

Discussant comments on “The Transformation of Manufacturing and the Decline in U.S. Employment” by Kerwin Charles, Erik Hurst, and Mariel Schwartz

Comments by Valerie A. Ramey, June 2018

1. Introduction

This fine paper by Charles, Hurst, and Schwartz (CHS) investigates the link between the post-2000 decline in manufacturing employment and the decline of the employment rate, and also analyzes the supporting roles played by transfer payments, geographic mobility, and opioid use. The paper is a particularly useful synthesis because it brings together threads from a number of other papers, including the authors’ own work, and it explores some competing explanations in a standardized empirical framework. The paper is a wonderful read because it tells a clear story based on an impressive marshalling of evidence.

The paper contains numerous findings that shed light on a variety of changes that occurred around the same time. The span of the analysis from aggregate to commuting zone level is particularly enlightening. Here are highlights of a just a few of the many interesting results:

- A 10 percentage point decline in the local manufacturing employment share of population reduced local employment rate by 3.7 percentage points for prime age men and 2.7 for prime age women.
- Manufacturing decline in local markets had similar effects on the prime age employment rate whether the decline was due to the China shock or other shocks.
- The correlation between manufacturing share changes and employment rate changes holds only in the post-2000 period, not the 1980-1990 period. Decreased geographical mobility appears to play a key role in this changing relationship.

In my discussion, I will begin by putting the changes post-2000 in historical context. I will then consider the validity of the shift-share instrument in light of the long-term effects of

initial conditions in manufacturing as well as in light of some new work on the econometrics of Bartik instruments. I will conclude by questioning whether one can draw aggregate inferences from cross-sectional data.

2. How New are these Trends?

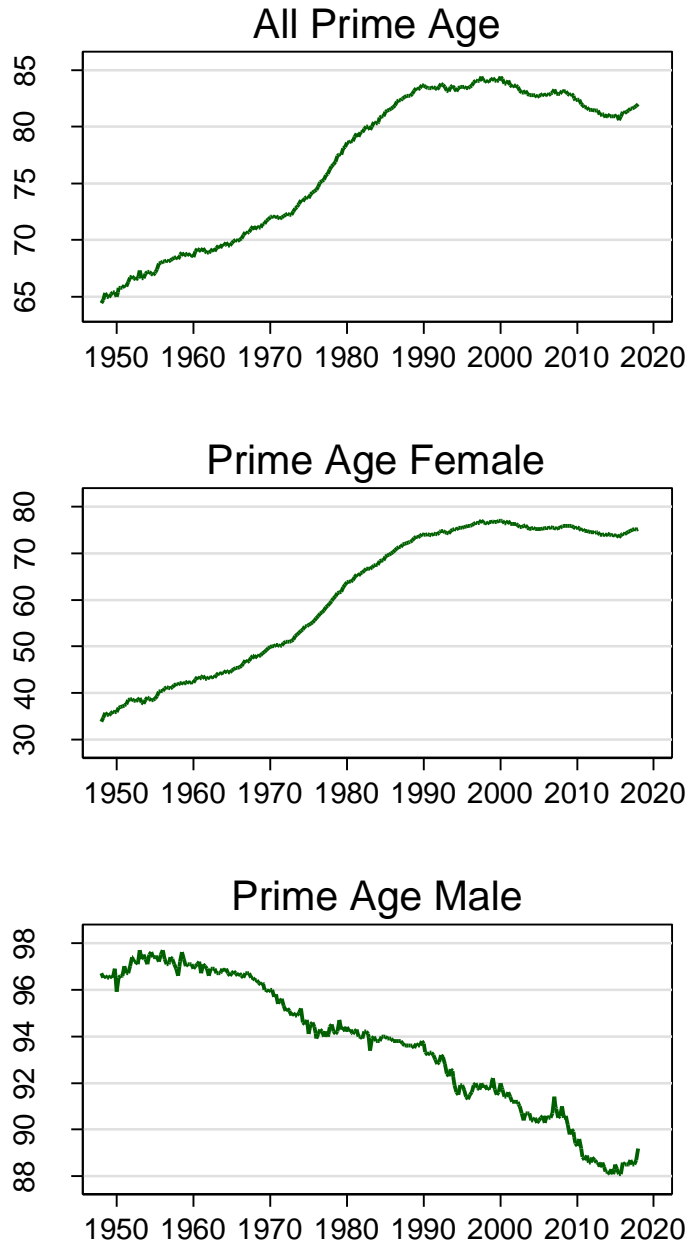
The authors focus on what they call “two profound changes in the U.S. economy” since 2000: the secular declines in manufacturing employment and in the employment rates of prime age individuals. In this section I will suggest that the movements post-2000 are largely a continuation of past trends, though I will highlight one key exception.

