ACS sample and include all adults between ages of 18 and 55.

Standard errors in all our analyses are clustered by state. Lastly, all regressions are weighted by MSA young adult population (age 18-33).

Since we mainly focus on the education choices of 18-25 year olds, we use the average labor market outcomes of non-college 18-25 year olds to measure opportunity costs. This group includes all individuals with just a high school degree (or equivalent) and high school dropouts. We measure both the employment rate and average wages for this group, averaging across individuals in each MSA-year. To compute individual wages, we divide the individual's reported annual earnings from the prior year by an estimate of their reported annual hours worked over the prior year.[29] To compute the skill premium within each MSA in each period t, we focus on the wages of 26-55 year olds for those with and without any college education. In the language of our model outlined above, this is our estimate of the individual's expected future college premium.[30] The Online Appendix that accompanies this paper provides a further description of the construction of all variables used in the paper.

Table 1 presents estimates of the effect of the housing boom on opportunity costs, using two different measures of what a young adult gives up in terms of labor market rewards by going to college in a given year. The table presents both OLS and TSLS results. Subsequent tables show only the preferred TSLS results; all corresponding OLS results for all other tables are presented in the Online Appendix. We show both sets of estimates in Table 1 to give a flavor of the pattern of results that we consistently find across the relationships we study: strongly significant TSLS estimates that are larger than their OLS counterparts, although the latter are consistently relatively large and generally statistically significant across all of our results.

The first column of Table 1 presents the results for the average prevailing employment rate among young adults without a college education. Both the OLS and TSLS results show that 2000-2006 growth in housing demand in an MSA raised employment among non-college educated young adults overall, and for men and women separately. The OLS results in the top panel suggest that in an MSA experiencing a 1 log point larger increase in housing demand between 2000 and 2006, mean employment rate was 3.1 percentage points higher among all 18-25 year olds and 3.0 and 3.2 percentage points higher among 18-25 non-college men and women, separately. The corresponding preferred TSLS estimates are 5.1, 5.9 and 4.3 percentage points. These effects, which are all strongly statistically significant, are relatively large given that the mean

[29] Specifically, to compute mean wages in an MSA during a given time period, we start with the same analysis samples described earlier. We then impose the following restrictions: (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 have exceeded $5,000, and (3) the individual must have worked at least 48 weeks in the prior year. Given these restrictions, our measure of average wages is for full-time workers with relatively few non-employment spells. To estimate annual hours worked last year, we multiply usual hours worked per week by the number of weeks worked last year.

[30] We also conduct robustness tests where the age range used to calculate opportunity costs is 18-33, and the age range for expected future returns was either 26-45, 33-45, or 33-55, and we found that our main results are robust to these alternative definitions.

employment rates for all non-college educated 18-25 year olds and for men and women separately, were 60.6, 64.3 and 56.5 percent, respectively. A one standard deviation change in housing demand across MSAs was 0.49. As a result, a one standard deviation change in housing demand was associated with a 2.9 and 2.2 percentage point increase in employment rates for 18-25 year old non-college men and women, respectively.

The estimated effect on housing demand shocks on wages for these young non-college workers are also relatively large and strongly significant. The TSLS results for log wage in the second column show that in an MSA experiencing a 1 percentage point larger increase in housing demand, young adults going to college forego 11.5 percent more in wages, with very similar effects for young men (11.1 percent) and young women (11.7 percent). Given the large increase in employment that the boom caused among non-college educated persons, some portion of this estimated wage effect may reflect compositional effects rather than increased returns from an hour of work.[31] Even with this caveat, both the OLS and TSLS results for employment and wages suggest that the boom substantially improved labor market opportunities for both young adult men and young adult women without college educations. In the third column of the table we present results for a summary measure of labor market conditions that we use elsewhere in the paper: the product of wages and the probability of employment. The estimates show that that a one standard deviation change in housing demand results in a 8.8 percent increase in wages adjusted by the probability of finding a job in the first place, for the pooled sample of men and women (0.179 * 0.49).

Were certain sectors particularly responsible for the improved labor market opportunities for young men and women presented in the first three columns? To the extent that people associate the housing boom with large increases in the building and renovation of houses, construction probably comes naturally to mind as a sector that ought to have been profoundly affected by the housing boom. As discussed above, the boom also involved massive changes in the volume of housing transactions – the amount of houses bought and sold. Many person performing the various tasks necessary for a sale to be consummated – things like advertising, listing, "showing", titling, insuring, procuring financing, etc. – would have been employed in the so-called "FIRE" sector of finance, insurance and real estate. Lastly, a broad set of sectors in retail and local services likely also responded to changes in housing demand through consumption increases coming from housing wealth effects or reduction of liquidity constraints (Mian and Sufi 2014).

Columns 4 and 5 of Table 1 present estimates of the effect of housing booms on the employment rate in the construction sector and the FIRE sector. The table shows both the point estimates for employment changes in the two sectors, and the ratio of those estimates divided by the overall employment effect from

---

[31] To assess the importance of composition effects, Online Appendix Table OA.8 reports results which use an MSA-specific wage growth measure that is residualized for age, race, gender, and marital status, following Shapiro (2006). We find very similar results with this alternative wage growth measure, which suggests that observable changes in composition do not account for the main results for average wages. In the same table, we also present estimates which adjust average wages for changes in average rents (to create estimate of change in real wages rather than nominal wages).

column 1. These ratios measure how much of the total employment effect from housing demand shocks for a given type of worker can be accounted for by changes in construction employment (column 4) and changes in FIRE employment (column 5).[32]

Focusing on the 2SLS results, our estimates show that 58.6% of the employment effect for young non-college men is concentrated in the construction sector while only 12.8% of the employment effect for young non-college women is in the construction sector. These results make intuitive sense given that young non-college men are much more likely to work in the construction sector. Conversely, our estimates show that 40.6% of the increase in employment for young non-college women can be traced to the FIRE sector (real estate agents, mortgage brokers, etc.). The comparable number for men is only 7.5%.[33]

To assess the robustness of the main 2SLS results, we report results from a wide range of alternative specifications in the Online Appendix, focusing on alternative control variables, alternative proxies for the change in local housing demand, and alternative ways of constructing the structural break instrument. We find similar results across different combinations of controls such as region fixed effects, controls for local manufacturing employment, routine employment, and local unemployment rate. We also find similar results using alternative proxies for changes in local housing demand, such as using the elasticity-weighted house price change instead of the sum of housing prices and housing permits. Lastly, we report results for alternative ways of constructing the structural break instrument, such as setting the structural break to 0 if the estimated structural break is not statistically significant. Overall, we find fairly similar results across all of our main labor market outcomes in all of these specifications.

In addition to the robustness analysis, we also tried to estimate whether the changes in opportunity costs were primarily driven by changes in housing prices or by changes in housing supply, since our primary housing demand measure combines them together, implicitly assuming that the employment and wage effects are similar. To assess whether this is a reasonable assumption, we carried out two exercises. First, we included both ΔP_k and ΔS_k as separate variable in our estimation of (5) and estimated this equation via OLS since we do not have separate instruments for each component. The coefficients on ΔP_k and ΔS_k were fairly similar, suggesting that both higher housing prices and the construction of more homes increased non-college employment and wages. Second, we explored whether housing demand changes had differential labor market effects in areas where housing supply is relatively elastic, using the local housing supply elasticities estimates from Saiz (2010), interacted with both our housing demand proxy and the structural break measure. Although precision is somewhat limited, we find no evidence that changes in local housing

[32] Note that the actual ratios, which are based on the division of the actual point estimates, sometimes differ slightly from the ratio of the rounded point estimates in the tables.

[33] We looked at several other sectors and found no meaningful effect of changes in local housing demand on employment in manufacturing, mining, and utilities. We therefore conclude that the remaining employment effect outside of construction and FIRE is accounted for by a broad range of jobs in the local retail and service sectors.

demand had differential effects in places where housing is inelastically supplied relative to places where housing is elastically supplied. We therefore conclude that our assumption of similar labor market effects of ΔP_k and ΔQ_k is a reasonable approximation that we carry through the rest of the analysis.

The results in Table 1 show clearly that the boom substantially increased the opportunity cost of college-going for both young men and women, although these increases came in different sectors. Table 2 explores how the boom changed the expected college earnings premium that a young adult could have expected to earn in the future. We estimate this by comparing the labor market outcomes of older, prime-aged persons (age 26-55) with and without any college education. We focus on the same labor market outcomes in the first three columns of Table 1, but the dependent variable is now the change over time in the *difference* in labor market outcomes between those with at least one year of college education to those without any college education.

Recall from the discussion in Section 2 that if housing booms raised the expected future college/non-college labor market premium, that effect would tend to offset any negative response to the opportunity cost changes presented in Table 1. The results in Table 2 argue strongly against this possibility. The TSLS point estimates indicate that local housing demand shocks *lowered* the employment rate gap between college and non-college working adults, strictly reducing the future college/non-college gain that a younger adult might reasonably have expected from getting a college education.

The estimates in the second column show that local housing booms did not meaningfully change the expected future college/non-college wage gains. In contrast to the employment rate results, local booms did not significantly increase or reduce the college/non-college wage gap among older working adults, with estimates generally close to zero. The results for the future wages weighted by the probability of finding employment (shown in Column 3) are similar to the employment rate results. In sum, the results in Table 2 show that a young adult during the boom years, trying to form a conjecture of how increasing housing demand in his local area would affect his future market returns from different education paths, would have reasonably concluded that the boom either had no effect on lifetime labor market streams from the college versus non-college path, or else potentially *reduced* the earnings and employment gain from becoming college educated. Nothing about expected future gain would have tended to militate against the effect of rising opportunity costs.

5. Housing Demand Shocks and College Attainment during the Boom

In this section, we present the paper's main results, which assess how local housing demand shocks during the boom affected young adults' college-going. We use a variety of methods and information about college-going from three different data sources. Moreover, as we show below, the particular limitations of each data source are strengths of at least one of the other two. Combining a range of estimation methods and different data sources therefore allows us to carry out a comprehensive investigation.

5.1 Census/ACS Estimates

Our first education results use self-reports of schooling attainment in the "same state" Census/ACS sample combining the 2000 Census and the 2006 ACS. Our primary measure of educational attainment is the fraction of individuals in a given age range with any college attainment regardless of degree completion. We refer to this measure as "Any College". Our second measure of educational attainment is the fraction of individuals in a given age range who completed at least a bachelor's degree. We refer to this measure as "At Least Bachelor's Degree". We calculate the mean educational attainment rates in an MSA among 18-25 year olds in 2000 and among 18-25 year olds in 2006. As we have shown, the latter group faced much higher opportunity costs of acquiring a college education when they made their college going decisions if they lived in an MSA experiencing a housing boom, so we would expect their average attainment levels to be lower in housing boom. We estimate first-difference regressions of the form

$$\Delta S_{kt} = \alpha_0 + \beta_{FD} \widehat{\Delta H_{kt}}^D + X_{kt}\Gamma + u_{kt}, \tag{6}$$

where ΔS_{kt} is change in average educational attainment among 18-25 year olds in an MSA between periods t and $t+s$. The coefficient β_{FD} is the first-difference estimate of how the growth in housing demand in an MSA affected the change in college attendance among young adults in that MSA.

Table 3 presents the TSLS results from estimating (6) using the structural break λ_k as an instrumental variable. The two columns show the results for the dependent variable defined for "Any College" and "At Least Bachelor's Degree", respectively. The first three panels show the results where the education measures are defined for all individuals in the 18-25 range, just males in this range, and just females in this range. The last panel shows results for a placebo specification where we look at changes in college attainment for all 26-33 year olds.

The results indicate that rising local housing demand during the national housing boom sharply lowered the fraction of 18-25 year olds with "Any College", with estimated effects that were very similar for men and women. The strongly statistically significant point estimates imply that a one percentage point increase in local housing demand reduced the fraction of 18-25 year olds who completed any amount of college training by about 2 percentage points – 1.8 for men and 2.4 for women. To further help interpret the magnitudes, a one-standard deviation increase in the local housing demand shock reduced the propensity of any college by roughly 1 percentage point for both men and women. As a benchmark, roughly 43 percent of men and 51 percent of women between the ages of 18 and 25 had any college attainment in 2000. By contrast, the second column in the table shows that the growth in local housing demand had no effect on the fraction of 18-25 year olds with at least a Bachelor's' degree. It is not only that the effects are statistically insignificant; the point estimates are tiny compared to the point estimates in column 1.

The point estimate of 0.021 for "Any College" attainment for me and women represents a 4.5 percent reduction in the fraction of 18-25 year olds with "Any College" training relative to the rate in 2000. The difference across the columns suggests that housing booms lower college attainment among 18-25 year olds almost entirely by decreasing the share of them completing some level of college less than a Bachelor's degree, with no effect on the share completing a Bachelor's degree. Because of the way education is coded in the Census/ACS, the "Any College" group consists of people who have completed an Associate's degree and persons with college training who have not completed any degree, but for whom we do not whether the college-training they received was at the Associates or Bachelor's level.[34] We nonetheless believe that is possible to say something about the likely implied effect of the point estimates in column 1 for just the portion of "Any College" sample receiving Associates-level training. To do this, we estimate what share of 18-25 year olds with "Any College" in 2006 were likely studying towards a Bachelor's degree.

As the table shows, among persons aged 18-25 in 2000, 46.9% of them had completed "Any College" and 10% had completed at least a Bachelor's degree. We also report in Table 3 average attainment in 2006 of persons who *had been* 18-25 in 2000. These means are 61.9 percent for "Any College" and 28.5% for at least a "Bachelor's". These numbers suggest that many 18-25 year olds in 2000 who would ultimately complete a Bachelor's degree had done not so as of 2000. Since Bachelor's degrees can take many years to complete, many of these people were probably studying at Bachelor's colleges in 2000 when they were between 18 and 25. Even though they did not complete the degree by this age, they were engaged in Bachelor's *training* in 2000. Assuming that all those completing Bachelor's degrees by 2006 had already started Bachelor's training by 2000, the share of 18-25 year olds with only Associate's level training would be

[34] In the Census/ACS IPUMS education codes, college-trained persons without degree are either "some college, but less than a year" (code 065) or "1 or more year of college credit, no degree" (code 071). There is no information about the *type* of college training these persons received.

no larger than 18.4% (46.9-28.5). If housing demand shock only affected Associates level attainment as the results hint might be the case, then the implied magnitude of the -0.021 point estimate relative to the share of people with Associate's level training before the boom would therefore represent at least 11.4 percent decline relative to the share of population studying at Associate's colleges. This implied effect is likely even larger since not all people who doing Bachelors' training in 2000 ultimately go on to complete a degree, which would reduce the share engaged in Associate's training even lower. Overall, these calculations suggest significant declines in Associate's training as a result of housing boom.

