HOUSING BOOMS, MANUFACTURING DECLINE, AND LABOR MARKET OUTCOMES*

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

July 2017

Abstract

We study the extent to which manufacturing decline and local housing booms contributed to changes in labor market outcomes during the 2000s, focusing primarily on the distributional consequences across geographical areas and demographic groups. Using a local labor markets design, we estimate that manufacturing decline significantly reduced employment between 2000 and 2006, while local housing booms increased employment by roughly the same magnitude. The effects of manufacturing decline persist through 2012, but we find no persistent employment effects of local housing booms, likely because housing booms were associated with subsequent busts of similar magnitude. These results suggest that housing booms "masked" negative employment growth that would have otherwise occurred earlier in the absence of the booms. This "masking" occurred both within and between cities and demographic groups. For example, manufacturing decline disproportionately affected older men without a college education, while the housing boom disproportionately affected younger men and women, as well as immigrants. Applying our local labor market estimates to the national labor market, we find that roughly 40 percent of the reduction in employment during the 2000s can be attributed to manufacturing decline and that these negative effects would have appeared in aggregate employment statistics earlier had it not been for the large, temporary increases in housing demand. (J21, E24, E32, R23)

* This is a revised version of a previously distributed NBER working paper (NBER Working Paper No. 18949). We thank John Bound, Tom Davidoff, Matt Gentzkow, Ed Glaeser, Erzo Luttmer, Frederic Vermeulen (Editor), and four anonymous referees for their detailed feedback, and we thank Hank Farber and Tom Lemieux for their comments as discussants. Additionally, we thank seminar participants at Columbia, Duke, Harvard, Maryland, Northwestern, Princeton, Tulane, University of British Columbia, University of Chicago, University of Illinois at Chicago, the AEA, Einaudi Institute, the NBER Summer Institute (Macro Perspectives), and the Atlanta, Chicago, Cleveland, and New York Federal Reserves for helpful comments. We are grateful to Loren Fryxell, 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 their financial support and hospitality while working on this project.

† kcharles@uchicago.edu, Harris School, University of Chicago and NBER
‡ erik.hurst@chicagobooth.edu, University of Chicago Booth School of Business and NBER
§ noto@northwestern.edu, Northwestern University and NBER

I. INTRODUCTION

The share of the employed population has fallen sharply since the peak of the last business cycle in 2007, with especially pronounced changes for those with less skill. For example, between 2007 and 2011, employment rates for men aged 21-55 with four-year college degrees fell from 89 percent to 84 percent, and decreased substantially from 83 percent to 74 percent for men aged 21-55 without a four-year college degree. What accounts for these changes? A number of recent papers have examined changes in employment outcomes since 2007, studying the role of factors like de-leveraging associated with falling housing prices (Mian and Sufi 2012), policy uncertainty (Bloom et al. 2012), unemployment benefit extension (Rothstein 2012), the expansion of government transfer programs (Mulligan 2012), and spatial and industry mismatch (Sahin et al. 2012). Yet, employment rates were actually decreasing *throughout* the 2000s, long before the start of the 2007-2010 recession.[1] Focusing on the two business cycle peaks before 2006, employment rates for prime-aged men declined by 1 percentage point between 1989 and 1999, and by an additional 2.5 percentage points between 1999 and 2006 -- both massive decreases, involving millions of workers.[2] These trends suggest that current patterns of employment may be partly attributable to economic forces that predate 2006, and that understanding current employment patterns requires a focus on a period spanning, at least, all of the 2000s.

This paper studies how employment during the entire 2000s was affected by two large changes in the national economy during the 2000s: the continuing decline of the manufacturing sector, and the national boom and bust in the housing market. We study both the separate effects of these two phenomena and how they interacted to affect employment for different population subgroups between 2000 and 2006 and over the entire 2000-2012 period. We focus on manufacturing decline and the housing boom/bust partly because of how large these phenomena were. In the two decades prior to 1999, U.S. manufacturing employment fell from roughly 18.2 million to 17.4 million. However, in the relatively short time between 1999 and 2006, U.S. manufacturing employment fell by an *additional 4 million jobs*. The decline continued through

[1] See Moffit (2012) for a discussion of this phenomenon.
[2] All numbers in this section come from the authors' calculations using the Current Population Survey (CPS). The sample was restricted to men between the ages of 21 and 55 (inclusive).

the 2006-2012 period, with an additional 2 million manufacturing jobs lost.[3] Changes in the housing market were equally dramatic: between 1997 and 2006, after decades of being relatively flat, housing prices surged by about 37 percent, before entirely collapsing over a couple of years.[4]

Beyond the scale of these changes, employment in manufacturing and in activities affected by changes in the housing market have historically been particularly important for less skilled persons – the sub-group experiencing the largest changes in employment since 2000. Figure 1 uses data from the Current Population Survey (CPS) to plot the share among all persons, whether working or not, of men and women aged 21-55 (henceforth, "prime-aged") without a four-year college degree (henceforth, "non-college") working in manufacturing and in construction. Increased housing demand should stimulate changes in construction activity and may also change demand for local labor services as household wealth increases from changes in housing prices. The patterns in Figure 1 for construction employment thus likely represent a lower bound on the total employment changes associated with changes in housing demand. Panel A of Figure 1 shows that fully 37 percent of all non-college men worked in one or the other of these sectors in 1977, and more than 20 percent of all such men continue to do so in 2012. Manufacturing employment for these men has declined sharply over time, falling from 27 percent in 1977 to 14 percent today. Construction employment among non-college men was fairly constant at about 10 percent between 1977 and 1997, then surged during the housing boom to 15 percent, before collapsing with the housing bust after 2006. Although lower than rates for non-college men, employment in manufacturing among non-college women has traditionally also been significant. These rates declined substantially during the early 2000s. Very few non-college women have historically worked in construction, a pattern which was unchanged over the course of the boom and bust in housing (Figure 1, Panel B).

Figure 1 offers suggestive hints that manufacturing decline and changes in the housing market may have played an important role in the evolution of employment since 2000. For example, the patterns suggest that between 2000 and 2006 the roughly five percentage point

[3] Data for changes in manufacturing employment are from the Bureau of Economic Analysis.
[4] There are two bodies of literature studying why these phenomena occurred -- something that is not the focus of our paper. For manufacturing decline, see Autor, Dorn, and Hanson (2013) for analysis of the role of import competition from China in explaining recent U.S. manufacturing declines and Pierce and Schott (2016) for a related analysis of the "surprisingly swift" decline in manufacturing employment coming from changes in trade agreements with China. For housing, see Mayer (2011) and the citations therein for a discussion of why house prices changed during the early 2000s and why they reverted during the late 2000s.

decline in the share of men working in manufacturing was roughly offset by the roughly five percentage point increase in the share of men working in construction. After 2006, the share of men working in either manufacturing or construction fell sharply as manufacturing continued to decline and the construction share reverted to its pre-housing boom level. Second, changes in construction employment during the 2000-2006 period did not offset the decline in the manufacturing share for non-college women. This result suggests that if the housing boom lifted the employment prospects of non-college women, it would likely be through sectors other than construction.

Moving beyond suggestive time series evidence, this paper studies in detail the effect of manufacturing decline and the temporary boom and bust in housing on employment, focusing on the distributional consequences across cities and demographic groups. The empirical work follows a local labor market strategy which exploits variation across metropolitan statistical areas (MSAs) during the 2000s in both the size of the manufacturing decline and in the size of the local housing demand change. To motivate this strategy, we develop a model of sectoral choice, employment, and wages, in the spirit of Roy's (1951) classic framework. This model is related to some of the sectoral models that have recently been developed to study discrimination and inequality (Hsieh et al. 2016, Adao 2016, Burstein et al. 2016), and allows for arbitrary number of sectors and demographic groups. The key insight from the model is that a shock in a single sector will affect wages and employment in that sector but will also affect overall employment (as workers move into non-working sector) as well as employment and wages in other sectors, with the magnitude of these responses governed by comparative advantage and structural of aggregate production function.

Turning to empirical implementation, to study manufacturing decline, we follow Bartik (1991) and Blanchard and Katz (1992) and construct a measure of the predicted change in manufacturing demand in an MSA given by the interaction between an MSA's initial industry mix and national changes in industry employment within narrowly-defined manufacturing industries. [5] The logic of this widely-used measure is that the national decline in the manufacturing sector differentially impacted MSAs because of *pre-existing* differences in the level and composition of manufacturing in the area and the fact that specific manufacturing

[5] Bound and Holzer (1993) employ a very similar method in their work showing a relatively sharp negative relationship between sectoral declines in manufacturing during the 1970s and 1980s and wage and employment outcomes for men.

industries experienced different trends over time. This measure is therefore likely to be systematically unrelated to any change specific to the MSA -- such as MSA-specific labor supply shocks during the 2000s -- that may also affect labor market outcomes. Reassuringly, we find that the measure of predicted local manufacturing change very strongly predicts *actual* changes in MSA manufacturing employment from 2000-2006, suggesting that the measure indeed captures changes in local manufacturing activity in our analysis.

To study changes in housing demand, we note that housing price changes were the most dramatic manifestation of housing demand changes over the 2000s, but there were also almost surely changes in the quantity (and/or quality) of housing which are less readily observed. Using a simple demand/supply framework, we derive a measure of changes in local housing demand that, in principle, captures both the price and quantity effect. Our predicted housing demand measure is a function of the observed price change in the local area and the change in the number of local building permits for new residential construction. We have used this measure in related work studying the effect of local housing booms and busts on educational attainment (Charles et al. 2016).

There is growing consensus that the large temporary changes in housing prices during the 2000s stemmed from factors like the expansion of credit to sub-prime borrowers, low interest rates, the rise of securitization instruments for mortgages in the financial sector, and investor speculative activity -- rather than from changes in household income, population, or construction costs (Mayer 2011; Sinai 2012). This suggests that most of the observed changes in housing demand during the 2000s may be independent of changes in traditional latent factors that also directly affect MSA labor market outcomes. Consistent with this interpretation, we find similar results from Two Stage Least Squares (TSLS) estimates where we instrument for the change in predicted housing demand. To do this, we use an instrumental variable that exploits structural breaks in the evolution in housing prices in an MSA, arguing that these "sharp," or relatively discrete, jumps in housing prices are exogenous with respect to any changes in latent confounds, like labor supply shocks or changes in labor demand, which likely evolve smoothly over time.[6] Across all of our main specifications, we find broadly similar effects for estimated housing demand changes in both the OLS and TSLS specifications, suggesting that variation in

[6] This instrumental variable is introduced and discussed in much greater detail in Charles et al. (2016).

MSA housing prices between 2000 and 2006 was not substantially confounded by unobserved labor supply shifts or other unobserved changes in labor demand.

We find that predicted 2000-2006 manufacturing decline in an MSA decreased employment, lowered wages, and reduced MSA population. The effects for employment and wages were substantial: a one standard deviation increase in the predicted decline in manufacturing in an MSA decreased the overall employment rate for prime-aged individuals in the MSA by 0.7 percentage points and reduced wages by 1.2 percent during the 2000-2006 period. The estimated effects on employment and wages were largest for non-college workers. Additionally, we find that positive shocks to housing demand in an MSA during 2000-2006 increased employment and increased wages. In particular, a one standard deviation increase in housing demand within an MSA increased the employment rate by 1.0 percentage points and increased wages by 1.4 percent points for all prime-age workers. The effect of the housing demand change was largest for non-college men and smallest for college women. Roughly two-thirds of the increase in employment for non-college men in response to the local housing demand increase was the result of increased employment in construction and FIRE (Finance, Insurance, and Real Estate). Non-college women also experienced a large increase in employment in response to the housing demand increase during the early 2000s, but virtually none of it resulted from increased construction employment. Positive housing demand changes increased employment of non-college women mainly through greater employment in the FIRE sector and in the retail and service sectors (Charles et al. 2016).

We next look at average wages by sector, focusing on the manufacturing sector, the construction and FIRE sectors (pooled together), and all other sectors. We find similar wage consequences of manufacturing decline and housing demand changes across each group. This implies that relative wages across sectors are not meaningfully affected by manufacturing decline and/or shifts in housing demand. We show theoretically that this is consistent with a general sectoral choice model that we develop to motivate our empirical analysis. In that model, the specific functional form assumptions regarding comparative advantage across sectors and regarding the aggregate production are sufficient to deliver a proposition which shows that average wages in different sectors are invariant to sector-specific shocks.

Interestingly, over the *entire* 2000-2012 period, we find that the effect of a change in housing demand in an MSA during the housing boom period was fairly small. This results from the fact that almost all of the MSAs experiencing large house price increases from 2000-2006 experienced similarly large reductions in housing prices from 2006-2012. The housing boom lifted local labor markets while the housing bust depressed them. These results contrast sharply with those for manufacturing decline, for which we estimate consistently large effects over the longer term.

According to our estimates, roughly 40 percent of the decrease in employment from 2000-2012 was attributable to declining manufacturing.[7] We show that a large portion of the manufacturing effect on employment was due to an increase in being out of the labor force rather than an increase in unemployment. Additionally, we find that most of our employment effect occurred prior to recent recession; manufacturing decline post-2006 accounted for roughly 12 percent of the decrease in employment during the 2006-2012 period. We find that between 2000 and 2006 the U.S. housing boom reduced the employment rate by roughly 1 percentage point. Over the 2000-2012 period, the housing boom explains very little of the change in employment because the subsequent housing bust undid the employment gains from the preceding housing boom.