Consider Figure 1 below, which shows the labor force participation rates for prime age individuals.¹ I use labor force participation rates instead of employment-population rates in order to remove the cyclicity of unemployment so that I can focus on the secular trends.

The “All Prime Age” graph in Figure 1 shows that the rise in the labor force participation rate slowed during the 1990s and then declined somewhat since 2000. It is this decline since 2000 that is the focus of CHS’s investigation. The graphs separated by sex show that, contrary to the emphasis on men and manufacturing in much of the paper’s discussion, the aggregate trend appears to be driven by the behavior of prime age female labor force trends. This feature is particularly striking once one takes account of the different scale of the graphs: the changes in female labor force participation rates swamp those of the males. Female labor force participation rates rose 43 percentage points from 1948 to the peak in 2000, and then declined 2 percentage points from 2000 to 2018. Prime age male labor force participation rates declined gradually, with the average annual rate of decline after 2000 similar to the rate from the late 1960s to 2000. Thus, what changed after 2000 was not the behavior of the prime age male labor force participation rate but rather the cessation of the decades-long rise in female labor force participation rates. This fact calls into question the role of manufacturing in explaining the aggregate participation rate since the authors’ own evidence suggests that manufacturing plays a much smaller role for female employment.

¹ Charles et al. define prime age as ages 21-55. I define prime age as 25-54 because it corresponds to the age category aggregate published by the BLS.


Figure 1. Trends in Labor Force Participation Rates



The second change the authors highlight is the decline in the share of manufacturing employment. I would argue that the decline in the *share* is part of a long-term trend, but the decline in the *level* is indeed unprecedented. The left panel of Figure 2 shows manufacturing employment as a percent of the working-age population, ages 20-64. The average annual loss rate in the share was 0.25 percentage points per year from 1970 to 2000 and 0.22 percentage

points per year from 2000 to 2018. Thus, the employment share in manufacturing is following a long-standing trend. On the other hand, as the graph in the right panel of Figure 2 shows, the decline in the *level* of employment in manufacturing since 2000 is unprecedented for the post-WWII period. Thus, to make the link to recent declines in the employment rate, it seems that explanations should focus on absolute levels of manufacturing employment rather than on shares. However, all of CHS’s key evidence, such as Figure 14 at the aggregate level and the estimates from commuting zone regressions in Table 2, is in terms of the manufacturing employment share relative to population, not changes in the levels. Thus, the estimates cannot inform us about the role of the actual variable that changed post-2000 – the level of manufacturing employment. Such an investigation might shed further light on the role of manufacturing.

Figure 2. U.S. Manufacturing Employment % of Population Ages 20-64



3. The Long Decline in Manufacturing and Why Less-Educated Workers were Adversely Affected

As my graphs in the previous section showed, the decline manufacturing employment relative to the working age population began just before 1970. A question that the current

paper's focus on manufacturing brings up is "what makes manufacturing special?" The U.S. economy has experienced other dramatic sectoral shifts in its history, such as the movement of workers out of agriculture. Why is the decline in manufacturing seen as such a negative shock to workers, particularly to the less-educated?