As a placebo specification, the last panel of results presents estimates for persons aged 26-33. The conceptual model emphasizes that older households should be less likely to respond to the housing boom. The results show, reassuringly, that a sample older than our 18-25 age group of interest did not respond to the housing demand shock. Additionally, the fact that we find no effect of housing boom on Bachelor's attainment in this older group suggest that the null effect we find for 18-25 year olds is not simply because they are too young to have completed their Bachelor's degrees.

Our finding that increases in local housing demand during the national boom lowered mean college attainment, with almost all of the effect coming from schooling that is less than a Bachelor's degree, is consistent with the predictions of our conceptual model emphasizing the role of opportunity costs. The fact that the massive Census samples allow us to precisely estimate means at the start and peak of the boom for relatively narrow birth cohorts by MSA is a unique advantage of this data source. This notwithstanding, there are some important limitation of the Census/ACS data.

One concern is that the Census/ACS education self-reports may be unreliable. This is an especially important concern because there is some evidence that the errors is self-reported education tend to be non-classical, with people claiming higher educational attainment than is suggested by other types of evidence (see Filmore 2014). A second concern, which has been already noted, is that Census/ACS data do not allow us to determine whether a person with college training but who has not finished a degree *had* been working towards a degree at a community college or a four-year university. This makes it very difficult to accurately characterize the type of college training received by an important part of the sample, calling into question any firm conclusions about the differential responses across different types of colleges. Third, the fact that the 2001-2004 ACS samples do not record the MSA of respondents prevents us from doing high frequency analysis that better allows us to exploit the differential timing of the boom across MSAs.

Finally, there is the important problem that because the Census/ACS data are pooled cross-sectional samples, we cannot definitively link people to their MSA at different points during the housing boom. Even with our sample restricted to persons living in their state of birth, we do not know their MSA at the start of the boom for persons who moved across MSAs *within* that state during the boom. The Census/ACS results may thus still be confounded by endogenous migration, even in the "same state" sample.

The next two data sets used to study schooling address these concerns and also permit a number of rich, alternate analyses that are not possible with the Census data. We discuss these two data sets, and results from them, in turn.

5.2 IPEDS ESTIMATES

Our second source of information on educational attainment is the Integrated Postsecondary Education Data System (IPEDS). The IPEDS is constructed from administrative data on enrollments reported annually by most of the colleges and universities in the U.S, including both community colleges and four-year colleges and universities.[35] The data tracks first-time, full-year enrollments, enabling us to identify persons enrolling in college for the first time during the boom. We match colleges and universities to MSAs by hand, and compute MSA-specific estimates of total first-time, full-year enrollments for different types of colleges and universities in each year between 1997 and 2006.

There are several important strengths of the IPEDS data. Because the enrollment information are from administrative records, they are likely less error-ridden than survey self- reported schooling attained in the Census/ACS. In addition, the IPEDS data specifically reports enrollment for different types of colleges separately, permitting a precise characterization of the type of college training for every enrolment (Associates' or Bachelor's' levels, 2 or 4 year). Something analogous is not possible with the Census data for persons whose completed schooling is less than a Bachelor's degree. Finally, the annual IPEDS reports provide high frequency information about college level training, which allows for some econometric tests that cannot be done with other data sources.

The main shortcoming of the IPEDS is that it is not an individual-based survey, but is rather a survey of enrollments in institutions. It is not possible to measure enrollment by birth cohort using the IPEDS, as we do with the Census/ACS. Indeed, we know very little about precisely who the first-time enrollees in the IPEDS are, including exactly where they are from. In the analysis we assume that they are from the local market that houses that college or university, but this will likely not be true some portion of enrollments if people move across MSAs for their college training.

For all of the IPEDS analysis, we use the per capita enrollment rate in the MSA, calculated by adjusting total first-time enrollment totals by the size of the 18-25 population in the MSA based on interpolations between the 1990 Census, the 2000 Census and the 2006 ACS. The analysis focuses on the 224 MSAs for

[35] Unfortunately, for-profit universities are underrepresented in IPEDS data, and they are growing fast during the housing boom period. In principle, we should be able to capture these educational investments in Census/ACS self-reported educational attainment variable.

which enrollment data are reported every year between 1997 and 2006. The sample of MSAs is thus smaller than the 275 MSAs used in our Census/ACS results above. We divide all IPEDS enrollments into two mutually exclusive and exhaustive categories: community college enrollments and 4-year colleges and universities. The community college category includes junior colleges and technical colleges, while the 4-year colleges and universities category includes all institutions that award bachelor's degrees.[36]

We perform two types of analysis with the IPEDS data. The first set of results follows the analysis Census/ACS data and focuses on changes in average enrollments during the 2001-2006 relative to average enrollments in the 1995-2000 period. The primary advantage of this "long difference" specification is it allows us to use our TSLS specification where we instrument for the housing demand change with our estimated structural break (λ_k). The second set of results exploits the higher frequency annual administrative data to examine whether the specific timing and magnitude of the structural break instrument lines up with the timing and magnitude of the enrollment change within each MSA.

TSLS Estimates for Changes in Per Capita Enrollment

For our first analysis, we show the results from a TSLS estimation of a first difference model of the effect of the 2000-2006 change in housing demand on the change in average annual per capita enrollment from the 1995-2000 period (when people made enrollment decisions before the boom began) to average annual per capita enrollment during the 2001-2006 period (when people made decisions during the boom). We instrument for the change in housing demand, $\widehat{\Delta H_k^D}$, using the structural break. This specification is identical to regression specification (6) used previously for the various first-difference Census/ACS results aside from the change in dependent variable.

As the first column Table 4 shows, the TSLS results indicate that a 1 log point increase in an MSA's housing demand from 2000 to 2006 lowered the five year average of per capita annual enrollment in 2-year colleges by about 1.1 percentage points.[37] This statistically significant effect was similar for male and female enrollment. The second column of Table 4 presents results for enrollment in 4-year colleges. The contrast with the results in the first column is very striking. We find that the 2000-2006 growth housing demand in

[36] Note that some 4-year colleges will award associate's degrees, but we cannot separately measure enrollments of students studying towards associate's degrees in the IPEDS data, so they are included in the 4-year college category. This means that the total community college enrollment in our analysis sample is likely an underestimate of total enrollment of students studying towards associate's degrees. The Data Appendix provides more detail on these categorizations and the construction of our data set.

[37] Because of space constraints, we show the OLS specification in the Online Appendix that accompanies the paper. All of the OLS results show a negative and statistically significant association between housing demand in an MSA and enrollment in 2 year colleges for both men and women, and no statistically significant relationship between housing demand in an MSA and enrollment in 4 year colleges.

an MSA had no statistically significant effect on per capital 4-year enrollment. For men, the 4-year enrollment point estimates are very small, less than one-tenth the size of the corresponding 2-year enrollment estimate. The estimated effect for women is larger but it is also imprecisely estimated. Given the large standard errors on the Bachelor's results, we cannot rule out whether women actually increased their bachelor's enrollments, perhaps because of the liquidity constraint mechanism discussed in Section 2.

For the third and fourth columns of Table 4, we use IPEDS data from several years before the boom to assess whether the results in the first two columns actually capture the causal effect of housing boom, or whether the regressions might simply be picking up the effect of pre-existing trends. In these placebo tests, we measure whether the 2000-2006 change in housing demand predicts the growth in annual average enrollment from a *previous* time period – specifically the change in average annual enrollment during the 1987-1990 period relative to average enrollment during the 1991-1996 period. Reassuringly, the results show that current booms do not predict previous changes in average annual per capita for either 2- or 4-year college enrollment. This suggests that the estimates in the first two columns are not simply capturing long-term trends and indeed capture causal effects attributable to the housing demand shock.

What are the implied magnitudes of the main results in column 1 of Table 4, and how do they compare to the TSLS Census results for completed schooling in Table 3? Although enrollment is a flow measure of schooling and the Census highest schooling completed variable studied in Table 3 is a measure of the stock of college training, the two constructs should offer the same basic picture of the effect of housing booms, since the years of college that a person has completed as of given year is necessarily a function of their enrollment decisions in several separate years before year in question. It is therefore reassuring that the two sets of results give the same qualitative picture, of a strongly significant negative effect of the booms on Associates-level type training, with much smaller effects for Bachelor's type training.

Relative to the five-year average of annual enrollment in 2-year colleges of 4.7 and 5.3 percent in 2000 for men and women, respectively, the point estimates in Table 4 imply that a one standard deviation increase in housing demand reduced per capital male and female enrollment in Associate's-granting colleges by 11 percent and 10 percent, respectively. Comparing these effects to the implied magnitudes from the Census/ACS results is difficult for several reasons. There is, first, the fact that the IPEDS results are based a sample of 224 MSAs rather than 275 used in the Census analysis. This does not guarantee that the results would be quantitatively the same even if the measures used in the two studies were identical. A second issue that frustrates easy comparison across the two sets of results is that we do not know the ages of enrollees in the IPEDS, whereas all of the Census/ACS results focus on schooling completed by persons in particular age bins. Another important difference between the data sources that makes comparison of the magnitudes difficult is that whereas the Census data are limited to persons born in the same state, there is no information about where IPEDS enrollee are from. Thus, IPEDS results include not only enrollment decisions of native-born persons from other states, but also immigrants. Finally, our IPEDS results are based on first-time, full year enrollment, from which it is impossible to translate into completion rates for different types of schooling. Despite these challenges, we think our conclusion that the likely effect of on standard deviation boom on Associate's-level highest completion was about one percentage points is highly consistent with the effect sizes for per capita enrollment rates in the IPEDS.

Difference in Difference Estimates for Per-Capita Enrollments

The second type of exercise we conduct with the IPEDS exploits the exogenous variation associated with the MSA-specific information about the timing and size of structural breaks in a difference in difference (DD) model. Specifically, using annual per capital enrollment in a given MSA during a given time period, e_{kt}, we estimate

$$e_{kt} = \alpha_k + \delta_t + \beta_{DD}\left(\left(\text{Post } t_k^*\right) \times \lambda_k\right) + X_k\Gamma + v_{kt} \qquad (7)$$

where $\left(\text{Post } t_k^*\right)$ is an indicator variable denoting time periods after the date of the MSA-specific structural break, t_k^* - that is, all t such that $t \geq t_k^*$. The variable λ_k is size of the structural break, and α_k and δ_t are, respectively, MSA and year fixed effects. The DD coefficient β_{DD} measures how per capita enrollment in an MSA in the years after the structural break differs from enrollment in the years before the break, with this post-break/pre-break difference weighted by the size of the structural break. A very appealing aspect of (7) is that it tests whether there is a change in MSA enrollment that coincides with the break its housing demand. Since this DD estimate controls very flexibly for time effects, and for fixed features of the MSA that affect enrollment, the interaction arguably yields unbiased estimates of our effects of interest, so long as the time breaks are random, as we have argued throughout.

Table 5 presents the DD results. Column 1 shows that there was a strongly statistically significant reduction in per capita enrollment in 2-year, Associates'-granting institutions in an MSA in the years after the MSA's break, compared to the years before the break. These enrollment declines occurred for both men and women. To figure out the implied magnitudes of the point estimates in this table, it is useful to know that the standard deviation of the structural break variable across the different MSAs is about 0.06. This means that a one standard deviation larger structural break meant annual enrollment in 2-year, Associate's-granting institutions during the post-break period was lower by 0.4 percentage points (0.066*0.06). Relative to annual average enrollment before the national boom, our results imply that immediately after the structural break in an MSA, and for the each subsequent year of the boom, per capita enrollment in Associate's-granting institutions in the MSA fell by about 8.2 percent (0.004/4.9) relative to enrollment in the years before the break. This approximately 8 percent decline was true for both men and women.

The results for per capita enrollment in 4-year colleges and universities in the second column are very different. They show that there was no statistically significant change in per capita enrollment in these Bachelor's degree-granting institutions in the years after the structural break, relative to enrollment before. Although not precisely estimated, the point estimates in column 2 do suggest that there may have been a modest *increase* in enrollment in Bachelor's-granting institutions in the years after the structural break. This is particularly true for women. However, none of the enrollment responses in Bachelor's granting institutions are statistically different from zero. We emphasize that this effect would be perfectly consistent with our conceptual framework, which argues that besides the opportunity cost mechanism that is this paper's main focus, housing booms may have eased liquidity constraints for some persons. To the extent that this effect exists, Bachelor's-granting institutions is precisely where one would expect to observe it, since these institutions are more expensive. The relatively small number of people in an MSA whose college-going decisions are immediately changed by increased in homeowner wealth also probably makes the effect difficult to precisely detect empirically. Why the effect, if it exist, should be bigger for women than it is for men as the point estimates suggest is unclear.

The DD results suggest that there were changes in MSA enrollments in the years after a structural break in the MSA's housing prices, and that how large that change was varied positively with the size of the structural break used by various TSLS analyses. Both the timing and magnitude of break therefore appear to be important for explaining trends in college enrollment for two-year colleges, but not for four-year colleges and universities.

To assess the robustness of the DD results, we conduct a randomization inference exercise which permutes assignment of year and magnitude of structural break in house prices across the various MSAs.[38] We carry this out 1,000 times, and each time we re-estimate the DD equation (7) on the permuted data. We then compare the main DD estimates in Table 5 to the empirical distribution formed by the 1,000 different estimates. The results are reported in Figure 11. On the left side of figure, the results for two-year colleges show that the main DD estimates from Table 5 are unusually large in magnitude compared to the empirical distribution and statistically significant for men, women, and all adults. The results are stronger for two-year colleges than four-year colleges and universities, as in Table 5. These results confirm that the precise timing

[38] We randomly permute the pair of the magnitude and year of structural break to preserve the weak correlation between them in each permutation. If we permute the magnitude and year of structural break independently, the results are very similar.

and magnitude of structural break in house prices is important for explaining timing and magnitude of changes in enrollment in two-year colleges.

5.3 Individual Panel Results from NLYS97

The third data source we use to study the effect of the boom on college attainment is the restricted use version of 1997 National Longitudinal Survey of Youth (NLSY97), for which we obtained special permission. This individual-level longitudinal panel data set initially surveyed a random sample of American youth aged 12-16 in 1997 and has followed them since.

The age range of the NLSY97 sample and the timing of the survey are ideal for our study: at 15-19 years old in 2000, these young people would have been making college-going decisions right around the time of the housing boom. Because the restricted-use NLSY97 provides information about respondents' MSA in each survey year, we identify exactly where a person lived at the beginning of the housing boom, and irrespective of whether they moved subsequently.[39] As noted, this is something that is impossible to do with the individual level Census/ACS data. More generally, uniquely among our data sources, the panel nature of the NLSY97 data allows us to track outcomes for particular persons through time. The NLSY97 is also the only one of the data sources we use that includes a rich set of control variables measuring family background and other demographic characteristics, including parental education, race, ethnicity, and family composition.

