

The zombie robot argument lurches on

There is no evidence that automation leads to
joblessness or inequality

Report • By [Lawrence Mishel](#) and [Josh Bivens](#) • May 24, 2017

The media are full of stories about robots and automation destroying the jobs of the past and leaving us jobless in the future; call it the coming Robot Apocalypse. We are also told that automation and technology are responsible for the poor wage growth and inequality bedeviling the American working class in recent decades, and that looming automation will only accelerate and ratchet up these problems. Recent research by economists Daron Acemoglu of MIT and Pascual Restrepo of Boston University is but the latest fuel for the automation media narrative (Acemoglu and Restrepo 2017a).

What is remarkable about this media narrative is that there is a strong desire to believe it despite so little evidence to support these claims. There clearly are serious problems in the labor market that have suppressed job and wage growth for far too long; but these problems have their roots in intentional policy decisions regarding globalization, collective bargaining, labor standards, and unemployment levels, not technology.

This report highlights the paucity of the evidence behind the alleged robot apocalypse, particularly as mischaracterized in the media coverage of the 2017 Acemoglu and Restrepo (A&R) report. Yes, automation has led to job displacements in particular occupations and industries in the past, but there is no basis for claiming that automation has led—or will lead—to increased joblessness, unemployment, or wage stagnation overall. We argue that the current excessive media attention to robots and automation destroying the jobs of the past and leaving us jobless in the future is a distraction from the main issues that need to be addressed: the poor wage growth and inequality caused by policies that have shifted economic power away from low- and moderate-wage workers. It is also the case that, as Atkinson and Wu (2017) argue, our productivity growth is too low, not too high.

Rather than debating *possible* problems that are more than a decade away, policymakers need to focus on addressing the decades-long crisis of wage stagnation by creating good jobs and supporting wage growth. And as it turns out, policies to expand good jobs and increase wages are the same measures needed to ensure that workers potentially displaced by automation have good jobs to transition to. For these workers, the education and training touted as

SECTIONS

1. The new narrative on robots • 4
2. The estimated impact of robots is small, and automation broadly defined does not explain recent labor market trends • 5
3. A&R's claim that robots reduced overall, national-level employment is unconvincing because it relies on unrealistic assumptions • 12
4. Maybe there has not been a robot apocalypse yet; but what about joblessness and inequality in the future? • 15
5. Conclusion • 20

About the authors • 21

Acknowledgments • 21

Endnotes • 21

References • 22

solutions in the mainstream robot narrative will be inadequate, just as they were not adequate to help displaced manufacturing workers over the last few decades.

Key findings

In this paper we make the following points:

Acemoglu and Restrepo’s new research does not show large and negative effects on overall employment stemming from automation

- A&R’s methodology delivers high-quality *local* estimates of the impact of one sliver of automation (literally looking just at robots). But their translation of these high-quality local estimates (for “commuting zones”) into national effects relies on stylized and largely unrealistic assumptions.
- Even if one takes the unreliable simulated (not estimated) national effects as given, they are small (40,000 jobs lost each year) relative to any reasonable benchmark. For example, our analysis shows that their estimated job losses from the “China trade shock” are roughly four times as large as their estimated job losses from growing robot adoption in the 2000s.
- While A&R’s report shows that “robots” are negatively correlated with employment growth across commuting zones, it finds that all other indicators of automation (nonrobot IT investment) are *positively* correlated or neutral with regard to employment. So even if robots displace some jobs in a given commuting zone, other automation (which presumably dwarfs robot automation in the scale of investment) creates many more jobs. It is curious that coverage of the A&R report ignores this major finding, especially since it essentially repudiates what has been the conventional wisdom for decades—that automation has hurt job growth (at least for less-credentialed Americans).
- The A&R results do not prove that automation will lead to joblessness in the future or overturn previous evidence that automation *writ large* has not led to higher aggregate unemployment.

Technological change and automation have not been the main forces driving the wage stagnation and inequality besieging working-class Americans

- **There is no historical correlation between increases in automation broadly defined and wage stagnation or increasing inequality.** Automation—the implementation of new technologies as capital equipment or software replace human labor in the workplace—has been an ongoing feature of our economy for decades. It cannot explain why median wages stagnated in some periods and grew in others, or why wage inequality grew in some periods and shrank in others.
 - Indicators of automation increased rapidly in the late 1990s and the early 2000s, a period that saw the best across-the-board wage growth for American workers in a generation.
 - Indicators of automation fell during two periods of stagnant (or worse) wage growth: from 1973 to 1995 and from 2002 to the present. In these periods,

inequality grew as wage growth for the richest Americans far outpaced wage growth of everyone else.

- During the long period of shared wage growth from the late 1940s to the mid-1970s (shared because all workers' wages grew at roughly the same pace), indicators of automation also increased rapidly.
- **There is no evidence that automation-driven occupational employment “polarization” has occurred in recent years, and thus no proof it has caused recent wage inequality or wage stagnation.**
 - First, numerous studies have documented that there was no occupational employment polarization—in which employment expands in higher-wage and lower-wage employment while hollowing out in the middle—in the 2000s. Employment has primarily expanded in the lowest-wage occupations. Yet wage inequality between the top and the middle has risen rapidly since 2000.
 - Second, wage inequality overwhelmingly occurs between workers within an occupation, not between workers in different occupations. So even if occupational employment polarization had occurred, it could not explain the growth of wage stagnation or inequality.
- **There is no evidence of an upsurge in automation in the last 10 to 15 years that has affected overall joblessness. The evidence indicates automation has slowed.**

Trends in productivity, capital investment, information equipment investment, and software investment suggest that automation has decelerated in the last 10 or so years. Also, the rate of shifts in occupational employment patterns has been slower in the 2000s than in any period since 1940. Therefore, there is no empirical support for the prominent notion that automation is currently accelerating exponentially and leading to a robot apocalypse.

The fact that robots have displaced some jobs in particular industries and occupations does not mean that automation has or will lead to increased overall joblessness.

- As noted above, data showing a recent deceleration in automation suggest that there is no footprint of an automation surge that can be expected to accelerate in the near future.
- Technological change and automation absolutely can, and have, displaced particular workers in particular economic sectors. But technology and automation also create dynamics (for example, falling relative prices of goods and services produced with fewer workers) that help create jobs in other sectors. And even when automation's job-generating and job-displacing forces don't balance out, government policy can largely ensure that automation does not lead to rising overall unemployment.
 - The narrative that automation creates joblessness is inconsistent with the fact that we had substantial and ongoing automation for many decades but did not have continuously rising unemployment. And the fall in unemployment from 10 percent to below 5 percent since 2010 is inconsistent with the claim that surging automation is generating greater unemployment.