George Borjas and I offered one possible answer to this question in a series of papers written in the 1990s (Borjas and Ramey (1994a,b, 1995). In that work, which explored the role of the decline in durable goods manufacturing employment in the rise in wage inequality, we presented time series evidence, evidence from a panel of cities, as well as a theoretical model that told the following story. At the end of WWII, U.S. manufacturing was dominant in the world, particularly in the durable goods industries such as automobiles and steel. U.S. companies enjoyed market power at home and abroad and earned significant rents as a result. American unions, which were strengthened by the New Deal legislation during the 1930s, compelled the companies to share their rents with their workers in the form of higher wages. Because the manufacturing workers who benefitted from these high wage premia were less educated, the education wage premium was lower in the 1950s and 1960s.

But circumstances began to change in the 1970s. Europe and Japan completed their recovery from WWII and their companies began to compete with U.S. firms. For example, Japan's success in the automobile industry beginning in the 1970s competed away many of the rents previously earned by U.S. automobile companies. Later in the 1990s and 2000s, China became a significant competitive force faced by U.S. companies and workers. The net effect was fewer rents earned by U.S. companies to share with the workers, and increased direct worker competition against less-educated U.S. workers. The result, we argued, was a rise in the education premium since fewer less-educated workers could find work in the high wage manufacturing industries.

In Borjas and Ramey (2000), we went on to investigate the long-term consequences for employment and wages in manufacturing industries. We found very high correlations of industry wage premia over decades, suggesting institutional barriers to lowering long-standing wage premia in the face of competition. However, we found that firms responded in other ways to the rise in competition. In particular, we found that industries that paid higher wage premia in 1960 experienced significantly slower employment growth over the subsequent decades relative to industries with lower wage premia. Further investigations found that the high initial wage

industries responded by replacing workers with capital and ended up experiencing faster growth in capital-labor ratios and faster labor productivity growth. In follow-up work, Shim and Yang (2018) showed that this process has continued up to the present, with the high initial wage premia industries investing much more in ICT (information and communication technology) and replacing their routine workers. Thus, it appears that the decline in manufacturing employment share is tied to its initial high rent/high wage status after WWII.

In sum, our story suggests that researchers and policymakers are correctly focused on the negative consequences of the decline in manufacturing for less educated workers. However, that decline was driven by forces that one should not expect to reverse. It also highlights the very long-term effects of the initial conditions. These effects might affect the exogeneity of CHS's instrument, as I will argue in the next section.

4. Comments on the Cross-Sectional Commuting Zone Estimates

An important part of the CHS paper is the extension of the previous research using variation at the level of the commuting zone to estimate causal effects of the decline in the manufacturing share. They use a Bartik-type shift-share instrument to estimate the differential decline in the employment rate of prime age individuals across commuting zones caused by differential declines in manufacturing employment shares. Their estimates (shown in Table 2 of their paper) indicate that a commuting zone that experienced a demand-driven one-percentage point larger decline in manufacturing than another commuting zone suffered a 0.37 percentage point larger decline in the employment rate of prime age men. The estimated effect for prime age women was 0.27. In both cases, the effects were larger for less educated individuals.

The commuting zone estimates are quite interesting and suggest important effects on individual worker outcomes. As CHS discuss, manufacturing tends to be highly spatially concentrated relative to other sectors. This implies that the decline in manufacturing might have larger effects on local labor markets than on aggregate ones. Decades ago, Blanchard and Katz (1992) showed that the departure of industries from local areas leads to a rise in the unemployment rate that falls only slowly afterward, and does so, not because other firms move in, but because the population eventually moves out. Their results raise a puzzle: why don't other firms move into the local area to take advantage of the weak labor market? This question deserves far more attention from the literature. The Blanchard and Katz finding, along with the

evidence that geographic mobility has declined, can perhaps explain why the long-run declines in manufacturing are now having more serious consequences for individuals. If firms do not move in to take the place of those that left, and if workers will no longer move out, then the result is local areas with large shares of prime age individuals who are no longer economically productive. As the CHS paper suggest, these areas may be more prone to ills such as opioid abuse.

My final comments consist of two observations about the cross-sectional estimates, one about the instrument and the other about the link between cross-sectional estimates and aggregate implications.