The big downside of the NLSY97 is its small sample size. Our NLSY sample consists of only approximately 6,000 men and women, which effectively renders credible comparisons of means across the 275 separate MSAs impossible. We therefore must collapse the cross-MSA variation in our housing demand proxy to an indicator variable for whether or not the individual lived in a "housing boom MSA" in 1997 (which is defined based on whether the MSA was in the top third of the distribution based on the structural break instrument).[40]

Our NLSY97 sample is restricted to individuals with non-missing data on employment, educational attainment, and control demographic variables. The final sample is 5,923 individuals (2,979 men and 2,944 women). We estimate a series of regression, relating people's outcomes in 2006 to the whether they lived in a "housing boom" MSA in 1997, controlling for a rich set of demographic controls.

[39] We use the individual's metropolitan area of residence in 1997 and assign the housing demand change of that MSA to the individual over the entire time period, even if that individual moves elsewhere during the sample period. Approximately 20 percent of the sample relocates during the 1997-2006 time period.

[40] Even without constraints imposed by sample size considerations, our data agreement with the NLS for the use of the restricted-use NLSY97 prevents us presenting results separately by individual MSA.

Table 6 presents the results.[41] The first outcome variable, in column 1, is whether the person is employed at age 20. We show the results for this variable because the NLSY is the only data set which allow us to directly assess not only whether people in boom markets faced situations where labor market prospects for young, non-college educated adults in general, but whether they *personally* were more likely to participate in labor market activity. The results show clearly that this was the case. We find that young adults living in "housing boom" MSAs just before the start of the housing boom were more likely to be employed at age 20, with a particularly pronounced effect for young men. Overall, being from a "housing boom" market increased the probability of being employed at age 20 by 5.8 percentage points. For men, the increase was a strongly statistically significant 7.1 percentage points. Though somewhat smaller than the effect for men, and marginally significant, the employed effect for young women from boom markets was also substantial at 4.1 percentage points. These effects represent 10.3 and 6.0 percent increases relative to mean employment rates, respectively.

The second and third columns of Table 6 use the available information in the NLYS97 about schooling attainment to measure the effect of being from a housing boom market on our two measures of schooling attainment defined as of 2006: whether the person had attended "any college" and whether they had received at least a Bachelor's degree. The results show that the effect of being from a boom MSA on having gotten "any college" training by 2006 was negative, substantial and strongly statistically significant. Overall we find that young adults in "boom" MSAs when the national housing boom began were 5.8 percentage points less likely to have attended college at all by 2006. Interestingly, at -4.9 compared to -6.8, the percentage point lower attainment effect for men was smaller than that for women, although the two are closer in terms of magnitude and not statistically different from each other at standard levels. The point estimates in the third column for effect of being from a housing boom MSA on Bachelor's degree attainment by 2006 are very small compared to the corresponding "any college" results, and none is statistically different from 0.[42]

[41] In Online Appendix Table OA.16, we report results from alternative specifications to assess robustness of the main NLSY results. In main results, we focus on comparing "housing boom MSAs" to other MSAs based on whether or not the MSA was in the top-third of structural break instrument. We alternatively define "housing boom MSAs" based on the estimated change in housing demand between 2000 and 2006 (change in prices plus change in permits). Both of these approaches are binary comparisons grouping cities in boom/non-boom groups. We prefer these results because of limited sample size in NLSY data set. However, we also estimate results which use continuous measure of both housing demand instrument and housing demand change proxy. In all cases, we find a similar pattern of results: living through housing boom causes statistically significant decline in share of population with "any college" by 2006 and no statistically significant change in share of population with "Bachelors degree." The patterns are similar for men and women, together and separately, across each of the outcomes.

[42] In Online Appendix Tables OA.4 and OA.5, we repeat the analysis in Table 1 and Table 3 using the same indicator variable used in the NLSY97 analysis instead of the main housing demand proxy. The point estimates are somewhat smaller in magnitude, but we find a broadly similar pattern of results, with statistically significant increases in employment and decreases in "Any College", and no statistically significant change in "At Least Bachelor's Degrees."

Because the NLSY analysis tracks individuals over time, even as they move across MSAs, the findings in Table 6 are free of any concerns about endogenous migration. The broad similarly between the NLSY results and those presented earlier in the paper suggest that endogenous migration is also unlikely to be the primarily explanation for the pattern of results in the "same state" Census/ACS sample, and suggests that, despite our ignorance about the identity of enrollees in the IPEDS data, the enrollment results too capture true causal effects of being from a housing boom market rather than the effect of migration.

5.4 DISCUSSION

We find a consistent pattern of results across different data sources and methods: a negative effect of local housing boom on college attainment (or college going) with the virtually all of the reduction coming college training below the Bachelor's degree level.

The patterns are consistent with the "single index" conceptual model of college choice presented in Section 2. Young adults deciding between Associate's level training and labor force participation should be particularly sensitive to prevailing labor market conditions for less-skilled persons, as we find. The very small (or modestly positive) effect on Bachelor's level training could stem either from the fact that labor market conditions for young unskilled persons workers are irrelevant to the decisions of person thinking of going to Bachelor's training because of the larger gains from this type of college, or because the degree to which housing booms relieve liquidity constraints, and counteract the force of opportunity costs, is particularly important at Bachelors-granting universities and colleges compared to the much cheaper community colleges.

Another potential explanation for our results is the possibility the various estimated effects do not reflect changes in the behavior of potential students, but rather how colleges respond over the course of a boom. In particular, we would find the same results if, instead, it was simply easier for Bachelor's than for Associate colleges to expand to deal with the increase in the number of students arising from population inflows into MSAs experiencing large booms. The best available evidence about the ease with which different types of colleges can expand, or even their differential *desire* to expand to accommodate interested students, suggests that this line of reasoning is very unlikely to explain our results. Exploiting exogenous variation in class size arising from cohort sizes Bound and Turner (2007) show that Associates and two year colleges are vastly more supply-elastic than their Bachelor's-granting counterparts. Moreover, they argue convincingly that Associates degree granting institutions clearly prioritize "access" as compared to Bachelor's-granting institutions which tend to emphasize maintaining student "quality". Consistent with this previous work, we find no evidence that changes in housing demand affect the average cost of two-year colleges in an MSA (Online Appendix OA.14).[43] We therefore conclude that the effects we estimate during the boom were the result of the decisions of potential students in response to changes in their labor market opportunities, and not due to supply-side responses of colleges and universities to the boom.

We conclude discussion with a rough calculation of what share of the "slowdown" in the share of young adults with any college attendance shown in Figure 2. Extrapolating the pre-1996 trend to 2006 suggests a roughly four percentage point decline relative to trend for both men and women. What share of this "slowdown" can be accounted for by our estimated effect of the housing boom? To answer this, we apply our local housing boom estimates to national housing boom. Aggregating the housing price and housing permits data, we estimate a national change in housing demand of 0.55 between 1997 and 2006. Applying the 2SLS estimates in Table 3 to this change yields a predicted decline by 2006 of 1.16 percentage points, or roughly 30 percent of the aggregate "slowdown" for both men and women.[44]

6. Effects During Bust and Full Housing Cycle, and Persistence of Boom Effects

Our results thus far have focused on changes during the 2000-2006 national housing boom. In this section, we study the effect of local housing demand changes during the massive national housing "bust" shown in Figure 1, and over the 2000-2012 interval spanning the entire boom and bust cycle. The goal of this section is to assess whether the housing boom had persistent effects on the educational attainment of individuals. As shown in our simple theory, a housing boom when an individual is young may have a greater effect on their human capital choice than a corresponding housing bust when an individual is older. However, new cohorts may respond to the housing bust in a symmetric way as prior cohorts did when the prior cohorts experienced an equally sized housing boom.

We begin with an analysis of how changes in housing demand during these periods affected the labor market opportunity costs of attending college faced by different generations of young adults. Using data from the Census/ACS, we estimate first difference models relating the change in average labor market

---

[43] The primary measure of average cost is calculated for each MSA-year by averaging across the two-year colleges (or four-year colleges and universities) in the sample, weighting by the college's enrollment in that year. The average cost measure is calculated as the net tuition revenue across all students divided by the full-time equivalent enrollment. The net tuition revenue is net of all grants (e.g., Pell, federal, and state grants), so that changes in average cost reflect both changes in prices as well as changes in the availability and eligibility of grants. We also find no evidence that changes in housing demand affect "direct costs", which is the sum of published in-state tuition and costs related to fees, books, and other supplies.

[44] Using either the results reported in Figure 4 and the IPEDS "difference-in-difference" results lead to similar percentages of between one-quarter and one-half of the "slowdown" explained by the housing boom. Alternatively, instead of accounting for slowdown in Figure 2, we can account for slowdown in "Any College" across cohorts (as shown in Figure 3). This is less straightforward since some individuals in older cohorts could still have been "treated" by the housing boom, but assuming that no cohorts before 1975 were treated by the boom, then one can fit a pre-1975 trend by gender and compare predicted effect of housing boom.

outcome among 18-25 year olds without a college education between 2006 and 2012 (the bust), and between 2000 and 2012 (the full cycle) to the 2000-2006 change in housing demand in the MSA. Throughout all these analyses, our housing demand shock is defined over the 2000-2006 period. These regressions, which we estimate by TSLS using the structural break as an instrument for the change in housing demand, assess how the size of the boom an MSA experienced affected labor market conditions in the MSA over the course of the bust, and from the beginning to end of the housing cycle. As we have already noted, the size of an MSA's boom is nearly perfect correlated with the size of the bust it later experienced. Consistent with all previous Census/ACS results, 2006 values are the average in the Census/ACS from 2005 to 2007, and the 2012 values are the Census/ACS means from 2011 to 2013.

Table 7 presents results for the three different measures we have used to reflect a young adult's opportunity costs of attending college: the change in mean employment rates of 18-25 year olds with no college education in the MSA (first set of columns), the change in average wages of 18-25 years old with no college educations in the MSA (second set of columns), and the product of these two measures (final set of columns). We find that, compared to young adults of the same age at the peak of the boom in 2006, 18-25 year olds making college-going decisions after the bust faced labor markets with much worse immediate opportunities for non-college educated young adults, the larger the preceding local boom. In term of the probability of being employed, we estimate that an MSA with a one point larger increase in housing demand during the boom saw a 9.2 percentage point larger decrease in the employment rate of young men over the course of the bust. The employment rate decrease among young adult women was not as massive as the decline among young adult men, but at 6.4 percentage points, it was also quite large. The effects were much larger than the estimates in Table 1 of the increases in employment rates these MSAs during the boom.

The large negative point estimates in the second set of columns suggest that, even with the well-known tendency for wages to exhibit "downward stickiness" (Card and Hyslop 1997), average wages for non-college-educated young adults may have fallen during the bust by at least as much as they rose during the boom, but none of the wage estimates are statistically significant. The results for the wages adjusted by the probability of employment offer a nice summary of how labor market circumstances confronting young adults changed from the peak of the boom and the end of the bust. Local markets with a one point larger increase in local housing demand during the boom experienced a 23.5 percent decline in adjusted wages for young men without college educations, and a 15.9 percent decline for their female counterparts by the end of the bust.

Across all three outcomes, our results show that, at the very minimum, the bust totally erased the favorable labor market conditions wrought by the boom. Indeed, all the point estimates suggest that the bust may have made labor market conditions for unskilled young adults slightly *worse* than they had been before the boom began. This would be consistent with the recent arguments about the effects of debt during the housing cycle emphasized by authors like Mian and Sufi (2014).

Table 8 presents TSLS first difference estimates of the effect of changes in housing demand on the 2006-2012 and the 2000-2012 changes in college attendance. Columns (1) and (2) show results for the share of 18-25 year olds in an MSA with "Any College" training as measured in the Census/ACS.[45] Although not precisely estimated, the point estimates suggest that young adult making schooling during the bust, when labor market conditions for non-college educated young adults had worsened substantially from the peak of the boom, were more likely to have attended college at all compared to similarly aged people at the peak of the boom. The results indicate that by the end of the bust, the size of the boom that an MSA had experienced during the 2000-2006 period had no effect on the share of 18-25 year olds with "Any College" training compared to what had been the case for young adults in that MSA before the start of the housing boom and bust cycle. Both of these Census results are consistent with the organizing theoretical framework about the effect of opportunity costs described above. As labor market conditions returned to normal (or slightly worsened) during the bust, young individuals started going back to college.

The remaining columns of Table 8 show results for changes in per capita enrollment using administrative IPEDS data. Following the specification in Table 4, the outcomes variable for the bust estimates is the difference between average annual enrollment during 2002-2006 and the average of annual enrollment during 2000-2012. Similarly, for the results from the start of the housing cycle to the end of the bust, the outcome variable is difference between average annual enrollment during 1996-2000 and the average of annual enrollment during 2000-2012. The results show that enrollment in two year colleges was higher during the years of the bust, by marginally statistically significant amounts, the bigger the size of the MSA's preceding boom. By the time that labor market conditions had essentially returned to levels seen before the start of the boom and bust cycle, annual enrollments were no different from what they had been in 2000, irrespective of the size of the preceding boom. This enrollment behavior is precisely what our conceptual model would predict.

The results in Table 8 are cross-generation comparisons: how one generation of young adults compares to another generation of young adults, making decisions at a different time. We study next whether the *specific* generation of young adults who, as we have shown, invested less in college during the boom, invested more in college training during the bust, when housing demand, and the associated labor market opportunities for non-college educated collapsed. Were the reductions in college training for this generation during the boom merely delays, or were the effects more permanent?

[45] As with all of the results from the boom period, we find that housing demand changes has no effect on "Bachelor's" degree during the bust.

Although we did not emphasize this point when we discussed Figure 4 in the Introduction, the patterns in that figure showing that cohorts from big boom markets, who made their college-going decisions during the years of the boom, had reduced attainment when we observe them relative to persons from other markets suggests that these results might be permanent. To see this, notice that among the cohorts in question (people born between 1978 and 1986), attainment rates for those in markets that had particularly large local booms remained lower as late as the 2013 ACS, when the earliest birth cohorts from this group were 33 years old and had presumably completed their college training.

Similarly suggestive evidence about the durability of reduced college attainment can be seen in first difference analyses from the Census/ACS. Using TSLS, we assess the effect of the change in local housing demand in an MSA on the difference in average attainment between persons aged 18-25 in 2006 and persons in the same MSA aged 24-31 in 2012. Notice, this latter group is drawn precisely from the same sample as people aged 18-25 in 2006, so the regression estimates how the size of a boom (and thus the later bust) affected the degree to which this particular birth cohort increased its schooling during the bust. The small and statistically insignificant effects we estimate suggest that this group, who we know had invested less in college during the boom years the larger their local housing demand change, did not increase their college attainment during the bust.[46]

While these results suggest that the effect of the decreased college investment during the boom was not reversed during the bust, the Census data do not track individual people over time, nor does the analysis account for endogenous migration across MSAs. Much sharper tests of whether there is persistence in reduced college attainment is provided by results using individual panel data from the NLSY97, where individual are tracked over time and where we know precisely where the person was at the start of the boom. In Table 9, we present results for the NLSY97 sample in 2012. The results in column 1 show that after the end of the bust, people from "housing boom" MSAs were no more likely to be employed. This contrasts sharply with the large employment increase we earlier document for these persons around the peak of the housing boom.