- As noted above, fluctuations in the pace of technological change have been associated with both good and bad labor market outcomes. So there is no reason to deduce that we should fear robots.

The American labor market has plenty of problems keeping it from working well for most Americans, and these are the problems that should occupy our attention.

- The problems afflicting American workers include the failure to secure genuine full employment through macroeconomic policy and the intentional policy assault on the bargaining power of low- and middle-wage workers; these are the causes of wage stagnation and rising inequality. Solving these actually existing problems should take precedence over worrying about hypothetical future effects of automation.

The new narrative on robots

Robots are everywhere in the media. In February 2017 *The New York Times Magazine* published an article titled, “Learning to Love Our Robot Co-workers” (Tingley 2017). An article in *The Washington Post* in March 2017 warned, “We’re So Unprepared for the Robot Apocalypse” (Guo 2017). And, in *The Atlantic* Derek Thompson (2015, 2016) paved the way in the summer of 2015 with “A World without Work,” followed in October 2016 with an article asking, “When Will Robots Take All the Jobs?”

The automation narrative told by these articles and other coverage is a story in which the inevitable march of technology is destroying jobs and suppressing wages and essentially making large swathes of American workers obsolete. The policy recommendations accompanying this flawed analysis range from weak-tea conventional wisdom urging workers to get better educated or retrained to wholesale utopian reconfigurations of the modern welfare state with some version of universal income grants. Given the flawed analyses behind them, neither the weak-tea education/training call nor the utopian universal basic income schemes make much sense as a response to developments in American labor markets in recent decades.

What is remarkable about the automation narrative is that *any* research on robots or technology feeds it, even if the bottom-line findings of the research do not validate any part of it. The most recent example is the new research by Acemoglu and Restrepo (2017a) on the impact of robots on employment and wages. Here is how *The New York Times* wrote about this research:

The paper adds to the evidence that automation, more than other factors like trade and offshoring that President Trump campaigned on, has been the bigger long-term threat to blue-collar jobs. The researchers said the findings—“large and robust negative effects of robots on employment and wages”—remained strong even after controlling for imports, offshoring, software that displaces jobs, worker demographics and the type of industry. (Miller 2017)

What Acemoglu and Restrepo actually find

As we explain in detail in the next section, A&R find a rather *small* not *large* impact of robots on overall employment. And they do not find that robots had a greater adverse impact on jobs than trade. A&R themselves find that U.S. trade with China displaced *three times* as many jobs as did robots. We show that in the 2000s, Chinese trade displaced *four* times as many jobs as robots.

Further, these simulated national effects do not persuasively demonstrate *any* negative impact of robots on *overall* employment that would lead us to question the established consensus that automation did not generate higher aggregate unemployment in the past. Unlike A&R's local estimates (discussed below), the findings on *overall* employment effects are not direct estimates of data but rather produced by a highly stylized model that uses quite unrealistic assumptions to derive its results. In short, the A&R findings on the national employment displacement caused by robot adoption are not reliable.

Additionally, contrary to claims that automation and technology are the primary threat to American workers' living standards, automation broadly defined could not plausibly have generated any of the wage stagnation or wage inequality experienced in the 2000s. Job polarization was not present and, in any case, is unrelated to wage trends. A&R's work, moreover, does not provide a basis to be fearful about *future* joblessness.

Finally, the media have been strangely selective in their coverage of A&R's recent research. News reports focus only on A&R's findings regarding one small sliver of potential automation: literally *robots*. News reports ignore A&R's finding that *other* capital investments have a neutral or even positive effect on employment and wages of working-class Americans. Given that the sheer scale of these nonrobot capital investments is almost surely *substantially larger* than the scale of investments in robots, it is reasonable to interpret A&R's results as showing that automation *writ large* has been positive for working-class Americans. This finding would seem to be newsworthy because it contradicts decades of claims that capital investments in information technology are the leading cause of labor market distress for these workers.

The estimated impact of *robots* is small, and *automation* broadly defined does not explain recent labor market trends

Acemoglu and Restrepo (2017a) estimate the employment and wage impacts of increased industrial robot use over the 1990–2007 period in local labor markets (which they call commuting zones, or CZs). They compare employment and wage changes between CZs in which a large share of employment is in industries with high robot use and compare employment and wage changes in those CZs with employment and wages changes in CZs with low shares of employment in high-robot-use industries. They then attempt to translate these direct estimates into an aggregate, national “robot” impact. Later in this report we

explain why the stringent assumptions one must accept to make this translation cast serious doubt on the claim that robot use reduces *overall* employment.

In this section we explain why—even if we accept A&R’s estimates of the national impact of robots—we find nothing in their report that establishes that automation broadly defined (including robots and nonrobot automation such as information technology) explains recent trends. There are two main reasons why this is so. First, their estimates of the national impact of robots alone are small compared with reasonable benchmarks. And second, their estimates do nothing to overturn previous research indicating no apparent links between the pace of automation and negative labor market outcomes, such as wage stagnation, wage inequality, and sluggish employment growth. The fact that broadly defined automation has probably decelerated in the last 10 to 15 years calls into question the narrative that robots or automation are having an accelerating (exponentially, some claim) impact.

The negative aggregate employment impact that A&R find is small

The New York Times and other outlets have described A&R’s national estimates of job loss due to robots as “large.” But A&R’s estimates of job loss due to robots are actually quite modest when scaled against reasonable benchmarks. For example, at the *upper* end of the range they find that 670,000 jobs were lost due to robots (Acemoglu and Restrepo 2017a, 36). Since these jobs were lost over a 17-year period, 1990–2007, this translates to about 40,000 jobs lost each year. (Note that 40,000 jobs equates to about 2 percent of jobs gained on net in recent years.) Further, A&R find that robot displacement of workers led to a 0.34 percentage-point decline in the share of the working-age population with a job. These two sample employment impacts pale in comparison with the employment impacts of trade with China, and with employment declines overall in recent decades.

A&R’s own estimate suggests that job displacements caused by trade with China in the 2000s were four times as large as their estimate of job loss due to robots

In footnote number 26 in their 2017 report, A&R note that their estimated employment decline due to the effect of robots is only a third as large as the negative employment impact of U.S. trade with China.¹ This finding flatly contradicts *The New York Times* reporting that A&R found larger effects stemming from robot adoption than from growing trade with China. We note that A&R estimated the impact of U.S. trade with China on U.S. employment for the entire 1990–2007 period. However, the impact of China trade on U.S. employment was mostly felt in the latter part of that period, when imports from China grew much faster than Chinese imports grew earlier in the period. This suggests that the “China trade effect” for the years from 1999 to 2007 would be even more than three times as large as the “robot effect” for those same years.