Our results suggest that the temporary housing price boom during the 2000-2006 period "masked" some of the adverse labor market effects of the sectoral decline in manufacturing, in the sense that the large employment effects caused by that sectoral decline would have otherwise been evident in the pre-recessionary period of 2000-2006. We emphasize three distinct dimensions to this masking. First, there was significant "cross-MSA" masking: many of the places experiencing large declines in manufacturing employment were different from the places experiencing large, positive housing demand changes. Second, there was "cross-individual" masking, in the sense that the effects of these sectoral changes affected different population sub-groups different (both within and between cities). For example, older workers were much more adversely affected by the decline in manufacturing than were younger workers, while younger workers were more likely to experience increased construction employment following increases in housing demand. Lastly, in related work we have documented significant "within-individual" masking, where the housing boom affects labor market outcomes of

[7] As we discuss below, our results are not substantially affected by accounting for the estimated migration response to the manufacturing and housing shocks when applying our local labor market estimates to a national context. We argue that, if anything, allowing for a migration response as well as other relevant general equilibrium considerations tends to increase the estimated importance of declining manufacturing in accounted for observed changed in employment.

individuals directly affected by manufacturing decline. Using detailed data from the Displaced Workers Survey (DWS), Charles et al. (2016) document that workers displaced from the manufacturing sectors during 2000-2006 were significantly more likely to end up in employment if they lived in a MSA in which housing demand increased sharply from 2000-2006.[8]

Beyond providing new evidence about the effects of arguably two of the largest market-wide phenomena of the past 20 years, our results speak to the ongoing debate about whether there is a structural component to the current high levels of non-employment in the U.S. The finding that the housing boom through 2007 masked systematically worsening labor market conditions from manufacturing decline suggests that changes in employment since 2007, the focus of much recent work, may overestimate the cyclical component in the U.S. labor market. Similarly, the result that manufacturing decline accounts, by itself, for 40 percent of the decrease in employment since 2000 suggests an important explanatory role for factors that are not purely cyclical. It is worth emphasizing that our results do not imply that cyclical forces do not matter importantly for high levels of employment. Indeed, the non-employment growth not accounted for by our estimates may be due to cyclical forces, labor supply responses to changing government policies, or to other structural forces such as spatial mismatch. Lastly, our results focus on short-to-medium run effects, which may overstate or understate longer run effects of manufacturing decline. For example, adverse employment effects of manufacturing decline may be ameliorated over the longer term as workers make adjustments like acquiring more formal human capital, training for new occupations, or moving to new locations.

The remainder of the paper proceeds as follows. In Section II we develop a model that uses the classic Roy framework to study changes in employment and wages in the presence of different sectoral shocks. We next discuss the empirical framework in Section III. Section IV

[8] Some descriptive evidence on the role of housing and manufacturing on aggregate employment is presented in Charles, Hurst, and Notowidigdo (2016). That paper presents time series evidence and cross-MSA regressions which are consistent with the "masking" that we study in detail in this paper; however, that paper only focuses on the employment of prime-age men without a college education, while this paper studies a broad range of demographic groups, as well as additional labor market outcomes such as wages, unemployment, and labor force participation. There are also several additional analyses in this paper that do not appear in Charles, Hurst, Notowidigdo (2016): this paper focuses on estimating the causal effect of local manufacturing shocks and local housing demand shocks using plausibly exogenous variation in manufacturing employment and local housing demand, and also uses the local labor markets estimates to quantify the role of manufacturing and housing in accounting for changes in aggregate employment (both overall and by demographic group).

discusses the data. Section V presents our main empirical results. In Section VI, we apply our local labor markets estimates to the national labor market to try to account for some of the national employment trends since 2000. We conclude in Section VII.

II. CONCEPTUAL FRAMEWORK

In this section, we develop a model of sectoral choice, employment, and wages, in the spirit of Roy's (1951) classic framework. This model is closely related to some of the sectoral models that have recently been developed to study discrimination and inequality (Hsieh et al. 2016, Adao 2016, Burstein et al. 2016).[9] The goal of the model is to provide predictions regarding the effect of sectoral shocks on employment and averages wages, both overall and by sector.

Graphical Model

Before presenting the full model, which allows for many sectors and many demographic groups, we begin with a simple graphical representation in the simplified case where there are only three sectors and a single demographic group. We assume that there are two sectors in which workers can be employed: manufacturing, *M*, and housing-related sectors, *H*. Extending the standard Roy framework, we assume workers have some reservation wage associated with allocating their time to employment sector instead of the non-employment sector, *N*. Workers with skill endowment e_s supply e_s efficiency units of labor in sector *s*. We assume that individual-specific productivity is perfectly negatively correlated so that $e_H = (1 - e_M)$.[10] A worker chooses non-employment if his reservation wage is larger than his highest wage across to two sectors; i.e., $r > \max\{e_M w_M, e_H w_H\}$, and will be employed otherwise. Workers have heterogeneous skill endowments and reservation wages, which are jointly distributed according to the joint distribution $F(e_M, r)$.

Aggregate market output is given by $Y = A_M L_M + A_H L_H$, where A_M and A_H are sector-specific shifters for *M* and *H*, and L_M and L_H are total labor supplies in the two sectors, denominated in

[9] The specific model developed in this paper is broadly similar to Hsieh et al. 2016, with the main difference that there is no endogenous human capital and no discrimination in the labor market. Additionally, we provide additional closed-form results in the case of a single demographic group.

[10] Given this, e_M represents the productivity of the worker in sector *M* relative to the worker's productivity in sector *H*, so that the individual-specific productivity in each sector is perfectly negatively correlated. This is relaxed in the more general model below.

efficiency units. This implies that wages per efficiency unit are pinned down by the demand shifters, so that $w_M = A_M$ and $w_H = A_H$; i.e., relative wages do not depend on relative supplies. With these assumptions, we can define total labor supplies and population shares in sectors M and H, which rely on a marginal worker with skill endowment e_M* who is indifferent between working in sector M and sector H at prevailing wages.[11]

In equilibrium, total labor supply across the sectors is determined by the endogenous self-selection of workers. Both the equilibrium of this simple model and comparative statics can be illustrated graphically. Figure 2 illustrates how workers, in equilibrium, self-select into sectors at all possible combinations of skill endowment and reservation wages, for different values of the productivity shocks. The y-axis in the figure is the reservation wage (r) and the x-axis is the relative skill endowment in manufacturing (e_M), with the entire plane representing all possible (e_M, r) combinations.

Panel A of Figure 2 depicts an initial equilibrium, with workers for whom $e_M > e_{M*}$ choosing to work in the manufacturing sector, M, as long as $e_M > r$. Workers with $s < s^*$ and $s^* > r$ will work in housing-related sectors, H. Workers with a high reservation wage or who have no relative skill advantage in either sector are more likely to be non-employed at any point in time. Panel B of Figure 2 illustrates the effect of a negative shock to manufacturing such as that studied throughout the paper. A negative manufacturing shock, represented by a fall in A_M, is predicted to lower the share of persons employed in manufacturing because of two margins of adjustment. As the figure illustrates, some workers switch from the manufacturing sector, M, to housing-related sectors, H, and other workers are predicted to leave manufacturing to enter non-employment, as represented by the area $M \rightarrow N$. Theory offers little guidance about the relative magnitude of these two effects, as they depend on the distribution of reservation wages and skill among workers. For example, if most workers have very low reservation wages, then a negative shock to one sector will mostly generate switching into the other sector, with little change in

[11] The marginal worker indifferent between sector M and sector H is implicitly defined by following expression:

$$A_M e_M^* = A_H (1 - e_M^*)$$

Total labor supplies in efficiency units are given by the following expressions:

$$L'_M = \int_{e_M^*}^{1} \int_0^{e_M A_M} e_M f(e_M, r) dr de_M \quad L'_H = \int_0^{e_M^*} \int_0^{(1-e_M) A_H} (1 - e_M) f(e_M, r) dr de_M$$

Population shares in each sector are given by the following expressions:

$$L_M = \int_{e_M^*}^{1} \int_0^{e_M A_M} f(e_M, r) dr de_M \quad L_H = \int_0^{e_M^*} \int_0^{(1-e_M) A_H} f(e_M, r) dr de_M \quad L_N = 1 - L_M - L_H$$

employment. This corresponds to a situation of inelastic labor supply, as in occupational choice models such as Autor, Levy, and Murnane (2003), where sector-specific shocks reallocate workers across sectors but do not change aggregate employment. Our various empirical results below suggest, by contrast, that many workers (especially the less-skilled) have reservation wages close to their market wages, since negative manufacturing shocks lead to substantial changes in overall employment in the short-to-medium run, consistent with the results from earlier decades reported in Bound and Holzer (1993).

Panel C and Panel D of Figure 2 illustrate the situation, such as what occurred in the early 2000s, where a negative manufacturing shock occurs simultaneously with a positive shock in the housing-related sector. In Panel C, we highlight only the adjustments along the employment margin. Panel D highlights the margin of substitution resulting from the movement of workers across sectors without the potential for a non-employment spell. The key result from Panel C is that the overall employment effect from a decline in manufacturing is attenuated, or "masked," for two reasons. First, there may be "within-person" masking. This is what occurs when individuals who would have otherwise entered non-employment because of decline in manufacturing are instead employed because of the temporary boom in housing. This area is represented by the diamond area *M->N->H*. In Charles et al. (2016), we use individual-level data from the Displaced Worked Survey to study the extent of within-person masking of manufacturing decline from the 2000-2006 housing boom. Panel C also highlights "across-person" masking that operates across different people, even perhaps across different cities. With this type of masking persons drawn out of non-employment because of growth in housing (*N->H*) are not the same as the persons who enter non-employment from manufacturing (*M->N*). In this paper, we focus primarily on the distributional consequences across geographic areas and demographic groups. This more aggregate notion of masking is a key input into our construction of counterfactual national employment estimates in the absence of the national housing boom and bust.[12]

General Sectoral Choice Model

[12] The model in this section can also be used to understand why it is empirically challenging to estimate the effects of sectoral shifts on wages, since the model reveals compositional shifts induced by shock. As a result, any observed change in wages will reflect both changes in wages for affected workers as well as composition effects. This is the interpretation given in Autor et al. (2013) for the somewhat puzzling pattern of wage effects that they estimate.

The graphical model has several limitations that make it difficult to use for empirical analysis. First, it only has two work sectors (manufacturing and housing-related sectors), so it does not allow us to study the other sectors that might be indirectly affected by manufacturing shocks and housing demand shocks. Second, the structure of the aggregate production function implies that wages are not affected by labor supplies, both at the sectoral level and in the aggregate. As a result, a shock to the manufacturing sector will only affect manufacturing wages, but not wages in other sectors. Since our empirical analysis will look separately at wages overall as well as sectoral wages, we develop a more general model to understand how sectoral shocks affect wages and employment overall and by sector. Lastly, we study outcomes for several demographic groups (age, education, gender), causing the distributional consequences of these shocks to vary both within and across groups. The more general model can accommodate all of these forces.

In the general model, there are M sectors (such as manufacturing and housing-related sectors) and G groups of individuals (which are intended to represent different demographic groups such as age cohorts, education groups, and genders). The wage per efficiency unit of labor supplied to sector m is given by w_m, and it is the same across demographic groups. There is a unit measure of individuals indexed by i, and individuals have sector-specific skills (measured in efficiency units) given by e_{im}. The share of individuals in group g is given by q_g. Each individual i in group g receives utility according in sector m according to $\log(w_m z_{gm} e_{im})$, where z_{gm} is a group-specific utility from working in occupation m.[13] This parameter captures persistent difference in sectoral choices across different groups that are unrelated to sectoral wage gaps.

Given this setup, individuals choose a potential work occupation based on the sector-specific wage (per efficiency unit), their group-by-occupation-specific utility term, and their own idiosyncratic comparative advantage efficiency terms.[14] Individuals then choose work or non-work (with non-working sector capturing leisure choices and/or home production) based on an individual-specific taste for non-work as well as a group-specific taste for non-work given by comparing utility in most preferred potential work occupation. To economize on notation, we

[13] Following Hsieh et al. (2016), this utility specification can be interpreted as arising from a utility function over income and occupation choice given by $U = \log(c_{im}) + \log(z_{gm})$, where consumption, c_{im}, is equal to income, w_m*e_{im}, and the utility over the group-specific taste for sector m is given by z_{gm}.
[14] Formally, each individual solves the following maximization problem: $m_i^* = \arg\max_m \log(w_m e_{im} z_{gm})$.

label the non-working sector to be sector $m = 0$, and we normalize $w_0 = 1$, so that the utility from choosing this sector as $\log(w_0 z_{g0} e_{i0})$. The taste for non-work can also be interpreted as an individual-specific reservation wage.[15]

The distribution of the efficiency terms and reservation wage for each individual is drawn from joint distribution, $F(e_{i0}, e_{i1}, ..., e_{im}, ..., e_{iM})$. The joint distribution allows individuals to have sector-specific comparative advantage. To highlight the role of self-selection, we assume that these skill endowments are exogenous characteristics of the individual, ruling out endogenous human capital investments as in Hsieh et al. (2016). Given these parameters and distributions, we can define the probability of choosing occupation m as P_{mg} and the probability of choosing work and non-working sector as P_g^W and $1 - P_g^W$, respectively.[16]

To complete the model, we assume an aggregate production function, $Y = Y(H_1, H_2, ..., H_M)$, and we assume that wages in each sector are equal to the marginal product of labor in each occupation (i.e., $w_m = \partial Y / \partial H_m$), where H_m is the total amount of labor (in efficiency units) supplied to each occupation.[17] Given this setup, an equilibrium in this multi-sector labor market is defined by the following conditions:

1. Individuals choose work or non-work to maximize utility, and if they choose to work then they work in the occupation that gives them maximum utility, given market wages, group-specific tastes for each sector, and their own idiosyncratic comparative advantage terms.
2. Wages are set equal to the marginal product of labor given the aggregate production function and labor supplies defined by the previous step.

To make this general model tractable, we make several parametric assumptions on the distribution of comparative advantage terms, the idiosyncratic taste for choosing non-working sector (i.e., the reservation wages), and the aggregate production function. First, we assume that

[15] While the model includes a non-working sector, it does not distinguish between voluntary non-participation and involuntary unemployment. Nevertheless, in the empirical analysis we estimate the overall effect of manufacturing decline and housing demand changes on non-employment, as well as separate effects on unemployment and non-participation.

[16] Given this general notation, these probabilities are defined as follows:

$P_g^W = \sum_{m=1}^{M} P_{mg}$; $1 - P_g^W = 1 - \sum_{m=1}^{M} P_{mg} = P_{0g}$

[17] The total amount of labor (in efficiency units) supplied to occupation m is given by the following expressions:

$H_m = \sum_{g=1}^{G} H_{gm}$; $H_{gm} = q_g * P_{gm} * E[e_{im} \mid \text{choosing } m]$

comparative advantage terms and reservation wages are joint distributed as independent Frechet random variables, as follows:

$$F(e_{i0}, e_{i1}, ..., e_{im}, ..., e_{iM}) = \exp\left[-\sum_{m=1}^{M} e_{im}^{-\theta}\right]$$

Note that this assumes a common shape parameter across each sector (including non-working sector). This is a critical assumption to be able to derive closed-form expression for population shares in each sector.

Second, we assume aggregate production function is CES with sector-specific demand shifters A_m and aggregate labor supplies (in efficiency units) as follows:

$$Y(H_1, H_2, ..., H_M) = \left[\sum_{m=1}^{M} (A_m H_m)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}$$

With this setup, we derive closed-form expressions for labor supplies and equilibrium wages for the case of a single demographic group (i.e., $G = 1$).[18] Using these expressions, it is straightforward to prove the following proposition in the single-group case:

Proposition: In the case with a single group ($G = 1$) and arbitrary number of sectors ($M > 1$), a negative shock to sector m (i.e., a reduction in A_m) reduces employment and average wages in sector m and increases the share of the population in the non-work sector and in the other work sectors. The ratio of averages wages across sectors is not affected by the shock, meaning that average wages in all sectors decline by the same proportion in response to any combination of sectoral shocks.