The second column shows that as of 2012, people from the top "boom" MSAs were still less likely to have attended college by about 5 percentage points. Their reduced education at the peak of the housing boom was about 6 percentage points, so there is arguably some evidence of very modest catch-up in Table 9. Even allowing for this, our results clearly suggest that the likelihood of ever has remained significantly depressed well after the housing boom ended.

Taken together, our various results across the different data sources indicate that there was a permanent "educational scarring" for the specific group of people who came of age during housing boom

[46] See the Online Appendix for these results.

and who were from markets with especially large booms: their college training was reduced during the boom and did not recover during the bust. Their schooling investment stands in sharp contrast to later generations of young adults in markets with large housing booms, who made college investments at rates identical to people from other markets after the bust had removed the large changes in labor market opportunities associated with the boom. The results suggest that the housing boom had a persistent effect on the human capital of younger individuals who experienced the boom.

7. CONCLUSION

This paper studies the effect of the recent housing boom and bust on the college-going decisions of young adults. We exploit cross-city variation in local housing booms, and we isolate plausibly exogenous variation in housing demand by building on the work of Ferreira and Gyourko (2011) by creating an instrument for local housing booms using the size of the structural break in housing prices during the early 2000s. We argue this instrument primarily captures speculative activity which generates increases in local housing demand, and we use this instrument to identify the causal effect of housing demand shocks on labor market outcomes, college attendance, and educational attainment.

Across several different (and complementary) data sources and empirical strategies, we consistently find evidence that the housing boom reduced college attendance and educational attainment. These effects are generally similar for men and women, and seem to be concentrated among students studying at two-year colleges or towards Associate's degrees. Applying our local labor market effects nationally, we find that the national housing boom can account for roughly 30 percent of the observed slowdown in college-going for young men and women.

We present a simple model of college attainment during a housing boom. The theory highlights the separate roles of (1) the opportunity cost of attending college, (2) the changing skill premium, and (3) the potential relaxation of liquidity constraints. Using detailed labor market data, we find that that the housing boom increased the employment and average wages of men and women without a college education, raising the opportunity cost of attending college. We find no evidence that housing boom altered the returns to going to college. This suggests that the estimated changes in educational attainment during the housing boom are likely coming through changes in opportunity costs rather than changes in returns to education.

Further evidence of the role of opportunity costs comes from the housing bust that followed the boom. We find that employment rates return roughly to pre-boom levels following the boom and bust in housing, and two-year college attendance in 2012 returns roughly to pre-boom levels, as well.[47] In contrast to these

[47] Although the labor market outcome themselves are not the primary focus of this paper, the labor market results indicate that both the housing bust *and* the housing boom affected aggregate employment, which suggests that the specific sectoral shifts

results, we find evidence of persistent declines in educational attainment for birth cohorts who were of college-going age during the boom.

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).[48] If true, then our results suggest a "scarring" effect of the housing boom for individuals who had the bad luck of being college-going age during the historically unprecedented boom and bust in housing.

Such "scarring" can also suggest part of the reason why growth has been so sluggish in the aftermath of housing boom and bust cycle. By forgoing schooling during the housing boom, there is a set of workers with lower marginal products than they would have had otherwise. The lower level of productivity for these workers can act as a drag on aggregate labor productivity, raising the question of whether our findings can help understand the decline in labor productivity within the U.S. in the aftermath of the Great Recession.

---

occurring during that housing boom had aggregate effects. Understanding why these particular sectoral shifts had such a large effect on aggregate employment represents an important area of future work.

[48] 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 (2015). "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.

Bhutta, Neil (2014). "The Ins and Outs of Mortgage Debt during the Housing Boom and Bust," *Journal of Monetary Economics*.

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.

Bound, John and Harry J. Holzer (2000). "Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s," *Journal of Labor Economics*.

Bound, John and Sarah Turner (2007). "Cohort Crowding: How Resources Affect Collegiate Attainment," *Journal of Public Economics*, 91(5-6): 877-899.

Burnside, Craig, Martin Eichenbaum and Sergio Rebelo (2015). "Understanding Boom and Busts in Housing Markets," *Journal of Political Economy*, forthcoming.

Cadena, Brian C. and Brian K. Kovak (2015). "Immigrants Equilibrate Local Labor Markets: Evidence from the Great Recession," *American Economic Journal: Applied Economics*, forthcoming.

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

Cameron, Stephen V. and Christopher Taber (2004). "Estimation of Educational Borrowing Constraints Using Returns to Schooling," *Journal of Political Economy*, 112(1), 132-182.

Card, David and Dean Hyslop (1997). "Does Inflation Grease the Wheels of the Labor Market," in *Reducing Inflation: Motivation and Strategy*. Christina Romer and David Romer, eds. Chicago: University of Chicago Press.

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.

Chen, Kaiji and Yi Wen (2014). "The Great Housing Boom of China," Working Paper.

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

Davidoff, Thomas (2015). "Supply Constraints are Not Valid Instrumental Variables for Home Prices Because they are Correlated with Many Demand Factors", *Critical Finance Review*, forthcoming.

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.

Glaeser, Edward and Charles Nathanson (2014). "Housing Bubbles". *NBER Working Paper 20426*.

Goldin, Claudia and Lawrence Katz (2010). *The Race Between Education and Technology*, Bekknap Press.

Hilger, Nathaniel (2014). "Parental Job Loss and Children's Long-Term Outcomes: Evidence from 7 Million Fathers' Layoffs," Working Paper.

Houthakker, Hendrick S. and Lester D. Taylor. 1970. *Consumer Demand in the United States*, 1929-1970. Cambridge: Harvard University Press.

Jorda Oscar, Moritz Schularick, and Alan M. Taylor (2014). "The Great Mortgaging: Housing Finance, Crises, and Business Cycles." *NBER Working Paper 20501*.

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

Manoli, Dayanand S. and Nicholas Turner (2014). "Cash-on-Hand & College Enrollment: Evidence from Population Tax Data and Policy Nonlinearities," *NBER Working Paper 19836*.

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 (2014). "What Explains the 2007–2009 Drop in Employment?" *Econometrica*, 82(6): 2197-2223.

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.

Notowidigdo, Matthew J. (2013). "The Incidence of Local Labor Market Shocks," Working Paper.

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.

Shapiro, Jesse (2006). "Smart Cities: Quality of Life, Productivity, and the Growth Effects of Human Capital." *Review of Economics and Statistics*, 88(2): 324-335.

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.

Shiller, Robert J. (2008). "Understanding Recent Trends in House Prices and Homeownership" in *Housing, Housing Finance and Monetary Policy*, Jackson Hole Conference Series, Federal Reserve Bank of Kansas City, 85-123.

Willis, Robert J. and Sherwin Rosen (1979). "Education and Self-Selection," *Journal of Political Economy*, 87(5).

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.

Figure 1: Fraction to Have Ever Attended College Among Persons Aged 18-29, 1980-2013

Panel A: Men

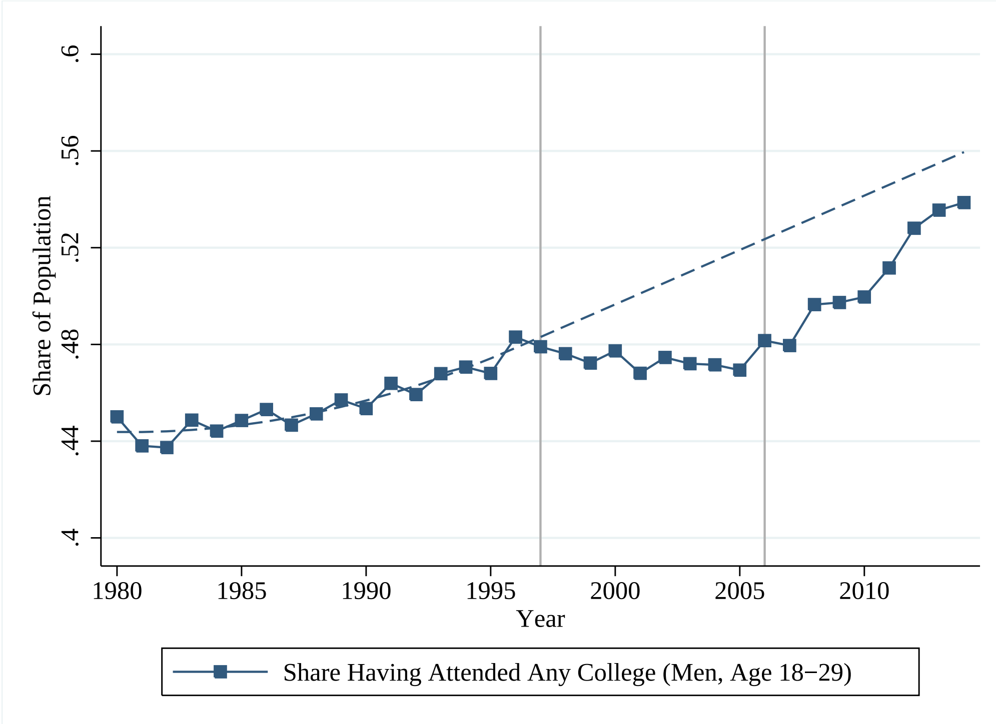

Panel B: Women

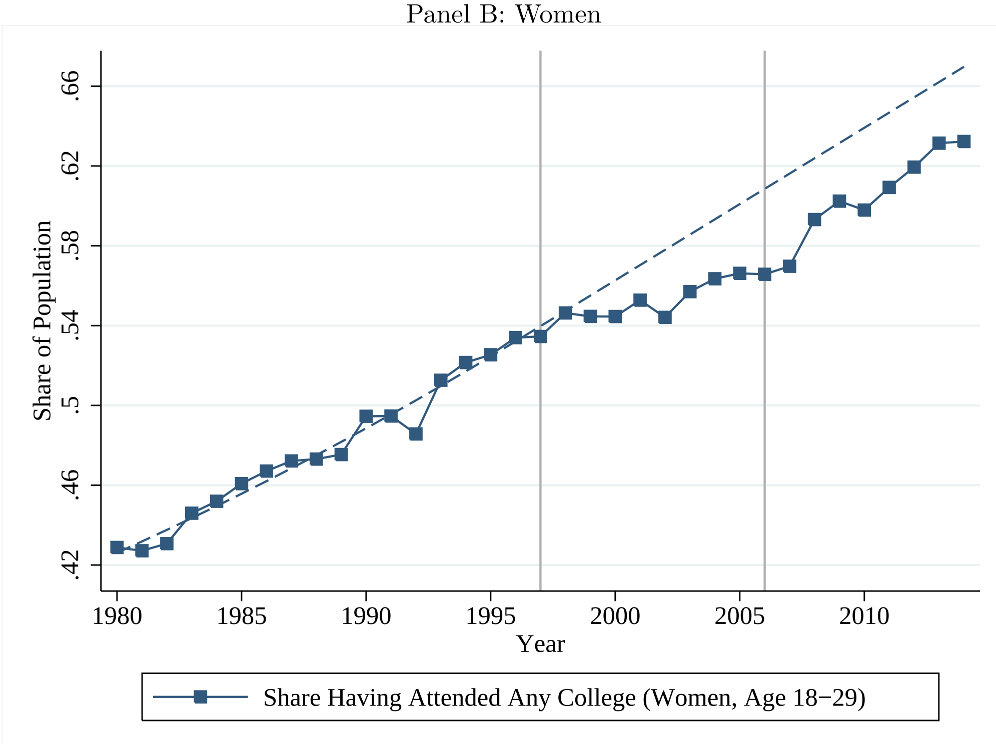

Notes: This figure reports trends in the share of men and women (age 18-29) who have attended at least one year of college. This series is constructed from the Current Population Survey (CPS) using CPS survey weights. The dashed line is the predicted college attendance rates based on a quadratic trend that is fit to the 1980-1996 period.

Figure 2: Fraction to Have Ever Attended College by Age 25,
By Year of Birth for Men and Women Born Between 1960 and 1989