To test this hypothesis, we examine the pace of growth of Chinese imports in the early

(1990–1999) and later (1999–2007) periods to allow us to compare the robot and Chinese import effects.² We find that the effect of Chinese imports on the employment rate over the 1990–1999 period was roughly equivalent to the impact of robots over the entire 1990–2007 period. Moreover, the impact of Chinese imports on the employment rate in the 1999–2007 period was *double* the impact of robots over the entire 17 years from 1990 to 2007. Thus A&R’s own findings indicate that Chinese imports had *four times* the annual impact of robots on the employment rate each year from 1999 to 2007 (U.S. Census Bureau 2017; BEA 2017).

It is also important to note that the impact of China trade is rightly described as a “shock,” as it represents a trend that accelerated in the 2000s relative to the 1980s or 1990s, and therefore partly explains the disappointing employment growth, manufacturing job loss, and broad-based wage stagnation of the 2000s.

The percentage-point decline in the employment rate in recent decades is nearly 12 times as great as A&R’s estimated employment rate decline due to robots

A&R find that robot displacement of workers led to a 0.34 percentage-point decline in the share of the working-age population with a job over the 1990–2007 period, based on their “preferred choice of parameters” (Acemoglu and Restrepo 2017a, 4). Specifically, their measure is the prime-age (age 25 to 54) employment-to-population ratio, or EPOP. But compare their estimated EPOP decline with changes in the prime-age EPOP in recent decades. From its peak in 2000 to 2016, the prime-age EPOP fell 4.1 percentage points, nearly 12 times the robot effect over the entire 1990–2007 period. Between its peak in the late 1990s boom and the trough of the 2001–2003 job-market recession (one often labeled “mild”), the prime-age EPOP fell 3.3 percentage points, nearly 10 times the employment rate decline due to robots. And between the 2006 peak and the trough of the Great Recession, the employment rate fell by 5.4 percentage points. Moreover, from its trough during the Great Recession, this measure has recovered by 3.5 percentage points (U.S. BLS various years.)

In short, even year-to-year movements of the EPOP stemming from macroeconomic conditions have utterly swamped any effect of robotic displacement. (As explained in more detail in the next section, even this small impact of robots cannot really explain why employment rates grew in the 1990s and have fallen since 2000, since robot use increased at the same pace in both the 1990s and 2000s (as can be seen in Figure 1 in the A&R paper).

Robot use may have accelerated post-1990, but there is no evidence that overall technological change and automation have had a larger or different effect in recent years

Technological change and automation have been major ongoing factors shaping the employment and wage structure for many decades. Demonstrating one particular technology's impact in only one period of time, as A&R (2017a) do with robots in the post-1990 period, cannot explain the onset of sluggish job growth or poor wage performance in recent decades relative to earlier decades. A convincing explanation for these adverse labor market trends requires identifying factors that are different now from those in the past, are large enough to affect the trajectory of the entire labor market, and map clearly onto periods of labor market change. So, for example, to say that automation is why we have had wage stagnation since the 1970s when we had strong wage growth in the 1950s and 1960s, or why we have had sluggish job growth, falling employment rates, and broad-based wage stagnation for both blue- and white-collar workers in the 2000s when we did not have these things in certain earlier decades, one needs to find that the ebbs and flows of automation were large and coincided with these labor market outcomes we are trying to explain. This does not describe the robot effect as described by A&R (2017a).

Why? A&R (2017a) show that the use of robots grew steadily over the 1990–2007 period. And we know that automation *writ large* (robot and nonrobot technology) has also been with us for many decades. Some of those decades had strong employment and median real wage growth, and some did not (Schmitt, Shierholz, and Mishel 2013; Schmitt 2013; Mishel 2013a, 2013b; Shierholz 2013). If automation had the same impact in each decade its cumulative impact would be large. However, a steady ongoing impact since the 1940s, let's say, could not explain why we had strong median wage growth in the 1950s and 1960s but median wage stagnation since the late 1970s—stagnation that receded in the late 1990s but accelerated again, with broad-based (blue- and white-collar) wage stagnation since 2002. It could not explain slow employment growth in the 2000s—including in the 2002-2007 recovery, even before the Great Recession. It could not explain the rise of wage inequality since the late 1970s. That is, the ongoing steady impact of automation cannot explain a *change* in wage or employment trends (Mishel and Bernstein 1994, 1998).

A&R (2017a) may argue that because robotization was essentially nonexistent at the start of their estimation period, their estimates are by definition an acceleration of technology's impact. Further, they may claim that robotization has been different from past technological changes in the degree or kinds of labor it displaces. But, crucially, they do not claim to be presenting any evidence that automation, broadly defined, had a greater overall impact in the 2000s than in the 1990s, 1980s, or 1970s. Robot use may be new but it does not necessarily just represent an addition to other forms of automation: rather, robots may well *replace* other forms of automation deployed in earlier decades. So, the examination of one new technology, robots, in one time period does not inform us of the

overall impact of automation in one period relative to another.

Automation broadly defined has actually been slower over the last 10 years or so

As Atkinson and Wu (2017) recently wrote, various technology-sector spokespeople and futurists frequently assert that “the pace of innovation is not just accelerating, but that it is accelerating ‘exponentially.’” This claim, pervasive in the media, is often followed by the sentiment that “we have not seen anything yet.” The empirical basis for these claims is absent because there is no evidence of any acceleration in workplace automation in recent years.

Automation is what occurs as the implementation of new technologies is incorporated along with new capital equipment or software to replace human labor in the workplace. Labor productivity and capital investment are both measures of automation in that they necessarily accompany the substitution of capital for labor. Thus if there were a recent surge of robots or automation we would expect to see the footprints in trends in productivity, capital investment, and software investment. In fact, a review of the indicators suggests that automation has decelerated, not accelerated, in the last 10 to 15 years (Mishel and Shierholz 2017). This directly contradicts the ongoing media narrative about robots.

Evidence for a deceleration of automation can be seen in **Figure A**, which presents trends in productivity, capital investment, and software investment for the earlier post–World War II years (1947–1973), the following period of slow productivity (1973–1995), the period of the IT-driven productivity boom (1995–2002), the period of slow employment growth in the 2002–2007 recovery, and the period since the onset of the Great Recession (2007–2016). Comparisons of these indicators across time periods is quite instructive.

The first indicator, productivity growth, has been slower in 2007–2016 than in any other period. The question then, frequently and appropriately asked, is, “If there is so much automation, then why has productivity growth slowed down?” And the recovery of 2002–2007 had slower productivity than the preceding seven years of the IT boom. It is hard to see the footprint of a surge in automation in the productivity data.


