ROBOTS, SKILL DEMAND, AND MANUFACTURING IN U.S. REGIONAL LABOR MARKETS

Nancey Green Leigh, PhD, FAICP*
Associate Dean for Research
College of Design
Professor of City and Regional Planning
Georgia Institute of Technology
ngleigh@gatech.edu

Benjamin Kraft
PhD Student
School of City and Regional Planning
Georgia Institute of Technology

Heonyeong Lee
PhD Student
School of City and Regional Planning
Georgia Institute of Technology

*Corresponding author

This research was supported by National Science Foundation Award #1637737, NRI: Workers, Firms, and Industries in Robotic Regions. We thank Professor Bill Drummond for his helpful comments and suggestions and Sian Llewellyn for her research assistance.

Introduction: Work and Technology in the 20th Century

Determining how technology changed the way we work over the second half of the 20th century is not a straightforward task, but the need to do so has increased with concerns over not just how work has changed, but whether work is substantially being eliminated. In contrast to Braverman's (1998) account of the disappearance of traditional craft skills as industrializing countries adopted Fordist production methods in the 20th century (which he labeled "deskilling"), mounting evidence suggests that—in measurable terms—workers' skill requirements increased. This "upskilling" is attributed to the information and communications technology revolution that spread to all sectors of the economy, and which also gave rise to robots and artificial intelligence. Skill Biased Technical Change (SBTC) is a concept employed by Autor, Levy, and Katz (2003) who found the increase in skill content of jobs over the second half of the 20th century was, controlling for industry, occupation, and education, directly related to the computerization of work. Within their "task model" framework, SBTC can be interpreted as a shift in demand away from routine job tasks, such as bookkeeping or factory assembly (i.e. tasks that can easily be automated) and toward complex tasks, such as tax planning or repairing industrial automation equipment.

However, upskilling was not the only notable trend in the labor market over the 20th century. Wage polarization is another, and the task model has done a poor job of accounting for the substantial growth in employment in "low-skill" service jobs (Acemoglu & Autor, 2011). To correct for this shortcoming, subsequent research incorporated changing consumer preferences and susceptibility to offshoring, and found that computerization is the primary explanation for redistribution away from middle skill work towards both the low and high skill ends of the labor market in the U.S. and Europe (Autor & Dorn, 2013; Goos, Manning, & Salomons, 2009).

While the econometric approach outlined above has been influential in policy debates about work and technology, it is only one of many possible approaches to examining the changing nature of work within the context of technology. Research from sociology, management, and organizational sciences highlights the social context of technological change by focusing on qualitative analyses, case studies, and "microdynamics" (Adler, 1992, p. 7), often at the individual workplace level. These approaches include critiques of the abstract conceptualization of technology and a re-emphasis on the actual physical nature of various workplace technologies—their "materiality"—and how they become incorporated into social systems (Leonardi, 2012; Orlikowski & Scott, 2008). This socially contingent perspective of technology generates nuanced sets of findings that differ from workplace to workplace, and that complicate the generalized trends described by statistical analysis.

For research seeking to determine technology's influence on work at meso and macro scales, a socially contingent perspective may imply machines are not the root "cause" of labor market problems in the first place. Rather, problems of skill bias, wage polarization, and technological unemployment may be fundamentally political—as opposed to technological—problems (Grint & Woolgar, 2013). For example, what is being interpreted as "upskilling" may actually be "upcredentialing," driven by increased policy emphasis on degree seeking (Keep & Mayhew, 2010) and increased employer selectiveness (Cappelli, 2012, 2014).

Even if machines are merely vehicles for politically induced labor market failures, the question of whether "this time is different" remains a pressing one. With the (notable) exceptions of Gordon (2012, 2014), who has been skeptical of our ability to innovate our way out of the long-term post-1973 economic stagnation, economists have generally been optimistic about the capacity of technological progress to generate a net surplus of jobs, as the virtuous cycle of rising wages and increased demand should spur the reallocation of labor to ever more useful functions. Only recently has it been suggested that we are at an “inflection point” whereby technological capabilities will outpace society’s ability to absorb the disruption that they cause, and radical changes in safety nets and education policy are in order (Brynjolfsson & McAfee, 2014).

With these complex and consequential issues in the background, this paper focuses more narrowly on the question of whether robots specifically, since the start of the recovery from the Great Recession, are exerting pressures on labor markets above and beyond that seen in previous rounds of industrial automation. We provide insight into this question using the case of industrial robots (those used in manufacturing operations), for two key reasons: 1) Even though the bulk of future disrupting technologies lie outside of manufacturing (e.g. drones, autonomous vehicles, personal care robots); they will largely be based on the technology pioneered in industrial robotics, and 2) industrial robots are already widely and systematically used in manufacturing, as opposed to service and consumer robots, which are in early stages of adoption.

This paper is structured as follows: First, we summarize the sparse literature on this topic, and discuss the main reason for its sparseness, that is, the limited availability of robotics data. We then develop our own indicator of robotics presence in a regional economy, and employ it in an empirical model using a novel data source, real time labor market information (RTLMI), that is derived from online job postings. We find that this new indicator performs better than others in explaining manufacturing job growth, and that the effect of robots is a slightly positive one.

Robots and Work: Actual Evidence

In contrast to the substantial literature on the general impacts of information and communications technology (ICT) (summarized briefly above), the body of work to date on the specific economic impact of robots is scarce (Graetz & Michaels, 2015) and has been conducted non-systematically, employing a variety of methods, time-frames, geographies, and data sources.

Throughout the 1980s—during the first major wave of industrial robot diffusion—there was some interest in predicting the diffusion and effects of robots in the U.S. Although this effort was largely abandoned by the 1990s, it reveals a noteworthy bullishness toward the uptake of industrial robots. Initial conservative scholarly and expert estimates of robot diffusion, reported in Hunt and Hunt (1983), predicted low-end ranges of 75,000 to 100,000 robots in operation by 1990 and a high-end of 150,000. However, even these self-described conservative forecasts were greater than the actual growth of robots through 1990. Indeed, as of 2004, the first year for which robot sales in the U.S. were directly reported to the International Federation of Robotics, the total number of robots in the U.S. was estimated to be less than 124,000 (2017a).

Predictions of robotics impacts were similarly inflated. Miller's "Impacts of Industrial Robotics" (1989) summarized several previous estimates showing relative agreement that displacement of labor by robots would be significant—on the order of between 33% and 55% between roughly 1990 and 2000—and confined to the handful jobs that involve tasks at which robots are especially adept. While this displacement rate estimate is probably too high, its recognition that displacement will be industrially localized is conceptually accurate, as robots are still intensively used in only a few manufacturing sectors. A separate estimate by Howell (1985) also predicted significant job losses due to robots, although the estimated range of 167,000 to 718,000 robot-displaced jobs did not consider jobs that may be created by potential increases in consumer demand for products made less expensive due to increased productivity from robot use.

Now that a time series of robot stock estimates exists (2004-2017), assessments of robot impact warrant higher degrees of confidence. We broadly categorize the current robot-impacts literature in terms of estimates of impacts to workers, distinguishing by whether robots a) have neutral effects on numbers of workers, or b) displace workers. The studies with neutral results, which we label "robots-as-status-quo" research, are Graetz and Michaels (2015) and Jäger et al (2015). Those that link labor displacement to robots, which we label "robots-as-displacers," are Acemoglu and Restrepo (2017) and Frey and Osborne (2013).

Table 1: Summary of Robot-Impact Research

	Year	Title	Employment Impacts	Other Results
Robots-as-status-quo				
Graetz & Michaels (G&M)	2015	Robots at Work	No effect on overall production hours worked; slight reduction in hours worked by lower-skilled workers	Increase labor productivity; Increase value added
Jäger et al	2015	Analysis of the Impact of Robotics Systems on Employment in the EU	No effect on number of production workers employed	Increase labor productivity; Decrease likelihood of offshoring production
Robots-as-displacers				
Frey & Osborne	2013	The Future of Employment: How Susceptible are Jobs to Computerization	47% of current occupations at high risk of automation	Not tested
Acemoglu & Restrepo (A&R)	2017	Robots and Jobs: Evidence from U.S. Labor Markets	One robot/thousand workers decreases employment by 3 to 6 workers and aggregate wages by .25% to .75%	Increased labor productivity (but of insufficient magnitude to mitigate displacement)

As Table 1 shows, the assertion that robots increase labor productivity is uncontroversial, putting robot-specific research in line with the SBTC consensus that machines generally make human work more productive.[1] Also in accordance with the current state of SBTC research is Graetz' and Michaels (2015) finding that the availability of work is essentially a zero-sum game *among workers*. That is, while there is no evidence that robots displace workers, there is evidence that they steer work in the direction of those with greater skill levels and away from those with less. This makes labor displacement an especially pernicious problem, because even when new jobs are created through increased consumer demand, they are not necessarily related geographically or functionally to the jobs that are lost. However, based on current research, this robot-specific skill biased pattern is no different or more severe than that associated with other types of automation technology. While rectifying these skill disparities has been a vexing policy problem for several decades, these findings at least suggest that the recent state of robotics technology does not portend the onset of mass technological unemployment. In fact, Jäger et al (2015) suggest that robots may be locally job preserving, as a high degree of robot use in an establishment is associated with a reduced likelihood of that establishment's production relocating to another country.

However, the "robots-as-displacers" results are more concerning because they suggest that rather than setting in motion a reallocation of skills within the human workforce, robotics (and AI more generally) is instead shrinking the size of the workforce—or at least the size of the portion of the workforce accessible to those with limited skills.

As previously noted, the four studies reviewed here are not perfectly comparable. Unlike the others, Frey and Osborne (2013) is a probabilistic forecast of threats to current occupations that can be expected in the near future based on mapping specific job tasks, as defined by the Standard Occupational Code (SOC), to experts' assessments of technological capabilities and

[1] Results suggesting otherwise would call into question the fundamental purpose of technological innovation, which is to enable a more efficient or liberating allocation of human labor, depending on one's philosophical attitude toward the "good life."

trajectories. It applies to AI *in general* rather than robots specifically. So, while the model predicts significant continued displacement to manufacturing jobs, it also includes logistics and administrative jobs that lie outside of the purview of traditional industrial robots. According to the model, most non-supervisory production occupations have at least a 50% chance of being "computerized," with the bulk having at least a 75% chance.

