

# The Employment Consequences of Robots: Firm-level Evidence

## **Abstract**

As a new general-purpose technology, robots have the potential to radically transform industries and employment. In contrast to prior studies at the industry level that predict dramatic employment declines, we find that investments in robotics are associated with increases in total firm employment, but decreases in the total number of managers. Similarly, we find that robot investments are associated with an increase in the span of control for managers remaining within the organization. We also provide evidence that robot adoption is not motivated by the desire to reduce labor costs, but is instead related to improving product and service quality. Our findings are consistent with the notion that robots reduce variance in production processes, diminishing the need for managers to monitor worker activities to ensure production quality. Decreases in managerial headcount may also arise from changes in workforce composition. We find that investments in robotics are associated with decreases in employment for middle-skilled workers, but increases in employment for low- and high-skilled workers, potentially changing managerial activities required by the firm. With respect to organizational change, we show that robots predict both centralization and decentralization of decision-making authority, but decision rights in either case are reassigned away from the managerial level of the hierarchy. This result contrasts with prior studies of information technology that have generally found decentralizing effects on decision authority within organizations. Overall, our results suggest that the impact of robots on employment and organizational practices is more nuanced than prior studies have shown.

# 1 Introduction

We explore how employment and organizations have changed in response to robot adoption. As robotics and artificial intelligence (AI) are increasingly used by firms as the next engine of innovation and productivity growth, their effects on labor, firm practices, and productivity have become a subject of growing importance. Anecdotal evidence in the popular press has documented extensively that robots reduce overall employment and exacerbate income inequality, as rapid advancements in vision, speech, natural language processing and prediction capabilities have achieved parity with or exceed human capabilities across a range of tasks. These technological advancements have shifted the comparative advantage from humans to machines for a growing list of occupations (Brynjolfsson and Mitchell 2017, Felten et al. 2019, Frey and Osborne 2017), potentially leaving human labor with substantially fewer activities that can add value (Brynjolfsson and McAfee 2014, Ford 2015). This technology-based labor substitution may displace a significant fraction of the overall workforce despite generating productivity gains (Acemoglu and Restrepo 2017, Autor and Salomons 2017, Ford 2015). If true, robots are likely to lead to significant changes in how firms organize production activities and manage their human capital (Bidwell 2013, Puranam, Alexy, and Reitzig, 2014, Zammuto et al. 2007).

Recent empirical studies using data at the industry or geographic region levels show that robots are associated with drastic declines in overall employment (Acemoglu and Restrepo 2017, Dinlersoz and Wolf 2018, Graetz and Michaels 2018, Mann and Püttmann 2017). However, robots have also been argued to be similar to past generations of general-purpose technologies (GPT) that ultimately increased labor demand. In this competing view, even as labor is displaced, the new jobs created will more than compensate for the jobs lost (Autor and Salomons 2017). Preliminary evidence using firm-level data supports this view, and finds that robot-adopting firms become more productive and ultimately increase total employment (Koch et al. 2019). These new jobs are likely to complement robots, suggesting a compositional change in labor within firms. As robots offer new capabilities that differ from prior information technology (IT) investments (Brynjolfsson and Mitchell 2017), changes in human capital and the organization of production activities may also differ from those caused by IT and reflect those that are complementary to robots.

Using comprehensive data on businesses in the Canadian economy spanning the years 2000-2015,<sup>1</sup> we show that robots are associated with increases in total employment, but the effect is not uniform across workers. Investments in robotics predict substantial declines in managerial employment, despite increases in non-managerial employment. This finding contrasts with prior information technologies that could not

---

<sup>1</sup> We use two main datasets for our empirical analysis spanning overlapping timeframes, described in more detail in the data section.