Figure A also tracks the trends in all capital investment and, specifically, in information technology (IT) equipment and software investments. It is reasonable to expect automation to be accompanied by investments in IT equipment and software, but rather than a surge, we see a deceleration of IT and software investments in the 2002–2007 and 2007–2016 relative to earlier times. We also see that capital investment in the workplace has grown more slowly since 2002 than in any other postwar period, a trend that many analysts have identified as a cause of the slow productivity growth. This is the opposite of what we would expect with a looming robot apocalypse based on automation accelerating exponentially.

Atkinson and Wu (2017) offer additional empirical evidence of a slowdown in automation


Figure A

Average annual growth rates of labor productivity, capital, and IT hardware and software, 1973–2016

Growth in labor productivity and the capital stock has decreased in recent periods



Capital investment in information technology has also slowed



Note: Using latest available data, 2016 measure includes data from 2015Q4–2016Q3.


Source: EPI analysis of [data \(xls\)](#) compiled by John Fernald of the Federal Reserve Bank of San Francisco

Economic Policy Institute

by tracking the rate at which occupational employment patterns have shifted in each decade from 1850 to 2010 and from 2010 to 2015. Automation (along with shifts in consumer demand and trade) would be expected to be a major factor in why employment in some occupations expands and employment in other occupations declines. Atkinson

Figure B

Change in occupational employment shares, by decade, 1940–2015



* Converted to decade rate by multiplying by two.

Source: Authors' analysis of data from Atkinson and Wu (2017)

Economic Policy Institute

and Wu compute two metrics of occupational change, and both show occupational employment shifts were far slower in 2000–2010 and 2010–2015 than in any decade in the 1900s.

Using the data graciously posted by Atkinson and Wu, we have computed another metric: the change in occupational employment shares (increases or decreases in occupations' shares of overall employment) in each decade. We examine the shares of total employment for each of the 250 occupations in the data for the beginning and end years of each decade and compute the changes in these shares. The sum of the occupational share gains will automatically be equal to the occupational share losses, so the metric of change in each decade is half the sum of the absolute change in employment shares. This metric adjusts for differences in the rate of employment growth in each decade and the absolute employment size of individual occupations. This metric of shifts in occupational employment measures the shares of total employment exchanged between occupations that gain and occupations that lose employment shares each decade.

The decadal rates of occupational employment shifts, starting in the 1940s, are shown in **Figure B**. The rate of change was fairly uniform over the 1940–1980 period and far more rapid than for any period since 1980. The period since 2000 has seen the lowest rate of change—half the rate of change of the 1940–1980 period. So, we see deceleration that corresponds to the slower growth in productivity and capital investments in IT discussed earlier.

Some of the high rates of occupational change in the 1940s and 1950s reflect the decline in seven farm occupations, which fell from 14.3 percent of total employment in 1940 to just 2.5 percent in 1960. However, even after adjusting the data to remove the impact of the erosion of employment in farm occupations in each decade, we find that rates of occupational change in the 2000s are still 45 percent lower than those in the 1940–1980 period. These are not turbulent times compared with the recent past.

A&R’s own results show that automation broadly defined has no effect on employment

The A&R estimates that we critiqued at the beginning of this section pertain only to the impact of robots. But, as evident in the discussion of Figure A above, there are a host of capital investments besides robots that can replace human labor in production. In fact, in decades past, many have blamed the labor market distress of working-class Americans on capital investments in IT. A&R include variables for these nonrobot investments in their CZ study and find that these investments have either a neutral or even a positive effect on employment and wages for these workers. Given that the sheer scale of these nonrobot capital investments is almost surely *substantially larger* than investments in robots, it is reasonable to interpret A&R’s results as showing that automation broadly defined has been positive for working-class Americans.

A&R’s claim that robots reduced overall, national-level employment is unconvincing because it relies on unrealistic assumptions

Acemoglu and Restrepo (2017a) estimate the employment and wage impacts of increased industrial robot use over the 1990–2007 period in local labor markets by comparing employment and wage-changes in CZs with large exposure to robot use with CZs that have not been exposed to as much robot use. Because these direct estimates of local effects appear to be of high quality, we focus here on their effort to estimate the aggregate, national impact of robots—and the media interpretation of the national estimate. A&R translate their direct, high-quality estimates of relative robot job displacements *across regions* into a *national* aggregate impact of between 360,000 and 670,000 jobs lost due to robots over 17 years. These aggregate employment projections are not statistical estimates. Rather, they combine the high-quality regression estimates on local data with a highly stylized model of national labor markets that relies on a number of stringent and likely unrealistic assumptions and simply rules out the possibility that some important channels of job creation offset jobs lost due to robots.

The bottom line is that A&R’s evidence is insufficient to overturn economists’ long-held,

evidence-based beliefs that automation, despite displacing jobs in particular occupations or industries, historically has *not* lowered aggregate employment or increased overall joblessness.

Automation displaces certain jobs but also creates other, offsetting jobs

Before tackling the unrealistic assumptions in the model it is worth discussing the many ways that automation can create jobs that potentially offset the jobs lost in industries affected by automation and thus leave overall employment and unemployment unchanged. The machinery used to replace human labor must be built and installed, creating jobs. Complementary jobs are created in activities that necessarily accompany the investment, such as programming and maintenance. The main dynamic, however, is that employers deploy automation because it will lower costs, and lowering costs will ultimately lower the prices of the goods and services produced (in addition to raising profits). These lower prices mean that consumers who purchase the less costly goods and services will have income left to purchase other goods and services—unless, of course, consumers run out of things they want to buy, which we doubt.

This increase in the demand for other goods and services will create jobs to generate those additional goods and services. Where these jobs will appear is unknowable, but history affirms that they do show up. The easiest illustration is what happened to agriculture. Due to automation, the share of the workforce employed in agriculture fell from around 50 percent in 1880 to just 20 percent in 1940 while still feeding the nation (Lebergott 1966, Table 2). Lower food prices allowed households to buy other goods, so employment rose in other sectors.

Despite many decades of ominous warnings, we have not seen evidence of automation leading to large-scale overall joblessness. That is, automation has been proceeding for many decades. If automation were creating increased joblessness we would have expected to see an ever-rising unemployment rate, but we have not seen that. The failure of past automation to cause increased overall joblessness has led economists to believe that automation displaces specific jobs but does not increase aggregate unemployment.³ The notion that automation has led to a rapid surge in joblessness is also hard to square with recent trends in unemployment: aggregate unemployment fell from 10 percent in late 2009 to under 5 percent in 2017.