Proof: See Appendix.

The Frechet assumption is a key condition behind this result. The sector m that receives negative shock will experience individuals shifting out of that sector and into other sectors (including non-working sector). As a result, wages in other sectors will fall due to increase in labor supply in other sectors (besides the sector receiving the shock). Although we do not have a formal proof and have not been able to find one in the literature, we believe that the Frechet

[18] In the case of a single group ($G = 1$), the key parameter restrictions for the existence and uniqueness of an equilibrium are that $\theta > 1$ and $\sigma > 0$. This is similar to (but somewhat weaker than) the conditions $\sigma > 0$ and $\theta > \sigma + 1$ in Eaton and Kortum (2002).

distribution is the unique distribution which satisfies the property where the overall wage effect is the same as the wage effect in each sector (in proportional terms). As a result, relative wages between occupations is not affected by sectoral shocks.[19]

In the more general case of multiple groups and multiple sectors (G > 1 and M > 1), we provide in the Appendix a proof for the same result regarding the ratio of average wages. Again, a negative shock to sector m will reduce average wages in all sectors by the same proportional amount. However, there is no longer an analogous result for employment shares. In numerical simulations of our model, we generally find that the employment share decreases in sector m in response to a negative for all groups and increases in the other unaffected sectors as long as the group-specific tastes within a sector are fairly similar across groups. However, this need not always be the case, and as a result, the clearest prediction from the model is that average wages should fall similarly across sectors (both overall and by group), but the changes in employment shares by sector and demographic group is ambiguous once there are multiple groups that have different group-specific tastes. As a result, it becomes an empirical question how employment of different demographic groups responds to sector-specific shocks.

In summary, the model provides the motivation for the empirical strategy described in the following section which relates local declines in the manufacturing sector and changes in local housing demand to changes in local labor market outcomes such as overall employment and average wages, as well as employment and average wages by sector and by demographic group.

III. EMPIRICAL FRAMEWORK

The empirical analysis focuses on comparisons across metropolitan statistical areas (MSAs), which we index by k. We assume that changes in labor market outcomes in a given MSA, ΔL_k, are determined, in part, by labor demand changes arising in three sectors: manufacturing (ΔD_k^M), the housing market (ΔD_k^H), and "other" sectors (ΔD_k^O). Labor market outcomes are also affected by unobserved labor supply changes, which we denote $\Delta\theta_k$. Observed changes in labor market outcomes in a given MSA can thus be written as the general function. We seek to estimate the

[19] Another key assumption behind this result is that wages must be set on the demand curve for each sector. This may not be an accurate approximation in manufacturing sector which has a meaningful union membership share and thus collective bargaining of wages.

effects of changes in the manufacturing sector ($d\Delta L_k / d\Delta D_k^M$) and the effects of changes in housing demand ($d\Delta L_k / d\Delta D_k^H$).

To do this, we construct measures for changes in local manufacturing demand and local housing demand. For local manufacturing demand changes, we use a variant of the widely-used measure in Bartik (1991) and Blanchard and Katz (1992).[20] Conceptually, this measure presumes that a national decline in the manufacturing sector differentially affects local manufacturing demand based on the importance and distribution of manufacturing employment in the local market at some time preceding the national change.

To derive a measure for the change in housing demand, we assume that the log of housing demand and housing supply in a market are given by the following expressions:

$$\begin{aligned} \log\left(H_k^D\right) &= \omega_k^D - \eta_k^{D,H} \log\left(P_k\right) \\ \log\left(H_k^S\right) &= \omega_k^S + \eta_k^{S,H} \log\left(P_k\right) \end{aligned} \tag{2}$$

In (2), ω_k^D and ω_k^S are, respectively, shocks that affect the demand and supply of housing at a given local housing price, P_k, while $\eta_k^{D,H}$ and $\eta_k^{S,H}$ are the price elasticities of housing demand and supply, respectively. Log differentiating the equilibrium condition $H_k^D\left(P_k\right) = H_k^S\left(P_k\right)$ and letting Δ denote log differences, the effect of a shock to housing demand can be expressed as:

$$\Delta\omega_k^H = \eta_k^{D,H} \Delta P_k + \Delta H_k^S. \tag{3}$$

This equation highlights that a change in housing demand produces two effects: a change in the equilibrium housing price and a change in the quantity of housing units supplied in the market. Both the effect on house prices and the change in the housing stock can affect local labor market outcomes, perhaps to different degrees. In particular, house price changes affect household wealth or liquidity and thus households' demand for goods and services produced in the local market (Mian and Sufi, 2012). Changes in the amount (or quality) of housing necessarily

[20] Specifically, we measure sectoral shifts in local manufacturing using:

$$\widehat{\Delta D_k^M} = -\sum_{j=1}^{J} \varphi_{k,j,2000} \left(\upsilon_{-k,j,2006} - \upsilon_{-k,j,2000} \right)$$

where $\varphi_{k,j,2000}$ is the share of relevant population employed in industry *j* in city *k* in the year 2000 and $\upsilon_{-k,j,t}$ is the national employment of industry *j* excluding city *k* in year *t*. The set of industries in *J* includes all 3-digit industries in manufacturing sector. See Autor and Duggan (2003), Luttmer (2005), and Notowidigdo (2012) for other examples of work using variants of this measure.

involves construction activity such as demolition, renovation, home improvements, or new construction. Our analysis does not disentangle the separate effects of household wealth and construction channels, but rather focuses on the combined effect of changes in housing demand. Under the assumption of no unobserved shocks to housing supply, equation (3) thus suggests $\widehat{\Delta\omega_k^D}$ as a natural empirical measure of a housing demand change, where $\widehat{\Delta\omega_k^D}$ is computed using observed changes in local house prices and changes in local housing supply, which we proxy for using housing permits data following Charles et al. (2016).

Given this derivation, we create the following empirical specification:

$$\Delta L_k = \beta_0 + \beta_1 \widehat{\Delta D_k^M} + \beta_2 \widehat{\Delta\omega_k^D} + \alpha X_h + \Delta D_k^O + \Delta\theta_k + \varepsilon_k, \tag{4}$$

where X_k is a vector of observable controls, ΔD_k^O and $\Delta\theta_k$ are unobserved labor demand and labor supply shocks, and ε_k is a mean-zero regression error. The parameters β_1 and β_2 measure, respectively, the direct effect of a predicted change in local manufacturing and of a change in local housing demand, holding the other variables constant. The total effect of either ΔD_k^M or ΔD_k^H consists of the sum of their relevant direct effect, plus any indirect effect operating through the effect of the variable in question on the other measure. We assume that changes in local housing demand do not directly affect local manufacturing activity predicted off national trends in manufacturing. The total effect of estimated housing demand changes on labor market outcomes, or $d\Delta L_k / d\widehat{\Delta\omega_k^D}$, is thus simply β_2. By contrast, standard spatial equilibrium models, such as (Roback 1982), suggest that housing demand is affected by changes in local labor supply and by changes in labor demand in any local sector. It therefore follows that our estimate of local housing demand changes may be written as:

$$\widehat{\Delta\omega_k^D} = \delta_0 + \delta_1 \widehat{\Delta D_k^M} + f(Z_k) + \gamma X_h + \Delta\theta_k + \nu_k, \tag{5}$$

Equation (5) includes several of the same variables as equation (4), along with v_k, which is a mean-zero error term, and Z_k, which represents factors that generate exogenous shocks to local housing demand, such as speculative activity in the housing market. Equations (4) and (5) jointly imply that the total effect of a manufacturing shock on labor market outcomes is therefore $d\Delta L_k / d\widehat{\Delta D_k^M} = \beta_1 + \delta_1\beta_2$. This combines both the direct effect of manufacturing on labor market outcomes as well as the indirect effect coming from the fact that declining manufacturing affects housing demand, which in turn affects local labor market.

In our main analysis, we report estimates of the total effect of changes in manufacturing and housing demand based on estimation of the parameters β_1, β_2 and δ_1. Our baseline estimates of these parameters are from a two-step OLS procedure. We first estimate (5) and retain the estimate δ_1. We then estimate (4) to recover estimates of β_1 and β_2. This regression consistently estimates the two direct effects so long as ΔD_k^M and $\widehat{\Delta\omega_k^D}$ are unrelated to any unobserved changes in other sectors or to unobserved changes in local labor supply. One of the key arguments justifying the use of the predicted manufacturing decline measure is precisely that a measure like ΔD_k^M is likely to be orthogonal to changes in local confounds because it is predicted off of national changes in manufacturing employment. By contrast, as (5) shows, estimated local housing demand changes may depend on changes in unobservable factors that also affect labor market outcomes. In addition, latent housing supply changes as well as measurement error in either ΔP_k or ΔH_k^S would introduce error into the measure of changes in housing demand, which would cause attenuation bias.

To address the possibility of bias in estimates $\widehat{\Delta\omega_k^D}$ from endogeneity and measurement error, we estimate equations (4) and (5) by Two Stage Least Squares (TSLS). To do this, we instrument for $\widehat{\Delta\omega_k^D}$ in the second step of the two-step estimation procedure, using instrumental variable Z_k that measures the degree to which the quarterly time series of housing prices in an MSA exhibited a sharply discontinuous structural break at some point between 2001 and 2005, rather than evolve smoothly over time. The presence and size of these structural breaks strongly predicts the predicted change in housing demand between 2000 and 2006 (i.e., $\widehat{\Delta\omega_k^D}$). As discussed in more detail in Charles et al. (2016), the economic justification for this instrument is that we are assuming that sectoral shocks or labor supply changes are smoothly incorporated into housing price changes. However, other housing demand shocks, such as those that might arise from speculative activity, can affect housing prices either smoothly or discontinuously. If these structural breaks are orthogonal to the effect of other latent confounds, then they are valid instruments for the change in housing prices in TSLS estimation of equation (4) and (5). As in Charles et al. (2016), we show in Appendix Figure OA.2 that the instrumental variable is uncorrelated with pre-existing levels and trends of a range of labor market variables. By contrast, Appendix Figure OA.3 shows that the instrument is strongly correlated with growth in the price-to-rent ratio and growth in the share of housing purchases made by "out of town" buyers, as measured by Chinco and Mayer (2016). These results are consistent with the instrumental variable primarily capturing variation across MSAs in speculative activity during this time period.

Throughout the analysis, we cluster standard errors by state. The analysis is conducted in first differences and thus implicitly accounts for time-invariant differences across MSAs. In most specifications, the X_k vector includes controls for the share of employed workers with a college degree, the share of women in the labor force, and the log of the MSA population. In the next section, we discuss the data used in the analysis in greater detail.

IV. DATA AND SUMMARY STATISTICS

The empirical analysis spans 2000-2012, which covers both the 2000-2006 housing boom and the 2006-2012 housing bust. We create a panel of MSAs using data from the 2000 Census and from various years of the American Community Survey (ACS) individual-level and household-level extracts from the Integrated Public Use Microsamples (IPUMS) database (Ruggles et al., 2004). Restricting attention to persons living in metropolitan areas, we compute mean wages, employment shares, employment shares in various occupations, and total population in each MSA. In 2000, these means are from the 2000 Census. For the 2006 numbers, we pool the ACS data from 2005 to 2007 to increase the precision of the MSA estimates. Similarly, we pool the 2011-2013 ACS for the 2012 numbers. Because of the large sample sizes, the various means can be reliably estimated for separate sex×education groups. The primary sample consists of non-institutionalized persons aged 21-55. Much of the analysis focuses on non-college men, but we also present results for non-college women and for college-educated men and women. We use 3-digit industry classifications for persons in the labor force in the Census and ACS data to construct the predicted manufacturing decline measure.

We compute local house prices using data from the Federal Housing Finance Agency (FHFA), which is a repeat-sales housing price index with data for most metropolitan areas. We mapped the FHFA metro areas to the Census/ACS metro areas by hand.[21] To mirror the ACS data, we construct average house price growth between 2000 and the average of house price in

[21] See the Charles et al. (2016) for details of this matching procedure.

the first quarter in 2005, 2006, and 2007. Similarly, when computing house price changes between 2000 and 2012, we use the pooled FHFA data for 2011, 2012, and 2013. To compute estimates of change in housing demand ($\widehat{\Delta\omega_k^D}$), we combine the change in house prices and change in housing permits, and we assume unitary elasticity of demand as in Charles et al. (2015) to implement the housing demand measure given by equation (3) above.[22]

Table 1 reports summary statistics of the housing market and manufacturing changes among the 275 MSAs with non-missing labor market and housing market data that constitute the main analysis sample. The top row of the table shows that over the boom period of 2000-2006, MSA house prices rose by roughly 50 percent on average. This increase is not driven by a few outlier MSAs. Prices rose sharply throughout the distribution, more than doubling at the 90th percentile MSA and increasing by 5.4 percent even at the 10th percentile.

The next two entries in the table are summary statistics for the two measures used in the paper to measure sectoral changes in housing and in manufacturing. As discussed earlier, housing price changes alone do not capture changes in local housing demand since there will, in general, be supply responses to these changes in demand. Our estimated housing demand measure is meant to account for both the price and supply effect. The table shows that during the boom the average MSA experienced a 60 percent increase in housing demand. The next entry in the table shows summary statistics for the predicted manufacturing change measure. From 2000-2006, the national decline in manufacturing was predicted to lower the share of all men and women employed in manufacturing by 1.5 percentage points.

A natural question about the two measures used in the paper is whether they are, in fact, strongly correlated with actual sectoral changes we contend they capture. Figure 3 shows that the predicted manufacturing measure is strongly correlated with actual changes in the share of the prime-aged population working in the manufacturing sector, suggesting that the predicted measure does capture local manufacturing demand shocks. Similarly reassuring is the strong positive association in Figure 4 between our estimated housing demand measure and the fraction

[22] In Charles et al. (2016), we report similar results using the MSA specific housing supply elasticity measures from Saiz (2010), who estimates local housing supply by MSA using detailed information on the amount of land available for development. The local housing supply elasticity is a substitute for the local housing permit data, because the observed change in price can be combined with local housing supply elasticity to construct an alternative estimate of overall change in housing demand.

of the total population in the MSA employed in construction -- an activity that would rise with positive local housing demand shocks.