easily replace managerial and professional work (Autor et al. 2006, Autor et al. 2003, David and Dorn 2013, Dustmann et al. 2009, Murnane et al. 1999). We find evidence that robots may affect managerial employment in two ways. First, robots may directly reduce the need to monitor and supervise workers when robots can substantially diminish human errors in the production process.<sup>2</sup> Given that supervision of workers accounts for a substantial portion of work done by managers (Hales 1986), demand for managerial labor to supervise workers may decline with robot adoption. Second, robots may also indirectly affect managerial employment through changing the types of workers needed. Although the total number of non-managerial employees increase with robot adoption, we also find that robot investments predict decreases in employment of middle-skilled workers and increases in employment of low- and high-skilled labor. These changes in labor composition may result in a reduction in managers (Malone 2003, Mintzberg 2013). Consistent with our findings of an increase in non-managerial employees and a decrease in the number of managers, we find that robot investments predict an increase in the span of control for managers remaining within the organization.

In examining the motivations for robot adoption by firms, we find that robot investment is not associated with the strategic importance of reducing labor costs, but is instead associated with an increase in the strategic importance of improving product and service quality. For the allocation of decision authority within organizations, we find that robot investments predict both centralization and decentralization of decision rights away from the managerial level of the hierarchy. This suggests that not only has managerial headcount decreased, but their decision authority is also diminished. This result differs from earlier studies that found IT generally led to decentralization of decision rights (Acemoglu et al. 2007, Bresnahan et al. 2002). Overall, our results show that changes in employment are related to complementary changes in organizational practices that are critical to the effective use of robots.

To the best of our knowledge, this study provides the most comprehensive evidence at the level of individual businesses on the employment and organizational effects of robot investments. The wide range of outcomes we examine—employment, labor composition, span of control, strategic priorities, and the allocation of decision rights—suggest that robots have a substantive effect on both employment and the organization of production in ways that differ from prior technologies. Our analysis also provides a deeper data-driven examination of how robots can change employment and organizational practices that are difficult to capture using country and industry level data (Raj and Seamans 2018). More broadly, our results suggest that detailed exploration at the level of individual organizations can provide useful insights in contributing to the important debate about the consequences of robots for labor and organizations.

---

<sup>2</sup> <https://www.indeed.com/career-advice/what-does-a-production-supervisor-do>

## 2 Theoretical considerations

The adoption of GPTs is often associated with substantial and widespread productivity gains across different sectors of the economy (Bresnahan and Trajtenberg 1995). To maximize the value of GPTs, firms must substantially reorganize their work activities and change the nature of work and human capital requirements (Autor et al. 2003, Bresnahan et al. 2002, Brynjolfsson et al. 2018). As a recent and rapidly proliferating GPT (Brynjolfsson et al. 2018, Cockburn et al. 2018), robots have the potential to transform employment, firm practices, and the economy (Agrawal et al. 2018, McAfee and Brynjolfsson 2017).

### 2.1 Robots and total employment

The effect of robots on employment remains an open question. Research examining the effect of robots on labor is still nascent, with only a few studies examining the substitutability of robots on work (Acemoglu and Restrepo 2017, Arntz et al. 2016, Frey and Osborne 2017, Mann and Püttmann 2017, Manyika 2017). However, most of these preliminary studies predict dire consequences of labor displacement from robot adoption. For example, Frey and Osborne (2017) find that up to 47% of all jobs in the United States may be displaced. Using a task-based approach breaking each occupation into a set of concrete tasks, OECD researchers find that 70% of tasks performed by labor could be automated (Arntz et al. 2016). Other studies using the task-based approach have concluded that more than 50% of work tasks are vulnerable to automation (Manyika 2017), leading to both labor displacement and wage reductions (Bessen et al. 2019). Using a measure of robot penetration at the industry level in the US, Acemoglu and Restrepo (2017) find that one robot can replace roughly six people. Graetz and Michaels (2018) also find that robot adoption is associated with a reduction in hours worked for low-skilled labor, using similar data on robot adoption for 17 countries.