Panel A: Men

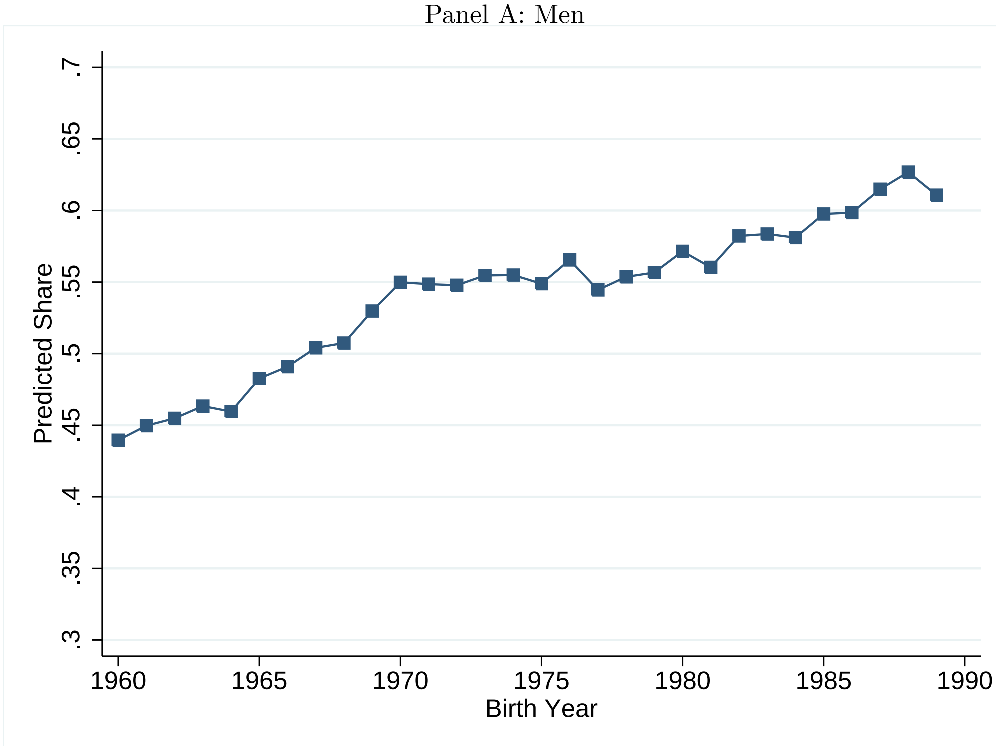

Panel B: Women

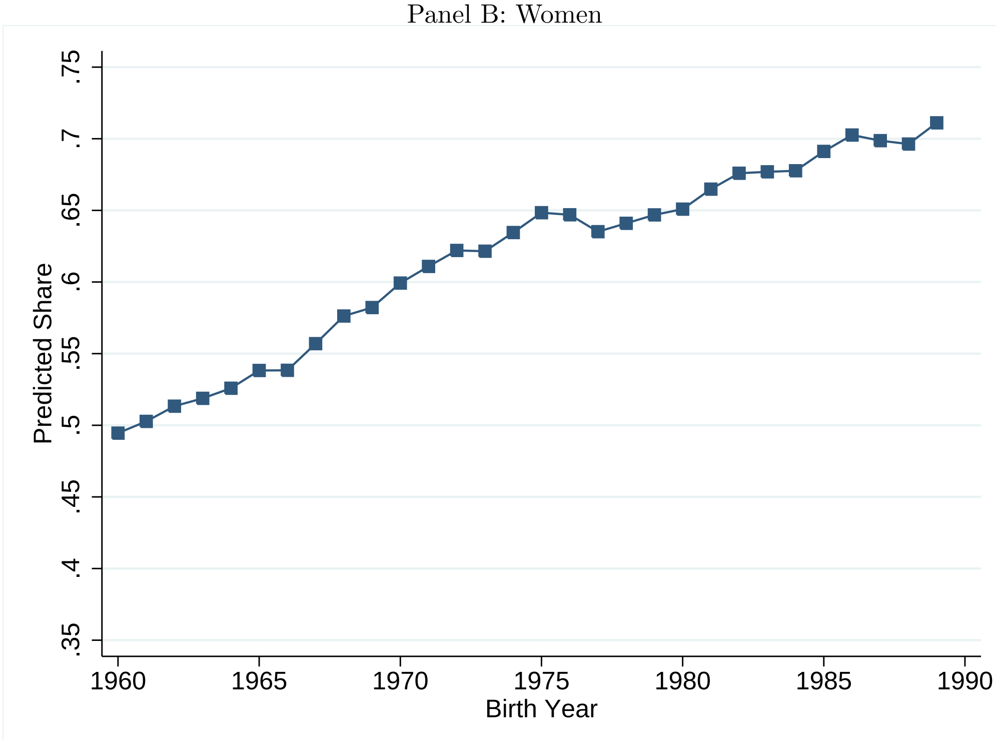

Notes: This figure reports estimated birth year (birth cohort) fixed effects in education for all men and women born between 1960 and 1990 (inclusive). The sample is all individuals between the ages of 25 and 54 (based on age in survey year), pooling CPS data sets between 1994 and 2014. The birth year fixed effects are recovered from an estimated model that regresses an indicator for whether individual has attended any college on a fourth-degree polynomial in age, birth year fixed effects, and normalized year fixed effects (setting the first and last year fixed effect equal to 0 and the sum of remaining year fixed effects equal to 0). The figures reported fitted values at age 25 using CPS survey weights. The sample is restricted to native-born men and women.

Figure 3: Share Any College Attendance for Individuals Born Between 1965 and 1987, by MSA House Price Growth

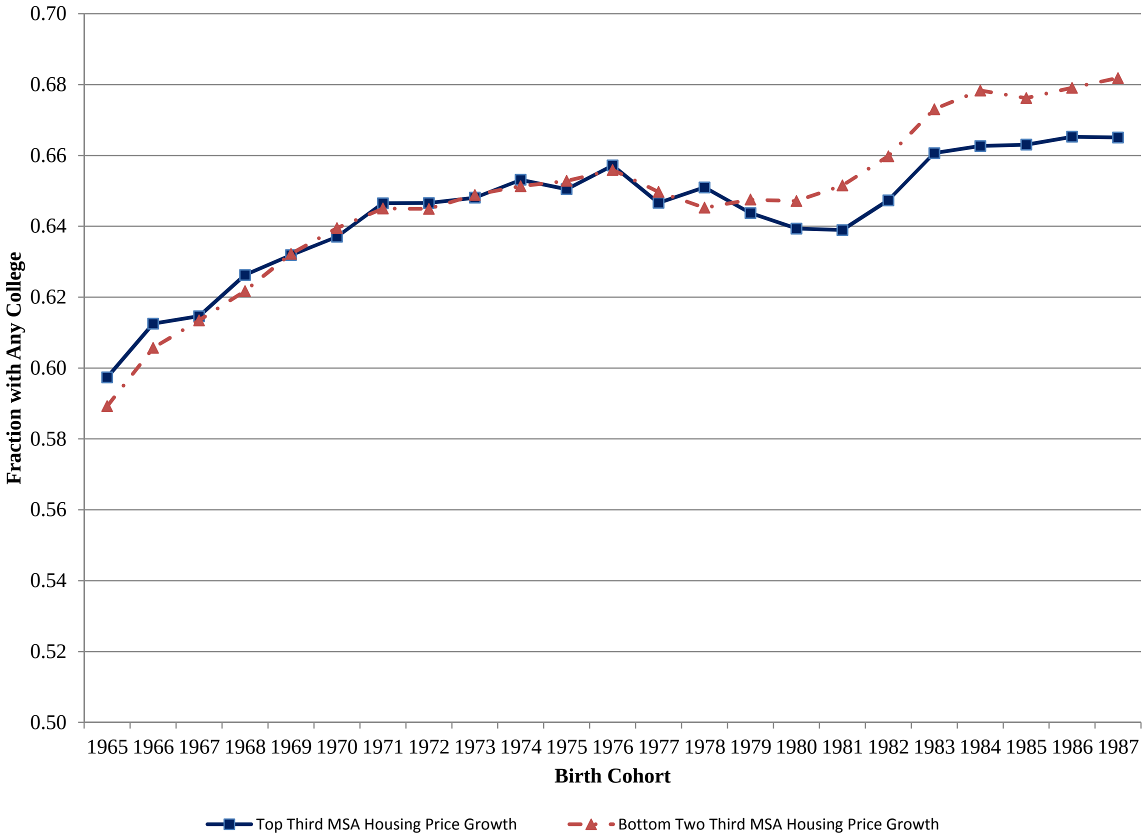

Notes: This figure reports education attainment by birth cohort for all men and women born between 1965 and 1987. The data come from the 1990 Census, the 2000 Census, and the 2005-2013 American Community Survey (ACS). The sample is restricted to all men and women between ages 25 and 54 in the survey year. The two lines are sub-samples of metropolitan areas based on whether or not the metropolitan was in the top tercile of distribution of house price changes between 2000 and 2006. We use FHFA house price data to compute MSA level house price changes. The data use Census/ACS survey weights.

Figure 4: Home Prices, Housing Permits, and Housing Transactions in the U.S., 1980-2012

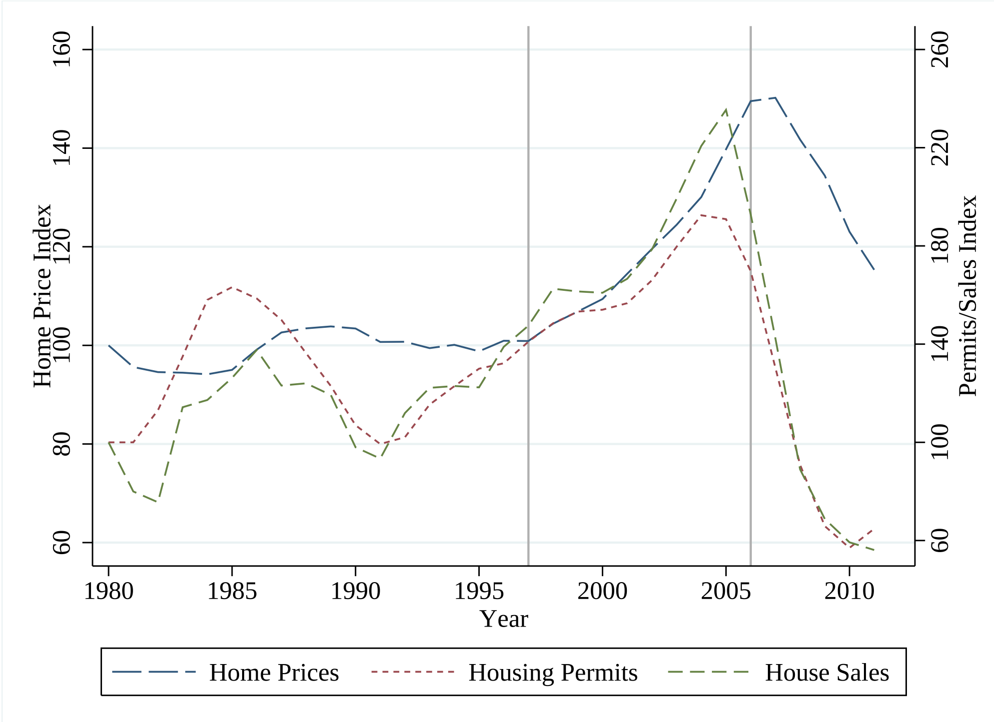

Notes: This figure reports trends in the FHFA National Home Price Index (1980 = 100), trends in the Census Housing Permits Index (1980 = 100), and trends in total new home sales from the Survey of Construction (1980 = 100). The FHFA series is a weighted, repeat-sales index that measures average changes in house prices across 363 metropolitan areas. The Census series is a building permits survey that estimates the number of new housing units (as authorized by building permits). The Survey of Construction series measures new house sales of single-family homes, whether or not building new homes in those areas requires a building permit.

Figure 5: An Economic Model of College Attendance

Panel A: College Attendance Decisions by Ability

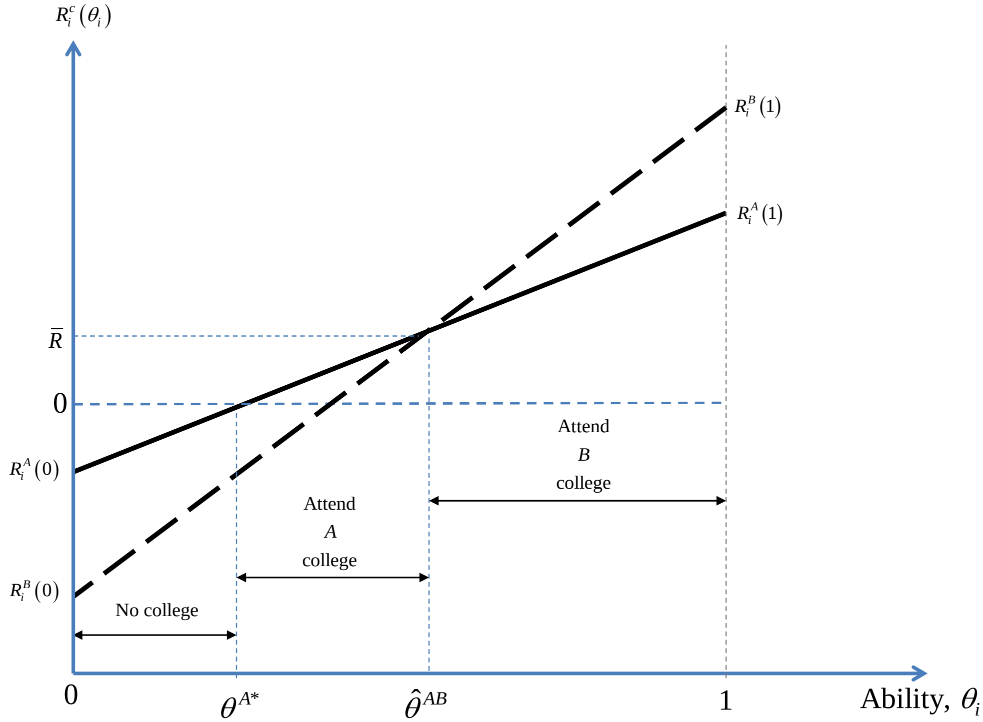

Panel B: Attendance Decisions Following Increase in Y_t^0 From Housing Boom

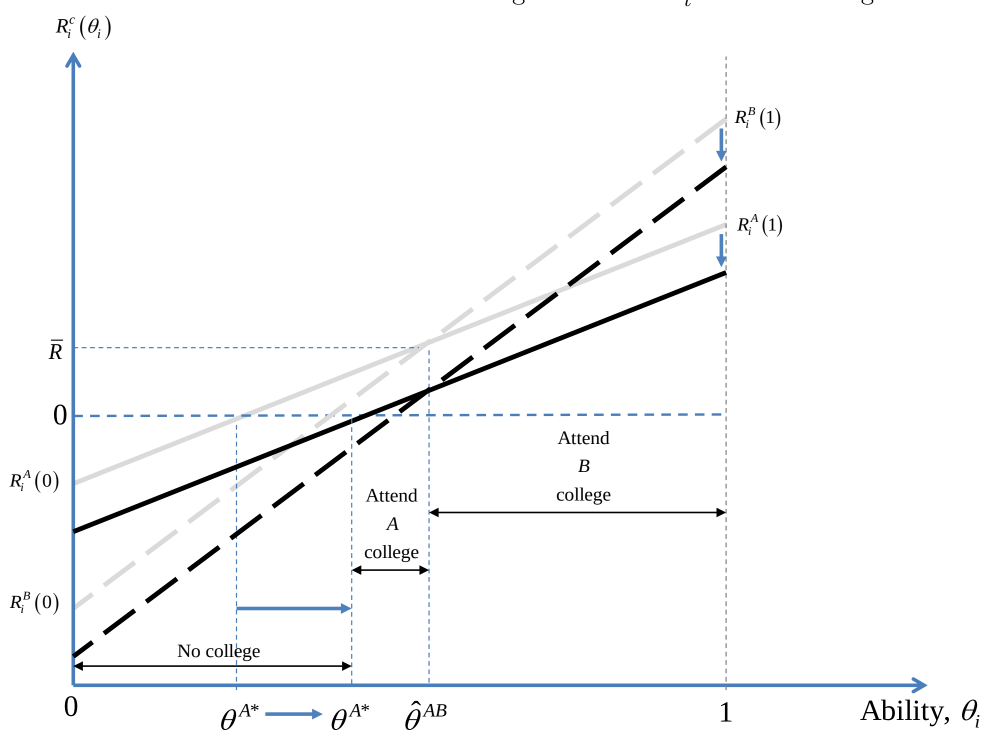

Notes: This figure shows the equilibrium college attendance decisions for individuals as a function of their underlying ability. The equilibrium shows individuals making choices of whether to attend college (and, if so, whether to attend type-A "Associate's" college or type-B "Bachelors" college). In the equilibrium, individuals sort based on comparative advantage, with low-ability individuals not attending college, middle-ability individuals attending type-A college, and high-ability individuals attending type-B college. In Panel B, there is a positive housing demand shock (housing boom), which raises income of all non-college-educated individuals by the same amount. In the new equilibrium, there is a reduction in share of population attending type-A college but no change in share attending type-B college.