Acemoglu and Restrepo (2017) build on Graetz and Michaels (2015) by deriving a subnational estimation of robot use (or "exposure" as they call it) from IFR's industry and country tabulations. Because the IFR only provides robot sales and stock data for countries and industries, Acemoglu and Restrepo map robot exposure by industry onto U.S. Commuting Zones (CZs), based on the concentration of each industry within each commuting zone. Their findings suggest that robots exert substantial downward pressure on local employment and wages, and that this pressure becomes intensified when interregional trade is taken into account.

Despite the divergent findings of the nascent robotics-impact literature, they are not necessarily contradictory. Differing results may be artifacts of the different methods, data sources, geographies, and time periods in question. For example, robots may simultaneously drive productivity growth with little job destruction when aggregated across the EU and OECD member states—as Jäger et al and Graetz and Michaels suggest—*and* displace workers and reduce salaries in the U.S. at the same time. While the geographic differential in robot uptake and impacts gets lost in the hyperbole of the media debate on the topic, Leigh and Kraft (2017) show that regions within the U.S. have substantial differences in their robotics knowledge bases and that these differences will likely affect the quantity and quality or local automation capabilities.

Robotics Data in Perspective

The unavailability of robotics data remains a significant impediment that accounts for the discrepancy in (and shortage of) robot-impact research. There is no North American Industrial Classification System (NAICS) code to single out robotics manufacturers or consulting firms. Nor is there a Standard Occupational Code (SOC) to identify employees who work directly with robots.[2] Consequently, there is no way to isolate robot users or makers in these commonly used, publicly available datasets.

One way to overcome the robotics data gap is to collect original data on robots. However, due in part to the time and expense involved in generating original datasets of this nature, these efforts have been scarce.

There are two examples to date of original robotics data collection. The first, a Fraunhofer Institute report to the European Commission (Jäger, Moll, Som, & Zanker, 2015) uses one question from the 2009 European Manufacturing Survey (EMS) about robot utilization to estimate several related impacts (summarized above). The question regarding robot utilization on the EMS asks respondents to rate their use of robots on a three-point "low-medium-high" scale, so exact robot-to-job or robot-to-output associations are not possible.

The second example, Leigh and Kraft (2017), generated a robotics "census" by mining proprietary business databases for robotics firms in the U.S. The resulting dataset demonstrates the uneven geography of robotics suppliers and service providers (called *integrators,* discussed below) across the country. While the census demonstrates that robotics-related employment is geographically correlated with the manufacturing sector, the actual *use* of robots cannot be

[2] The Occupational Information Network (O*NET) taxonomy is based on the SOC taxonomy, and contains two expanded occupations that identify robot-related work: robotics engineers (17-2199.08) and robotics technicians (17-3024.01). However, neither the U.S. Census nor the Bureau of Labor Statistics (BLS) collects data at this level of disaggregation.

determined from these data. Therefore, the census cannot be used to make definitive links to changes in employment, wages, or productivity.

The other two studies (G&M and A&R) use data from The International Federation of Robotics (IFR), which has been tracking robot sales since the early 1990s. The IFR sells a longitudinal dataset of robot sales and robot stock (the latter is estimated by applying depreciation rates to previous years' sales) by country and industry for each year from 1993 to the present.

While the IFR data are the only of their type available, they come with several problems. The first is that the concept of "robot stock" is an approximation based on annual robot sales. These stocks are inferred by depreciating accumulated robot sales to each country based on the assumption of a 12-year working life of a robot—a number which the IFR itself acknowledges is uncertain (2017b). Robotics is a rapidly developing technology, so it is reasonable to assume that the rate of obsolescence is increasing.[3]

In fact, robot sales in North America (the bulk of which the U.S. is responsible for) experienced an abrupt discontinuity beginning in 2010, three years after the end of both Graetz and Michaels' (2015) and Acemoglu and Restrepo's (2017) time series (see Figure 1). In North America, annual sales records were set every year from 2010 to 2016 (Robotics Industries Association, 2017). However, the extent to which this period represents decisions to replace old robots or to add new robots where none had previously existed is unknown. Do these sales increases coincide with increases in robotics technology, declining costs, deepening of pent-up capital after the Great Recession, or some combination thereof? This sharp break in robotics sales patterns presents researchers with a difficult choice: limiting analyses to pre- or post-Great

[3] We note that Graetz and Michaels (2015) find no meaningful impacts to their results from the application of different rates of depreciation to baseline 1993 levels of robot stocks.

Recession robot diffusion trends satisfies an important condition for internal validity, but the results may not be generalizable to subsequent periods.

Figure 1: Orders and Shipments of Robots in North America, 1999-2017

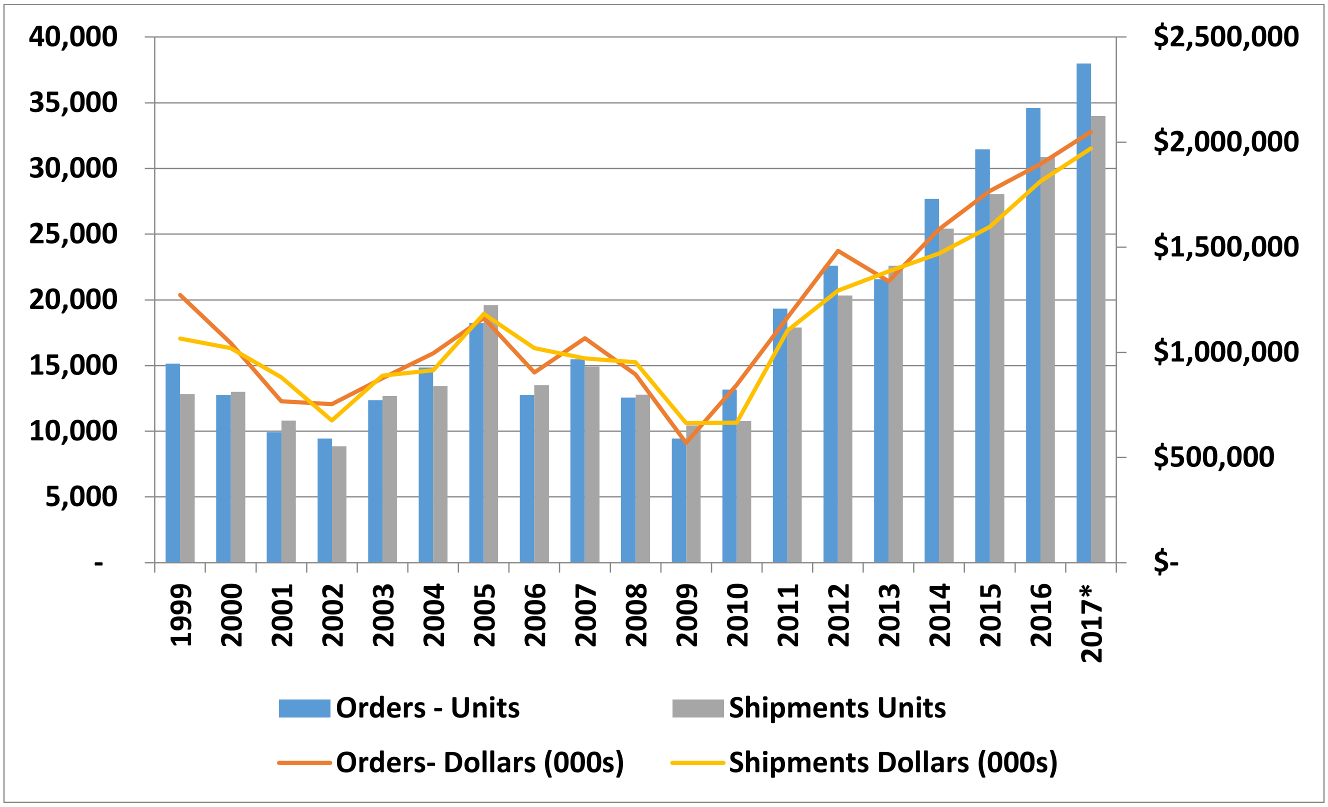

Chart Adapted from Robotics Industries Association; Includes U.S., Canada, and Mexico
**2017 projected*

Further, reporting sources have changed over the course of the IFR's data collection efforts, so annual robot stock figures are not fully comparable over time. Most significantly, no sales data were reported for North America until 2004 (International Federation of Robotics, 2017b), so any estimates of U.S. robots prior to that year must be derived from other nations' diffusion patterns.

Alternative Constructs of Robot Exposure and Their Limitations

Studies of the impacts technology defined broadly have been able to sidestep data problems by using expansive datasets in terms of both observations and time periods (e.g. Autor, Katz, and Kearney, 2003). These investigations are also afforded some leeway in conceptualizing the primary variable, technology, because "computerization" applies in the same general sense to all industries, from accounting and finance to manufacturing and agriculture, and for at least the second half of the 20^{th} century, computers have been so ubiquitously interwoven into daily life in both work and leisure that it makes little sense to try to theorize one overarching mechanism for how they "affect" a person's employment.

However, researchers interested in specific technological niches such as robots cannot resort to this rationale, and robot-specific impact studies must be more carefully operationalized. A&R's term, "Robot Exposure," is an intuitive and useful name for the concept at stake here, reflecting a concrete, epidemiological notion of exposure. However, as discussed above, available robot data do not necessarily provide ready analogues to epidemiological vectors.

Two approaches to measuring robot exposure have so far been used: Graetz and Michaels' "robot density," and A&R's aforementioned "robot exposure." Both derive worker-robot contact measures by creating a ratio of robots to human labor in each industry. Both use the IFR's robot stock estimates as the "robot" input. G&M's robot density uses "millions of hours worked" as the labor input, while robot exposure uses "thousands of workers."

In addition to underlying problems with robot data, both of these constructs, when operationalized empirically, employ questionable assumptions that threaten their construct validity. The first of these assumptions is that robots are uniformly diffused and employed, and the second is that their uses and impacts outside of the manufacturing sector are comparable to those within the sector.