CHS use a shift-share Bartik instrument as a demand instrument for the local manufacturing share. This type of instrument has been used in a number of recent papers. They construct their instrument as the sum over all of the local industries of the product of two terms: (i) the share of prime age individuals in commuting zone k working in manufacturing industry j in the year 2000; and (ii) the change in aggregate employment (less commuting zone k) in manufacturing industry j from 2000 to 2016. Essentially, the first term measures the exposure of the commuting zone to the aggregate shock measured in the second term.

A recent paper by Goldsmith-Pinkham, Sorkin and Swift (2018) analyzes the econometrics of Bartik instruments and reveals a number of previously unknown or underappreciated aspects. They summarize one of their key findings as: “Bartik is formally equivalent to a GMM estimator with the industry shares as instruments. Hence, we argue that the identifying assumption is best stated in terms of the industry shares – the national growth rates are simply a weight matrix.” (p. 25) This insight implies that CHS’s identifying assumption requires that the initial distribution of manufacturing industries across commuting zones in 2000 have no effect on the change in the employment-population ratio from 2000-2016 except through the demand-channel effect on the change in the manufacturing share in the commuting zone. It is at this point we can link back to my discussion in Section 3 of the unique history of the manufacturing sector in the U.S. and the long-term effects of the initial high rent/high wage premia in U.S. manufacturing in the couple of decades after WWII. It is very possible that those long term features also affect the employment-population ratio in other ways, and perhaps might be able to explain the reluctance of other industries to move into local areas.

Goldsmith-Pinkham, Sorkin and Swift (2018) also analyze the Autor, Dorn, and Hanson (2013) commuting zone estimates of the effect of China, and one of the findings might be relevant to the CHS paper since the instruments share similarities. Goldsmith-Pinkham et al. show that despite the fact that Autor et al. have over 700 Bartik instruments, only a few instruments are actually driving their estimates. In particular, it turns out that the highest-weight instrument in their sample is electronic computers in the 2000s, which implies that their estimator is “comparing outcomes in locations with high and low shares of the electronic computers industry.” (Goldsmith-Pinkham et al., p. 3). If this feature is also true for the closely-related CHS shift-share instrument, then the interpretation of their estimates would change dramatically. In particular, since they are estimating only *relative* effects across regions, their results might be reinterpreted as an explanation for why the employment rate of prime age individuals rose in Silicon Valley.

This last point brings me to my second comment regarding the link between cross-sectional estimates and aggregate implications. In several places in the paper, CHS use their commuting zone estimates to make inferences about aggregates. Cognizant of the leap required, they always preference these extrapolations with the adjective “naïve.” As they recognize, one cannot identify aggregate effects from cross-sectional estimates (since the constant term nets out aggregate effects). The additional identifying assumptions required to translate cross-sectional estimates to the aggregate are very strong, a crucial point overlooked by some of the recent literature. Muendler’s (2017) recent WTO report discusses the pitfalls and potentially faulty inference that can result when naively extrapolating cross-sectional estimates. These issues should be kept in mind when assessing how much the fascinating cross-commuting zone estimates of CHS’s paper can tell us about the aggregate trends.

Supplemental References not included in Charles-Hurst-Schwartz References

- Borjas, George J. and Valerie A. Ramey, "Time Series Evidence on the Sources of Trends in Wage Inequality," with George J. Borjas. *American Economic Review* 84 (May 1994): 10-16.
- , "The Relationship between Wage Inequality and Trade" in *The Changing Distribution of Income in an Open U.S. Economy*, eds. J.H. Bergstrand, T.F. Cosimano, J.W. Houck and R.G. Sheehan. Amsterdam: North Holland, 1994.
- , "Foreign Competition, Market Power and Wage Inequality.," *Quarterly Journal of Economics* 110 (November 1995): 1075 - 1110.
- , "Market Responses to Interindustry Wage Differentials," NBER Working Paper 7799, July 2000.
- Muendler, Marc-Andreas, "Trade, Technology and Prosperity: An Account of Evidence from a Labor Market Perspective," World Trade Organization, Staff Working Paper ERSD-2017-15, November 2017.
- Shim, Myungkyu and Hee-Seung Yang, "Interindustry Wage Differentials, Technology Adoption, and Job Polarization," *Journal of Economic Behavior and Organization*, 146 (2018), pp. 141-160.