The findings from these initial studies stand in stark contrast to earlier generations of technologies that have been found to increase employment in conjunction with productivity, ultimately leading to labor's share of productivity remaining constant. Instead of reducing employment, robots may also positively affect employment through 1) productivity increases from labor substitution inducing demand for other goods and services that require non-automated tasks; 2) capital deepening that increases the effectiveness of robots, which can increase productivity without further reducing labor; or 3) the creation of new tasks or increased demand for existing tasks that are complementary to robots (Acemoglu and Restrepo 2018, Brynjolfsson et al. 2018). Initial results from surveys of Spanish manufacturing firms suggest that organizations that adopt robots experience both productivity and employment gains (Koch et al. 2019).

These differing results are in part due to difficulties in observing these countervailing effects in an entire economy using data at the industry and geographic region levels. Studies at these levels of analysis cannot clearly examine how firms use robotics to substitute or complement labor. As prior literature

examining the link between IT and productivity has shown, analysis at more aggregated levels can often lead to markedly different conclusions from empirical studies conducted at the firm level (Bresnahan et al. 2002, Brynjolfsson and Hitt 1996). These differences can arise due to the substantial heterogeneity in productivity growth across firms that cannot be clearly observed at the industry level or other aggregated levels of analysis (Syverson 2004). For example, robot adopting firms may experience productivity and employment gains while non-adopting firms in the same industry experience employment and productivity losses. If true, even if robots are observed to cause employment losses at the industry level, it remains unclear whether robots displace workers within robot-adopting firms or if workers are instead displaced in non-adopting firms due to a decrease in competitiveness. Without clarity behind these underlying mechanisms, meaningful inferences become particularly challenging, with similar empirical issues hampering early attempts to understand the effects of IT investment on organizations. Ultimately, more precise measurement of both IT and organizational capabilities at the firm level was critical to resolving the IT-productivity paradox that earlier studies discovered, and to uncovering the factors explaining the heterogeneous effects of IT on firm outcomes (Brynjolfsson et al. 2002). With a firm-level measure of robot investments for the population of firms in Canada, we empirically investigate the competing hypotheses of whether robot-adopting firms increase or decrease employment in firms.

**H1a: Robot investments are associated with increases in total employment.**

**H1b: Robot investments are associated with decreases in total employment.**

## **2.2 Robots and non-managerial employment**

Regardless of the effect on total employment, workforce composition is likely to change with robot adoption as demand for different skills changes within the firm, similar to prior generations of skill-biased technical change. For example, the rise of IT in the late 1990s led to a reduction in the demand for low- and middle-skill occupations as routine tasks became automated, and a corresponding increase in demand for nonroutine and cognitively challenging tasks including managing employees (Autor et al. 2006, Autor et al. 2003, Card and DiNardo 2002, Murnane et al. 1999). Similar to these studies, we define low-skilled workers as those working in occupations requiring a high school degree or less; middle-skilled workers as those working in occupations requiring vocational/trades accreditation or an associate degree; and high-skilled workers as those working in occupations requiring at least an undergraduate university degree. Although nonroutine and cognitively challenging tasks have been argued to be difficult to automate (Autor et al. 2003, Murnane et al. 1999), the increasing sophistication of robots is likely to automate tasks that were previously unaffected by automation.

With advances in vision, speech, and prediction capabilities, robotics has advanced beyond automating simple routine tasks and become capable of performing more cognitively complex work, as well as tasks

involving specific types of manual dexterity. Middle-skilled workers are more likely to perform these tasks, which robots are more able to automate. For example, in the healthcare and pharmaceutical industries, robots have been used to handle and prepare materials, follow complex protocols to prepare and analyze samples, and help coordinate patient care without human intervention (Gombolay et al. 2018). Firms with significant warehousing operations have also experienced similar effects. Robots have automated a large range of warehousing logistics activities by effectively transporting objects between locations without human intervention. By relieving humans from lifting and handling awkward, heavy objects during inventory management, robots not only avoid injuries but also provide consistency in product quality and decrease overall delivery time.<sup>3</sup> In manufacturing, industrial robots can substantially reduce variance in product quality. Machine vision enables robots in the automotive industry to consistently install and weld parts onto car bodies with a high degree of precision, minimizing errors in the production process.<sup>4</sup> This can involve difficult manual manipulations such as 360-degree multi-arm rotations with many repetitions. Robots can be programmed to perform these tasks precisely over a long period of time. Accordingly, robots can substantially reduce both unintended human errors such as those arising from fatigue, as well as deliberate actions such as gaming production quotas that have previously impeded productivity and effective management (Helper and Henderson 2014).