Not All Robots are “Created” Equal

Even if we accept that robot stock data are, on average, reasonable estimates of the number of robots in use each year, the more problematic assumption is that *the number of robots in use is in fact the key determinant of other employment-related outcomes*. Robots have a wide variety of specifications and are used in a wide variety of applications. For example, some perform lightweight repetitive tasks at high speeds, while others may manipulate, weld, or cut large pieces of material more methodically. The number of possible combinations of relationships between robots, other automation equipment, workers, and output is essentially limitless. Condensing the complexity of the human-robot relationship to one simplified parameter is an outgrowth of the main problem of the task model upon which this research is based: generalizations that divide work—whether performed by humans or machines—into a few tractable categories break down at all but the most abstract and aggregate levels of analysis.

Another problematic abstraction arising from the “robot stock” variable is that the primary mechanism for industrial robot diffusion, robot *integration*, is overlooked. Industrial robotic systems are planned, engineered, installed, and often sold by robot integrators. Integration may be performed by in-house engineers and technicians, field engineering teams from robot suppliers, independent consultants, or some combination of these categories. Either way, integration is a substantial endeavor: Leigh and Kraft (2017) show that in the U.S., integrators at specialized automation engineering firms account for two-thirds of overall robotics-related employment.

Estimates for integration costs, however, are difficult to attain for several likely reasons. For one, integrators may not want to disclose costs publicly because they must routinely bid against other firms for new jobs. Also, the process of integration is highly variable and complex, preventing the establishment of standard pricing. Differing production processes, automation goals, and factory capabilities ensure that engineering challenges presented to integrators remain novel.

Robots, in and of themselves, are far from the sole determinant of the costs and consequences of robotic processes. An encouraging aspect of A&R's analysis is that their results stand up well to controls for capital expenditures—both those from IT specifically and the BLS's account of general capital stocks—and Leigh's and Kraft's (2017) measure of local robot-related employment, suggesting that the purchase of robots is independent, at the industry level, from other capital spending. However, it is far from clear how integrator costs are reported. Some of these costs are for software and computing (e.g. machine vision), while some are for hardware (e.g. physical robots and end-of-arm-tools). Some of these expenses are simply for the hourly costs of engineering services.[4]

The high variability in robotic automation processes—and thus diffusion and depreciation—confounds attempts to apply generalized robotic diffusion patterns over time to geographically specific places. A&R's model, which infers sub-national robot exposure patterns from national-level data weighted by commuting zone industrial mixes, is particularly vulnerable to this limitation. The "robot exposure" measure assumes that the nature and effects of robot use in a large auto plant in Michigan is essentially the same as it is in a large auto plant in South

[4] While the recent development of more user-friendly "collaborative robots" suggests movement toward more of a "plug-and-play" model, specialized engineering services will be required for most industrial automation for the foreseeable future.

Carolina, *and* that small and medium-sized auto suppliers use robots in a fashion similar to their larger counterparts across the country. This is problematic not only because of the variability in the integration process itself, but also because prior research suggests that geography is an important and more fundamental factor in firm integration of technology. For example, intra-metropolitan differences in locations of metalworking firms are highly predictive of whether they chose to adopt programmable automation technology (Harrison, Kelley, & Gant, 1996). Additionally, national culture strongly shapes attitudes towards technology and technicians within firms, thereby influencing the ultimate performance and productivity of the technology (Gertler, 1995).

Impacts of Robots Used Outside of Manufacturing

As it is applied in A&R, the "robot exposure" measure includes manufacturing subsectors at the three-digit NAICS level, and the two-digit Construction/Extraction, Utilities, and Education/Research sectors. However, outside of manufacturing, robots have minimal presence and are unlikely to have the same "effects" on workers as they do in manufacturing, where they have been systematically used for several decades. For example, in 2007, the IFR reported that 50 robots were used in the U.S. in the construction and extraction industry,[5] while workers in the industry, according to EUKLEMS, were about 7.8 million. Thus, we arrive at a robot exposure figure of .0064 robots per thousand workers.

The assumption that this tiny number of robots meaningfully affects labor market outcomes is untenable. The error compounds when the exposure factor is applied to all local U.S. labor markets. Fifty robots obviously cannot be distributed across 709 commuting zones; most

[5] There are only 20 robots officially specified as being used in this industry, but if we distribute the remaining "unspecified" robots among the various industries at levels proportional to the distribution of "specified" robots, as A&R do, there should be roughly 30 more in this category.

will have no construction robots—and thus no exposure. Because construction employment is a locally traded industry and relatively constant as a percentage of employment across labor markets—whereas manufacturing varies significantly—the "effect" of these fifty robots nationwide will be disproportionately large in most labor markets, and disproportionately small in the one or several that actually have these robots. The same problem applies to the utilities sector.

Moreover, the "education and research" sector should be excluded from the exposure measure for a different reason altogether, which is that educational robots do not fall within the same theoretical framework. Little, if any, robot related displacement would be expected in the sector. Educators increasingly use robots as instructional tools as the demand for robotics skills in the workplace rises. While a limited number of robots may be used in biomedical research labs in place of technicians, technology researchers conduct research on the robots themselves, suggesting that increased numbers of robots in research would be driven by more researchers and teams engaged in this area.

To date, the most impactful and ubiquitous types of non-manufacturing robots are those used in warehouses and distribution centers. However, these robots were not commercially available by the end of previous studies' time series. Further, their pattern of diffusion today is highly skewed. E-commerce giant, Amazon, bought the first-mover firm in warehouse robots, Kiva, for its exclusive use. Consequently, Amazon's estimated 100,000 warehouse robots (Wingfield, 2017) are still not included in IFR statistics, since Amazon does not sell its robotic systems on the open market. Other warehouse robot makers are starting to emerge in the market and should eventually part of IFR statistics, though the results will continue to under-represent warehouse robots until Amazon robots are incorporated.

Data and Methods

We use two approaches to address the problems with IFR robot stock data outlined above. The first approach modifies A&R's robot exposure variable, confining both the construction of the variable and the estimated impacts to the manufacturing industry. We call this the Modified Robot Exposure variable (MRE). The second approach uses an alternative data source, labor market data, to create an indicator for robot exposure.

While these approaches in this paper may be compared on conceptual grounds to those in the previous studies discussed above, they are not replications. A primary reason for this is that our labor market data (labeled RTLMI and explained below) begin in 2010, preventing the investigation of prior time periods. Rather, this paper should be viewed primarily as a comparison of two new measurement constructs of robotics' impacts.

Modified Robot Exposure

Modified Robot Exposure (MRE) limits the robot stock-derived exposure variable to the manufacturing industry, thereby eliminating the outsized influence of sectors that have minimal robot use but significant employment. Otherwise, it works similarly to the exposure measure created by A&R, with national-level industry concentrations of robots being applied at the local level. Rather than using commuting zones as labor market boundaries, we use U.S. Census defined Core Based Statistical Areas (CBSAs). CBSAs include metropolitan and micropolitan statistical areas, but exclude counties that are not part of an urban cluster of 10,000 people.[6] While commuting zones are more representative of actual labor markets and include more members of the U.S. workforce, we use CBSAs for two reasons. The first is that nonmetropolitan

[6] We use the terms "metro" and "metro area" interchangeably with CBSA, even though CBSAs include "micro" areas (those with urban clusters of at between 10,000 and 50,000 residents). Since for analytical purposes we limit observations to those with a minimum level of job postings, we are effectively only considering metropolitan areas.

and nonmicropolitan counties are subject to substantial suppression of employment data at the level of manufacturing subsectors, meaning that missing data would undermine accurate estimates for small counties. The second is that our labor market data are also biased toward urban areas (see subsequent section on Data Limitations); excluding nonmetropolitan areas maintains relatively consistency in the universes across independent and dependent variables.

Robotics Skill Demand Index

Using labor market data allows us to adjust, conceptually, the mechanism whereby jobs are impacted. With skill demand, the theorized impact no longer stems directly from the robot. Rather, it comes from the skills required by the presence of a robot. This slight shift in where the causal burden is applied adds a new dimension to existing robot impact research. We call the resulting measure the Robotics Skill Demand Index (RSDI).

The input for the RSDI comes from a novel source called Real Time Labor Market Information (RTLMI) to develop our robotics skills indicator. RTLMI is labor market data that are extracted from internet job postings on a continual basis. Of the several RTLMI vendors, we chose Burning Glass Technologies (BGT) for our study. Using proprietary algorithms, BGT collects online job postings on a daily basis, cleans the raw data, and formats it into a structured dataset. Although it has primarily been used by human resources departments and workforce and economic development agencies, it has emerged as a useful tool for academic researchers. Recent work has analyzed labor market skill demands (Deming & Kahn, 2017; Keith Wardrip, Stuart Andreason, & Zeeuw, 2017; Mason et al., 2016), the adoption of labor-saving technologies (Hershbein & Kahn, 2016), and job recruiting and searching activities (Banfi & Villena-Roldán, 2016; Modestino, Shoag, & Ballance, 2016).

We utilize a variable created by BGT called "skill clusters" to derive our composite robot indicator. These 559 skill clusters were defined by BGT by subjecting the text of job ads to k-means cluster analysis and additional qualitative scrutiny (Taska, 2017b). Each job ad is, in turn, associated with one or more skill clusters based on the text of the ad. Out of BGT's 559 clusters, there is one specifically called "robotics." Although the robotics skill cluster appears in 91,214 job ads in our dataset, this frequency is insufficient for regression analysis. Thus, to achieve enough robot-related skill observations for the RSDI, we developed a composite measure of skill clusters related to and including robotics. Skill clusters were selected for the RSDI by calculating the magnitude of correlation between the robotics skill cluster and all other skill clusters in the database for job ads in the manufacturing industry. The RSDI is comprised of the fifteen skill clusters with the highest correlation coefficients (>.04) (See Table 1).