V. MAIN RESULTS

Graphical Results

We begin our analysis of masking with some graphical evidence. We first characterize MSAs that experienced especially large housing demand changes, as those in the top tercile (one-third) of the distribution of the housing demand change measure, $\widehat{\Delta\omega_k^D}$. We refer to these MSAs as "housing boom MSAs." We then plot the relationship between predicted decline manufacturing between 2000 and 2006 (ΔD_k^M) in an MSA and the change in the share of non-college men in non-employment during the same time period, separately by "housing boom MSAs" and all other MSAs.

Figure 5 presents the 2000-2006 plots. In the figure, "housing boom MSAs" are represented with triangles, and the remaining two-thirds of MSAs are shown with circles. The gray line is the bivariate regression line for MSAs with housing price changes in the bottom 2/3 of the sample. The large and precisely estimated negative slope coefficient (-1.19, s.e. 0.25) implies that predicted manufacturing declines sharply increase non-employment among non-college men. Most of the triangles in the figure lie below the regression line, implying that MSAs with especially large housing demand changes experienced larger increases in employment rate among non-college men than did other types of MSAs with similar predicted changes in manufacturing. Formally, housing boom MSAs systematically had 2.4 percentage point higher employment growth for any given manufacturing decline than non-housing boom MSAs (standard error of the difference = 0.7 percentage points).

Figures 6 through 8 are analogous to Figure 5, with the change in non-employment replaced with the change in construction employment, average wages, and manufacturing employment. In Figures 6 and 7, the results show that housing boom MSAs had systematically higher increases in construction employment and average wages. For manufacturing employment, there is no clear difference between the housing boom MSAs and other MSAs (Figure 8). This is consistent with the identifying assumption that there is no direct effect of housing boom on manufacturing employment.

Lastly, in Figure 9 the change in employment for non-college men is defined over the 2000-2012 period. The results show that the temporary housing demand shock during 2000-2006 had no lasting effects on employment over the entire 2000-2012 period. This can be seen from the fact that "housing boom" MSAs are distributed evenly around the regression line for the other MSAs. Formally, there is no difference in the intercept of the regression line based on the MSAs that did and did not experience a housing boom between 2000 and 2006 (intercept difference = -0.001 with a standard error of 0.016).

Overall, these results suggest that there was significant masking during the 2000-2006 period both within and between MSAs. In MSAs that experienced a large decline to manufacturing demand, those that also experienced a large housing boom had smaller increases in employment during the 2000-2006 period. The masking results were undone as the housing bust occurred. Over the entire 2000-2012 period, MSAs that experienced a large decline in manufacturing had similar levels of employment regardless of what happened to housing prices in that MSA during the 2000-2006 period. These results focus on non-college men, but similar patterns can be seen in analogous figures for other demographic groups (not shown).

Employment Estimates: 2000-2006

Panel A of Table 2 presents the OLS estimates of the joint estimation of equations (4) and (5), using the two-step OLS estimator described in Section 3. To interpret the magnitudes, the rows below the estimated coefficients are re-scaled to represent a one standard deviation change.[23] The point estimates in the first column of the top panel of Table 2 imply that a one standard deviation larger predicted manufacturing decline decreased employment among non-college men by 0.8 percentage points during 2000-2006. Likewise, over the same period, a one standard deviation increase in housing prices increased the employment of non-college men by 1.7 percentage points. Column 2 presents results for college-educated men. The standardized effects are quite small relative to those for non-college men -- less than half the size in the case of predicted manufacturing decline and about one-fifth the size for estimated changes in housing demand. As columns 3 and 4 show, whereas the effects of manufacturing and housing demand shocks on employment for non-college women are comparable to the effects for non-college

[23] The coefficients are always standardized by the cross-city standard deviation in magnitude of the manufacturing shock or the housing shock during the time period analyzed.

men, there was little effect on the employment of college educated women. Employment effects for the entire population of men and women aged 21-55 are shown in column 5. These results are closer to the results for persons without a four-year college degree which is not surprising given that this sub-sample is roughly two-thirds of the overall sample population.

How much of these changes in employment from housing demand increases can be attributed to changes in construction employment? Panel B of Table 2 presents results analogous to those in Panel A, but with the change in the share of individuals in the MSA working in construction and FIRE as the dependent variable. The standardized effect of the housing demand change in Panel B divided by the standardized effect of the housing demand change in Panel A measures how much of the employment effect is from construction and FIRE. For example, a one standard deviation increase in the housing demand for non-college men increased their construction and FIRE employment by 1.1 percentage points, which accounts for roughly 70 percent (1.1/1.7) of the decline in employment of non-college men in response to a housing demand. Notice that for non-college women, roughly 30 percent of the reduction in employment to the housing demand change comes from increased construction and FIRE employment. Most of this increase is due to FIRE and not construction, suggesting that the effect of the housing boom on employment for women operated through increased employment in sectors other than construction.

These results are broadly consistent with the aggregate time series patterns in Figure 1, showing a large increase in construction employment for non-college men but none for non-college women during 2000-2006. The estimates in Table 2 also illustrate the important limitation of using only construction to measure the effect of housing demand increases on employment during the 2000s. For non-college men, housing demand changes strongly affected construction employment, but there were also employment effects outside of construction. For non-college women, virtually none of the labor market response to the housing demand increase occurred via increased construction employment. Through a local spillover mechanism, changes in local housing demand affected employment through other channels -- most likely in local retail and services. Panel B of Table 2 also highlights the local spillover effects of manufacturing decline on employment in the construction sector. Across all individuals, a one standard deviation decline in manufacturing demand reduced construction employment by 0.2 percentage points.

As manufacturing declines in a locality, housing demand also falls (Blanchard and Katz 1992). Given our joint estimation of (5) and (6), the effect of a manufacturing decline on employment that we report includes both the direct effect as well as the indirect effect through changes in local housing demand. Using all estimated parameters in the equations, the direct and indirect effects can be reported separately, as shown in Appendix Table A.1. These results suggest that a large share of the overall negative effect of manufacturing decline on employment comes from indirect effect of declining housing demand. However, the direct effects are economically significant for non-college men and women, although not precisely estimated for non-college men. These results potentially provide a way to re-interpret the results of the recent work studying manufacturing decline. For example, our results suggest that an important part of the large negative employment effects of China trade in Autor et al. (2013) may come from the indirect effects of manufacturing decline on housing demand, and are thus broadly consistent with the large role of housing demand on overall employment estimated in Mian and Sufi (2012) as well as the important role of construction employment in accounting for trends in aggregate employment in both the United States and Germany (Hoffman and Lemieux 2016).

Effect on Different Sub-Populations: 2000-2006

One interesting question is whether the effects in Table 2 differ by other key demographic traits. For example, one might imagine that a sectoral decline affects workers differently based on their age, since industry-specific human capital grows as workers age. Table 3 presents results for non-college men (columns (1) and (3)) and for all workers (columns (2) and (4)) separately by two age-groups: ages 21-35 and ages 36-55.[24] We find that changes in estimated housing demand produced broadly similar employment effects for both older and younger workers. By contrast, declines in manufacturing decreased employment among older workers by nearly twice as much as was true for younger workers.

We also explored the degree to which the results -- particularly for housing -- differ across native workers and immigrants. To this end, we have re-estimated the models in Table 2 only on a sample of workers who were born in the U.S. These results are presented in the last two columns of Table 3. Among native workers, the manufacturing results are very similar to those

[24] To conserve space, some of our future tables only highlight the results for non-college men and for all workers. However, in the Online Appendix we provide analogous tables showing the effects for non-college women, college men, and college women.

reported in Table 2. However, the effect of the housing demand shock on employment is roughly 40 to 60 percent smaller in the sample of native workers. For example, for native workers, a one-standard deviation increase in housing demand increased employment of non-college men by 1.0 percentage points, as opposed to 1.6 percentage points in the full sample. This suggests that employment of immigrants was particularly responsive to changes in housing demand during this time period.

In summary, the manufacturing and housing demand changes experienced during the 2000s had differential effects across sub-groups based on gender, education, age, and immigrant status. In particular, the manufacturing decline hit older workers harder than younger workers and housing demand changes affected native workers somewhat less than immigrants. Both manufacturing decline and the housing boom seem to have caused larger changes for non-college workers (compared to college workers).

Effects on Average Wages, Overall and by Sector: 2000-2006

The model in Section 2 suggests that sectoral declines in manufacturing or increases in housing demand affects labor market outcomes via changes in labor demand. If this reasoning is correct, then falling labor demand in manufacturing sector in an MSA should be accompanied by declining local wages. Likewise, housing demand increases in an MSA should be associated with rising local wages. The wage effects should also be largest for those groups that had the largest employment response to the sectoral shift.

The regressions in Panel A of Table 4 explores these ideas. These regressions are analogous to the regressions in Tables 2 and 3, except that the dependent variable is now the growth in average log wages in the MSA for a given group during a given time period.[25] As Table 4 shows, a one standard deviation manufacturing decline reduced wage growth for non-college men between 2000 and 2006 by 2.1 percentage points. For all workers, the wage response between 2000 and 2006 to the manufacturing decline was smaller at 1.2 percentage points. With respect to a one-standard deviation housing demand increase, the wage response between 2000

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

and 2006 was 2.0 percentage points and 1.4 percentage points for non-college men and all individuals, respectively. These results are consistent with our interpretation that these sectoral shifts affect local labor markets through their effect on labor demand.

The general model in Section 2 predicts that average wages by sector should not be differentially affected by either of the sectoral shocks. We test this in Table 5 by replacing average wages (overall, across all sectors) with average wages in specific sectors. We look at manufacturing sector (all industries in manufacturing grouped together), housing-related sector (which we continue to define as construction and FIRE, or Finance, Insurance, and Real Estate), and a catch-all “other” sector. In each case, the change in average wages overall is very similar to the sector-specific wages, consistent with the prediction from the general model in Section 2. We are not able to statistically distinguish these wage responses. This suggests that the reallocation of workers across sectors (and into non-employment) in response to manufacturing and housing demand changes may be consistent with the patterns of comparative advantage described by the model in Section 2. Of course, we would not necessarily expect this pattern to hold everywhere, but in our setting it appears that the standard sectoral choice model (in terms of functional form for comparative advantage terms and aggregate production function) describes the pattern of wage adjustment fairly accurately.

Longer Run Employment Effects: 2000-2012

The results in Tables 2 and 3 show the shorter run effect of the sectoral changes during the 2000-2006 period. How long-lasting were these effects? In Table 6, we examine the effect of manufacturing and housing demand changes over the entire 2000s. Columns 1 and 2 re-display the corresponding results for non-college men and all workers from panel A of Table 2. In columns 3 and 4, we assess whether the 2000-2006 sectoral changes had persistent employment effects over the entire 2000-2012 period. The results indicate that the effects of predicted manufacturing decline during the 2000-2006 period for both non-college men and for the overall population were, in fact, fairly durable, although the standard errors increase substantially for both groups in these specifications, so that only the results for all men and women remain statistically significant. The results for the employment effects of housing demand changes, however, differed sharply over 2000-2006 and the longer 2000-2012 period. In particular, we find that changes in estimated housing demand during the housing boom period (2000-2006) had no significant long-term effect on employment of either non-skilled men or of the entire population during the 2000-2012 period.

What accounts for this pattern? We believe that the key explanation has to do with the nature of transitory housing price variation over the 2000s. Since there was a strong correlation between the magnitude of a MSAs housing price growth during the housing boom and its subsequent price decline during the years of the housing bust, for most MSAs there was little change in estimated housing demand over entire decade. This point can be seen quite dramatically in Online Appendix Figure OA.1. This figure plots a MSAs housing price reduction between 2006 and 2012, against its price increase from 2000 to 2006. The line in the figure is a 45-degree line. The figure shows clearly that for the overwhelming majority of MSA, price increases during the boom were nearly exactly offset by declines during the housing bust. Although not shown in Table 6, we estimate that there was a very strong relationship between housing demand declines during the housing bust of 2006-2012 and local labor market outcomes during the bust. The estimated magnitudes were nearly identical to the estimates during the boom period.[26]

Employment, Labor Force Participation, and Unemployment

The results in Tables 2 and 3 focused on the effects of housing demand changes and manufacturing decline on the employment to population ratio. While the sectoral choice model emphasizes workers freely choosing between paid work and the non-working sector, some of the increase in non-employment may be due to involuntary unemployment. As a result, in Table 7, we decompose the employment effects of housing demand changes and manufacturing decline into changes in non-participation and changes in unemployment. The results in Table 7 show that roughly half of the overall estimated employment effect is accounted for by changes in non-participation for non-college men, and this is similar for both housing demand changes and manufacturing decline. For non-college women and for the overall sample of all prime-aged men and women, the estimates show a similar pattern of results for manufacturing decline, but for housing demand change estimates, there is a somewhat larger role for unemployment

[26] Our results during the bust period are similar to recent research by Mian and Sufi (2012) and Midrigan and Philippon (2011). Both papers show that during the recession, places with large house price declines had larger increases in employment. Our results, however, suggest that in the pre-recessionary period, places that had housing booms also had large declines in employment. Over the decade as a whole, the housing boom/bust cycle had very little impact on local labor markets.

(relative to non-participation) in accounting for overall change in employment. This is consistent with changes in non-participation from housing booms as being relatively more important for non-college men than other demographic groups.

Migration Effects

In Panel B of Table 8, we estimate whether local changes in manufacturing and housing result in migration across MSAs. As one location receives a negative shock to labor demand, previous work suggests that some individuals respond in part by migrating elsewhere (Blanchard and Katz 1992; Notowidigdo 2013). We find that in response to a one standard deviation manufacturing decline (housing demand increase) change during the 2000-2006 period, the MSA population of prime age non-college men fell by 1.8 percentage points (increased by 1.6 percentage points) during that same period. The results are nearly identical for all prime age men and women. The migration response to the manufacturing decline was actually larger over the longer 2000-2012 period while the response to housing demand increases was smaller. This is not surprising given that the 2000-2012 period witnessed the continuing decline in manufacturing, and the growth and reversal of the housing boom.

Robustness to TSLS Estimation of Housing Demand Change

As expression (6) shows, the measure of housing demand changes may be endogenous in OLS regressions. Additionally, since the housing demand change measure in empirical analysis is constructed with the assumption that there are no housing supply shocks, it may be an error-ridden version of true housing demand changes. We address both of these potential concerns using Two Stage Least Squares (TSLS) analysis described above.

Table 8 reports TSLS results analogous to the main results in Table 2.[27] Across all five columns, we consistently find that the estimated effects of house price booms during the 2000-

[27] In Online Appendix Table OA.1, we report estimate of first stage behind instrumental variables estimates. Columns 1 and 2 of the table show that the size of a structural break in a city's quarterly price series strongly predicts the size of the city's 2000-2006 change in housing prices. The large and strongly statistically significant point estimates are robust to the inclusion of the set of controls used previously and to controlling for the predicted manufacturing decline measure. The final column in the table examines how employment in manufacturing among non-college men is affected by the structural break variable. We find no relationship between these measures suggesting that the instrument is orthogonal to changes in manufacturing demand. Additionally, the F-statistic on the structural break measure is always around 30, which suggests that there is no "weak instrument" concern.