These illustrative examples suggest that robots can automate certain complex tasks that were primarily the responsibility of middle-skilled workers, including technicians, machinists, and operations personnel from a variety of industries that are responsible for following complex protocols to ensure production quality. These tasks may also involve certain types of manual dexterity that require significant learning over time for humans. With robots, many of these tasks can be automated using algorithms, eliminating human errors and the need to provide training for these skills. By reducing production quality variance, robots can decrease the demand for middle-skilled work, given the vulnerability of these tasks to robot-based automation.

## **H2: Robot investments are associated with decreases in middle-skilled employment.**

However, investments in robotics may also create demand for human labor and tasks that complement robots. While demand for middle-skilled work may decrease through direct substitution, demand for complementary work, either lower or higher-skilled, may increase with robot adoption. For firms that redesign their production processes to leverage the capabilities that robots can offer, productivity may increase, ultimately leading to increases in employment for specific types of workers. Despite recent technological advances, robots are often unable to fully automate most production processes; for many of these “residual tasks,” human labor remains a more efficient and cost-effective solution (Autor et al. 2003,

---

<sup>3</sup> <https://www.nytimes.com/2017/09/10/technology/amazon-robots-workers.html>

<sup>4</sup> <https://blog.robotiq.com/bid/69722/Top-5-Robotic-Applications-in-the-Automotive-Industry>

Brynjolfsson and Mitchell 2017). For example, Elon Musk famously scaled back investments in automation in Tesla's factory and reintroduced human workers after too much automation slowed the production of the Model 3 electric vehicle and delayed its market launch.<sup>5</sup> In order to effectively utilize robots, human capital must also be reorganized and reassigned to aid production. As an example, Amazon significantly redesigned work in its warehouses to effectively utilize its Kiva Robotic systems. Robots are used to travel between locations within the warehouse, but human workers pick and pack products delivered by the robots. In this case, instead of having middle-skilled workers managing inventory by walking from shelf to shelf to examine and handle products, robots and algorithms can automate this process and bring inventory directly to human workers, who then pick them up and place them into shipping boxes. Researchers have also systematically matched occupations to what machine learning can do and find that many of the manual skills performed by low-skilled labor cannot be easily replaced using technology (Brynjolfsson and Mitchell 2017, Felten et al. 2019). While machine learning is not identical to robot technologies, robotics rely heavily upon on machine learning to make inferences and thus can serve as a useful indicator of how robots may affect work.

Current evidence suggests that although robots can increase manual dexterity for certain tasks, they cannot yet effectively perform many manual tasks that humans can easily do. Accordingly, productivity increases from robot investments will lead to increases in demand for low-skilled workers doing these residual tasks.

### **H3: Robot investments are associated with increases in low-skilled employment.**

Demand for high-skilled workers may also increase with robot adoption. Similar to the example of how Amazon reorganized warehouse work activities after robot adoption, the majority of productivity gains from technology adoption come from the complementary redesign of work (Bresnahan et al. 2002, Hammer 1990). Implementing the necessary process improvements and work reorganization requires highly skilled professionals (Bresnahan et al. 2002, Hammer 1990, Helper and Henderson 2014, Huselid and Becker 1997, Ichniowski et al. 1997), some of whom are needed to program, repair, customize and work with robots (Acemoglu and Restrepo 2017, Autor and Salomons 2018, Brynjolfsson and Mitchell 2017). However, demand for high-skill workers may also increase for those that do not directly work with robots, as automating certain routine tasks may free up resources to engage in more cognitively complex tasks. For example, when hospitals adopt robots to lift patients out of beds, nurses are not only relieved of the physical strain from tasks that are prone to injuries, but also given more time to interact with patients and participate in clinical treatment (Gombolay et al. 2018). Similarly, by algorithmically providing pills and other medications directly to the patients (Bepko Jr et al. 2009), nurses can spend more time ensuring compliance