The RSDI is further divided into "broad" and "narrow" definitions, because the skills contained in the index are themselves general and not uniquely applicable to robots. The broad definition includes all 23 skill clusters *and* the robotics skill cluster (24 total clusters), while the narrow definition includes only the top five (shown in gray) and the robotics skill cluster (6 total clusters). Dividing the RSDI into these two versions enables us to gauge their relative performance in empirical tests, and thereby to assess whether the RSDI classification scheme is diluted by casting too wide of a net.

The actual RSDI used in the regression analyses is calculated by dividing the number of job postings containing at least one skill cluster represented in Table 1 by the total number of manufacturing industry job postings within each Census-defined core based statistical area (CBSA) in the U.S. (separate RSDIs are calculated for broad and narrow definitions).

Table 2: Robotics Skill Demand Index Clusters Ranked by Correlation with Robotics Skill Cluster

Rank	Skill Clusters	Correlation	Frequency	% Posting
1	Computer-Aided Manufacturing	0.1351	487,373	5.0%
2	Mechanical Engineering	0.0954	560,776	5.7%
3	Welding	0.0878	373,363	3.8%
4	Drafting and Engineering Design	0.0858	703,780	7.2%
5	Engineering Software	0.0751	148,989	1.5%
6	Electrical Construction	0.0745	168,960	1.7%
7	Engineering Activities	0.0568	517,778	5.3%
8	Basic Electrical Systems	0.0567	373,693	3.8%
9	Schematic Diagrams	0.0532	239,218	2.4%
10	Electrical and Computer Engineering	0.0515	661,015	6.7%
11	Industrial Engineering	0.0492	380,644	3.9%
12	Machine Tools	0.0464	708,631	7.2%
13	Manufacturing Processes	0.0455	849,328	8.6%
14	Process Engineering	0.0437	405,092	4.1%
15	Machinery	0.0433	498,566	5.1%
16	Electrical Power	0.0401	89,161	0.9%
17	Equipment Repair and Maintenance	0.0392	389,998	4.0%
18	Circuitry	0.0370	163,410	1.7%
19	Industrial Engineering Industry Knowledge	0.0363	334,997	3.4%
20	Lean Manufacturing	0.0363	1166,877	11.9%
21	IT Hardware	0.0359	203,750	2.1%
22	Manufacturing Design	0.0309	43,161	0.4%
23	Signal Processing	0.0304	61,668	0.6%

Source: BGT
N=9,856,829
Robotics Skill Cluster is also included in both narrow and broad definitions

The low correlations between the skill clusters are a methodological artifact. That is, since the k-means clustering process is used to establish unique groupings, high correlation coefficients are not expected. This is particularly the case given the significant internal variation in the dataset stemming from its 559 skill clusters and 9,856,829 total manufacturing job posting observations.

Despite low statistical correlations, the skill clusters included in both the broad and narrow robot skill demand indices are qualitatively related to industrial robot use and representative of advanced manufacturing in general.

Data Limitations

As we have emphasized, no existing dataset is ideal for examining the impact of robotics. While RTLMI correlates over time with the U.S. Bureau of Labor Statistics' (BLS) Job Openings and Labor Turnover Survey (JOLTS)[7], it contains inherent biases. RTLMI data are not derived from standardized or mandatory surveys, so they are subject to changes in posting strategies by employers over time. Since RTLMI only contains online postings, it provides better coverage of jobs with higher skill and education requirements (Carnevale, Jayasundera, & Repnikov, 2014), as well as jobs located in urbanized areas. As such, it misses word-of-mouth jobs, and most significant for the current study, internally advertised union positions.

We employ several strategies to minimize the effects of these biases, including designing an RSDI measure that internalizes changes in overall job postings over time, and adding a unionization variable to the model. We also limit our empirical test to include only CBSAs with at least 1,000 manufacturing job postings from 2010 to 2015, to avoid potential bias caused by one or several large employers accounting for an outsized number of postings where online recruiting has limited penetration.

[7] Although they are not exactly comparable, JOLTS is the public time series that most closely resembles BGT. From January 2010 to January 2017 the two time series correlate closely, with a coefficient of 0.89 (Taska, 2017a).

Geography of Robot Skill Demand

The geographic distribution of robotics demand defined by RSDI corresponds broadly with other indicators of robot penetration. Figures 2a and 2b show that robotics job demand is concentrated on the coasts and the Midwest stretching south along the Interstate 75 corridor.

Higher numbers of robotics job demand are on the coasts, but the higher concentrations are in the interior. This greater intensity of robot skill demand in the Midwest and Southeast reflects the skill needs of the more manufacturing-centric labor markets in these areas.

Figure 2a: Robotics Skill Demand (narrow definition) in CBSAs, 2015 (Magnitude)

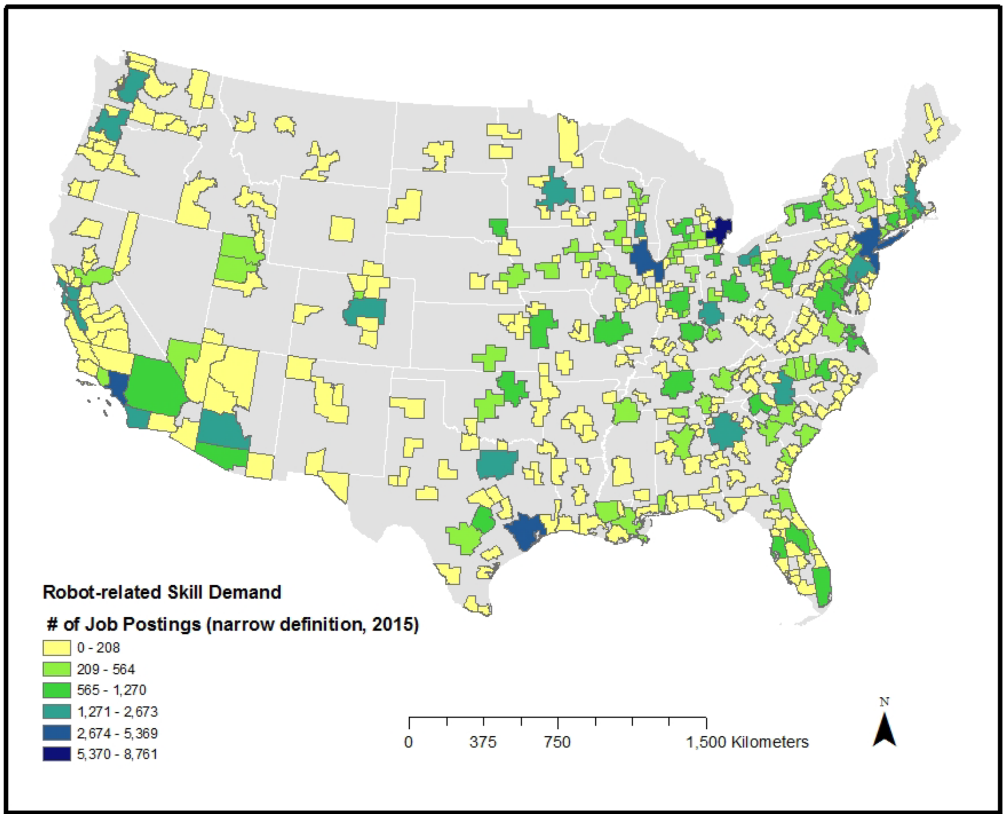

Source: Authors' calculation from BGT

Figure 2b: Robotics Skill Demand (narrow definition) in CBSAs, 2015 (Concentration)

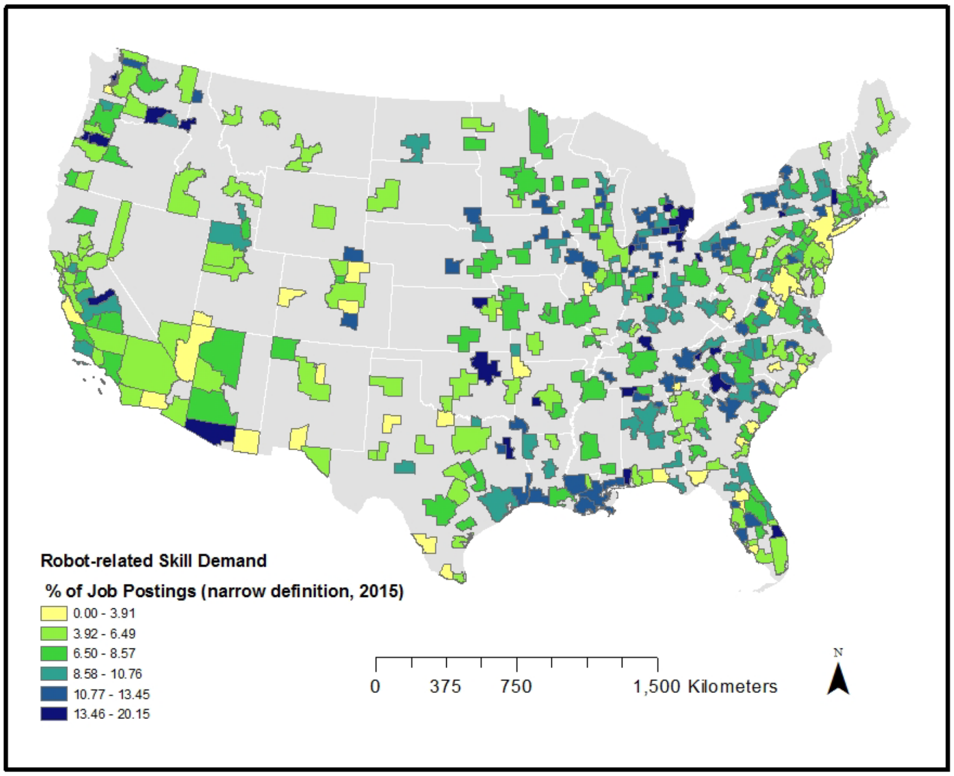

Source: Authors' calculations of BGT

Empirical Models

We employ two primary models to estimate the impact of robotics skill demand on regional manufacturing labor markets. The first, which we call the "initial conditions (IC) model" because its dependent variables are the starting conditions in metro areas, follows Glaeser et al (1992), where the associations between initial conditions of—in this case, robot skill demands—and outcomes (employment and wage growth) of regional labor markets, are tested. The model is based on a production function of output at a time t given by $A_t f(l_t)$, where A_t represents total productivity resulting from technological progress, such as increasing industrial robotics in the workplace, and l_t is a labor input. Assuming that the marginal product of labor is equivalent to wages, the equation can be rewritten in terms of growth rates, yielding equation (1), where w represents wages at time t. By setting $f(l) = l^{1-\alpha}, 0 < \alpha < 1$, we obtain our first empirical models for employment and wages presented in (2) and (3), respectively (Glaeser et al., 1992).