2006 period are similar to our OLS results. The point estimates are generally slightly larger than the corresponding OLS results, which is consistent with the idea that either some of the variation in house price changes was actually the result of changes in unobserved labor demand or labor supply or that there is some measurement error in our housing demand estimates. However, the broad similarity between the OLS and TSLS results suggests that most of the variation in housing prices at the MSA level between 2000 and 2006 was not significantly confounded by omitted variables or by housing supply shocks. In the Online Appendix, we present analogous TSLS for all of the results Tables 2 through 8. Across all of these specifications, we consistently find broad similarities between the OLS and TSLS results. In particular, we find broadly similar distributional consequences across education, age, and immigrant status, and we continue to find similar changes in sectoral wages to both manufacturing decline and changes in housing demand.

VI. AGGREGATE MASKING AND COUNTERFACTUAL EMPLOYMENT ESTIMATES

In this section, we use the estimated effects of manufacturing and housing demand changes to conduct counterfactual analyses of aggregate national employment during the 2000-2012 period. This analysis also provides, in essence, an estimate of the sum of cross-individual and the within-individual masking illustrated above.

To perform the counterfactual exercise, we combine the main point estimates from Table 2 with national time series changes in the employment rate, housing demand changes, and manufacturing employment shares to compute the separate contributions of declining manufacturing and housing demand changes on aggregate employment. Panel A of Table 10 reports the exercise for all prime age men and women. The share of all prime age men and women employed in manufacturing declined by 3.2 percentage points between 2000 and 2006. Using the estimates in column 5 of Table 2, this is predicted to decrease employment by 2.1 percentage points.[28] With respect to housing demand, the mean change over the 2000-2006

[28] For the national trends in employment and manufacturing over the 2000-2006 period and the 2000-2012 period, we use data from the CPS. These are the same data used in Figure 1. We use data from the CPS rather than the Census/ACS because the employment rates in the 2000 Census are systematically high relative to both the 2000

period of 0.6 (see Table 1) and the point estimates in Table 2 imply a decline in employment of -1.1 percentage points. Together, the two types of shocks were thus predicted to decrease non-employment for all prime aged men and women by 0.9 percentage points between 2000-2006. The actual increase in non-employment for all prime age men and women was 1.9 percentage points. Therefore, these two sectoral changes we study are estimated to jointly explain roughly one-half of the observed changes in employment during the early-to-mid 2000s. Notice that we would have predicted a 30 percent larger increase in employment during 2000-2006 had there been no housing demand changes. This is thus a measure how much, in total, the temporary housing demand changes masked the effect of manufacturing decline during the period. The next row of Panel A examines the entire 2000-2012 period. The results show that the predicted change in employment attributable to manufacturing decline is 3.0 percentage points, or 40 percent of the actual increase during the longer time period. As we have discussed at length, housing demand changes explain none of the changes in employment over the longer term. The results for non-college men in Panel B are broadly similar to the results in Panel A for all men and women. In particular, while the estimated effects imply greater absolute increases in employment for non-college men, the percentage of overall employment growth accounted for by manufacturing decline and housing demand changes is very similar.

Collectively, the results indicate that a non-trivial portion of the decrease in employment of both non-college men and all workers can be attributed to the continuing decline in the manufacturing sector, suggesting that structural forces account in part for the recent weak performance of the U.S. labor market. The results also imply that without the temporary boom in housing, and the masking associated with it, 1.3 million prime age workers would have been non-employed as early as 2006. The negative effects of structural manufacturing decline were masked in aggregate statistics during the early to mid-2000s. Importantly, for both non-college men and all workers, roughly two-thirds to three-quarters of effect of manufacturing decline pre-dated the 2008 recession, as a comparison across rows in column 3 of Table 10 shows. Although we argue that structural forces associated with manufacturing decline appear to have clearly mattered importantly for employment, it should be emphasized that the results do rule out a key role for cyclical forces.

CPS to the 2001 ACS. This fact has been carefully documented Clark et al. (2003). Using the CPS data guarantees provides a more consistent time series trend.

While we find these counterfactual exercises useful to assess the role of manufacturing decline and the housing boom on aggregate employment, there are various concerns associated with applying "local" MSA-level estimates to the national labor market. One issue is migration. In Table 8, we showed that one-standard deviation changes in manufacturing and housing shocks generate migration responses of about 2 percentage points. Given that the two sectoral changes both affect employment by roughly 1 percentage point it is possible to bound how much endogenous migration could affect the counterfactual estimates. For one bound, we assume that all of the migrants would have been non-employed had they not moved. In this case, the aggregate employment rate in response to a one standard deviation manufacturing shock would have increased by an additional 2 percentage points, from 1 to 3. The counterfactual estimates above would thus be severely underestimated. If we assume instead that all migrants would have been employed had they not moved, the estimated response to a one-standard deviation manufacturing shock would fall by roughly 0.02 percentage points, from 1 to about 0.98. This effect is so small because the number of people migrating out of the MSA in response to manufacturing shock is very small relative to the number of people who are employed in the MSA. Therefore, assuming that migrants are either more employable than the average non-migrant or roughly similar to the average non-migrant has a negligible effect on our results. If, however, the marginal migrant is much less employable, then our counterfactual estimates are quite conservative.

A second potential concern is that the counterfactual results ignore potential general equilibrium and feedback effects. For example, changes in house prices may have a direct effect on U.S. manufacturing demand. Mian and Sufi (2011) show that households that experienced large increases in housing prices increased their purchase of both local services and nondurable expenditures because of either a wealth or liquidity effect. Local housing booms can thus affect the national demand for manufacturing goods. As with migration, this type of feedback would again cause us to underestimate the extent of masking during the 2000-2006 period, since the decline in manufacturing in those years would have been even greater had there been no housing boom in the U.S. which effectively "propped up" manufacturing demand. A similar type of potential feedback is the possibility that manufacturing decline during 2000-2006 could have been one of the proximate causes of the housing boom. This channel seems highly implausible, since we find that local declines in manufacturing put downward pressure on local housing demand. This suggests that an important feedback effect of manufacturing decline is the housing market, and our empirical specification is designed to try to capture this, so that our local estimates of manufacturing decline capture both the direct effect of manufacturing decline and the indirect effect coming through declining housing demand.[29]

VII. CONCLUSION

This paper studies how manufacturing decline and housing booms affect labor market outcomes, with a particular emphasis on employment among the two-thirds of workers without a college degree. We estimate a variety of cross-MSA models which exploit the variation in both the magnitude of the negative shock to manufacturing as well as the sudden and dramatic increases in housing demand.

We find that roughly 40 percent of the decrease in employment during the 2000-2012 period can be attributed to the decline in manufacturing. These employment effects were very large for non-college men, but we find that local manufacturing shocks significantly decreased employment for other groups as well, such as non-college women. The large adverse labor market effects of manufacturing decline are present during the housing boom (2000-2006), during the collapse in the housing market (2006-2012) and over the entirety of the 2000s, over 2000-2012. We also find that increases in housing demand sharply raised employment during 2000-2006, especially among non-college men and women. The reversal of the housing market during 2006-2012 among cities experiencing unusually large increases in housing demand during 2000-2006 implies that, over the entire 2000-2012 period, local housing booms did not significantly contribute to longer run changes in labor market outcomes.

[29] Finally, for reasons similar to the general equilibrium effects during the housing boom, our analysis may overstate the effect of manufacturing decline during the housing bust period if falling housing prices dampened demand for manufactured goods during the 2006-2012 period. In this case, the change in manufacturing between 2006 and 2012 on which the counterfactuals are based would be too large. We do two things to address this concern. First, we redo our counterfactuals assuming that the trend in manufacturing between 2000 and 2006 continued through 2012. This assumption strikes us as reasonable, given the relatively steady, 40-year decline in manufacturing in the U.S. Linearly extrapolating the trend in manufacturing through 2012, we find nearly identical results to those in Table 10, since the actual decline in manufacturing employment between 2006 and 2012 is very close to the linear extrapolation based on the 2000-2006 trend. Second, we re-estimate the model ignoring the decline in manufacturing during the recession and focusing on the manufacturing decline before the recession. Again, we find sizable effects of the manufacturing decline during 2000-2006 on current employment prospects in the U.S.

The results imply that the positive labor market effects of the temporary housing boom "masked" the negative effect of sectoral decline in manufacturing that would have otherwise been more evident in the mid-2000s. The collapsing of the housing market during 2006-2012 not only had an independent adverse effect on labor market outcomes for some sub-groups but also "unmasked" the negative manufacturing effect that would have been seen earlier. A key conclusion from our analysis is that there are important distributional consequences, with different labor market effects by education, age, gender, and immigrant status. While we find significant “masking” in the aggregate, there are meaningful differences across geographic areas and demographic groups.

Sectoral booms and busts are often linked to aggregate business cycle dynamics. All else equal, a sectoral boom will increase wages and employment during the expansion and result in wages and employment falling during the contraction. Our results, however, highlight that sectoral booms and busts have very different aggregate employment dynamics when another sector in the economy is experiencing consistent, ongoing decline. In this case, negative labor market effects are muted during the boom and very large during the bust. The behavior of the employment rate since the early 1980s suggests the potential importance of a mechanism like this in the U.S. labor market. Since 1980, the employment rate of men in the U.S. has been relatively stable during expansions and has adjusted sharply around contractions. This point has been emphasized recently by Jaimovich and Siu (2014), and our results suggest that booms and busts in other sectors combined with a sectoral decline in manufacturing partly generate these patterns.

Finally, we think that our results may inform the current policy debate about how best to stimulate employment. The type of employment we have identified is the result of the longer run sectoral decline in manufacturing. Temporary boosts to labor demand from hiring subsidies or infrastructure investments are unlikely to have permanent effects on the employment prospects of non-college individuals, since labor demand for these workers will remain depressed once these subsidies expire because of the decline in the manufacturing sector. In this sense, our paper documents a significant role for structural forces in explaining the current high level of employment in the U.S. As noted above, over longer periods of time, non-employed workers (as well as subsequent generations of workers) may find it beneficial to invest in human capital accumulation. Therefore, addressing barriers to skill acquisition may have most lasting effect on improving the employment prospects of those workers who leave the labor force as a result of the ongoing decline in the manufacturing sector.

REFERENCES

Adao, Rodrigo (2016). “Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence from Brazil,” Working Paper.

Autor, David, David Dorn, and Gordon Hanson (2013). "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*.

Autor, David and Mark Duggan (2003). "The Rise in the Disability Rolls and the Decline in Unemployment", *Quarterly Journal of Economics*, 118(1), 157-206.

Autor, David, Frank Levy, and Richard Murnane (2003). "The Skill Content Of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 118(4): 1279-1333.

Bartik, Timothy (1991). "Who Benefits From State and Local Economic Development Policies", W.E. Upjohn Institute for Employment Research: Kalamazoo, Mich.

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

Bound, John and Harry J. Holzer (1993). "Industrial Shifts, Skills Levels, and the Labor Market for White and Black Males," *The Review of Economics and Statistics*, 75(3): 387-96.

Burstein, Ariel, Eduardo Morales, and Jonathan Vogel (2016). “Changes in Between Group Inequality: Computers, Occupations, and International Trade,” Working Paper.

Charles, Kerwin, Erik Hurst, and Matthew Notowidigdo (2015). “Housing Booms and Busts, Labor Market Opportunities, and College Attendance,” NBER Working Paper #21587.

Charles, Kerwin, Erik Hurst, and Matthew Notowidigdo (2016). “The Masking of Declining Manufacturing Employment by the Housing Bubble,” *Journal of Economic Perspectives*, 30(2): 179-200.

Chinco, Alex and Christopher Mayer (2016). “Misinformed Speculators and Mispricing in the Housing Market,” Review of Financial Studies, 29(2): 486-522.

Clark, Sandra Luckett, John Iceland, Thomas Palumbo, Kirby Posey, and Mai Weismantle (2003). "Comparing Employment, Income, and Poverty: Census 2000 and the Current Population Survey." Report for the Housing and Household Economic Statistics Division of the U.S. Census.

Eaton, Jonathan and Samuel Kortum (2002). “Technology, Geography, and Trade,” *Econometrica*, 70(5): 1741-1779.

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.

Hoffman, Florian and Thomas Lemieux (2016). “Unemployment in the Great Recession: A Comparison of Germany, Canada, and the United States,” *Journal of Labor Economics* 34:S1: S95-S139.

Hsieh, Chang-Tai, Erik Hurst, Chad Jones, and Peter Klenow (2016). “The Allocation of Talent and U.S. Economic Growth,” Working Paper.

Jaimovich, Nir and Henry Siu (2014). “The Trend is the Cycle: Job Polarization and Jobless Recoveries,” Working Paper.

Luttmer, Erzo (2005). "Neighbors as Negatives: Relative Earnings and Well-Being", *Quarterly Journal of Economics*, 120(3), 963-1002.

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

Mian, Atif and Amir Sufi (2012). "What Explains High Unemployment? The Aggregate Demand Channel", University of Chicago Working Paper.

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.

Midrigan, Virgiliu and Thomas Philippon (2011). "Household Leverage and the Recession", NBER Working Paper 16965.

Moffit, Robert (2012). "Th U.S. Employment-Population Reversal in the 2000s: Facts and Explanations", Brookings Papers for Economic Activity, forthcoming.

Mulligan, Casey (2012). The Redistribution Recession: How Labor Market Distortions Contracted the Economy. Oxford University Press, New York.

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

Rothstein, Jesse (2012). "Unemployment Insurance and Job Search in the Great Recession", Brookings Papers on Economic Activity.

Roy, A. D. (1951). "Some Thoughts on the Distribution of Earnings," Oxford Economic Papers, 3, 135-46.

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

Sahin, Aysegul, Joseph Song, Giorgio Topa, and Giovanni Violante (2012). "Mismatch Unemployment", New York University Working Paper.

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

Sinai, Todd M. (2012). "House Price Moments in Boom-Bust Cycles," NBER Working Paper #18059.