Figure 6: Trends in Housing Demand Over Time and Across Metropolitan Areas

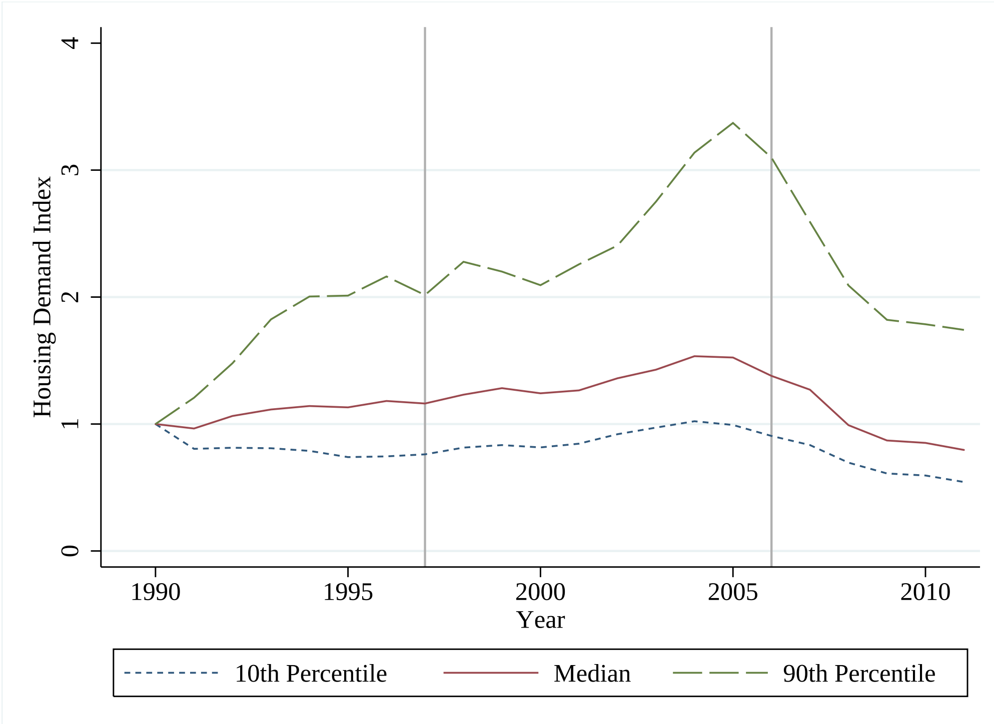

Notes: This figure reports trends in our constructed housing demand index (which is the log sum of the prices and permits indexes) at the 10th percentile, median, and 90th percentile, normalized to 1980 values within each percentile.

Figure 7: Variation in Structural Break Across Cities

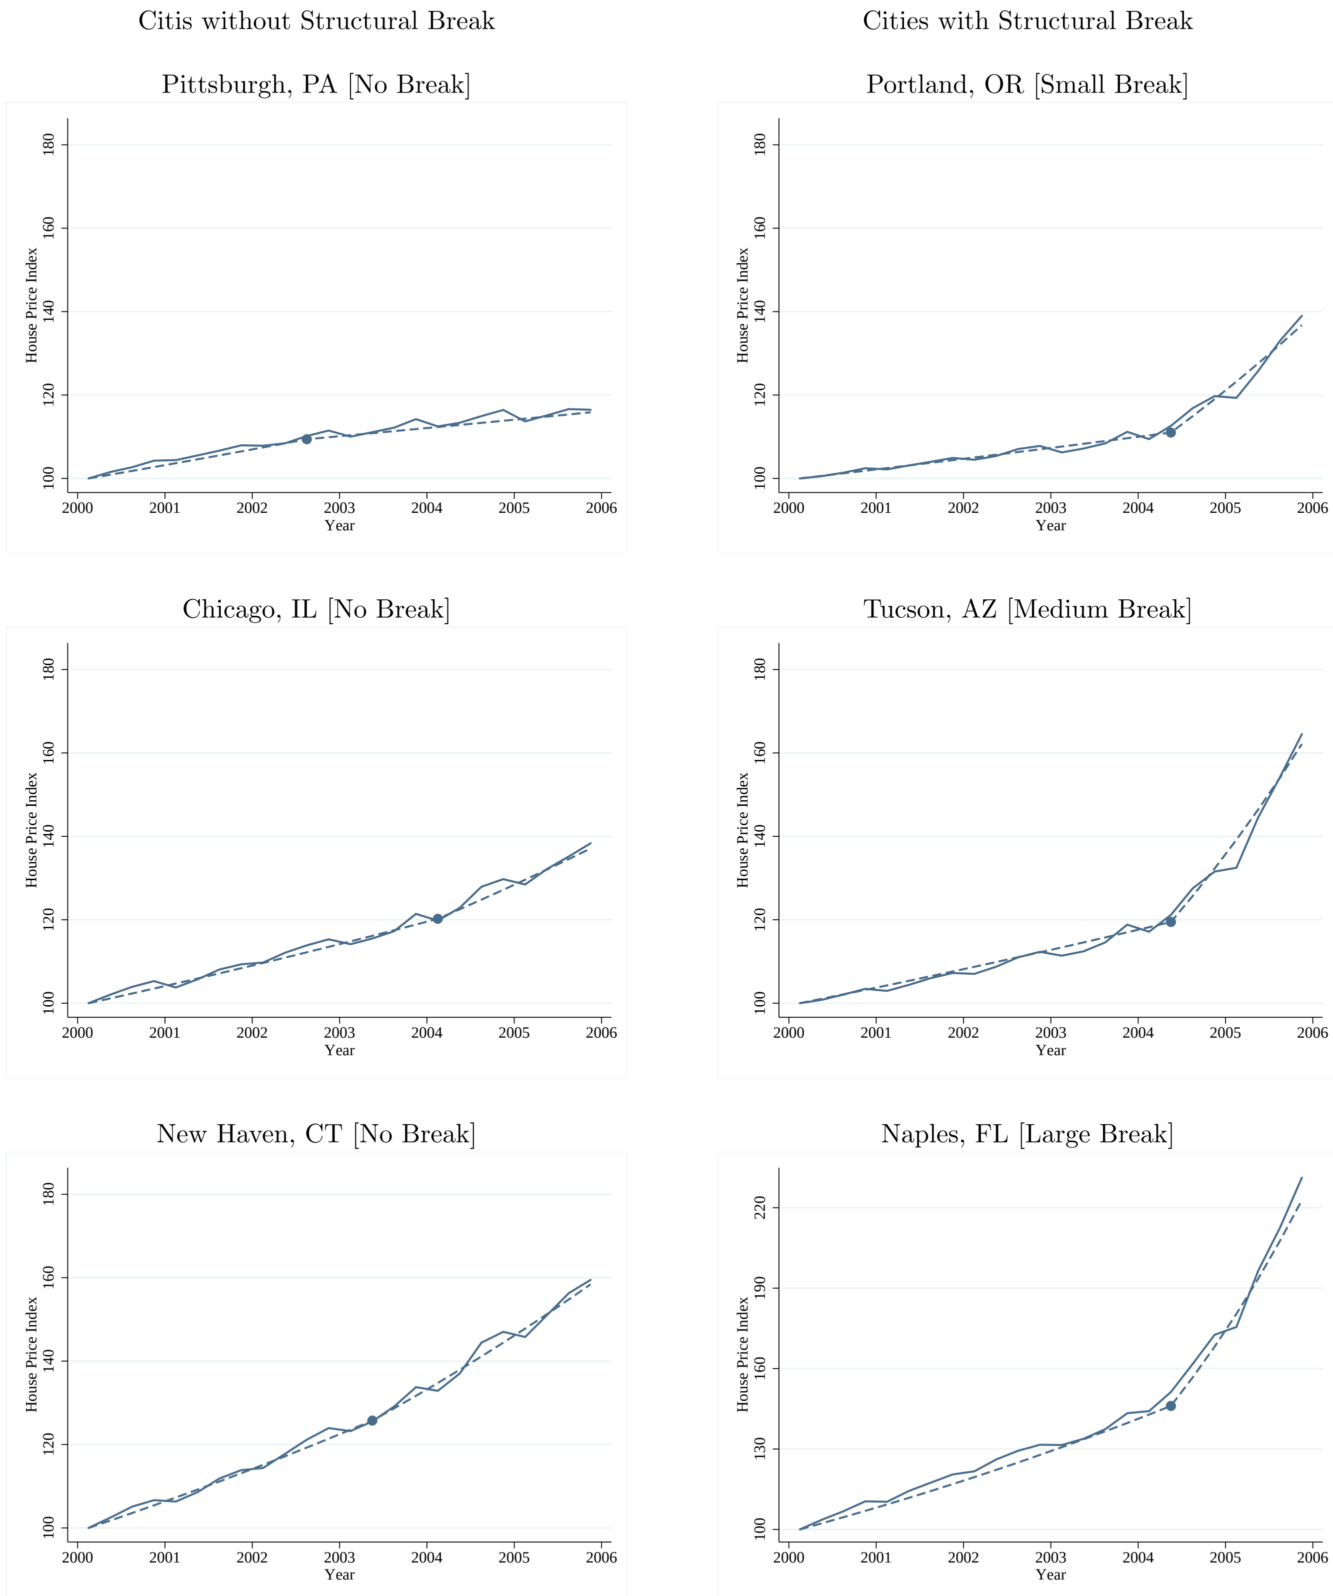

Notes: This figure shows graphs of quarterly house price data for six MSAs. The house price index for each city is normalized so that Q1, 2000 = 100. The solid lines report the house price series, while the dashed lines reports the structural break estimates, with a solid dot indicating the estimated quarter of the structural break. The MSAs in first column have small estimated structural breaks, and the MSAs in the second column have relatively larger estimated structural breaks. The rows group MSAs based on overall house price growth up until the estimated structural break.

Figure 8: 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, 2001 and Q4, 2005) and the estimated housing demand change across 2000-2006. 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 9: Lack of Correlation Between Structural Break Instrument and ...

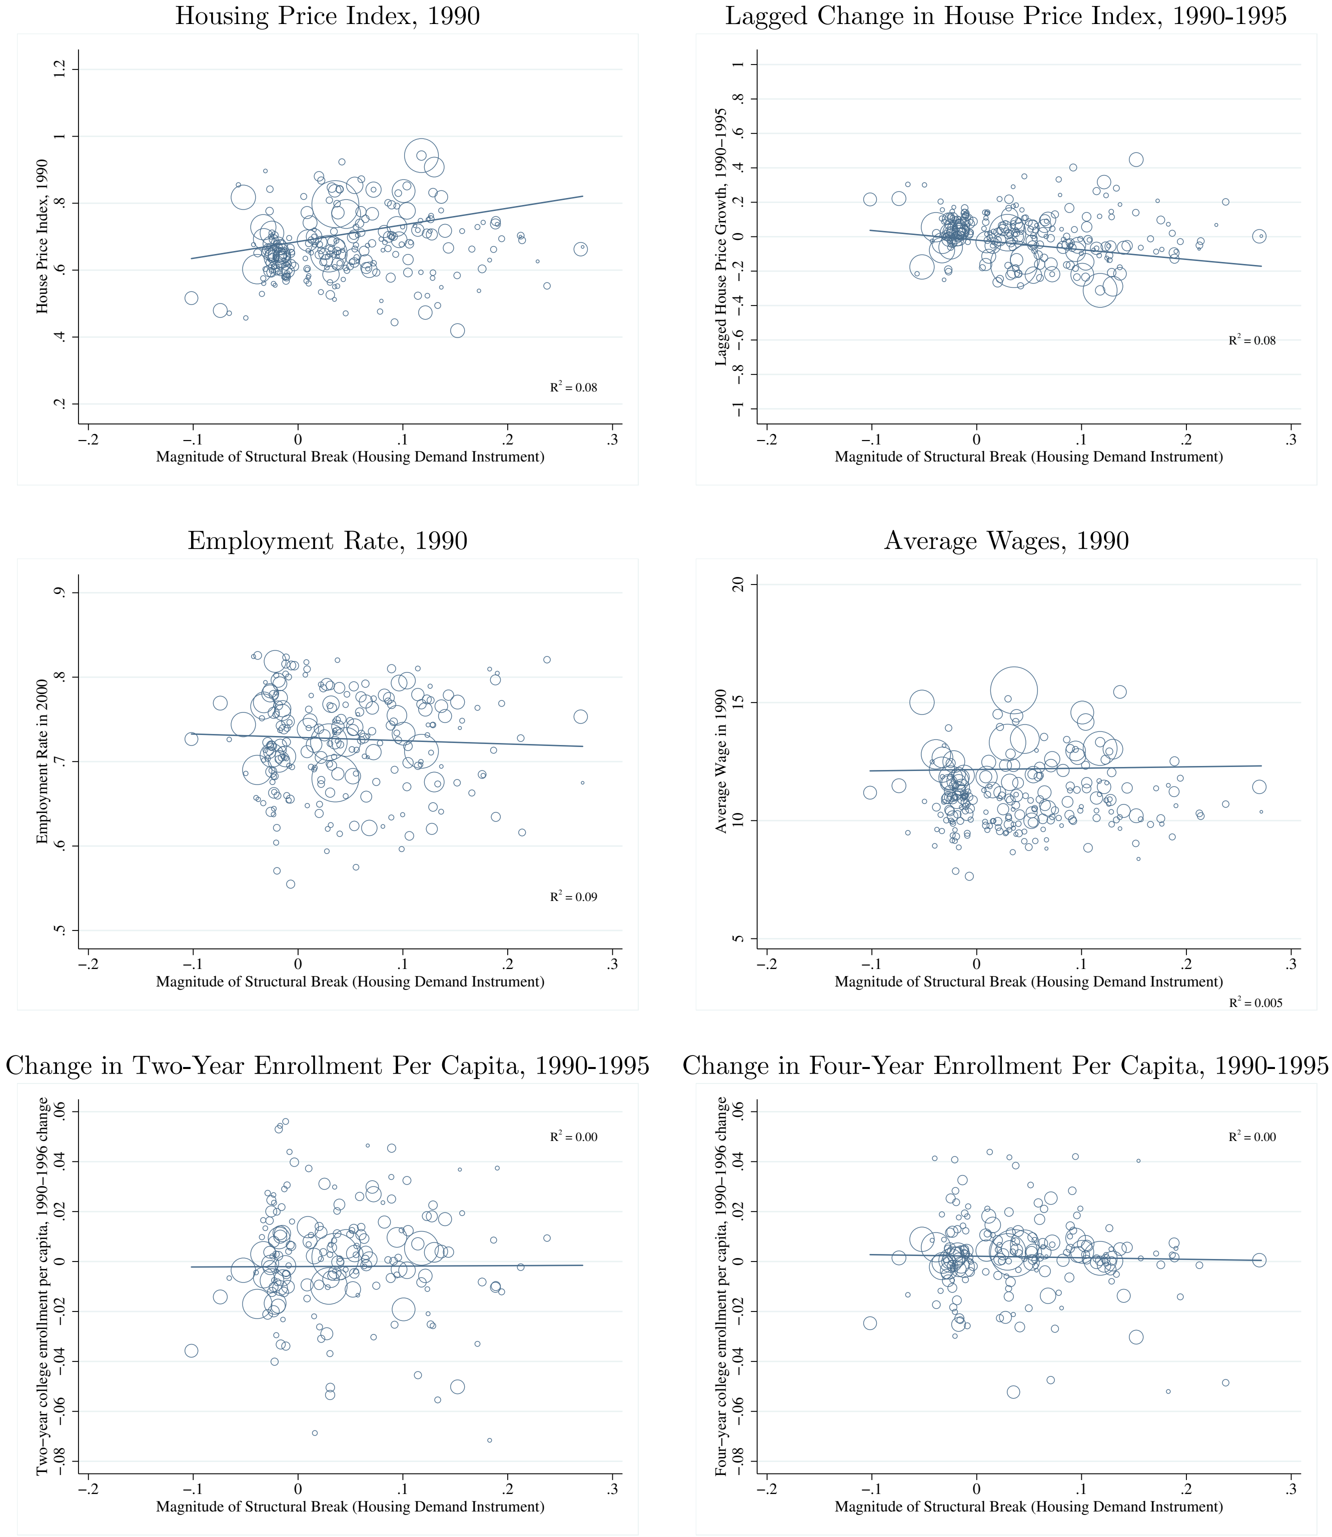

Notes: This figure reports the correlation between the structural break instrument used in the IV specifications and the following: lagged house price levels, lagged (1990-1995) house price growth, lagged employment rate and average wages from the 1990 Census, and lagged changes in two-year and four-year college enrollment (per capita) from IPEDS.