$$\log\left(\frac{A_{t+1}}{A_t}\right) = \log\left(\frac{w_{t+1}}{w_t}\right) - \log\left[\frac{f'(l_{t+1})}{f'(l_t)}\right] \tag{1}$$

$$\alpha\log\left(\frac{l_{t+1}}{l_t}\right) = -\log\left(\frac{w_{t+1}}{w_t}\right) + \log\left[\frac{A_{t+1}}{A_t}\right] \tag{2}$$

$$\log\left(\frac{w_{t+1}}{w_t}\right) = -\alpha\log\left(\frac{l_{t+1}}{l_t}\right) + \log\left[\frac{A_{t+1}}{A_t}\right] \tag{3}$$

A metropolitan area with an initial condition of a high concentration of industrial robots, A_t, is expected to grow faster in subsequent years than an area with a lower concentration in, due to productivity effects. This leaves open the possibility that employment, wages, or both, would also increase. However, it is also possible that employment would decrease due to strong substitution effects of robots, or that wages would stagnate for structural reasons not related directly to factory technology.

To test the hypotheses that employment and wages grow, we use RSDI and the "competing" measures of robot inputs as variables to explain the change in the ratio of manufacturing productivity in 2016 to 2010.[8] Our control variables for the models include Census division dummies, average employment and wage growth in metropolitan areas,[9] the size and concentration of metropolitan manufacturing industry,[10] the composition of durable manufacturing,[11] and a "right-to-work" legislation dummy,[12] since we expect union membership to exert a competing influence on our dependent variables.

[8] from Bureau of Economic Analysis
[9] from U.S. Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW)
[10] from QCEW
[11] from QCEW

Our final specifications, which predict employment and wage changes in manufacturing within Core Based Statistical Areas (CBSAs) are shown in equations (4) and (5), respectively, where X_i represents a vector of explanatory and control variables for a metropolitan area i.

$$\log\left(\frac{l_{i,t+1}}{l_{i,t}}\right) = \beta_0 + \sum \beta_k X_i + \varepsilon_i \tag{4}$$

$$\log\left(\frac{w_{i,t+1}}{w_{i,t}}\right) = \beta_0 + \sum \beta_k X_i + \varepsilon_i \tag{5}$$

There are two key limitations in the model specification. First, it does not allow for disaggregation of employment changes into their constituent parts. For example, it cannot determine the extent to which a net employment gain in a metro area due to productivity effects may be moderated by simultaneous substitution effects. Second, because the model framework only accounts for initial conditions of metropolitan areas observed in 2010, it cannot identify which non-technological aspects of industrial restructuring may have affected subsequent employment. For example, a regional transition into newer and more competitive industries may boost employment independently of whether existing firms upgrade their technology.

With these limitations in mind, we also test a first difference model (which we call the "difference model"), shown in (6) and (7). The models explain the log-transformed employment ratio as they correspond to the change in the intensity of industrial robots in regional production processes (note that equations (6) and (7) are identical to (4) and (5) except for the "delta" preceding the *X* variable, indicating the change in the robot skill demand index). By following changes over time in metro areas, these models should capture changes in industrial structure that may be concurrent with changes in process technology. The Leigh and Kraft (2017) local

[12] Right-to-work (RTW) laws present barriers to union organizing as part of a "business friendly" strategy. With the exception of Kentucky's county-specific laws, they are passed on a statewide basis. Since some metropolitan areas straddle more than one state, the core city was used to designate whether the CBSA was RTW or not during any part of the time frame, although it is acknowledged that statewide RTW status may be a factor in intra-regional competition. RTW status was derived from Peck (2016) and the National Conference of State Legislatures Right-to-work resource webpage (2017).

robot industry explanatory variable is not included in this model because it was only observed for a single time period. Control variables include the change in metropolitan employment and wage and a dummy representing change in right-to-work status from the same sources as the variables in the previous model.

$$\log\left(\frac{l_{i,t+1}}{l_{i,t}}\right) = \beta_0 + \sum \beta_k \Delta X_i + \varepsilon_i \qquad (6)$$

$$\log\left(\frac{w_{i,t+1}}{w_{i,t}}\right) = \beta_0 + \sum \beta_k \Delta X_i + \varepsilon_i \qquad (7)$$

Results

We display two main sets of results: employment outcomes from both the "Initial Conditions" (IC) and the "difference" models. Only the models with employment outcomes are shown, because none of the robot indicators register as statistically significant predictors of wage outcomes.

While the initial sample of metro areas includes 356 CBSAs based on the criterion of having at least 1,000 job listings over the study period in the labor market data, four additional outlying CBSAs were removed from the analysis because they were detected as unusual and influential cases based on several criteria, including Leverage (observations with extreme predictor values) and Cook's Distance (influence of a case on the predicted mean), and Studentized residual values (difference between predicted and observed).

The outliers warrant this brief discussion because they are all small metro areas, with relatively undiversified economies and less representation in RTLMI. Factors like a lack of industrial diversity or plant openings or closings will have outsized effects on input metrics for these communities. In general, our models were more accurate in predicting outcomes for large metros than for small metros.

Initial Conditions Model

For this sample of 352 CBSAs, the average Modified Robot Exposure (MRE) value is slightly over 13 robots per thousand workers. The average metro-wide value for the broad Robotics Skill Demand Index (RSDI) is 22.7% of job postings, while the more restrictive narrow RSDI is 9.3% of postings. With an average location quotient of 1.26, and an average percentage of employment in durable manufacturing of 55.4% in 2010, this is a relatively manufacturing intensive group of CBSAs, which is by design, as previously explained.

Table 3: Descriptive Statistics, Initial Conditions (IC) Model

Variables	Mean	Std. Dev.	Min.	Max.
ln (Manufacturing employment ratio (2016 / 2010))	0.070	0.123	-0.233	0.419
Relative employment size of local robot industry[a]	2.980	10.613	0.000	123.213
MRE[b] 2010	13.345	6.744	2.783	42.720
RSDI[c] Broad	22.697	8.829	0.000	58.824
RSDI[c] Narrow	9.309	5.824	0.000	32.787
ln (CBSA employment in 2016 / CBSA employment in 2010)	0.105	0.075	-0.159	0.371
ln (average wage in manufacturing in 2010)	10.850	0.214	10.250	11.865
Average years of education (2010)	26.466	8.449	10.912	57.503
Employment size (in millions, 2010)	0.027	0.064	0.001	0.898
Location Quotient in 2010	1.256	0.691	0.127	4.691
Percent durable manufacturing employment in 2010	55.422	17.764	7.195	93.716
Right to work legislation dummy	0.429	0.496	0.000	1.000

N = 352
a: Leigh and Kraft (2017) local robot industry variable
b: Modified Robot Exposure (robots per 1,000 workers)
c: Robot Skill Demand Index; Unit of measurement is percent

For the Initial Conditions model, the modified robot exposure variable (MRE) and the Robot Skill Demand Index (RSDI) both have significant, positive effects on employment (see Table 4). While the narrow RSDI has a slightly stronger relationship to the manufacturing employment ratio than does MRE, the broad RSDI has no discernible relationship. Because the dependent variable is in natural log form, a unit change in either the narrow RSDI or MRE can be interpreted to predict a .2% increase in manufacturing employment ratio. In other words, a CBSA with either a one percent above-average RSDI, or one more robot per thousand workers than average in 2010 can expect approximately .2% more manufacturing employment growth than average by 2016 (with the understanding that growth is defined relative to size in 2010).

Table 4: Initial Conditions Model; Dependent Variable = ln(Manufacturing Employment in CBSA 2016/Manufacturing Employment in CBSA 2010)

Variable	RSDI Broad skill definition					RSDI Narrow skill definition				
	Coef.	**Std. Err.**	**T-value**	**Sig.**	**Beta**	**Coef.**	**Std. Err.**	**T-value**	**Sig.**	**Beta**
Constant	0.956	0.322	2.970	***		0.893	0.324	2.760	***	
Relative employment size of local robot industry[a]	0.000	0.001	-0.220		-0.010	0.000	0.001	-0.370		-0.017
MRE[b]	0.002	0.001	2.430	**	0.121	0.002	0.001	2.390	**	0.120
RSDI[c]	0.002	0.001	2.850	***	0.136	0.001	0.001	1.400		0.065
Census division dummies (New England is reference)										
Middle Atlantic	0.008	0.033	0.240		0.019	0.013	0.033	0.390		0.031
East North Central	0.110	0.031	3.530	***	0.332	0.111	0.032	3.510	***	0.334
West North Central	0.052	0.035	1.490		0.119	0.050	0.036	1.420		0.114
South Atlantic	0.028	0.034	0.820		0.093	0.024	0.035	0.700		0.079
East South Central	0.070	0.035	1.980	**	0.148	0.070	0.035	1.960	*	0.148
West South Central	-0.028	0.038	-0.740		-0.074	-0.038	0.038	-0.980		-0.099
Mountain	0.067	0.035	1.940	*	0.162	0.062	0.035	1.780	*	0.149
Pacific	0.078	0.033	2.360	**	0.205	0.085	0.033	2.560	**	0.223
Ln (CBSA employment in 2016 / CBSA employment in 2010)	0.737	0.079	9.350	***	0.449	0.732	0.080	9.200	***	0.446
Ln (average wage in manufacturing 2010)	-0.095	0.032	-3.000	***	-0.164	-0.090	0.032	-2.830	***	-0.156
Average year of education (2010)	-0.003	0.011	-0.260		-0.015	0.000	0.011	-0.010		0.000
Employment size (in millions, 2010)	-0.150	0.090	-1.660	*	-0.078	-0.146	0.091	-1.600		-0.076
Location Quotient 2010	-0.009	0.010	-0.910		-0.049	-0.004	0.010	-0.470		-0.025
Percent of durable manufacturing employment 2010	0.000	0.000	-0.230		-0.012	0.000	0.000	-0.300		-0.015
Right to work legislation dummy	0.004	0.019	0.220		0.017	0.012	0.019	0.640		0.049

Number of Observations: 352; R-squared (adjusted): 0.3671 (0.3329), 0.3581 (0.3234)
a: Leigh & Kraft (2017) local robot industry variable
b: Modified Robot Exposure (robots per 1,000 workers)
c: Robot Skill Demand Index (RSDI)

There is also a notable geographic difference in the regional manufacturing employment response to increased robotics skill demand in 2010. The East North Central census division, which is largely contiguous with what is thought of as the historical U.S. manufacturing heartland, or the "rustbelt," sees the largest employment increases when holding 2010 robotics skill demand constant. There are likely regionally specific omitted variables contributing to these differences. The post-recession bailouts of the "Big 3" American auto companies and their subsequent employment rebounds (McNulty & Wisner, 2014) likely had some influence on the outsized manufacturing performance of the Midwest (East North Central), but would not have had as direct or significant of an effect on the East South Central and Pacific regions, each of which also showed strong manufacturing growth.