Table 1
Descriptive Statistics for Changes in Housing Demand and Manufacturing Decline

Change defined across following years:	2000-2006					
			Std.	Percentiles		
	N	Mean	Dev.	10th	50th	90th
Change in Housing Prices (*dP*)	275	0.475	0.388	0.054	0.408	1.207
Change in Housing Permits (*dQ*)	275	0.130	0.227	-0.154	0.113	0.343
Change in Housing Demand (*dP* + *dQ*)	275	0.605	0.515	0.000	0.510	1.395
Predicted Change in Manufacturing Employment	275	-0.015	0.008	-0.024	-0.013	-0.007

Notes: This table reports the summary statistics for the baseline sample of 275 metropolitan areas (MSAs) studied in the regressions that follow. The Housing Demand Change is constructed by adding the log change in housing prices (from FHFA house price index) to the log change in the number of housing permits for new construction (from Building Permits Survey). This procedure creates a proxy for the change in housing demand in an MSA. The Predicted Change in Manufacturing Employment is the negatve value of the Manufacturing Decline variable used in main regressions. It is constructed using the 2000 Census, the 2005-2007 American Community Survey, and the 2011-2013 American Community Survey following the procedure in Bartik (1991) and described in more detail in the main text. All of the reported sample statistics are computed using the 2000 population of prime-aged men and women in the MSA (from Census) as weights, since these weights are used in the regressions that follow.

Table 2
Employment and Construction Employment Share Response to
Housing Demand Change and Manufacturing Decline

Sample:	Non-College Men	College Men	Non-College Women	College Women	All Men and Women
	(1)	(2)	(3)	(4)	(5)
Panel A: Dependent Variable is Change in Employment to Population Ratio, 2000-2006					
Housing Demand Change	0.029	0.009	0.014	0.004	0.018
	(0.008)	(0.004)	(0.002)	(0.003)	(0.003)
	[0.000]	[0.013]	[0.000]	[0.271]	[0.000]
Manufacturing Decline	-0.747	-0.392	-0.769	-0.373	-0.657
	(0.270)	(0.125)	(0.165)	(0.149)	(0.155)
	[0.008]	[0.003]	[0.000]	[0.016]	[0.000]
Standardized (1σ) effects:					
Housing demand change	0.017	0.005	0.008	0.002	0.010
Manufacturing decline	-0.008	-0.004	-0.008	-0.004	-0.007
R^2	0.72	0.28	0.67	0.13	0.78
Panel B: Dependent Variable is Change in Share Employed in Construction and FIRE, 2000-2006					
Housing Demand Change	0.020	0.007	0.004	0.004	0.010
	(0.005)	(0.003)	(0.001)	(0.003)	(0.002)
	[0.000]	[0.028]	[0.000]	[0.140]	[0.000]
Manufacturing Decline	-0.438	-0.244	-0.153	0.128	-0.224
	(0.266)	(0.143)	(0.098)	(0.119)	(0.141)
	[0.107]	[0.095]	[0.125]	[0.287]	[0.119]
Standardized (1σ) effects:					
Housing demand change	0.011	0.004	0.002	0.002	0.006
Manufacturing decline	-0.005	-0.003	-0.002	0.001	-0.002
R^2	0.45	0.12	0.19	0.07	0.45
N	275	275	275	275	275
Include baseline controls	y	y	y	y	y

Notes: This table reports results of estimating equations (5) and (6) by OLS for various demographic groups. A 0.1 unit increase in the Housing Demand Change represents a 10 log point increase in housing demand, while a 0.01 unit decrease in Manufacturing Decline variable corresponds to a 1 percentage point decrease in predicted share of population employed in manufacturing. The baseline controls include the initial (year 2000) values of the share of employed workers with a college degree, the share of women in the labor force, and the log population in the MSA. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each state, are in parentheses and p-values are in brackets.

Table 3
Employment Response to Housing Demand Change and Manufacturing Decline, by Age Group and Immigration Status

Dependent Variable is Change in Employment to Population Ratio, 2000-2006						
Restriction:	Age 21-35		Age 36-55		Drop Immigrants	
Sample:	Non-College Men	All Men and Women	Non-College Men	All Men and Women	Non-College Men	All Men and Women
	(1)	(2)	(3)	(4)	(3)	(4)
Housing Demand Change	0.031	0.021	0.027	0.016	0.018	0.011
	(0.010)	(0.003)	(0.006)	(0.004)	(0.005)	(0.003)
	[0.004]	[0.000]	[0.000]	[0.000]	[0.001]	[0.002]
Manufacturing Decline	-0.436	-0.433	-0.883	-0.774	-0.805	-0.743
	(0.215)	(0.138)	(0.226)	(0.169)	(0.135)	(0.103)
	[0.048]	[0.003]	[0.000]	[0.000]	[0.000]	[0.000]
Standardized (1σ) effects:						
Housing demand change	0.017	0.012	0.015	0.009	0.010	0.006
Manufacturing decline	-0.009	-0.007	-0.011	-0.006	-0.013	-0.009
R^2	0.61	0.71	0.71	0.74	0.63	0.70
Include baseline controls	y	y	y	y	y	y

Notes: N=275 in all columns. This table reports OLS estimates analogous to columns (1) and (5) in Table 2 for alternative samples of either non-college men or all prime-aged men and women, using the same set of baseline controls. See Table 2 for more details. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each state, are in parentheses and p-values are in brackets.

Table 4
Wage Response to Housing Demand Change and Manufacturing Decline

Sample:	Non-College Men	College Men	Non-College Women	College Women	All Men and Women
	(1)	(2)	(3)	(4)	(5)
	Dependent Variable is Change in Average Log Wage, 2000-2006				
Housing Demand Change	0.035	0.019	0.018	0.007	0.024
	(0.005)	(0.010)	(0.005)	(0.007)	(0.005)
	[0.000]	[0.066]	[0.001]	[0.305]	[0.000]
Manufacturing Decline	-2.009	-0.660	-1.064	-0.637	-1.209
	(0.369)	(0.433)	(0.249)	(0.304)	(0.286)
	[0.000]	[0.135]	[0.000]	[0.042]	[0.000]
Standardized (1σ) effects:					
Housing demand change	0.020	0.011	0.010	0.004	0.014
Manufacturing decline	-0.021	-0.007	-0.011	-0.007	-0.012
R^2	0.43	0.14	0.45	0.11	0.41
N	275	275	275	275	275
Include baseline controls	y	y	y	y	y

Notes: This table reports results of estimating equations (5) and (6) by OLS for various demographic groups. A 0.1 unit increase in the Predicted Housing Demand Change represents a 10 percent increase in housing demand, while a 0.1 unit change in Predicted Manufacturing Decline variable corresponds to a 10 percentage point change in predicted share of population employed in manufacturing. The baseline controls include the initial (year 2000) values of the share of employed workers with a college degree, the share of women in the labor force, and the log population in the MSA. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each state, are in parentheses and p-values are in brackets.

Table 5
Sectoral Wage Responses to Housing Demand Change and Manufacturing Decline

Sample:	Average Log Wage, All Sectors	Sectoral Wage Responses: Manufacturing Only	Construction and FIRE Only	All Other Sectors, Excl. Manuf. and Construction/FIRE
	(1)	(2)	(3)	(4)
Panel A: Sectoral Wage Responses for Non-College Men, 2000-2006				
Housing Demand Change	0.035	0.032	0.033	0.031
	(0.005)	(0.006)	(0.005)	(0.006)
	[0.000]	[0.000]	[0.000]	[0.000]
Manufacturing Decline	-2.009	-1.849	-2.013	-1.796
	(0.369)	(0.403)	(0.373)	(0.399)
	[0.000]	[0.000]	[0.000]	[0.000]
Standardized (1σ) effects:				
Housing demand change	0.020	0.018	0.019	0.018
Manufacturing decline	-0.021	-0.019	-0.021	-0.019
R^2	0.43	0.42	0.40	0.38
Panel B: Sectoral Wage Responses for All Men and Women, 2000-2006				
Housing Demand Change	0.024	0.022	0.021	0.020
	(0.005)	(0.005)	(0.005)	(0.005)
	[0.000]	[0.000]	[0.000]	[0.001]
Manufacturing Decline	-1.209	-1.151	-1.104	-1.025
	(0.286)	(0.304)	(0.306)	(0.322)
	[0.000]	[0.000]	[0.001]	[0.003]
Standardized (1σ) effects:				
Housing demand change	0.014	0.013	0.012	0.011
Manufacturing decline	-0.012	-0.012	-0.011	-0.011
R^2	0.41	0.40	0.39	0.38
N	275	275	275	275
Include baseline controls	y	y	y	y

Notes: This table reports results of estimating equations (5) and (6) by OLS for various demographic groups. A 0.1 unit increase in the Predicted Housing Demand Change represents a 10 percent increase in housing demand, while a 0.1 unit change in Predicted Manufacturing Decline variable corresponds to a 10 percentage point change in predicted share of population employed in manufacturing. All Average Wage dependent variables are calculated as the average of the log of the wage. The baseline controls include the initial (year 2000) values of the share of employed workers with a college degree, the share of women in the labor force, and the log population in the MSA. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each state, are in parentheses and p-values are

Table 6
Employment Response to Housing Demand Change and Manufacturing Decline: Longer Run Results

Dependent Variable is Change in Employment to Popuplation Ratio				
Change defined across following years:	2000-2006		2000-2012	
Sample:	Non-College Men	All Men and Women	Non-College Men	All Men and Women
	(1)	(2)	(3)	(4)
Predicted Housing Demand Change, 2000-2006	0.029	0.018	0.013	0.006
	(0.008)	(0.003)	(0.013)	(0.006)
	[0.000]	[0.000]	[0.350]	[0.318]
Predicted Manufacturing Decline, 2000-2006	-0.747	-0.657	-0.461	-0.481
	(0.270)	(0.155)	(0.415)	(0.213)
	[0.008]	[0.000]	[0.273]	[0.029]
Standardized (1σ) effects:				
Housing demand change	0.017	0.010	0.007	0.004
Manufacturing decline	-0.008	-0.007	-0.005	-0.005
R^2	0.72	0.78	0.55	0.62
Include baseline controls	y	y	y	y

Notes: N=275 in all columns. This table reports OLS estimates analogous to columns (1) and (5) in Table 2 for alternative sample periods for dependent variable (but keeping right-hand side variables the same). See Table 2 for more details. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each state, are in parentheses and p-values are in brackets.

Table 7
Decomposing Employment Responses into Non-participation and Unemployment

Dependent variable:	Change in Employment to Population Ratio, 2000-		Change in Non-participant/Population		Change in Unemployed/Population Ratio, 2000-2006	
Sample:	Non-College Men	All Men and Women	Non-College Men	All Men and Women	Non-College Men	All Men and Women
	(1)	(2)	(3)	(4)	(5)	(6)
Housing Demand Change	0.029	0.018	-0.012	-0.009	-0.018	-0.009
	(0.008)	(0.003)	(0.004)	(0.002)	(0.005)	(0.002)
	[0.000]	[0.000]	[0.003]	[0.000]	[0.001]	[0.000]
Manufacturing Decline	-0.747	-0.657	0.418	0.316	0.329	0.341
	(0.270)	(0.155)	(0.136)	(0.089)	(0.195)	(0.101)
	[0.008]	[0.000]	[0.004]	[0.001]	[0.098]	[0.002]
Standardized (1σ) effects:						
Housing demand change	0.017	0.010	-0.007	-0.005	-0.010	-0.005
Manufacturing decline	-0.008	-0.007	0.004	0.003	0.003	0.004
R^2	0.72	0.78	0.49	0.57	0.69	0.73
Include baseline controls	y	y	y	y	y	y

Notes: N=275 in all columns. This table reports OLS estimates analogous to columns (1) and (5) in Table 2 for alternative dependent variables, allowing the overall employment effect to be decomposed into a change in unemployment rate and change in labor force participation rate. See Table 2 for more details. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each state, are in parentheses and p-values are in brackets.

Table 8
Population Response to Housing Demand Change and Manufacturing Decline

Sample:	Non-College Men	College Men	Non-College Women	College Women	All Men and Women
	(1)	(2)	(3)	(4)	(5)
	Dependent Variable is Change in Population, 2000-2006				
Housing Demand Change	0.029	0.057	0.020	0.059	0.036
	(0.025)	(0.016)	(0.025)	(0.014)	(0.021)
	[0.266]	[0.001]	[0.411]	[0.000]	[0.089]
Manufacturing Decline	-1.752	-1.484	-1.845	-1.374	-1.730
	(0.744)	(0.833)	(0.737)	(0.846)	(0.718)
	[0.023]	[0.082]	[0.016]	[0.112]	[0.020]
Standardized (1σ) effects:					
Housing demand change	0.016	0.032	0.012	0.034	0.020
Manufacturing decline	-0.018	-0.015	-0.019	-0.014	-0.018
R^2	0.11	0.15	0.20	0.22	0.13
N	275	275	275	275	275
Include baseline controls	y	y	y	y	y

Notes: This table reports results of estimating equations (5) and (6) by OLS for various demographic groups. A 0.1 unit increase in the Predicted Housing Demand Change represents a 10 percent increase in housing demand, while a 0.1 unit change in Predicted Manufacturing Decline variable corresponds to a 10 percentage point change in predicted share of population employed in manufacturing. The baseline controls include the initial (year 2000) values of the share of employed workers with a college degree, the share of women in the labor force, and the log population in the MSA. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each state, are in parentheses and p-values are in brackets.

Table 9
Employment and Wage Response to Housing Demand Change and Manufacturing Decline: Instrumental Variable Estimates using Magnitude of Structral Break in House Prices

Sample:	Non-College Men	College Men	Non-College Women	College Women	All Men and Women
	(1)	(2)	(3)	(4)	(5)
	Dependent Variable is Change in Employment to Population Ratio, 2000-2006				
Housing Demand Change	0.029	0.008	0.005	-0.002	0.013
	(0.011)	(0.003)	(0.006)	(0.005)	(0.005)
	[0.005]	[0.003]	[0.452]	[0.687]	[0.005]
Manufacturing Decline	-0.499	-0.323	-0.728	-0.389	-0.549
	(0.218)	(0.109)	(0.137)	(0.149)	(0.119)
	[0.027]	[0.005]	[0.000]	[0.012]	[0.000]
Standardized (1σ) effects:					
Housing demand change	0.017	0.005	0.003	-0.001	0.007
Manufacturing decline	-0.005	-0.003	-0.008	-0.004	-0.006
	Dependent Variable is Change in Average Log Wage, 2000-2006				
Housing Demand Change	0.048	0.033	0.036	0.038	0.040
	(0.012)	(0.015)	(0.008)	(0.015)	(0.011)
	[0.000]	[0.031]	[0.000]	[0.010]	[0.000]
Manufacturing Decline	-1.602	-0.382	-0.761	-0.317	-0.870
	(0.322)	(0.396)	(0.264)	(0.381)	(0.283)
	[0.000]	[0.340]	[0.006]	[0.409]	[0.004]
Standardized (1σ) effects:					
Housing demand change	0.027	0.019	0.020	0.022	0.023
Manufacturing decline	-0.017	-0.004	-0.008	-0.003	-0.009
First stage F-statistic	28.40	28.40	28.40	28.40	28.40
N	275	275	275	275	275
Include baseline controls	y	y	y	y	y

Notes: N=275 in all columns. This table reports IV estimates analogous to Table 2 using the Instrumental Variable in Table 7. The standardized effects rescale the coefficient by a one standard deviation change using the cross-MSA standard deviation. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each state, are in parentheses and p-values are in brackets.