Figure 10: Significant Correlation Between Structural Break Instrument and ...

Growth in Price-Rent Ratio

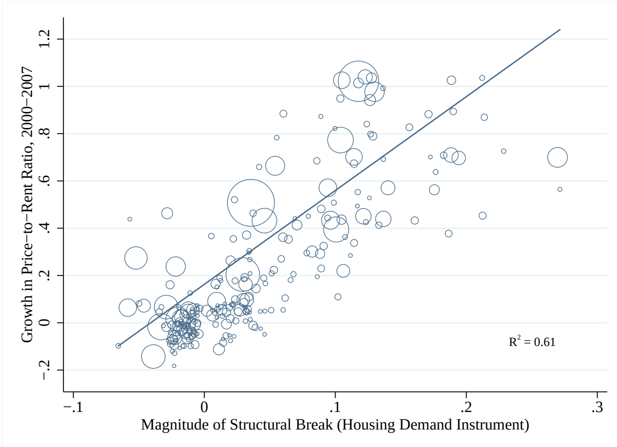

Change in “Out-of-Town” Buyers Share

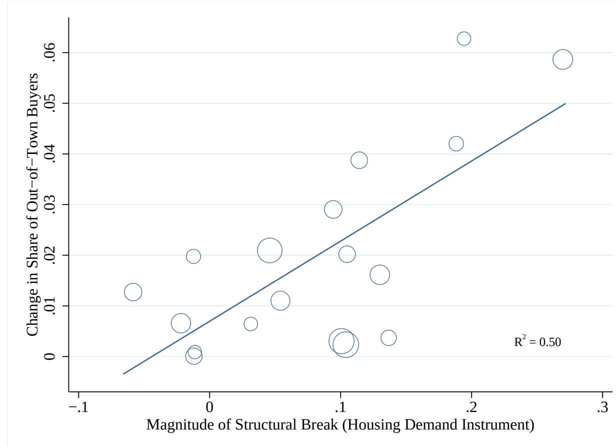

Notes: This figure reports the correlation between the structural break instrument used in the IV specifications and the growth in the price-rent ratio and the change in the share of “out-of-town” buyers, which can be interpreted as proxies for speculation. See text for details of the price-rent ratio calculation and the source of the “out of town” buyer share .

Figure 11: Randomization tests of college enrollment results exploiting timing and magnitude of local housing boom

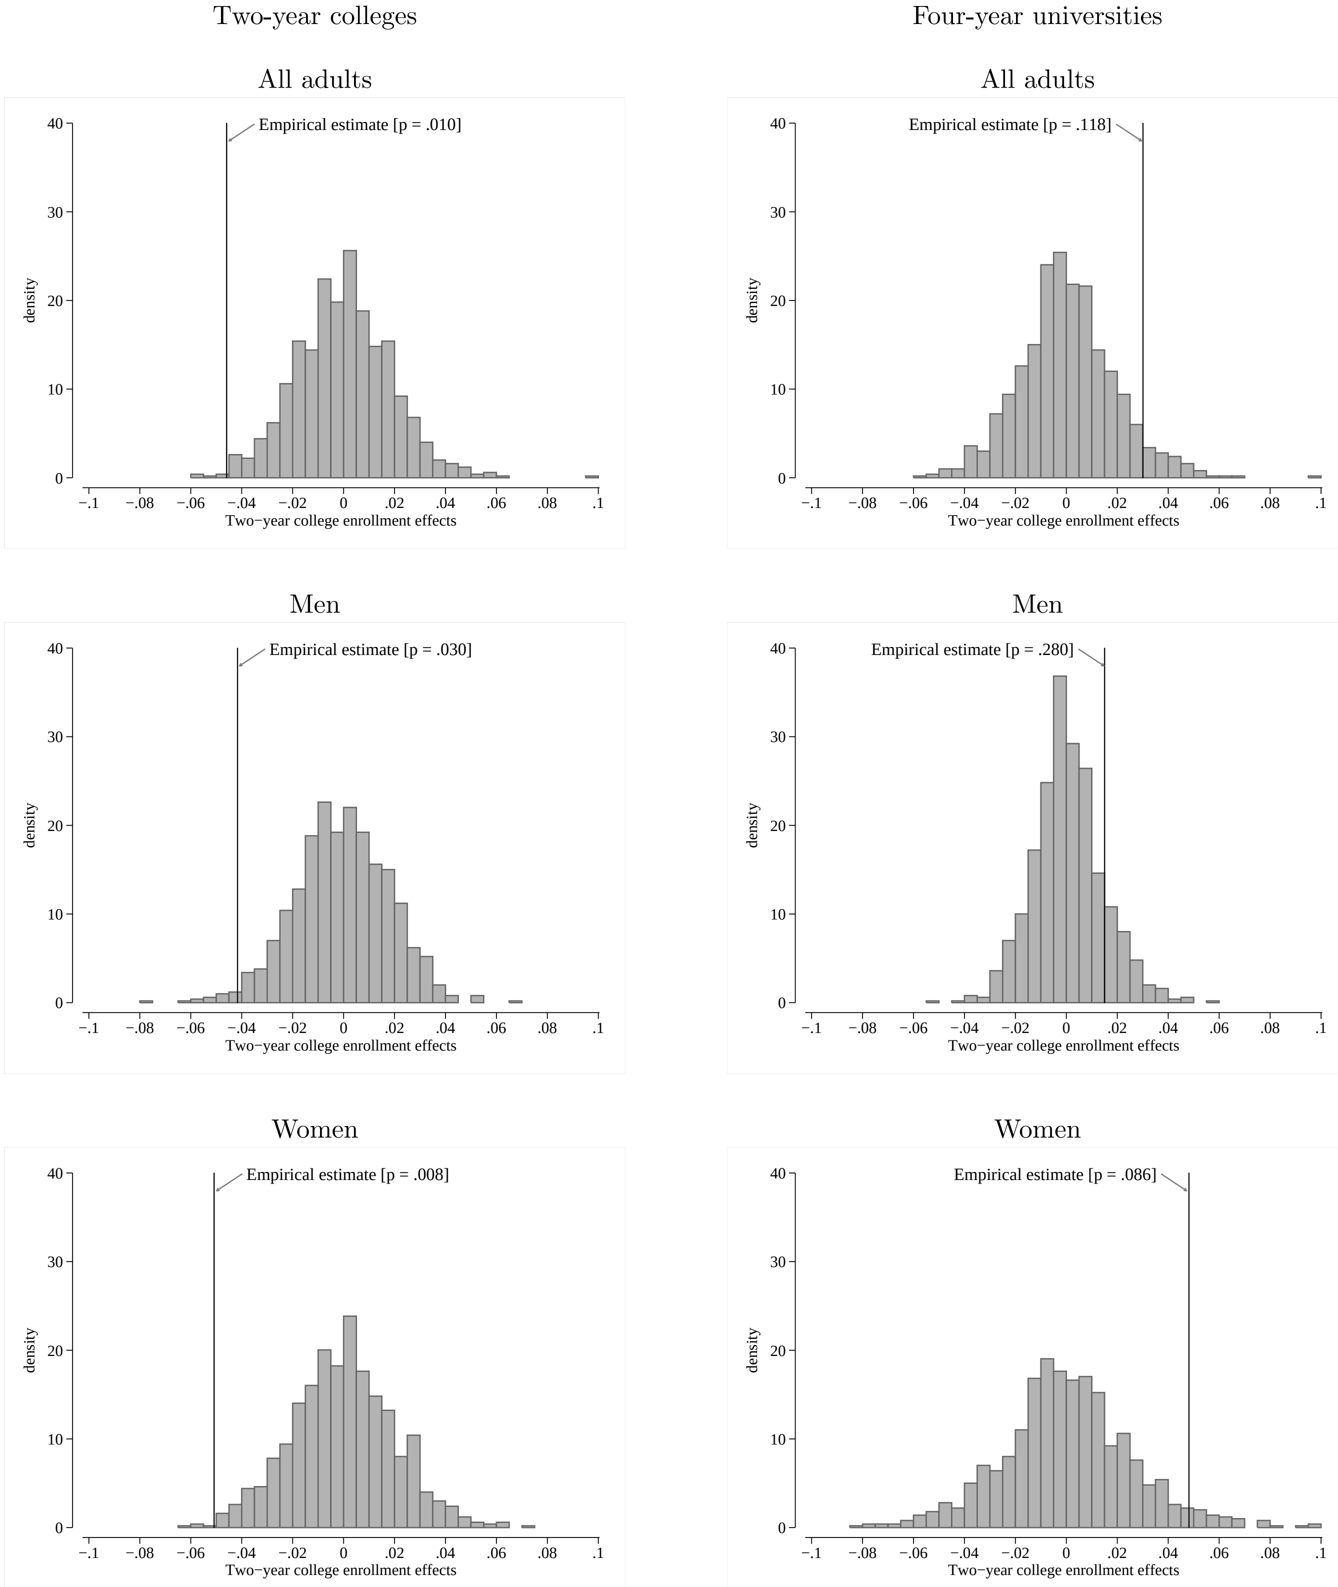

Notes: This figure shows histograms of distribution of estimated effects of housing boom on two-year and four-year college enrollment per capita for 1,000 permutation samples which permute the magnitude and year of structural break in local house prices across each city. The vertical lines indicate the corresponding estimates from the true data shown in Table 5.

Table 1

Housing Booms and Labor Market Opportunities for Adults Without Any College Education

Dependent Variable is 2000-2006 Change in:	Employment Rate (1)	Average Wages (2)	Emp. Rate * Average Wage (3)	Share Employed in Construction (4)	Share Employed in FIRE (5)
			Panel A: OLS Estimates		
All Adults Age 18-25					
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	0.031	0.079	0.115	0.016	0.002
	(0.008)	(0.011)	(0.023)	(0.003)	(0.001)
Share of Total Employment Change				51.6%	5.0%
R^2	0.31	0.17	0.33	0.13	0.12
Men Age 18-25					
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	0.030	0.080	0.119	0.026	0.001
	(0.008)	(0.011)	(0.023)	(0.004)	(0.001)
Share of Total Employment Change				86.6%	1.8%
R^2	0.20	0.14	0.25	0.13	0.03
Women Age 18-25					
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	0.032	0.085	0.115	0.005	0.002
	(0.009)	(0.014)	(0.023)	(0.001)	(0.002)
Share of Total Employment Change				16.8%	6.5%
R^2	0.22	0.09	0.23	0.10	0.10
			Panel B: 2SLS Estimates		
All Adults Age 18-25					
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	0.051	0.115	0.179	0.021	0.010
	(0.018)	(0.025)	(0.049)	(0.005)	(0.004)
Share of Total Employment Change				41.4%	20.1%
Men Age 18-25					
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	0.059	0.111	0.202	0.034	0.004
	(0.020)	(0.029)	(0.060)	(0.008)	(0.004)
Share of Total Employment Change				58.6%	7.5%
Women Age 18-25					
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	0.043	0.117	0.152	0.006	0.018
	(0.018)	(0.032)	(0.050)	(0.002)	(0.009)
Share of Total Employment Change				12.8%	40.6%
First stage F-statistic	38.86	38.86	38.86	38.86	38.86
N	275	275	275	275	275
Include baseline controls	y	y	y	y	y

Notes: This table reports OLS and 2SLS estimates. All samples are from Census/ACS data, have been restricted to ages 18-25, have been restricted to individuals who live in same state where they were born, and excludes individuals in group quarters. Additionally, all individuals have no college education, which includes high school dropouts and high school graduates with no reported college attendance. The baseline controls included in all columns are the following: log of MSA population in 2000, share of employed adults with a college degree, the share of adults who are foreign born, and the share of women in the labor force. The average 18-25 employment rate in 2000 is 0.61 for adults, 0.64 for men, and 0.57 for women. Standard errors are shown in parentheses and are clustered by state.

Table 2
Housing Booms and the Lifetime Returns to Education

Variable:	Employment Rate (1)	Average Wage (2)	Employment Rate * Average Wage (3)
	Dependent Variable is 2000-2006 Change in Difference Between Any College and No College		
2SLS Estimates for Adults Age 26-55			
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	-0.026	0.002	-0.060
	(0.006)	(0.011)	(0.019)
2SLS Estimates for Men Age 26-55			
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	-0.020	0.008	-0.041
	(0.006)	(0.014)	(0.024)
2SLS Estimates for Women Age 26-55			
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	-0.030	-0.014	-0.080
	(0.007)	(0.009)	(0.021)
First stage F-statistic	38.86	38.86	38.86
N	275	275	275
Include baseline controls	y	y	y

Notes: This table reports 2SLS estimates for alternative gender and education groups. All samples are restricted to ages 26-55 and have been restricted to individuals who live in same state where they were born and excludes individuals in group quarters. All individuals with no college education represents high school dropouts and high school graduates with no reported college attendance; all individuals with "any college" reported attending college for at least part of one year (which includes college graduates and college dropouts). The dependent variables are the difference in the change in labor market outcomes for those with "any college" relative to the same labor market change for those with "no college". A negative coefficient means the labor market outcomes of those with "no college" improved relative to those with "any college" during the housing boom. The baseline controls are described in Table 1. Standard errors are shown in parentheses and are clustered by state.