CBSAs with greater overall employment growth over the period had strong positive associations with robot skill demand (CBSA employment 2016/CBSA employment 2010), while those that paid higher initial manufacturing wages (average wage in manufacturing 2010) experienced depressed growth. The presence of right-to-work (i.e. anti-union) legislation showed no significant effect.

Difference Model

In contrast to the Initial Conditions (IC) model, which uses only baseline 2010 values of explanatory robotics variables as predictors, the Difference model takes into account how these variables changed over the study period.

These changes were significant. The average change in robot exposure (MRE) nearly doubled over the time period, increasing from 13.3 robots per thousand workers in 2010 to 26.3

robots per thousand workers in 2015 (an increase of nearly 13 robots per thousand workers; see Table 5). This change significantly outpaced the RSDI, which grew by 5.8% and 2.0% for broad and narrow definitions respectively over the same time period.

Table 5: Descriptive Statistics, Difference Model

Variable	Mean	Std. Dev.	Min.	Max.
Change in MRE[a] (2010-2015)	12.981	6.458	-4.178	53.248
Change in RSDI[b] (broad, 2010-2015)	5.783	9.452	-22.745	42.308
Change in RSDI[b] (narrow, 2010-2015)	1.987	6.317	-20.960	23.896
Ln (CBSA employment in 2016/2010)	0.105	0.075	-0.159	0.371
Ln (CBSA average wage in 2016/2010)	0.130	0.041	-0.060	0.294
Change in years of education	0.213	0.149	-0.360	0.716
Change in RTW legislation (from non RTW to RTW state)	0.105	0.307	0.000	1.000

N = 352
a: Modified Robot Exposure (robots per 1,000 workers)
b: Robot Skill Demand Index; Unit of measurement is percent

Results for the Difference model (see Table 6) show less agreement between the two primary explanatory variables. MRE's effect on manufacturing employment change is neither statistically significant nor of any measurable magnitude. The change in the narrow RSDI is the only robot-related variable that has an acceptably significant positive impact on growth in manufacturing employment in the Difference model.[13] The magnitude of this effect in the Difference model, .2%, is roughly equivalent to its magnitude in the IC model.

Change in overall metro employment and the passage of right-to-work legislation have positive effects on employment, although only three Midwestern states—Michigan, Wisconsin, and Indiana—changed their right-to-work status (from non-RTW to RTW) during the study period. As mentioned above, other regional idiosyncrasies, such as post-recession auto industry

[13] The broad RSDI, not shown, is significant at the 0.1 level, but has a nearly negligible coefficient of 0.0010072.

bailouts or concentrations of growing manufacturing sectors, may also have contributed to the above average employment growth. As the Initial Conditions model suggests, a state's RTW status in 2010 had no impact on manufacturing employment five years later.

Table 6: Difference Model, Narrow RSDI definition; Dependent Variable = ln(Manufacturing Employment 2016/Manufacturing Employment 2010)

Variables	**Coefficient**	**Std. Err.**	**T-value**	**Sig.**	**Beta**
Constant	-0.019	0.025	-0.750		
Change in MRE[a] (2010-2015)	0.000	0.001	-0.540		-0.026
Change in RSDI[b] (narrow, 2010-2015)	0.002	0.001	2.430	**	0.113
Ln (CBSA employment in 2016/2010)	0.694	0.081	8.540	***	0.423
Ln (CBSA average wage in 2016/2010)	0.176	0.151	1.160		0.058
Change in years of education	-0.064	0.039	-1.630		-0.077
Change in RTW legislation (from non RTW to RTW state)	0.081	0.019	4.260	***	0.202

Number of Observations: 352; R-squared (adjusted): 0.2621 (0.2493)
a: Modified Robot Exposure (robots per 1,000 workers)
b: Robot Skill Demand Index; Unit of measurement is percent

DISCUSSION

These results broadly indicate that for the post-recession time period of 2010-2016, robots and robotics-competent workers helped to bolster regional manufacturing employment. But we may ask what exactly does a .2% increase in the 2016-2010 manufacturing employment ratio mean? We can ground and contextualize these numbers by looking at the marginal effects of the explanatory variables in several hypothetical and real examples.

To do so, first we examine what happens to the theoretical mean case by applying coefficients from the Initial Conditions model and increasing the initial stock of robots and demand for robot skills. Table A2 in the appendix presents the descriptive statistics for the 352 metro areas included in the analysis, including average level of manufacturing employment, number of robots per thousand workers, and percent of jobs requiring robot-related skills in 2010. The mean CBSA would have begun the period with 26,929 manufacturing workers, 359 robots, and demanded robotics skills in 22.7% of its advertised manufacturing jobs. The model predicts that this mean metro area would have added 1,670 manufacturing employees by the end of 2015, for a 2016-2010 employment ratio of 1.081. A metro area identical to the mean in all other aspects but with one more initial robot per thousand workers (about 27 robots in this case) should see an increased employment ratio of 1.083 (1.081*1.002), or about 65 additional manufacturing workers (Table A3).

A one-standard deviation initial advantage in robot stock (equal to a total of 541 robots for a metro area of this size) predicts an employment advantage in 2016 of 436 jobs (Table A3). Even though the unstandardized coefficients for MRE and RSDI are essentially equivalent, advantages in robot skill demand (RSDI) have a slightly better job payoff, as suggested by its higher *beta* coefficient (see Table 4). A one-standard deviation initial advantage in robot skill demand (32% of metro-wide manufacturing job listings) should result in 487 more manufacturing jobs overall in 2016 (Table A4). Thus, this one-standard deviation above average metro area would have generated 922 more jobs (435+487) in five years due to its initial robotics advantage.

Note that these are the additional employees that are considered to be “due” to robots. In fact, the model suggests that promoting *general* regional economic growth, indicated by the all-sector employment ratio variable, remains the most impactful way to strengthen the manufacturing sector. In contrast, higher initial manufacturing wages depress future employment, a result not entirely unexpected as firms seek to locate in areas where labor costs are low. However, regions may not be able to rely on low wages as a business attraction feature for long, as the Difference model suggests that wage *growth* has a positive but not quite statistically significant effect on manufacturing employment growth (Table 6). More competitive—and more highly robotized—regional manufacturing sectors may require more highly skilled and better-compensated workers.

In reality, there is no such thing as a mean or average metropolitan area. Our model cannot capture all of the socio-political determinants that drive metro employment growth or decline. Metro areas that did exceedingly well or poorly in stewarding their manufacturing economies may have done so for reasons that our models fail to reflect because of omitted variables or because the assumptions of linearity and normality in traditional ordinary least squares (OLS) modeling do not apply to regional economic growth.

One way to account for these latent influences and problematic assumptions is to conduct a quantile regression, which fits data to specified quantile breakpoints of the dependent variable (2016-2010 employment ratio) rather than the mean, without discarding observations from other quantiles. Quantile regressions for the Initial Conditions model at the .75 and median (.5) levels generally agree with the OLS version of the model (see Table A5). The agreement is strongest for the .75 model, where the primary explanatory variables maintain high levels of statistical power and see increases in magnitude over the general OLS Initial Conditions model. However, MRE does not appear to be an influential factor for those in the bottom quarter of manufacturing growth. Instead, RSDI (broad) maintains strong or moderate statistical power (>10%) for all quantiles. This difference across categories and the low correlations between the two indicators (see Table A6 in appendix) suggests that they are indeed measuring different aspects of robotics use.

Finally, we test our Initial Conditions model against the actual employment changes in our sample metro areas. Using residuals, or differences in predicted versus observed values, we investigate qualitative patterns in the model's accuracy. It is reassuring that the model is fairly accurate for many large metro areas that we would expect to be influenced by robotics. Some of the largest metro areas with important manufacturing economies are within the top 20 percent of overall error (smaller residual absolute value = smaller error) (see Table A7). The model is very accurate in estimating San Francisco's[14] and Cincinnati's growth and does reasonably well for other important manufacturing hubs such as Los Angeles, Philadelphia, Seattle, New York, and Boston.[15]

To ground these predictions in terms of robots and employees, we use the Cincinnati metro area, for which our model predicts employment outcomes with a high degree of accuracy. Here, its 2010 advantages of 2.5 robots per thousand workers and 0.75% higher robotics skill demand are compared to a counterfactual case where both indicators were equal to the sample average. Doing so suggests that Cincinnati's initial robot stock (MRE) accounted for an additional 622 manufacturing workers and its initial skill demand (RSDI broad) accounted for another 162, for a total of 783 manufacturing workers over the subsequent six years.

Specialized manufacturing metros are also on the list of CBSAs that fit closely to the model, including Milwaukee, Dayton, Rochester (NY), and Syracuse. The three primary U.S. robotics research hubs—Boston, San Jose, and Pittsburgh—also have small differences between actual and predicted values. These 20 "well-fit" metro areas represented about 29% of total U.S. manufacturing employment in 2015.