As stated earlier, A&R's main contribution is to estimate the employment impact of robots by comparing employment trends in geographic areas (CZs) with heavy robot use with trends in CZs with less robot use. But one cannot simply extrapolate these local estimates into national ones, because this extrapolation would not capture many of the jobs generated by automation, as described above. To make this concrete, note that one of the biggest potential offsets to technology-induced displacement of jobs in one particular locale is the job boost from new technologies in other locales. Specifically, new technologies should boost productivity and hence lower the prices of the goods and services being newly produced with robots. Much of this effect could operate *outside* the

CZs that have seen increasing use of robots.

Because it highlights how the A&R model falls short, it is worth restating the two ways that technology-induced price declines can be expected to increase demand. First as industries adopt robots, prices for these industries' products fall, increasing demand. Second, as consumers pay less for the robot-using industries' output, they can direct their purchasing power to other industries, boosting demand for *all other industries' output* and thereby creating jobs. But if these other industries are not located in the same CZs as the robot-adopting industries, A&R's *direct* estimates of displacement at the CZ level will not pick this up. Further, while the spread of robots may displace jobs in particular CZs, the robots do need to be designed and built and then eventually replaced. Because these robot production activities probably occur in other CZs, the associated jobs also would not be picked up by A&R's direct regional estimation.

In short, one cannot use the cross-geographic area estimates of robots' impact on jobs to simply extrapolate the overall, national impact.

A&R's stylized model is based on unrealistic assumptions

Acemoglu and Restrepo recognized that their estimates of the impact of robots across geographic areas cannot be taken as a national, aggregate effect because these estimates do not capture many of the jobs generated by automation. This recognition led them to generate model-based computations of national employment effects that do aim to capture these job-creation effects. But this computation requires accepting a number of stylized and likely untrue assumptions about the economy. If one does not accept these assumptions, one should not rely on A&R's national job-loss number, and should not use this job-loss number to reject the longstanding and widespread consensus that automation does not lead to aggregate job loss or higher unemployment. Even A&R themselves note that their national employment computations should not be seen as the last word on this subject, due to the very reasons we cite above, writing:

We view our paper as a first step in a comprehensive evaluation of how robots will affect, and are already affecting, the labor market equilibrium. It is indeed only a first step, because our methodology directly estimates only the effect of robots on employment in a commuting zone relative to other commuting zones that have become less exposed to robots. We then used the structure of a model of trade between commuting zones to infer the aggregate effects of robots. Alternative strategies of estimating the aggregate implications of the spread of robots, either by focusing on cross-country comparisons or by directly or indirectly measuring the flows of goods across local labor markets, would be obviously highly complementary to this approach. (A&R 2017a, 36)

The first and largest stylized and likely untrue assumption is that the economy is in "equilibrium," or full employment, at both the beginning and the end of their sample period. This means that wages and employment are simply set by the intersection of labor demand and supply curves and thus there is no scope for changes in aggregate demand (spending by households, businesses, and government) to affect employment growth.

There is no particular reason to believe that the U.S. economy was in such an equilibrium in 1990 or even 2007. In many years in this period, unemployment rates were notably lower than they were at the beginning and end of this period. Given this, it is impossible to justify their claim that only changes in technology, not in aggregate demand, are driving relationships between wages, employment, and robot-use in those years.

The second major assumption is that labor markets are perfectly competitive, with wages and the employment set simply by the intersection of supply and demand curves, with no role for institutions in setting wages. An example shows why this casts doubt on A&R's inferences. Take a situation in which robotization disproportionately displaces jobs in industries with high unionization rates. If this leads to wage declines through reduced bargaining power, this breaks the mechanical link between wage and employment growth posited in A&R's stylized model of competitive labor markets. Specifically, if wages fall due to erosion of union effects, this could reduce labor supply, independently of the effect of robots.⁴

One potential sign that A&R's estimates should not be strictly read as representing equilibrium effects from a well-behaved "competitive" labor market model can be found in Table A4 in their report. This table shows that robotization increases unemployment in high-use CZs and that this effect is significant. But unemployment represents potential workers *actively searching for jobs yet not finding them*. In short, it represents a disequilibrium outcome not represented by the intersection of competitive labor supply and demand curves. Because even their directly estimated, within-CZ effects show signs of being out of equilibrium, it seems highly likely that the overall labor market is out of equilibrium as well.

We should be clear that adding in nonequilibrium effects could make the employment effect of robotization worse. For example, if increased use of robots depresses wages but boosts the incomes of robot owners, robotization represents a transfer of income to upper-income households that have higher propensities to save rather than spend. This shift of income to high-income savers could slow aggregate demand growth and increase unemployment. But slow demand that is driven by a shift of income to high-income savers would have policy implications that are very different from the education and training policies that are generally discussed when the labor market distress is blamed on robots. Specifically, the economic drag from regressive distribution of income potentially caused by robotization would simply call for macroeconomic policymakers to use levers of monetary or fiscal policy to boost aggregate demand growth. And to be clear, the A&R-estimated within-CZ effects that include years of the Great Recession are notably higher than the estimates that do not include these years. This suggests that the relationship between robot use and job loss is strengthened by labor market slack. But the labor market may well have been slack in many years between 1990 and 2007 (if not in a Great Recession-level crisis).⁵

Maybe there has not been a robot apocalypse yet; but what about

joblessness and inequality in the future?

There are two main worries about future robot automation and its effect on American workers: that it will lead to automation-induced unemployment—an increase in overall joblessness—and that it will lead to greater inequality. The evidence to date strongly argues against worrying much about either one of these potential outcomes.

Do not expect higher joblessness due to robots

A&R in their conclusion seem to acknowledge that their estimates of job loss due to robots are small, referring to their estimates as “limited.” They also state that there are relatively few robots in the U.S economy (and by the way, A&R show that the number of robots in the United States is considerably smaller than the number of robots in the European economy). They then note that the implications could be large if robot use significantly expands, as various experts have claimed. The potential future impact of robots is also emphasized in a recent A&R article:

So far, there are relatively few robots in the US economy, and so the number of jobs lost due to robots has been limited to between 360,000 and 670,000 jobs. If the robots spread as predicted, future aggregate job losses will be much larger. (Acemoglu and Restrepo 2017b)

However, their research does not actually provide a convincing direct estimate of a negative *aggregate* employment effect (as argued earlier), so worrying about future joblessness is unnecessary.

A&R cite studies widely interpreted as projecting that future automation will drive joblessness (Acemoglu and Restrepo 2017a, 3). But these studies only address potential *occupational*-specific displacements due to automation. Specifically, they provide projections of how many occupations are “susceptible ... to automation,” according to A&R. These studies make *no* claims about aggregate job losses due to automation and do not assess any of the potential job-creating effects of automation, reviewed earlier, that could offset the displacements. That is, studies of how many occupations can be automated do not tell us whether automation will lead to future joblessness unless they also estimate the positive impact of future technologies. It is also worth noting that many credible studies, such as the 2016 study by the Organization of Economic Cooperation and Development (Arntz 2016), call into question whether automation will affect a near majority of jobs.