---

<sup>5</sup> <https://www.theverge.com/2018/4/13/17234296/tesla-model-3-robots-production-hell-elon-musk>

and making other clinical decisions. In the context of manufacturing, when much of the routine production process is done by robots and low-skilled labor, this can free up time and resources for high-skilled professionals to design and market new products and optimize production processes (Felten et al. 2019). Programmable robots can also increase a firm's flexibility in serving different types of orders and in providing a greater range of products. This can further increase the demand for higher-skilled workers who can design a greater range of products. Consistent with these arguments, Autor and Dorn (2009) find that investments in computer technologies over the last several decades contributed to the widespread increase in high-skilled jobs engaged in creative, problem-solving, and coordination tasks. Similarly, Felten et al. (2019) find that investments in artificial intelligence are correlated with increased employment of high-skilled workers such as software engineers. Thus, we expect employment of high-skilled workers to also increase after robot adoption.

#### **H4: Robot investments are associated with increases in high-skilled employment.**

### **2.3 Robots and managerial employment**

Managerial employment may also change significantly with robot adoption. When production is automated using robotics, human errors are substantially reduced and variance in production quality decreases (Verl 2019). Unlike humans, robots can precisely perform the same complex protocol repeatedly for long periods of time without experiencing fatigue, leading to both productivity increases as well as fewer errors in the production process. Agency problems arising from information asymmetries also do not exist with robots, as robots do not operate for self-interested gains the way humans might in work settings (Eisenhardt 1989, Hong et al. 2019, Jensen and Meckling 1976). Given the substantial costs of employee monitoring for firms (Dickens et al. 1989, 1990) and considerable time spent by managers monitoring employee activities (Hales 1986, 1999), the adoption of robots in production can substantially reduce the need to closely monitor work effort and quality. Through both a reduction in variance in the production process and the lack of agency costs associated with managing robots, the level of monitoring required to ensure production quality is likely to decline. Given that monitoring and control constitute a significant portion of managerial activities (Kolbjørnsrud et al. 2016), the demand for managerial labor is likely to decrease after robot adoption.

While robots can reduce demand for managers through decreasing the need to monitor employees in production, they may also affect managerial work by changing the composition of non-managerial employees within the organization. If robot adoption is associated with a decline in middle-skilled workers and an increase in high- and low-skilled workers, managerial activities may change for the newly transformed workforce. Managing low-skilled workers can differ substantially from managing other types of employees as low-skilled work is typically more standardized, and consequently easier to monitor and



evaluate than their higher skilled counterparts ( Mintzberg 1980, Perrow 1967). Also, an individual manager can potentially supervise many more employees if digital tools automate aspects of monitoring standardized work. For example, technology can be used to organize and report the output of simple routine tasks and even make predictions about work outcomes (Aral et al. 2012), especially for standardized work where inputs and outputs can be specified and clearly measured (Brynjolfsson and Mitchell 2017). In the case of Amazon, the productivity of warehouse workers is tracked in real time and an automated system generates recommendations for employee warnings and terminations when productivity targets are not met.<sup>6</sup> Having an objective measure of productivity recorded using automation technology also reduces disruptive conflicts between managers and subordinates, since objective productivity measures are more difficult to dispute (Scully 2000, Wu 2013). As the fraction of low-skilled workers in the organization's workforce increases, fewer managers may be needed within the organization.