---

[14] We use central city names to designate metro areas. For example, "San Francisco" represents the entire San Francisco-Oakland-Hayward, CA statistical area.

[15] Boston did not make the "top 20%" list (appendix Table A6) but would be the next metro in line if the list were expanded.

The most notably absent metros (i.e. metros where growth is poorly explained by the model) are four traditional “rust belt” communities: Buffalo, Cleveland, Detroit, and Chicago. These omissions are confounding because based on previous research into metro-level robotics-related employment (Leigh & Kraft, 2017), one could consider the stretch of Lake Erie shore from Buffalo to Detroit somewhat of a “Robotics Crescent” due to its high concentration of robotics-related business establishments and employment. Although not geographically contiguous to the Lake Erie group, Chicago is second overall in number of robotics-related establishments. Including the smaller metro areas between Buffalo and Detroit, the Lake Erie robotics crescent and Chicago together account for 65% of robotics-related employment (calculations based on unpublished data from Leigh and Kraft, 2017).

The fact that these key manufacturing economies do not conform to the model explains to a large extent why the local robotics industry (LRI) variable derived from Leigh and Kraft (2017) does not exert any statistical power in predicting employment changes. However, the large errors for these metros remain confounding. None of them have exceedingly high or low robot exposure or robotics skill demand index values, nor did they experience drastic changes to their manufacturing economies during the period. The model overestimated impacts in Detroit and Buffalo and underestimated Cleveland and Chicago, but no systematic pattern is apparent for why the model misses high on some metros and low on others. These discrepancies underscore the need for further geographically informed qualitative research into the economic impact of robotics.

CONCLUSION

This paper's unequivocal support for robots as supporting employment growth places it into the "robots-as-status-quo" category, because it supports the traditional theory that technology leads to net job growth. Its other contribution has been to critique and improve upon previously tested constructs for robot exposure. Three new constructs for measuring robot exposure were tested, and two of them—a robot stock variable modified from Acemoglu and Restrepo (2017) and a Robot Skill Demand Index derived for online job postings—are associated with employment gains since 2010. Estimates of these gains can be quantified to the effect of an approximate 900 manufacturing employee gain over six years for a one-standard deviation above average level of robot use and demand robotics-competent workers in manufacturing.

While these metrics are measuring different aspect of robot utilization, they fall short of both a complete accounting of robots and a rigorously valid operationalization of the concept of robot exposure. Until better robotics data exist, the economic impact of robots will have to be estimated via inference and triangulation, as well as more rigorously informed by how robots are used and diffused *in addition* to economic theory. Finally, because of the limited diffusion of commercial robotics and other smart machines, geography must be closely considered before making generalizations. Regional economies in the U.S. demonstrate significant variation in their response to robotization.

This paper contributes to efforts to identify the economic impacts of robots by a) restricting analysis to the manufacturing sector, which is the only sector where robots are historically and systematically used, b) triangulating robot stocks and skill demand to derive a more holistic measure of robotics influence, and c) qualitatively disaggregating impacts according to geography. Data limitations confine the analysis to a unique six-year period during which the U.S recovered from an historic recession. Thus, whether the findings are unique to this period, or indicators of a long term positive relationship between employment and robotics, cannot be answered at this point in time. However, the findings clearly indicate the need for more research, data, and theory to answer critical questions about the future of work and all of the ensuing implications of that future for local and regional economies and society.

References

Acemoglu, D., & Autor, D. (2011). Skills, Tasks, and Technologies: Implications for Employment and Earnings. *Handbook of Labor Economics, 4*.

Acemoglu, D., & Restrepo, P. (2017). Robots and Jobs: Evidence from U.S. Labor Markets. *Working Paper*.

Adler, P. S. (1992). Introduction. In P. S. Adler (Ed.), *Technology and the Future of Work*. Oxford: Oxford University Press.

Autor, D. H., & Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *The American Economic Review, 103*(5), 1553-1597. doi:10.2307/42920623

Banfi, S., & Villena-Roldán, B. (2016). Do high-wage jobs attract more applicants? Directed search evidence from the online labor market. *CEA WP, 327*.

Braverman, H. (1998). *Labor and monopoly capital: The degradation of work in the twentieth century*: NYU Press.

Brynjolfsson, E., & McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W. W. Norton & Company.

Cappelli, P. (2012). *Why good people can't get jobs: The skills gap and what companies can do about it*: Wharton Digital Press.

Cappelli, P. (2014). *Skill gaps, skill shortages and skill mismatches: evidence for the US*. Retrieved from

Carnevale, A. P., Jayasundera, T., & Repnikov, D. (2014). Understanding Online Job Ads Data. *Georgetown University, Center on Education and the Workforce, Technical Report (April)*.

Deming, D., & Kahn, L. B. (2017). *Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals*. Retrieved from

Gertler, M. S. (1995). "Being There": Proximity, organization, and culture in the development and adoption of advanced manufacturing technologies. *Economic Geography, 71*(1), 1-26. doi:10.2307/144433

Glaeser, E. L., Kallal, H. D., Scheinkman, J. A., & Shleifer, A. (1992). Growth in cities. *Journal of Political Economy, 100*(6), 1126-1152.

Goos, M., Manning, A., & Salomons, A. (2009). Job Polarization in Europe. *The American Economic Review, 99*(2), 58-63. doi:10.2307/25592375

Gordon, R. J. (2012). *Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds*. NBER Working Paper Series. National Bureau of Economic Research. Cambridge, MA. Retrieved from http://www.nber.org/papers/w18315.pdf?new_window=1

Gordon, R. J. (2014). The Demise of U.S. Economic Growth: Restatement, Rebuttal, and Reflections. *National Bureau of Economic Research Working Paper Series*.

Graetz, G., & Michaels, G. (2015). *Robots at work*. London School of Economics Centre for Economic Performance. London School of Economics. Retrieved from http://cep.lse.ac.uk/pubs/download/dp1335.pdf

Grint, K., & Woolgar, S. (2013). *The machine at work: Technology, work and organization*: John Wiley & Sons.

Harrison, B., Kelley, M. R., & Gant, J. (1996). Innovative Firm Behavior and Local Milieu: Exploring the Intersection of Agglomeration, Firm Effects, and Technological Change. *Economic Geography, 72*(3), 233-258. doi:10.2307/144400

Hershbein, B., & Kahn, L. B. (2016). *Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings*. Retrieved from

Howell, D. R. (1985). The Future Employment Impacts of Industrial Robots: An Input-Output Approach. *Technological Forecasting and Social Change, 28*, 297-310. doi:http://dx.doi.org/10.1016/0040-1625(85)90032-0

Hunt, H. A., & Hunt, T. L. (1983). *Human resource implications of robotics*. Kalamazoo, Mich.: W.E. Upjohn Institute for Employment Research.

International Federation of Robotics. (2017a). *World Robotics Statistics*.

International Federation of Robotics. (2017b). *World Robotics: Industrial Robots 2017*. Retrieved from

Jäger, A., Moll, C., Som, O., & Zanker, C. (2015). *Analysis of the impact of robotic systems on employment in the European Union*. Retrieved from Luxembourg: http://publica.fraunhofer.de/dokumente/N-366982.html

Keep, E., & Mayhew, K. (2010). Moving beyond skills as a social and economic panacea. *Work, Employment & Society, 24*(3), 565-577.

Keith Wardrip, Stuart Andreason, & Zeeuw, M. d. (2017). *Uneven Opportunity: Exploring Employers' Educational Preferences for Middle-Skills Jobs*. Retrieved from

Leigh, N. G., & Kraft, B. R. (2017). Emerging robotic regions in the United States: insights for regional economic evolution. *Regional Studies*, 1-13. doi:10.1080/00343404.2016.1269158

Leonardi, P. M. (2012). Materiality, Sociomateriality, and Socio-Technical Systems: What do These Terms Mean? How Are They Different? Do We Need Them? In J. K. Paul M. Leonardi, Bonnie A. Nardi (Ed.), *Materiality and Organizing: Social Interaction in a Technological World* (pp. 25-48). Oxford: Oxford University Press.

Mason, J. L., Johnston, E., Berndt, S., Segal, K., Lei, M., & Wiest, J. S. (2016). Labor and skills gap analysis of the biomedical research workforce. *The FASEB Journal*.

McNulty, P. A., & Wisner, J. D. (2014). The impacts of the 2008 government bailouts on the US automobile industry. *The Journal of Human Resource and Adult Learning, 10*(2), 20.

Miller, S. M. (1989). *Impacts of industrial robotics : potential effects on labor and costs within the metalworking industries*. University of Wisconsin Press: Madison, Wis.

Modestino, A. S., Shoag, D., & Ballance, J. (2016). Downskilling: changes in employer skill requirements over the business cycle. *Labour Economics, 41*, 333-347. doi:http://dx.doi.org/10.1016/j.labeco.2016.05.010

National Conference of State Legislatures. (2017). Right-to-work Resources. Retrieved from http://www.ncsl.org/research/labor-and-employment/right-to-work-laws-and-bills.aspx

Orlikowski, W. J., & Scott, S. V. (2008). Chapter 10: Sociomateriality: Challenging the Separation of Technology, Work and Organization. *Academy of Management Annals, 2*(1), 433-474. doi:10.1080/19416520802211644

Peck, J. (2016). The Right to Work, and the Right at Work. *Economic Geography, 92*(1), 4-30. doi:10.1080/00130095.2015.1112233

Robotics Industries Association. (2017). Industry Statistics. Retrieved from https://www.robotics.org/Industry-Statistics

Taska, B. (2017a, August 16). [Chief Economist, Burning Glass Technologies, Personal Communication].

Taska, B. (2017b, November 15). [Chief Economist, Burning Glass Technologies. Personal Communication].