As noted, we don’t find A&R’s evidence of an aggregate negative employment impact of automation to be comprehensive and convincing. Worries about large-scale unemployment driven by future automation seem unwarranted. A rapid pace of automation could overwhelm the offsetting positive employment impacts, but that is speculation. Those warning that automation will cause large-scale future joblessness have failed to articulate why “this time will be different.”

That does not mean we should ignore the plight of displaced workers. We have no quarrel with anyone who wants to strengthen the safety net and provide education and training to workers displaced by automation or any other factor. But the biggest challenge displaced workers face is not getting training, but getting training that leads to a good job. That is why we find no comfort in the conventional and limited call for education and training as a cure-all, such as that espoused by Fred Hiatt in *The Washington Post*, when he wrote, “What government should be doing in response to automation and other forms of economic dislocation [is]... education and training, infrastructure, research” (Hiatt 2017).

We would note that the lack of enough good jobs is what has prevented the labor market from successfully absorbing displaced manufacturing workers. Our biggest challenge is ensuring that the jobs of the future pay well, including jobs that do or do not require a college degree; the quality and pay of both blue-collar and white-collar jobs has degraded rapidly in recent decades (Mishel et al. 2012; Bivens et al. 2014). Therefore, even if one seeks to address a far-off potential robot apocalypse, the appropriate policy response is the one needed to address today’s real problems: wage stagnation and rising inequality. These appropriate policies include ensuring genuine full employment; raising the minimum wage, combatting wage theft by employers, and taking other actions to strengthen labor standards; and rebuilding collective bargaining. These policies are articulated in EPI’s 2016 [Agenda to Raise America’s Pay](#).⁶

Automation broadly defined has not in recent years induced wage inequality via occupational employment polarization

Will robots cause wage inequality? We think not. Contrary to conventional wisdom, there is no basis for believing that technology or robots have been generating wage inequality in recent years and not much reason to worry about the future impacts. For example, “occupational job polarization” is often offered as evidence that automation causes inequality, but recent data show that polarization cannot plausibly explain the inequalities we have been experiencing.

Studies of occupational polarization examine whether technological change is shrinking middle-wage occupations and expanding high-wage and low-wage occupations, resulting in a hollowed middle class, greater wage inequality between earners at the top and the middle of the wage distribution, and less inequality between the middle and the bottom. The implied theory is that future robotization could accelerate an ongoing trend wherein technological change and automation broadly defined reduce the demand for “routine” tasks (i.e., more readily automatable) and hence reduce employment and thus wages of many American workers, particularly those without a four-year college degree. For noncollege workers displaced from middle-wage “routine” occupations, the only option is to find work in a “manual” low-wage sector that sees expansion following the hollowing-out of the middle class.

There are three key problems with this chain of inferences linking occupational

employment patterns and wage inequality.

1. Job polarization cannot explain the doubling of the top 1.0 and 0.1 percent income and wage shares

As we have written elsewhere (Bivens and Mishel 2013), wage and income shares of the top 1 percent of taxpayers doubled from 1979 to 2007. For taxpayers at the 99.5 percentile, wage and income shares grew even more. The enormous growth of top incomes and wages was due to the expansion of the financial sector, higher pay in the financial sector, and the tremendous growth in executive and managerial compensation. Robots or automation more generally were not factors, and the job polarization literature has never even addressed top income and wage growth.

2. Job polarization has not been occurring since 1999, so is not relevant for predicting the future and cannot possibly explain recent trends


Job polarization cannot explain recent wage inequality in the bottom 99 percent of the wage structure simply because job polarization has not been present. For example, Acemoglu and Autor (2011) show that occupational employment polarization did not occur between 1999 and 2007, since *all* of the employment expansion in this period was in lower-wage occupations (the bottom 25 percent of occupations arrayed by wage levels) and employment did not expand, as a share of total employment, in the upper 75 percent of occupations arrayed by wage levels (Autor 2010). Later work by Autor first presented in 2014 shows that job polarization did not occur in the 2007–2012 period either: there was very little or no employment expansion in the upper 75 percent of the occupational wage structure (there was mild expansion in the upper 17 percent) (Autor 2014). As such, job polarization did not appear in occupational employment patterns post-1999, and therefore cannot be considered a serious contributor to the continuing rapid growth of wage inequality between the upper 10 percent of wage earners and middle-wage workers.

The key chart from Autor’s most recent analysis is replicated in **Figure C** here. It shows that there has been little or no polarization since 1999. The analysis combines all occupations into 100 occupation units and ranks them by wage level. The chart shows the expansion of employment (as a share of total employment) on the vertical axis and occupations ranked by wages from lowest to highest on the horizontal axis. Any portion of the dotted line above (or below) the horizontal line indicates occupations that increased (or decreased) their share of employment. The figure shows that “polarization” occurred in 1989–1999 because employment expanded in the low-wage occupations (the bottom 10 percent) and the high-wage occupations, with the greatest expansion among the highest-wage occupations. But significantly, the chart shows that employment grew only among low-wage occupations in 1999–2007 (orange line) and primarily in low-wage occupations in 2007–2012 (green line).

Autor is not the only researcher to find there has been no polarization in the 2000s. Acemoglu and Restrepo (2017a) overlooked many studies published before Autor’s 2014

Figure C

Smoothed employment changes by occupational skill percentile, 1979–2012



Source: Reprinted from David Autor, "Polanyi's Paradox and the Shape of Employment Growth," MIT Economics (Autor 2014)

Economic Policy Institute

paper that also found that polarization has not occurred in the 2000s. These studies include Beaudry, Green, and Sand (2013); Schmitt (2013) and Shierholz (2013), both in January 2013 with a fuller analysis in November 2013 (Schmitt, Shierholz, and Mishel 2013); and Mandelman (2013) at the Atlanta Federal Reserve Board (Chart 2) in July 2013.

3. Occupational employment patterns do not generate or explain wage inequality

Equally important, patterns of occupational employment growth—whether exhibiting polarization or not—do not drive wage stagnation and are largely irrelevant for the study of wage inequality. Wage inequality is a phenomenon that occurs overwhelmingly *within* occupations. Further, the expansion and decline of an occupation's employment shares are unrelated to whether wages in that occupation rise or fall relative to other occupations (Schmitt, Shierholz, and Mishel 2013). Therefore, the mechanism by which occupational employment patterns would be expected to drive wage trends is absent.