In addition to differences in managing low-skilled work, managing highly skilled professionals is also likely to differ from managing middle-skilled workers. High-skilled workers often engage in more cognitively challenging tasks which provide higher added value, such as product design and production optimization. Managing these types of workers is likely to differ substantively from managing workers doing routine manual tasks (MacDuffie 1997, Parker and Slaughter 1988). Supervising low- and middle-skilled workers primarily involves ensuring that employees arrive on time, verifying compliance with rules and regulations, monitoring their work procedures and output, issuing commands, and training them to do the job properly (Helper and Henderson 2014, Taylor 1977). By contrast, employees doing more cognitively complex work are often experts themselves in dealing with problems outside of routine operations, and can resolve production problems better than their managers (Helper et al. 2000, Kenny and Florida 1993). These employees are often empowered to make more decisions because they are more capable of solving relevant problems than their managers (Huselid and Becker 1997, Ichniowski et al. 1997). As a consequence, managing these employees may entail less direct issuing of commands and more advising and empowerment of employees to solve problems (Malone 2003, Mintzberg 1973, Mintzberg 2013).

While we expect span of control to increase for managing low-skilled workers, the expected change is ambiguous when subordinates are high-skilled workers. If workers require more advising and coaching from managers, managerial span of control may decrease (Malone 2003, Malone 2004). High-skilled workers have also been argued to pose unique challenges to the efficiency of organizational hierarchies due to greater needs for communication and conflict resolution, which can be mitigated by decreasing span of control (Bell 1967, Meyer 1968). However, effective utilization of high-skilled labor often leads to granting employees greater autonomy (Bresnahan et al. 2002), potentially leading to increases in span of control

---

<sup>6</sup> <https://www.inc.com/suzanne-lucas/amazon-fires-hundreds-via-computer-algorithm-im-okay-with-that.html>

(Simon 1946). Prior literature examining the relationship between skill composition changes and span of control in the presence of technology adoption has been limited, but available evidence generally finds net positive effects on span of control (Scott et al. 1994).<sup>7</sup> If decreases in the demand for managerial labor from reduced monitoring requirements and skill composition changes dominate potential increases due to productivity gains, demand for managerial labor may ultimately decline. Given these arguments, we expect that managerial employment will decrease with robot adoption.

**H5: Robot investments are associated with decreases in managerial employment.**

### 3 Data and Measures

#### 3.1 Data

To measure robot investment at the firm level, we use data capturing the purchases of robots imported by Canadian firms provided by the Canadian Border Services Agency (CBSA) from 1996 to 2017. Global production of robotics hardware is highly concentrated in relatively few countries including Japan, Germany, the United States and increasingly China. By contrast, Canada does not produce a meaningful quantity of robotics hardware domestically and consequently must import robots from foreign producers, allowing us to exploit data on import transactions to measure robot adoption by firms. For all import transactions, the CBSA classifies goods according to Harmonized System (HS) codes, and classifies industrial robots separately from other types of technologies, machinery, and equipment.<sup>89</sup> In addition to the HS code, the name of the exporting firm, product country of origin, name and address of the importing firm, business number of the importing firm (a unique government-issued identifier for Canadian businesses) and value of the transaction are recorded. As an additional validity check of our measure of robot investment, we benchmark our measure to data reported by the Robotics Industry Association (RIA), and find both measures are comparable, showing similar trends over time (see detailed discussion in Appendix section S1).

Since we are using import data, the definition of robots is ultimately based upon what types of import transactions are being classified as “robots.” As a starting point, we note that the International Federation of Robotics (IFR) defines industrial robots as having the characteristics of being 1) automatically controlled, 2) reprogrammable, 3) a multi-purpose manipulator in three or more axes, and 4) used in industrial automation applications. The IFR provides a number of examples of robots and what their primary

---

<sup>7</sup> Bloom et al. (2014) also find a positive relationship between IT investment and span of control, but were not able to observe corresponding changes in employee skill composition.

<sup>8</sup> Industrial robots are a separate classification at the ten digit HS code level recorded by the CBSA.