Table 10
Accounting for the Effect of Housing Demand Change and Manufacturing Decline on National Trends in Non-Employment

	Actual Change (1)	Predicted Change due to Housing Demand Change (2)	Predicted Change due to Manufacturing Decline (3)	Residual Change, (1) - (2) - (3) (4)	Share of Actual Change Explained by Manufacturing + Housing (5)
Panel A: Accounting for National Non-employment Trends for All Men and Women					
2000-2006	0.019	-0.011	0.021	0.009	52.9%
2000-2012	0.073	0.000	0.030	0.043	40.5%
Panel B: Accounting for National Non-employment Trends for Non-College Men					
2000-2006	0.022	-0.018	0.036	0.004	82.6%
2000-2012	0.108	0.000	0.053	0.055	49.1%
Panel C: Accounting for National Nonemployment Trends for Non-College Women					
2000-2006	0.027	-0.009	0.020	0.016	42.4%
2000-2012	0.081	0.000	0.031	0.050	38.0%

Notes: This table reports counterfactual estimates of predicted changes in aggregate non-employment for different demographic groups. The coefficient estimates from Table 2 and Table 5 are used to compute the predicted values. Actual changes in non-employment, housing prices, and manufacturing employment are taken from the CPS.

Figure 1: Trends in Employment in Manufacturing and Construction for Non-College Men and Non-College Women, 1977-2011

Panel A: Non-College Men

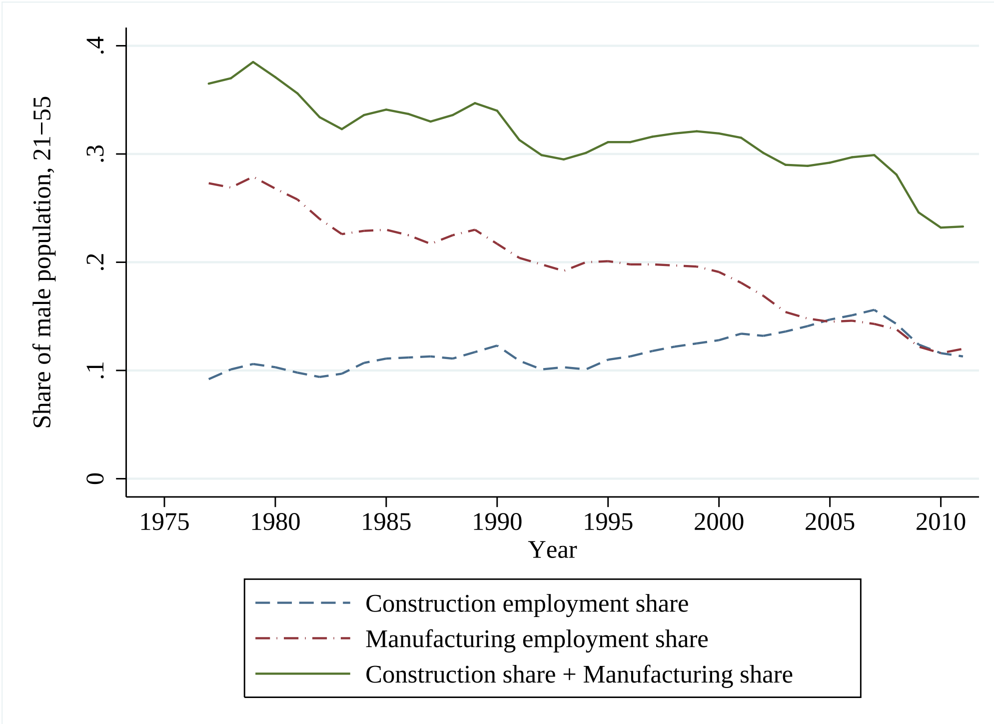

Panel A: Non-College Women

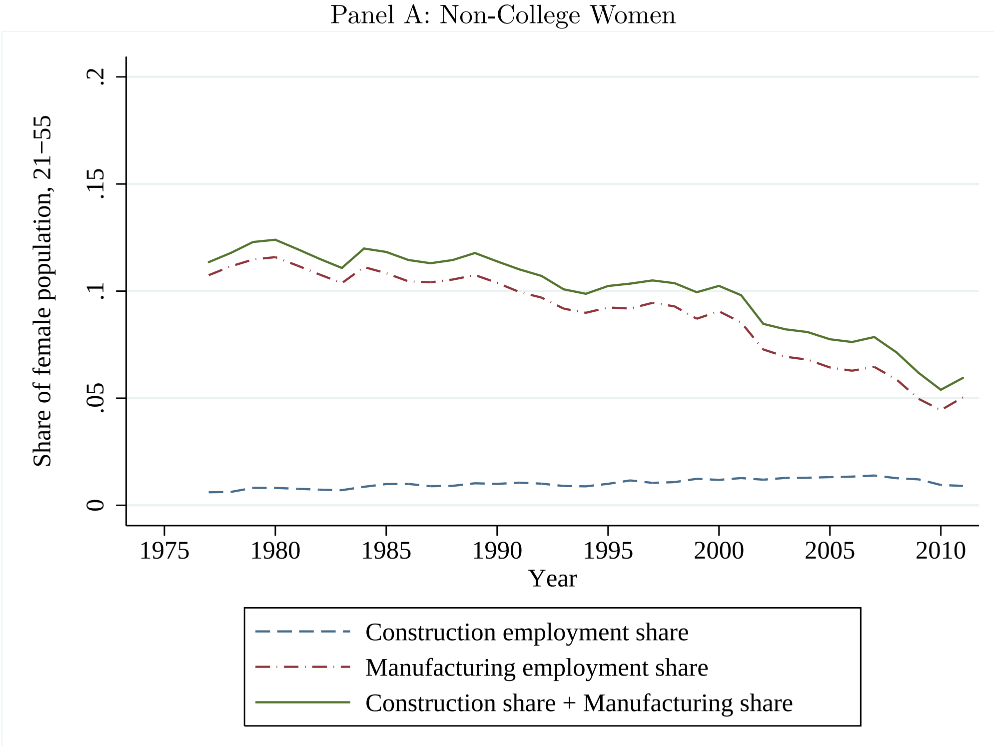

Notes: These figures use data from the March CPS. The sample includes all men and women without a four-year college degree, age 21-55. All employment shares are calculated using individual-level survey weights.

Figure 2: Graphical Solutions of Sectoral Choice Model

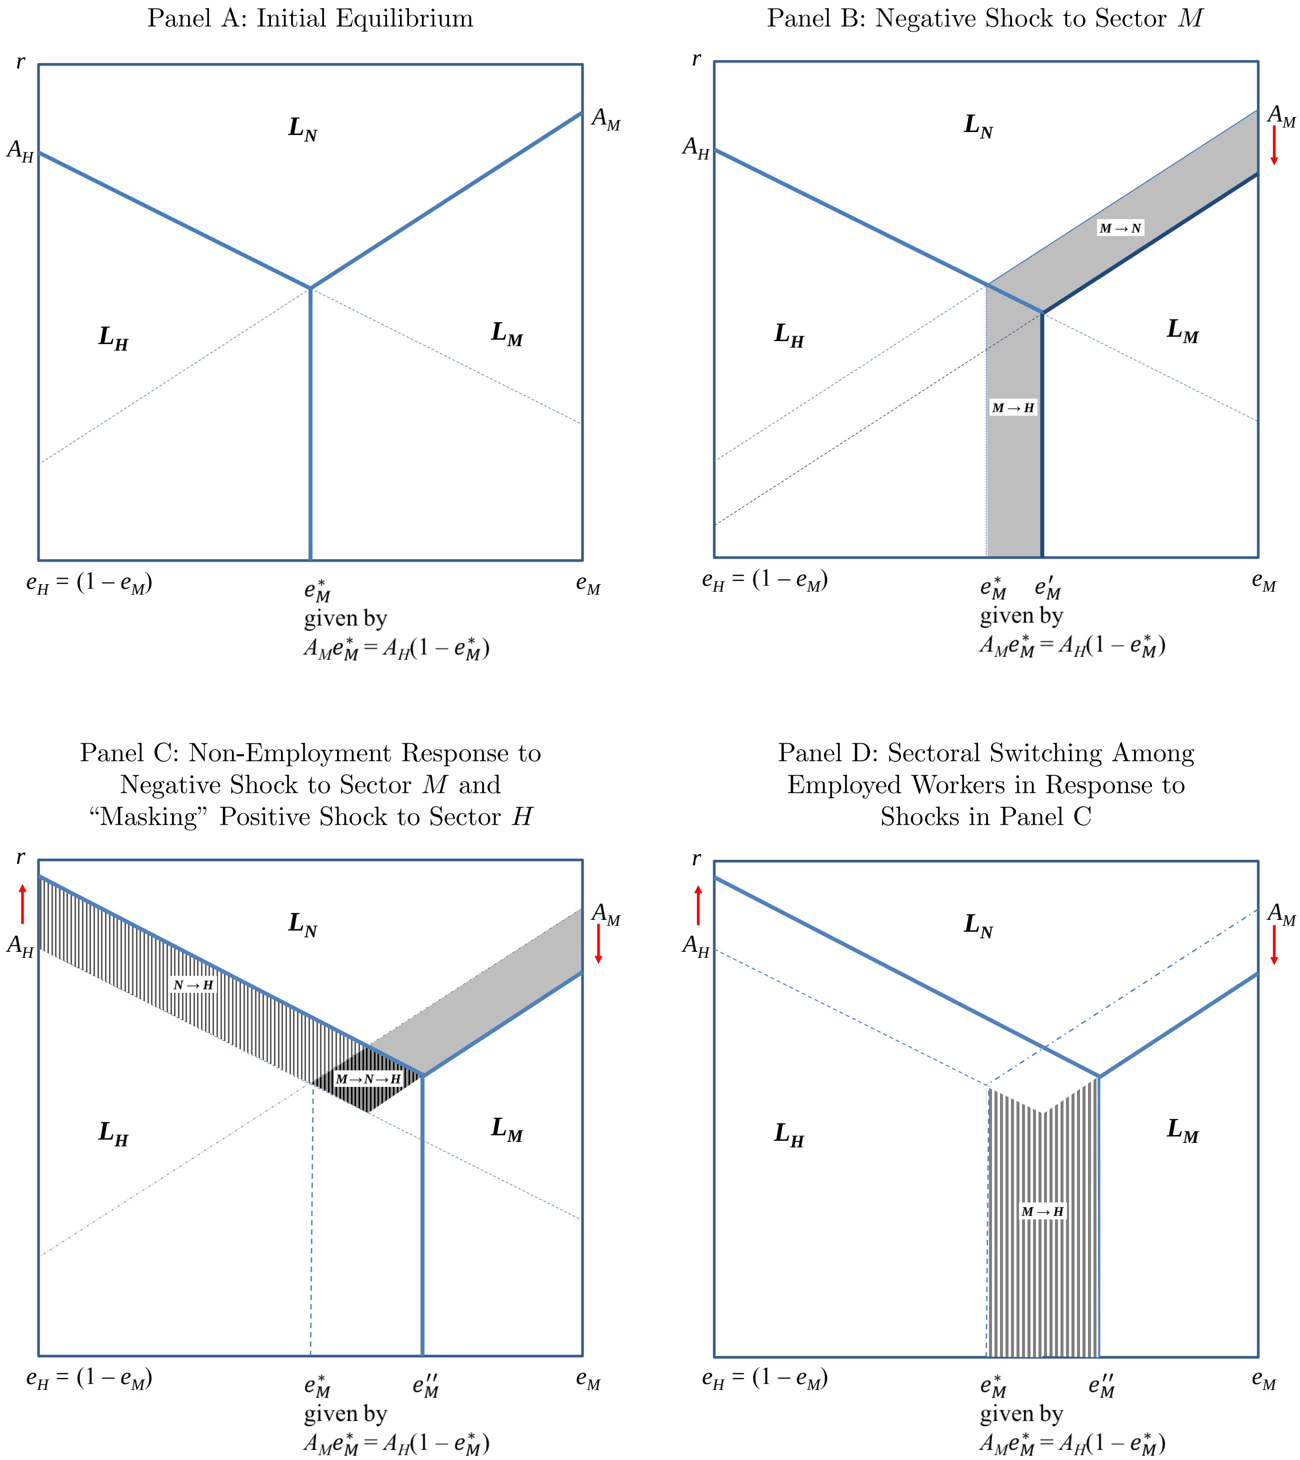

<u>Notes:</u> These figures show the graphic solutions of the model. In Panel A, we show the initial equilibrium, which shows the combination of e_M and r) parameters determine how workers self-select into sectors (or into non-employment, N). Panel B shows how the equilibrium responds to a negative shock to sector M; workers leave sector M for either sector H or enter non-employment (sector N), with the relative importance of these two channels depending on the mass of workers along each margin. Lastly, Panels C and D show how the equilibrium responds a "masking" positive shock to sector H. In this case, some workers who would have entered non-employment in Panel B instead remain employed and enter sector H (center diamond in Panel C).

Figure 3: Predicted Manufacturing Decline and Manufacturing Employment, 2000-2006

Notes: This figure reports the correlation across cities between the predicted change in manufacturing employment and changes in manufacturing employment between 2000 and 2006. The manufacturing decline variable is constructed following Bartik (1991); see main text for details. The change in manufacturing employment is defined as the change in the share of the total population of men and women age 21-55 employed in manufacturing. Each circle represents a metropolitan area, and the size of the circle is proportional to the prime-age population in the metropolitan area as computed in the 2000 Census.