To summarize, automation-induced job polarization just should not rank highly in our concerns about wage inequality because this polarization hasn't happened since 1999 or 2000 and would not affect wage trends even if it did happen. This can readily be illustrated by examining what happened to the wages of low-wage workers in the 2000s.

The lowest-wage occupations expanded in the 2000s while middle- and higher-wage occupations, for the most part, had declining or stable employment shares, as reviewed above. This would lead one to believe that low-wage workers would have seen the fastest wage growth, as employers competing to hire low-wage workers would need to increase their wage offers. Yet wages at the 20th and 30th percentiles grew about 2 percent during 1999–2016 (when low-wage occupations greatly expanded) while wages in the middle (40th, 50th, and 60th percentiles) grew about 4 percent. During this same period, wages at the top (90th and 95th percentiles) grew, respectively, by 17 and 23 percent.⁷ In short, analysts interested in understanding wage trends should not be focused on occupational employment patterns.

To summarize, take the metaphor of a murder mystery. If the victim were the labor market prospects of working American, automation-induced occupational job polarization could not have been the cause of death. For one, as a weapon it was not present at the crime scene (polarization has not been present since 1999). For another, it largely shoots blanks (i.e., occupational employment polarization does not lead to wage inequality or wage stagnation). Finally, we know who the major perpetrators were and how they committed the crime (the top 1 percent rigged the rules of the economy to claim a wildly disproportionate share of income growth, as explained in Bivens (2016)).

Conclusion

Maybe one day robots or some other technological advance really will become a primary threat to American living standards. Anything can happen. But scarce resources and attention today should be focused on the current threat to these living standards: policies that have shifted bargaining power from working Americans to capital owners and corporate managers. Alleviating the struggles of working Americans will thus require a redistribution of economic leverage and bargaining power. Making this happen will be a huge political project, and will require a clear focus.

Claims that the struggles of working Americans are instead the sad byproduct of bloodless trends in technology obscure this clear focus. The overwhelming evidence still strongly suggests that rapid technological advances are associated with better, not worse, economic outcomes for working Americans. To be clear, rapid technological advances alone are just a necessary, not a sufficient, condition to boost wages across the board. American workers also need policies that support their ability to claim an equitable share of productivity growth.

Adopting policies to lift labor standards, broaden collective bargaining, and maintain genuine full employment would help ensure that there are good-quality jobs available for workers displaced by technology. This, more than training and education, is what is needed to respond to any potential wave of automation. These policies would also confront the wage stagnation already upon us. These policies are identified in our *Agenda to Raise America's Pay* (EPI 2016). If even a quarter of the attention paid to robots were instead shifted to these policy changes, we could have a much more fruitful debate about economic policy.

About the authors

Lawrence Mishel, a nationally recognized economist, has been president of the Economic Policy Institute since 2002. Prior to that he was EPI's first research director (starting in 1987) and later became vice president. He is the co-author of all editions of *The State of Working America*. He holds a Ph.D. in economics from the University of Wisconsin at Madison, and his articles have appeared in a variety of academic and nonacademic journals. His areas of research are labor economics, wage and income distribution, industrial relations, productivity growth, and the economics of education.

Josh Bivens joined the Economic Policy Institute in 2002 and is currently the director of research. His primary areas of research include macroeconomics, social insurance, and globalization. He has authored or co-authored three books (including *The State of Working America, 12th Edition*) while working at EPI, edited another, and has written numerous research papers, including for academic journals. He often appears in media outlets to offer economic commentary and has testified several times before the U.S. Congress. He earned his Ph.D. from The New School for Social Research.

Acknowledgments

The authors would like to acknowledge the helpful advice of John Schmitt, Dean Baker, Jared Bernstein, Ben Spielberg, and those of EPI colleagues at an internal seminar. Teresa Kroeger and Jessica Scheider both provided able research assistance. Thanks also to Lora Engdahl for her editing.

Endnotes

1. Their footnote says, "These magnitudes can be compared to the size of the effects from exposure to imports from China. Table A5 in the Appendix shows that in our long-differences specification, the exposure to China variable has a negative but insignificant coefficient. However, in the stacked-differences specification, which corresponds to the one used in Autor, Dorn and Hanson (2013), it is negative and significant. Using this latter estimate, the implied quantitative magnitude from the national rise in Chinese imports for a commuting zone with the average exposure to Chinese imports (compared to a no exposure commuting zone) is a decline of roughly 1 percentage point in the employment to population ratio, which is about *three times* the 0.37 percentage points implied decline from the adoption of robots [emphasis added]."
2. Chinese import growth is measured as a share of gross domestic product, and rose from 0.25 percent in 1990 to 0.85 percent in 1999 to 2.22 percent in 2007. That means that 70 percent (1.37 percentage points) of the 1.97 percentage-point growth in Chinese imports occurred from 1999 to 2007. A&R (2017a) indicate that the impact of China imports over 1990–2007 was about a 1 percentage-point reduction in the employment rate, roughly three times greater than A&R's estimated impact of robots of 0.37 percentage points. The annual impact of Chinese imports on the employment rate was .088 per year (1.0*0.7 over eight years) from 1999 to 2007. In contrast, the annual impact of robots on the employment rate was 0.022 (0.037 over 17 years). Chinese imports data are from the Census; GDP data are from BEA's national income and product accounts

(U.S. Census Bureau 2017; BEA 2017).

3. If automation leads to price declines in directly affected industries and price declines do not lead to increased consumption spending overall, policymakers can also simply raise the level of aggregate demand through expansionary fiscal or monetary policy to make sure job losses do not result.
4. Union effects refers to research (e.g., Mishel et al. 2012) showing that all else equal, workers in unions earn higher wages than their nonunion peers. Other research (e.g., Rosenfeld, Denice, and Laird 2016) has shown that declining unionization lowers the wages of all workers.
5. Before closing, it is useful to mention two other assumptions in the A&R model that make the national estimates unreliable. First, they assume the economy starts from a position of essentially having no robots. This could well describe their dataset, but it does make inferences using this method about the future break down. Second, they assume that technology and the cost of robots (and robot installation) are invariant across commuting zones. The analysis of A&R estimates would benefit from some actual data, which they do not present.
6. To be clear, the agenda of boosting workers' bargaining power to ensure a higher number of "good" jobs and give workers displaced by robots (or anything else) decent work to transition to is a viable strategy as long as (1) bargaining power and institutions matter for wage setting, and (2) automation has not made workers obsolete at living wages. We think the evidence clearly indicates these two things are true. If instead all that mattered for labor market outcomes were well-defined supply and demand curves, and were the labor demand curve for some types of workers (generally the less-credentialed) so shifted by technology that these workers were unemployable at living wages, then policies like a universal basic income could be the only ones able to provide income to these Americans. Again, we find no evidence that this dynamic has begun.
7. Calculations made from wage-by-percentile data presented in the State of Working America Data Library (EPI 2017).