Wingfield, N. (2017, September 10). As Amazon Pushes Forward With Robots, Workers Find New Roles. *New York Times*. Retrieved from https://www.nytimes.com/2017/09/10/technology/amazon-robots-workers.html

APPENDIX

Table A1: Outliers

CBSA	2010 Manufacturing Employment	Residual
Bloomington, IL	4,565	-0.471563
Bangor, ME	3,922	-0.370401
Gulfport-Biloxi-Pascagoula, MS	20,665	-0.369981
Elizabethtown-Fort Knox, KY	5,239	0.2969948

Table A2: Average Metropolitan Area Characteristics Based on Sample of 352 CBSAs, excluding outliers

Statistic	Value
Manufacturing Employment 2010	26,929
Manufacturing Employment 2016	29,110
Manufacturing Employment Change	2,181
Manufacturing Employment Ratio 2016-2010	1.081
MRE (Robots/Thousand Mfg Employees)	13.345
Number of Robots	359
MRE Standard Deviation	6.744
RSDI (Percent Job Ads Demanding Robot Skills)	22.697
RSDI Standard Deviation	8.829

Table A3: Employment Predictions for Advantages in Modified Robot Exposure (MRE) Coefficients based on Average CBSA

	1 Unit Advantage	.5 Standard Deviation Advantage	1 Standard Deviation Advantage
Robots/Thousand Mfg Employees	14.345	16.717	20.089
Number of Robots	386	450	541
Marginal Effect on Employment Ratio	.0024	0.0081	0.0162
Predicted Employment Ratio	1.083	1.089	1.097
Predicted 2016 Employment	29,180	29,329	29,476
Predicted Change in Employment 2010-2016	2,245	2,399	2,618
Predicted Increase over Average	65	218	436

Table A4: Employment Predictions for Advantages in Robotics Skill Demand Index (RSDI) Coefficients based on Average CBSA

	1 Unit Advantage	.5 Standard Deviation Advantage	1 Standard Deviation Advantage
Pct Job Ads Demanding Robot Skills	23.697	27.112	31.526
Marginal Effect on Employment Ratio	0.0021	0.0091	0.0181
Predicted Employment Ratio	1.0831	1.090	1.099
Predicted 2015 Employment	29,166	29,355	29,598
Predicted Change in Employment 2010-2015	2,237	2,426	2,669
Predicted Increase over Average	55	245	487

Table A5: Median, 0.25, and 0.75 Quantile Regression Results with RSDI (broad)

Variable	**Median**			**0.25 Quantile**			**0.75 Quantile**		
	Coef.	**Std. Err.**	**Sig.**	**Coef.**	**Std. Err.**	**Sig.**	**Coef.**	**Std. Err.**	**Sig.**
Constant	1.075	0.382	***	0.911	0.474	*	1.147	0.429	***
Relative employment size of local robot industry[a]	-0.001	0.001		0.000	0.001		-0.001	0.001	
MRE[b]	0.004	0.001	***	0.001	0.001		0.005	0.001	***
RSDI[c] (Broad)	0.001	0.001	*	0.002	0.001	**	0.002	0.001	*
Census division dummies									
Middle Atlantic	-0.009	0.039		0.033	0.048		0.015	0.044	
East North Central	0.095	0.037	**	0.126	0.046	***	0.142	0.042	***
West North Central	0.050	0.042		0.075	0.052		0.070	0.047	
South Atlantic	0.033	0.041		0.047	0.051		0.077	0.046	*
East South Central	0.050	0.042		0.063	0.052		0.116	0.047	**
West South Central	-0.046	0.045		-0.008	0.056		0.039	0.051	
Mountain	0.075	0.041	*	0.031	0.051		0.148	0.046	***
Pacific	0.068	0.039	*	0.069	0.049		0.096	0.044	**
Ln (CBSA employment in 2016 / 2010)	0.758	0.093	***	0.931	0.116	***	0.837	0.105	***
Ln (average wage in manufacturing in 2010)	-0.104	0.037	***	-0.106	0.047	**	-0.085	0.042	**
Average year of education (2010)	-0.006	0.013		0.000	0.016		-0.023	0.014	
Employment size (in millions, 2010)	-0.152	0.107		-0.078	0.133		-0.217	0.120	*
Location Quotient in 2010	-0.001	0.011		0.004	0.014		-0.024	0.013	*
Percent of durable manufacturing employment in 2010	0.000	0.000		0.001	0.001		0.000	0.000	
Right to work legislation dummy	0.001	0.023		0.013	0.028		-0.010	0.026	

Number of Observations: 352

Table A6: Correlation Matrix for Robotics Variables

	LRI	MRE	RSDI (Broad)	RSDI (Narrow)	Change in MRE	Change in RSDI (Broad)	Change in RSDI (Narrow)
Local Robotic Index (LRI)	1.000						
Modified Robotic Exposure (MRE)	0.051	1.000					
Robotic Skill Demand Index (Broad)	-0.032	-0.019	1.000				
Robotic Skill Demand Index (Narrow)	-0.004	-0.035	0.561	1.000			
Change in MRE	0.055	0.247	0.104	0.054	1.000		
Change in RSDI (Broad)	-0.022	0.112	-0.560	-0.288	0.058	1.000	
Change in RSDI (Narrow)	0.007	0.076	-0.321	-0.668	0.030	0.531	1.000

Table A7: Top 20% Residual Values (i.e. top 20% best most accurately predicted CBSAs)

	CBSA	Residual		CBSA	Residual
1	Yakima, WA	1.03175E-05	36	San Jose-Sunnyvale-Santa Clara, CA	0.011785216
2	San Francisco-Oakland-Hayward, CA	0.000189208	37	Albany, OR	0.011820175
3	Cincinnati, OH-KY-IN	0.000513808	38	Syracuse, NY	0.012002562
4	Texarkana, TX-AR	0.001188924	29	Longview, WA	0.012249806
5	San Luis Obispo-Paso Robles-Arroyo Grande, CA	0.001496651	40	Odessa, TX	0.012304634
6	Ann Arbor, MI	0.002391054	41	Sioux City, IA-NE-SD	0.012550895
7	College Station-Bryan, TX	0.002544946	42	Niles-Benton Harbor, MI	0.012762969
8	Sioux Falls, SD	0.003116937	43	Watertown-Fort Drum, NY	0.012857568
9	San Antonio-New Braunfels, TX	0.003280213	44	Denver-Aurora-Lakewood, CO	0.013225475
10	Winston-Salem, NC	0.003332793	45	Amarillo, TX	0.013607227
11	El Paso, TX	0.003455844	46	Birmingham-Hoover, AL	0.013759348
12	Kennewick-Richland, WA	0.003461991	47	Staunton-Waynesboro, VA	0.013829897
13	Burlington, NC	0.003698296	48	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	0.014453323
14	Chattanooga, TN-GA	0.003859612	49	Des Moines-West Des Moines, IA	0.014817915
15	Altoona, PA	0.004337472	50	Lawrence, KS	0.015194151
16	Abilene, TX	0.004532591	51	Greeley, CO	0.015499496
17	Ames, IA	0.004608882	52	Bay City, MI	0.015677979
18	Providence-Warwick, RI-MA	0.005144149	53	Gainesville, FL	0.015926976
19	Cape Girardeau, MO-IL	0.005387467	54	Seattle-Tacoma-Bellevue, WA	0.015951125
20	Augusta-Richmond County, GA-SC	0.00562124	55	Rochester, NY	0.016237415
21	Greenville, NC	0.00657293	56	La Crosse-Onalaska, WI-MN	0.016431658
22	Worcester, MA-CT	0.006616897	57	Lansing-East Lansing, MI	0.017301565
23	Scranton--Wilkes-Barre--Hazleton, PA	0.006959064	57	Portland-Vancouver-Hillsboro, OR-WA	0.017497735
24	Fargo, ND-MN	0.007693685	59	Jefferson City, MO	0.017784689
25	Auburn-Opelika, AL	0.008083295	60	New York-Newark-Jersey City, NY-NJ-PA	0.017909208
26	Lafayette-West Lafayette, IN	0.0085481	61	Dayton, OH	0.017957604
27	Los Angeles-Long Beach-Anaheim, CA	0.00887201	62	Jonesboro, AR	0.018643463
28	Pittsburgh, PA	0.008976536	63	Spokane-Spokane Valley, WA	0.019244101
29	Asheville, NC	0.009470141	64	Milwaukee-Waukesha-West Allis, WI	0.01958066
30	Lewiston-Auburn, ME	0.009529696	65	Winchester, VA-WV	0.01965549
31	McAllen-Edinburg-Mission, TX	0.009705436	66	Tampa-St. Petersburg-Clearwater, FL	0.019716749
32	Appleton, WI	0.009782419	67	Orlando-Kissimmee-Sanford, FL	0.019751681
33	Jacksonville, FL	0.010018547	68	Pittsfield, MA	0.020242767
34	Charlottesville, VA	0.010373693	69	South Bend-Mishawaka, IN-MI	0.02053442
35	Lancaster, PA	0.011180566			

GLOSSARY OF KEY TERMS AND ABBREVIATIONS

Difference Model
Model in which primary explanatory robot indicator variables reflect changes from 2010 to 2015.

Employment Ratio
Measure of manufacturing employment change in a CBSA, equivalent to Manufacturing employment in 2016 / Manufacturing employment in 2010

Initial Conditions (IC) Model
Model in which primary explanatory robot indicator variables reflect initial conditions in 2010.

Modified Robot Exposure (MRE)
Indicator of robot exposure based on robot exposure indicator created by Acemoglu and Restrepo (2017) using International Federation of Robotics robot stocks data. MRE is limited to the manufacturing sector, while A&R's original robot exposure variable includes all sectors of the economy.

Real Time Labor Market Information (RTLMI)
Data derived from online job postings, provided by Burning Glass Technologies.

Robotics Skill Demand Index (RSDI)
Index created by measuring percentage of online job ads compiled in RTLMI within a core-based statistical area that list robotics or robotics-related skills as qualifications.

- *Narrow Definition* – restrictive measure of robotics skill demand, including the robotics skill cluster and additional five skill clusters most highly correlated with robotics (skill clusters created by Burning Glass Technologies).
- *Broad Definition* – Expanded measure of robotics skill demand, including six skill clusters in narrow definition and the next 18 skill clusters ranked in order of correlation with the robotics skill cluster.