Table 3
Housing Booms and Educational Attainment: 2SLS Estimates, Census/ACS Data

	Dependent Variable is 2000-2006 Change in Share With College Education	
College Education Definition:	Any College	Bachelors Degree or Higher
	(1)	(2)
2SLS Estimates for All Adults Age 18-25		
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	-0.021	0.003
	(0.007)	(0.004)
Average for Adults Age 18-25 in 2000		
Average in 2000	0.469	0.100
Average in 2006	0.619	0.285
2SLS Estimates for Men Age 18-25		
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	-0.018	-0.002
	(0.008)	(0.004)
Average for Men Age 18-25 in 2000		
Average in 2000	0.426	0.083
Average in 2006	0.572	0.250
2SLS Estimates for Women Age 18-25		
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	-0.024	0.008
	(0.009)	(0.006)
Average for Women Age 18-25 in 2000		
Average in 2000	0.511	0.118
Average in 2006	0.664	0.319
[Placebo Specification]		
2SLS Estimates for All Adults Age 26-33		
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	0.008	0.004
	(0.010)	(0.007)
Average for Adults Age 26-33 in 2000		
Average in 2000	0.609	0.269
Average in 2006	0.626	0.301
First stage F-statistic	38.86	38.86
N	275	275
Include baseline controls	y	y

Notes: This table reports 2SLS estimates for alternative gender and age groups. All samples are restricted to ages listed in panel heading have been restricted to individuals who live in same state where they were born, and excluded those in group quarters. All individuals with "any college" reported attending college for at least a portion of one year. The baseline controls are described in Table 1. Standard errors are shown in parentheses and are clustered by state.

Table 4
Housing Booms and College Enrollment: 2SLS Estimates, IPEDS Data

Dependent Variable is the Change in Average Annual Enrollment Per Capita				
Change defined betwee following years:	2000 and 2006		1990 and 1996	
Enrollment outcome:	2-year colleges	4-year colleges and universities	2-year colleges	4-year colleges and universities
	(1)	(2)	(3)	(4)
2SLS Estimates for Men and Women				
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	-0.011	0.001	-0.003	0.002
	(0.005)	(0.004)	(0.011)	(0.002)
Average level at start of period	0.050	0.064	0.053	0.056
2SLS Estimates for Men Only				
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	-0.011	-0.002	-0.001	0.003
	(0.005)	(0.003)	(0.011)	(0.002)
Average level at start of period	0.047	0.058	0.048	0.052
2SLS Estimates for Women Only				
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	-0.011	0.004	-0.005	0.002
	(0.005)	(0.007)	(0.012)	(0.002)
Average level at start of period	0.053	0.071	0.058	0.060
First Stage F-statistic	38.71	38.71	26.80	26.80
N (Number of Metropolitan Areas)	224	224	191	191
Include baseline controls	y	y	y	y

Notes: The unit of observation is the metropolitan area, and the data come from the IPEDS data set. The dependent variable are long differences across years reported in column headings. Each endpoint is average annual enrollment during the preceding five years. The enrollment data are matched to metropolitan areas by county, using 2000 metropolitan area definitions. Two-year colleges are defined to be any college that does not offer a four-year degree. Some 4-year colleges may offer two-year degrees but they will be included in columns (2) and (4). This table reports 2SLS estimates for alternative demographic groups. The baseline controls are described in Table 1. Standard errors are shown in parentheses and are clustered by state.

Table 5
Housing Booms and College Enrollment: Difference-in-Difference Estimates Exploiting Timing and Magnitude of Housing Boom, IPEDS Data

Dependent Variable is Enrollment Per Capita, Annual Data 1990-2006

Enrollment outcome:	2-year colleges (1)	4-year colleges and universities (2)
Panel A: OLS Estimates for All Adults		
Interaction between magnitude and timing of structural break, $\lambda_k \times (\text{Post } t_k^*)$	-0.066	0.019
	(0.022)	(0.020)
Mean of dependent variable	0.049	0.069
Panel B: OLS Estimates for Men Only		
Interaction between magnitude and timing of structural break, $\lambda_k \times (\text{Post } t_k^*)$	-0.062	0.012
	(0.020)	(0.013)
Mean of dependent variable	0.045	0.061
Panel C: OLS Estimates for Women Only		
Interaction between magnitude and timing of structural break, $\lambda_k \times (\text{Post } t_k^*)$	-0.071	0.028
	(0.025)	(0.028)
Mean of dependent variable	0.053	0.076
N	3853	3625
Number of Metropolitan Areas	223	223
Metropolitan Area FEs and Year FEs	y	y

Notes: The unit of observation is the metropolitan area-by-year and come from the IPEDS data set. The enrollment data are matched to metropolitan areas by county, using 2000 metropolitan area definitions. Two-year colleges are defined to be any college that does not offer a four-year degree. Some 4-year colleges may offer two-year degrees but they will be included in columns (2). This table reports OLS estimates for alternative demographic groups. All regressions include MSA and year fixed effects. The baseline controls from previous tables are not included because they are not identified when metropolitan area fixed effects are included. The right-hand side variable is interaction of structural break variable and indicator for whether the year is after the estimated year of structural break. Standard errors are shown in parentheses and are clustered by state.

Table 6
Housing Booms, Employment, and Educational Attainment:
Evidence from Individual-Level Panel Data from NLSY

Dependent variable:	Employed at Age 20	Has Attended Any College, Year 2006	Has Bachelors Degree, Year 2006
	(1)	(2)	(3)
Panel A: OLS Reduced Form Estimates for Men and Women			
Living in Housing Boom MSA in 1997	0.058	-0.058	-0.014
	(0.018)	(0.020)	(0.025)
Mean of Dependent Variable	0.684	0.569	0.233
N	5362	5362	5362
Panel B: OLS Reduced Form Estimates for Men Only			
Living in Housing Boom MSA in 1997	0.071	-0.049	-0.020
	(0.026)	(0.028)	(0.035)
Mean of Dependent Variable	0.689	0.522	0.191
N	2697	2697	2697
Panel C: OLS Reduced Form Estimates for Women Only			
Living in Housing Boom MSA in 1997	0.041	-0.068	-0.009
	(0.026)	(0.028)	(0.024)
Mean of Dependent Variable	0.679	0.616	0.274
N	2665	2665	2665
Include Baseline Controls (Metropolitan Area)	y	y	y
Include Additional Individual-Level Controls	y	y	y

Notes: The unit of observation is individual, and the assignment of housing demand change (between 2000 and 2006) is based on where the individual was living in 1997 at start of the NLSY97 sample. This table reports OLS estimates for alternative demographic groups, and each column reports results for a different dependent variable. The key independent variable is a dummy variable indicating whether the estimated structural break instrument was in the top tercile across MSAs. The baseline controls are the same as the controls in Table 1, and the additional individual-level controls are the following: age, demographic indicators for black, hispanic, mixed race, non-black; separate indicators for father's and mother's education (missing, high school dropout, high school graduate, some college, and Bachelors or greater). Standard errors are shown in parentheses and are clustered by state.

Table 7

Housing Booms, Housing Busts, and Labor Market Opportunities, Census/ACS Data

Dependent variable:	Employment Rate		Average Wage		Employment Rate * Average Wage	
Change defined between following periods:	2006 and 2012	2000 and 2012	2006 and 2012	2000 and 2012	2006 and 2012	2000 and 2012
	(1)	(2)	(3)	(4)	(5)	(6)
2SLS Estimates for All Adults Age 18-25						
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	-0.081	-0.029	-0.046	0.069	-0.203	-0.024
	(0.018)	(0.022)	(0.034)	(0.044)	(0.048)	(0.063)
2SLS Estimates for Men Age 18-25						
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	-0.092	-0.033	-0.043	0.068	-0.235	-0.033
	(0.018)	(0.024)	(0.047)	(0.052)	(0.049)	(0.072)
2SLS Estimates for Women Age 18-25						
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	-0.066	-0.023	-0.034	0.083	-0.159	-0.007
	(0.026)	(0.025)	(0.048)	(0.041)	(0.068)	(0.063)
First stage F-statistic	38.86	38.86	38.86	38.86	38.86	38.86
N	275	275	275	275	275	275
Include baseline controls	y	y	y	y	y	y

Notes: This table reports 2SLS estimates for alternative demographic groups. All samples are restricted to ages 18-25 and have been restricted to individuals who live in same state where they were born. Additionally, all individuals have no college education, which includes high school dropouts and high school graduate with no reported college attendance. The Census/ACS sample and baseline controls are described in Table 1. Standard errors are shown in parentheses and are clustered by state.

Table 8

Housing Booms and Housing Busts: Educational Attainment and College Enrollment, Census/ACS and IPEDS Data

Dependent variable:	Share With Any College		Two-Year College Enrollment	
Change defined over following periods:	2006 and 2012	2000 and 2012	2006 and 2012	2000 and 2012
	(1)	(2)	(3)	(4)
2SLS Estimates for All Adults Age 18-25				
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	0.016	-0.005	0.008	-0.002
	(0.010)	(0.010)	(0.005)	(0.007)
2SLS Estimates for Men Age 18-25				
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	0.018	-0.000	0.008	-0.003
	(0.011)	(0.012)	(0.005)	(0.008)
2SLS Estimates for Women Age 18-25				
Housing Demand Change 2000-2006, $\widehat{\Delta H_k^D}$	0.015	-0.009	0.008	-0.003
	(0.011)	(0.012)	(0.005)	(0.008)
First stage F-statistic	38.86	38.86	38.47	38.19
N	275	275	230	227
Include baseline controls	y	y	y	y

Notes: This table reports 2SLS estimates for alternative demographic groups. Columns (1) and (2) report results using Census/ACS data that are analogous to results in Table 3 for alternative years. Columns (3) and (4) report results using IPEDS data that are analogous to results in Table 4 for alternative years. The baseline controls are described in Table 1. Standard errors are shown in parentheses and are clustered by state.

Table 9
The Persistent Effects of Housing Booms on Educational Attainment, NLSY Data

Dependent variable:	Employed, Year 2012	Has Attended Any College, Year 2012	Has Bachelors Degree, Year 2012
	(1)	(2)	(3)
Panel A: OLS Reduced Form Estimates for All Individuals			
Living in Housing Boom MSA in 1997	0.004	-0.048	-0.022
	(0.022)	(0.025)	(0.018)
Mean of Dependent Variable	0.747	0.613	0.302
N	5362	5362	5362
Panel B: OLS Reduced Form Estimates for Men Only			
Living in Housing Boom MSA in 1997	0.008	-0.049	-0.029
	(0.016)	(0.029)	(0.025)
Mean of Dependent Variable	0.778	0.562	0.263
N	2697	2697	2697
Panel C: OLS Reduced Form Estimates for Women Only			
Living in Housing Boom MSA in 1997	0.006	-0.049	-0.015
	(0.022)	(0.032)	(0.026)
Mean of Dependent Variable	0.716	0.664	0.342
N	2665	2665	2665
Include Baseline Controls (Metropolitan Area)	y	y	y
Include Additional Individual-Level Controls	y	y	y

Notes: The unit of observation is individual, and the assignment of housing demand change (between 2000 and 2006) is based on where the individual was living in 1997 at start of the NLSY97 sample. This table reports OLS estimates for alternative demographic groups, and each column reports results for a different dependent variable. The controls and key idenpendent variable are described in Table 6. Standard errors are shown in parentheses and are clustered by state.

Appendix Table A1
First Stage for Housing Demand Change Using
Magnitude of Structural Break in House Prices as an Instrumental Variable

Dependent Variable is 2000-2006 Change in:	Housing Demand, $\widehat{\Delta H_k^D}$	House Prices, *dP*	Housing Permits, *dQ*
	(1)	(3)	(5)
Magnitude of Structural Break in House Prices	3.780	3.292	0.704
[Housing Boom Instrument]	(0.606)	(0.484)	(0.382)
First-stage F-statistic	36.68		
R^2	0.70	0.72	0.34
N	275	275	275
Include baseline controls	y	y	y

Notes: This table reports OLS estimates of the first stage underlying the 2SLS regressions reported in Tables 1-4, 7, and 8. The baseline control variables are described in Table 1. 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, 2001 and Q1, 2005 (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. Standard errors are shown in parentheses and are clustered by state.

Appendix Table A2
List of Metropolitan Areas with Largest Structural Breaks

Rank	Metropolitan Area	Housing Demand Change, 2000-2006	[Housing Demand Instrument] Magnitude of Structural Break	Timing of Structural Break
1	Yuma, AZ	1.396	0.271	2004 Q2
2	Phoenix, AZ	1.009	0.270	2004 Q2
3	Biloxi - Gulfport, MS	0.461	0.250	2005 Q2
4	Boise City, ID	0.868	0.238	2005 Q1
5	Fort Walton Beach, FL	1.193	0.229	2003 Q3
6	Visalia - Tulare - Porterville, CA	1.681	0.214	2003 Q3
7	Lakeland - Winterhaven, FL	1.743	0.212	2004 Q3
8	Naples, FL	1.055	0.212	2004 Q2
9	Albany, GA	0.180	0.207	2005 Q3
10	Las Vegas, NV	1.154	0.194	2003 Q2
11	Fort Myers - Cape Coral, FL	2.128	0.190	2004 Q2
12	Bakersfield, CA	2.075	0.189	2003 Q2
13	Orlando, FL	1.103	0.188	2004 Q2
14	Pensacola, FL	0.821	0.187	2004 Q1
15	Reno, NV	0.930	0.183	2003 Q2
16	Ocala, FL	1.748	0.177	2004 Q2
17	Tucson, AZ	0.909	0.176	2004 Q2
18	Flagstaff, AZ-UT	1.131	0.173	2004 Q1
19	Melbourne, FL	1.644	0.171	2003 Q3
20	Odessa, TX	0.887	0.166	2005 Q1
21	Daytona Beach, FL	1.670	0.157	2004 Q1
22	Wichita Falls, TX	0.806	0.155	2005 Q3
23	Jacksonville, NC	0.954	0.154	2005 Q1
24	Salt Lake City - Ogden, UT	0.456	0.152	2005 Q1
25	Honolulu, HI	1.515	0.143	2003 Q3

Notes: This table reports the top 25 MSAs in the main sample by the magnitude of structural break instrumental variable. See notes to Appendix Table 1 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 0.039 and the average estimated housing demand change is 0.539.

Appendix Table A3
Distribution of Estimated Structural Breaks

Timing of Structural Break	Number of Metropolitan Areas	Percent
2001 Q1	14	4.9%
2001 Q2	6	2.1%
2001 Q3	14	4.9%
2001 Q4	16	5.7%
2002 Q1	1	0.4%
2002 Q2	5	1.8%
2002 Q3	27	9.5%
2002 Q4	7	2.5%
2003 Q1	9	3.2%
2003 Q2	32	11.3%
2003 Q3	22	7.8%
2003 Q4	5	1.8%
2004 Q1	47	16.6%
2004 Q2	14	4.9%
2004 Q3	6	2.1%
2004 Q4	3	1.1%
2005 Q1	55	19.4%

Notes: This table reports the distribution of quarter of estimated structural break for the 275 cities in main sample. See notes to Appendix Table 1 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 housing price data covers 2000 Q1 through 2005 Q4, but structural breaks are limited to 2001 Q1 through 2005 Q1 to follow the 15% "trimming" recommended by Andrews (1993).