References

- Acemoglu, Daron, and David Autor. 2011 "Skill, Tasks and Technologies: Implications for Employment Earnings." In *The Handbook of Labor Economics* vol. 4b, O. Ashenfelter and D. Card, eds. Amsterdam: Elsevier.
- Acemoglu, Daron, and Pascual Restrepo. 2017a. *Robots and Jobs: Evidence from US Labor Markets*. NBER Working Paper No. 23285. National Bureau of Economic Research.
- Acemoglu, Daron, and Pascual Restrepo. 2017b. "Robots and Jobs: Evidence from the US" (article based on working paper). *Vox*: Centre for Economic Policy Research, April 10.
- Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. 2016. "The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis." *OECD Social, Employment and Migration Working Papers*. Organisation of Economic Co-operation and Development.
- Atkinson, Robert D., and John Wu. 2017. *False Alarmism: Technological Disruption and the U.S. Labor Market, 1850–2015*. Information Technology & Innovation Foundation, May 8.
- Autor, David. 2010. "The Polarization of Job Opportunities in the U.S. Labor Market: Implications for

- Employment and Earnings.**” Center for American Progress and the Brookings Institution: The Hamilton Project.
- Autor, David. 2014. “**Polanyi’s Paradox and the Shape of Employment Growth.**” MIT Economics.
- Beaudry, Paul, David A. Green, and Benjamin M. Sand. 2013. “**The Great Reversal in the Demand for Skill and Cognitive Tasks.**” NBER Working Paper No. 18901. National Bureau of Economic Research.
- Bivens, Josh. 2016. *Progressive Redistribution Without Guilt: Using Policy to Shift Economic Power and Make U.S. Incomes Grow Fairer and Faster.* Economic Policy Institute.
- Bivens, Josh, Elise Gould, Lawrence Mishel, and Heidi Shierholz. 2014. *Raising America’s Pay: Why It’s Our Central Economic Policy Challenge*, Economic Policy Institute, Briefing Paper no. 378,
- Bivens, Josh, and Lawrence Mishel. 2013. “The Pay of Corporate Executives and Financial Professionals as Evidence of Rents in Top 1 Percent Incomes.” *Journal of Economic Perspectives* vol. 27, no. 3, 57–78.
- Bureau of Economic Analysis (BEA), U.S. Department of Commerce. 2017. “**Table 11.5: Gross Domestic Product**” [data table]. In *National Income and Product Accounts* [database]. Accessed April 10, 2017.
- Economic Policy Institute (EPI). 2016. “**The Agenda to Raise America’s Pay.**”
- Economic Policy Institute (EPI). 2017. “Wages by Percentile.” In *State of Working America Data Library.*
- Guo, Jeff. 2017. “**We’re So Unprepared for the Robot Apocalypse.**” *Washington Post: Wonkblog*, March 30.
- Hiatt, Fred. 2017. “**How Democrats Can Be Progressive without Being Irresponsible.**” *Washington Post*, April 9.
- Lebergott, Stanley. 1966. “**Labor Force and Employment, 1800–1960.**” In *Output, Employment, and Productivity in the United States after 1800*, Dorothy S. Brady, ed. Cambridge, MA: National Bureau of Economic Research.
- Mandelman, Federico S. 2013. “**The Asymmetric Impact of Offshoring and Low-Skill Immigration on Employment and Wages.**” Federal Reserve Bank of Atlanta.
- Miller, Claire Cain. 2017. “**Evidence That Robots Are Winning the Race for American Jobs.**” *New York Times: The Upshot.*
- Mishel, Lawrence. 2013a. “Occupation Employment Trends and Wage Inequality: What the Long View Tells Us.” *Working Economics* (Economic Policy Institute blog), January 15.
- Mishel, Lawrence. 2013b. “Timing Matters: Can Job Polarization Explain Wage Trends?” *Working Economics* (Economic Policy Institute blog), January 14.
- Mishel, Lawrence, and Jared Bernstein. 1994. *Is the Technology Black Box Empty? An Empirical Examination of the Impact of Technology on Wage Inequality and the Employment Structure.* Washington, D.C.: Economic Policy Institute.
- Mishel, Lawrence, and Jared Bernstein. 1998. “**Technology and the Wage Structure: Has Technology’s Impact Accelerated Since the 1970s?**” *Research in Labor Economics* vol. 17.

- Mishel, Lawrence, Josh Bivens, Elise Gould, and Heidi Shierholz. 2012. *The State of Working America, 12th Edition*. Washington, D.C., and Ithaca, N.Y.: Economic Policy Institute and Cornell University Press.
- Mishel, Lawrence, and Heidi Shierholz. 2017. “[Robots, or Automation, Are Not the Problem: Too Little Worker Power Is](#).” Economic Policy Institute.
- Rosenfeld, Jake, Patrick Denice, and Jennifer Laird. 2016. *Union Decline Lowers Wages of Nonunion Workers: The Overlooked Reason Why Wages Are Stuck and Inequality Is Growing*. Economic Policy Institute report.
- Schmitt, John. 2013. “[Job Polarization in the 2000s?](#)” *Working Economics* (Economic Policy Institute blog), January 13.
- Schmitt, John, Heidi Shierholz, and Lawrence Mishel. 2013. “[Don’t Blame the Robots: Assessing the Job Polarization Explanation of Growing Wage Inequality](#).” EPI-CEPR Working Paper: Economic Policy Institute and Center for Economic and Policy Research, November 13.
- Shierholz, Heidi. 2013. “Are the Job Polarization Data Robust?” *Working Economics* (Economic Policy Institute blog), January 18.
- Thompson, Derek. 2015. “[A World Without Work](#).” *The Atlantic*, July/August.
- Thompson, Derek. 2016. “[When Will Robots Take All the Jobs?](#)” *The Atlantic*, October 31.
- Tingley, Kim. 2017. “[Learning to Love Our Robot Co-Workers](#).” *New York Times Magazine*, February 23.
- U.S. Bureau of Labor Statistics (BLS). Various years. Current Population Survey labor force statistics. Labor force statistics from the Current Population Survey used in this paper are available from the Bureau of Labor Statistics through three primary channels: [annual data \[Tables 8, 14, 19, 20, and 21\]](#), [monthly household data \[2015 and 2016\]](#), and [through series reports \[LNS12032194, LNS12032195, LNS12032196, and LNS12005977\]](#).
- U.S. Census Bureau. 2017. *Trade in Goods with China* [database]. Accessed April 10, 2017.