Figure 4: Predicted Housing Demand Change and Construction Employment, 2000-2006

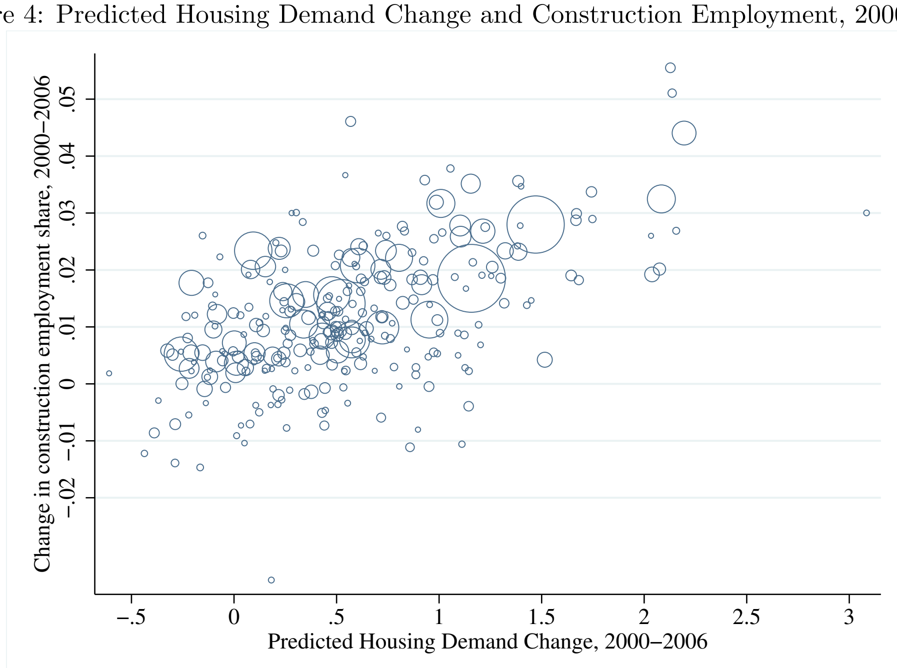

Notes: This figure reports correlation across cities between the 2000-2006 change in share of population employed in construction and the change in housing demand over the same time period. Each circle represents a metropolitan area, and the size of the circle is proportional to the number of prime-age men and women in the metropolitan area as computed in the 2000 Census.

Figure 5: Manufacturing Decline and Non-Employment of Non-College Men Across MSAs with Different Housing Demand Shocks, 2000-2006

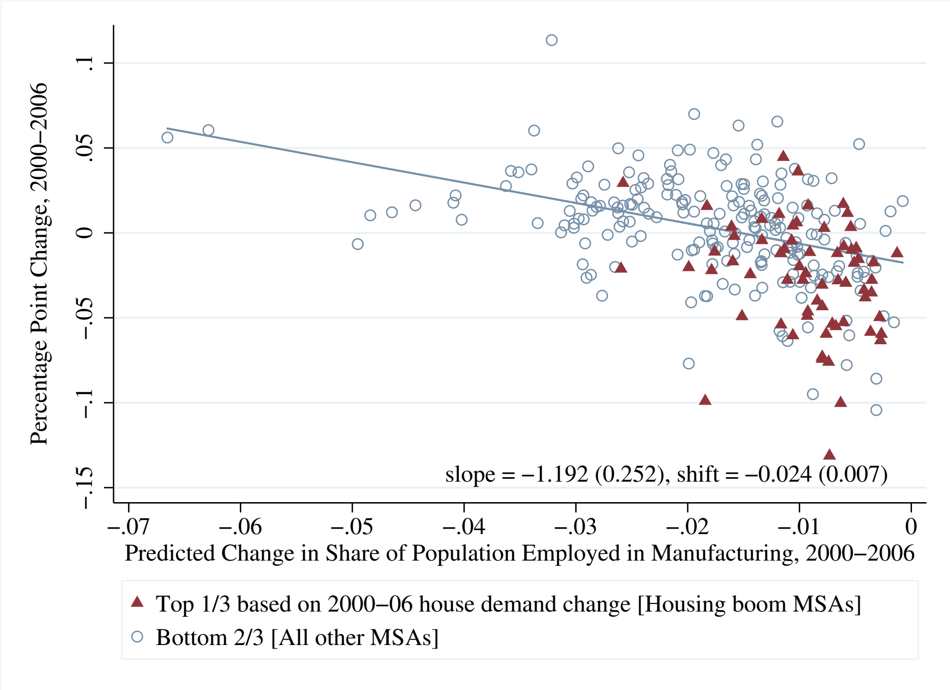

Figure 6: Manufacturing Decline and Construction Employment of Non-College Men Across MSAs with Different Housing Demand Shocks, 2000-2006

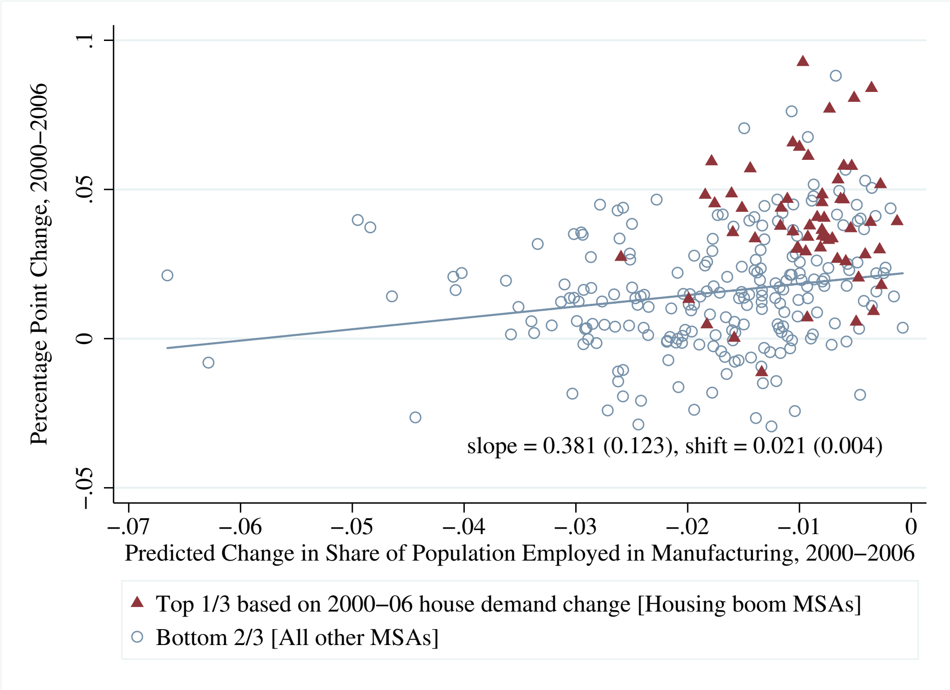

Notes: These figures report the correlation across cities between the predicted change in manufacturing emplyoment and the change in the non-employment rate and construction employment share of non-college men (age 21-55) between 2000-2006. The sample is divided based on the Housing Demand Change in the metropolitan area between 2000 and 2006. The bottom two-thirds of the metropolitan areas based on this residualized measure are shown in light-colored circles; the top one-third are shown in dark-colored triangles. The solid line represents the OLS regression line that is computed based on the bottom two-thirds sample. The slope of this line is reported along with the average difference between the regression line and the top one-third "housing boom MSA" sample.

Figure 7: Manufacturing Decline and Average Wages of Non-College Men Across MSAs with Different Housing Demand Shocks, 2000-2006

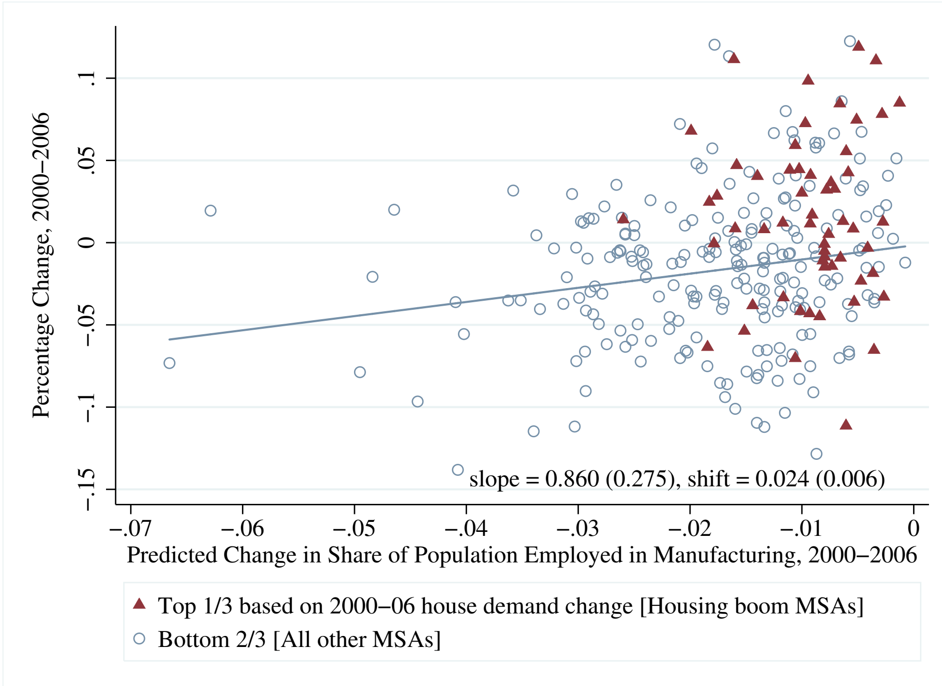

Figure 8: Manufacturing Decline and Manufacturing Employment of Non-College Men Across MSAs with Different Housing Demand Shocks, 2000-2012

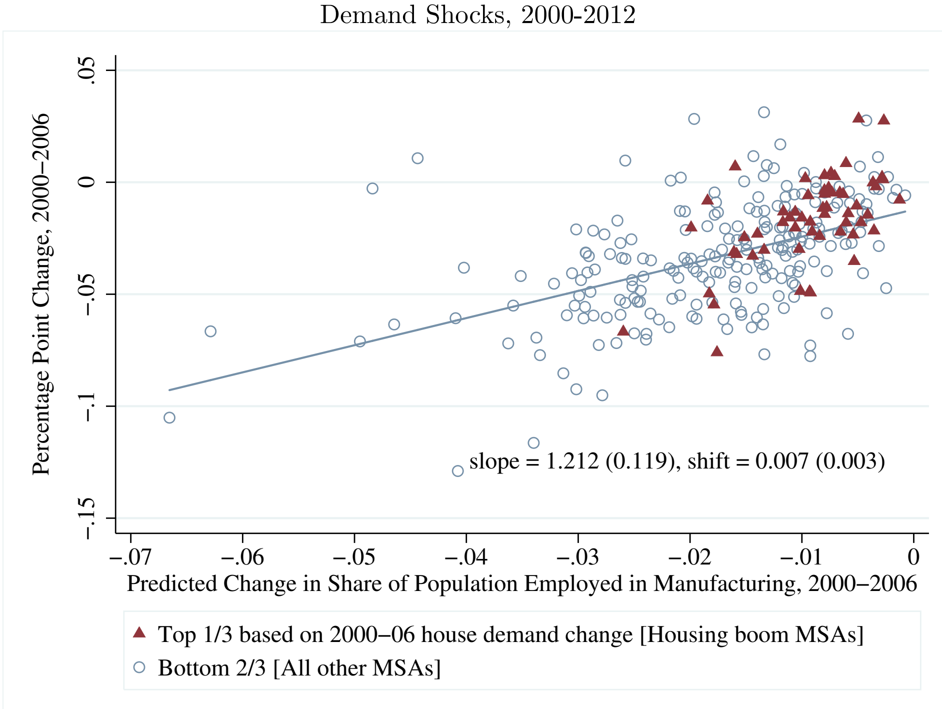

Notes: These figures report the correlation across cities between the predicted change in manufacturing emplyoment and the change in the average wage and manufacturing employment share of non-college men (age 21-55) between 2000-2006. The sample is divided based on the Housing Demand Change in the metropolitan area between 2000 and 2006. The bottom two-thirds of the metropolitan areas based on this residualized measure are shown in light-colored circles; the top one-third are shown in dark-colored triangles. The solid line represents the OLS regression line that is computed based on the bottom two-thirds sample. The slope of this line is reported along with the average difference between the regression line and the top one-third "housing boom MSA" sample.

Figure 9: Manufacturing Decline and Non-Employment of Non-College Men Across MSAs with Different Housing Demand Shocks, 2000-2006

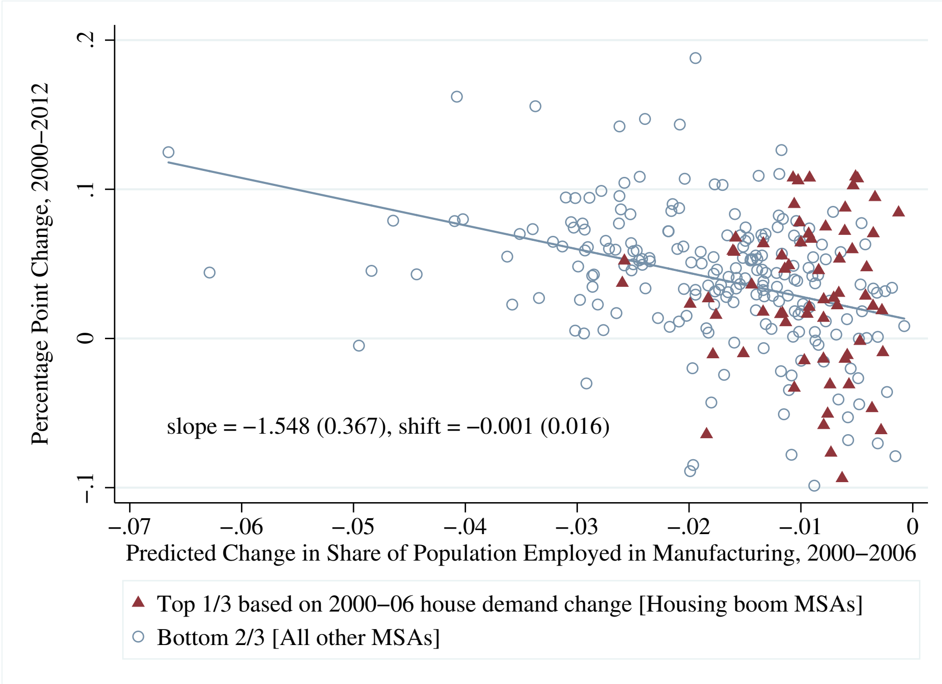

Notes: These figures report the correlation across cities between the predicted change in manufacturing emplyoment and the change in the non-employment rate of non-college men (age 21-55) between 2000-2012 (i.e., this figure is sample as Figure 5 exception that non-employment rate change is extended beyond 2000-2006). The sample is divided based on the Housing Demand Change in the metropolitan area between 2000 and 2006. The bottom two-thirds of the metropolitan areas based on this residualized measure are shown in light-colored circles; the top one-third are shown in dark-colored triangles. The solid line represents the OLS regression line that is computed based on the bottom two-thirds sample. The slope of this line is reported along with the average difference between the regression line and the top one-third "housing boom MSA" sample.