<sup>9</sup> The classification details several different types of robots as distinct HS codes, which can be grouped into two consistent categories across the time period of our data: 1) robots for automotive assembly lines and 2) all other types of industrial robots. Our measure of total robot investment is the sum of these two categories of robots.

functions are in both their published material and on their website, which include activities such as assembly, welding, painting, packaging, picking and placing, and handling materials for metal casting. In principle, firms that are members of IFR-affiliated industry associations are likely to define “robots” to be consistent with the IFR definition.

To examine our measure of robotics investment in greater detail, we manually conducted searches in the public domain for transactions accounting for 95% of the total value of robot purchases in our data. Members of IFR-affiliated industry associations (e.g., Robotics Industry Association, Japan Robot Association) accounted for 58.4% of the total value of imports in our data. Firms that were not robotics association members but that advertised selling the same type of robots accounted for another 13.3% of the value. Most often these firms specialized in installing and integrating robots actually produced by association members. For an additional 2% of the total value of transactions, the exporting firms were not affiliated with a robotics industry association but manufactured robots for scientific laboratories. Imported mainly by firms in the healthcare industry in our data, these robots automate a variety of repetitive tasks in biology and chemistry research, such as pipetting.

An additional 19.0% of the total value was attributable to importing firms in industries that use robots intensively: primarily the automotive industry, but also machine tools and plastics manufacturing. Some firms in these industries are members of robotics industry associations, but our data has more comprehensive coverage of firms that invest in robots. Given the well-documented prevalence of their use in these industries and from examining the types of robots used by the importing firms in these transactions, we were able to infer that these transactions reflected investments in robotics similar to transactions involving robotics association members. For the remaining 2.3%, firm websites confirmed robots being utilized across a variety of activities, including performing repairs, handling materials in hazardous environments such as pipelines or nuclear power plants, as well as in construction and demolition.

We merge our robot investment data with two datasets maintained by Statistics Canada containing measures of firm characteristics: 1) the National Accounts Longitudinal Microdata File (NALMF), a panel dataset that contains measures of aggregate firm-level employment and economic inputs derived from tax filing data from 2000-2015; and 2) the Workplace and Employee Survey (WES), developed and administered by the Business and Labour Market Analysis Division and the Labour Statistics Division at Statistics Canada.<sup>10</sup> The WES consists of both an employer component which contains comprehensive information on employment and management practices at the organizational level, and a linked employee component measuring individual-level job characteristics and activities. The employer survey sample is a random stratified sample in a panel structure, representative of the population of business establishments in

---

<sup>10</sup> Datasets are linked by organization-year using unique identifiers provided by Statistics Canada.

the Canadian economy in each year.<sup>11</sup><sup>12</sup> For the employee sample, individual employees were randomly chosen within each organization and surveyed for two consecutive years, with Statistics Canada resampling individuals from each organization after each 2 year cycle was completed. The WES employer survey data we use spans the years 2001-2006, while the employee survey data we use follows employees during the years 2001-2002 and 2003-2004.<sup>13</sup>

We make several adjustments to both our NALMF and WES samples to more precisely capture those firms of sufficient size that purchased robots with the intention of implementing them as an end user for production. Here, we only include firms with at least ten employees, and removed those firms in the finance and insurance (NAICS code 52) and real estate rental and leasing sectors (NAICS code 53), as firms in these sectors were found to be primarily involved in leasing robots to other firms and comprised a negligible percentage of total robot imports into Canada. We also removed firms in service industries that were engaged in programming imported robots for the purpose of reselling them to other firms (NAICS codes 5413, 5414, 5415, 5416), and firms in the wholesale trade sector (NAICS code 41). In our final data used for analysis, our NALMF sample contains 168,729 firms in total, our WES employer sample contains 3,981 businesses establishments, and our WES employee sample contains 7,958 individual employees.

### **3.2. Robot capabilities**

Based on our examination of individual robot import transactions, we found that robots are especially active in the automotive and machinery and equipment assembly sectors (see Appendix Figure A4), plastic processing industries, and in metal and manufacturing. In automotive manufacturing, robots are usually organized along a structured assembly line to fetch and position parts; fasten, rivet or weld parts together; and apply coatings and/or paint to the assembled parts. Robots are also prominent in the electronics assembly industry, where “pick and place” robots select circuits and place them on circuit boards or silicon wafers. They handle small, delicate parts with precision, selecting among different types and pressing them on to circuit boards. They can also visually inspect circuit boards, as well as test the connections, and can also be involved in etching circuit boards. Robots may also be involved in packaging finished products. In addition to improving quality, one of the main motivations for adopting robots in the electronics industry

---

<sup>11</sup> We note that the WES employer survey data is at the establishment level (approximately 95% of firms in the Canadian economy have only one establishment), and we conduct additional checks with Statistics Canada to ensure our results are robust to this issue.

<sup>12</sup> An important strength of the WES is that responding to the survey was mandatory under Canadian law, which resulted in regular response rates of approximately 90 percent, mitigating concerns of non-response bias in our analysis.

<sup>13</sup> In the NALMF data, there are 5,180 firms that adopt robots during the time period of the data (1.4% of all firms). In the WES data, there are 48 organizations that adopt robots during the time period of the data (0.85% of businesses in the sample)

is the increase in flexibility in serving different orders, switching from large volume orders to smaller batches.

Robots are also used extensively in the processing of plastics, where they primarily perform secondary machine tending roles. They also apply labels and move parts to other areas where they are further modified or packaged for shipment. In the injection molding of plastic parts and packaging materials, they are also used to pick items and as label applicators. Overall, in the plastic processing industry, robots can replace a substantial part of repetitive manual labor.

In minerals and metals manufacturing, robots are involved in loading and unloading metal blanks into computer numerical control (CNC) machine tools, repositioning semi-finished parts during the machining process, and deburring afterwards. A primary motivation for robot adoption by firms in die-casting industries is the improvement of worker safety. Foundries are dangerous work environments where robots (or workers) are subjected to intense heat and toxic fumes. The molded parts would then need to be cooled, modified, and inspected. Robots can control for quality in all these steps. When the quality of molded parts depends upon the skill of individual workers, robots offer much greater consistency. Individuals working with robots are also able to work much more safely and efficiently.

In addition to these industry specific applications, palletizing is a ubiquitous application that robots can facilitate across many industries. Robots can recognize, pick up, orient and stack packages on pallets. They can move easily between packages of different sizes and varieties of different quantities. Combined with the ability to control for quality, robots can efficiently place items in packages, seal them, and label them with machine readable codes. This not only increases efficiency and precision, but also reduces injuries associated with palletizing large objects.

### **3.3 Measures**

Here, we describe the measures used for our main baseline tests.

*Robot investment.* Using our data capturing imports of robotics hardware, we create a measure of robot capital stock by adding all robot purchases by each firm recorded in each year. To adjust our robot capital stock measure for economic depreciation, we assume a useful life of 12 years based upon stated guidance given by the International Federation of Robotics (IFR).

*Employee count, hiring, and departures.* To measure the total number of employees within the firm, we use the total count of employees provided in the NALMF data for each firm-year, recorded from payroll deduction remittance forms submitted by all Canadian firms to the Canada Revenue Agency (CRA). Total numbers of managerial and non-managerial employees are recorded as responses in each year of the WES

employer survey.<sup>14</sup> The total number of new employee hires and departures are also recorded for each year of the survey data for both managerial and non-managerial employees. Non-managerial employee headcount is also reported by skill type, including middle-skilled, low-skilled, and high-skilled workers.<sup>15</sup>

*Strategic importance of labor cost reductions and quality improvements.* To measure the strategic importance of labor cost reductions and quality improvements to the firm, we exploit a section of the WES employer survey which asks respondents to “please rate the following factors with respect to their relative importance in your workplace general business strategy” for the years 2001, 2003, and 2005. Respondents are asked to choose the importance of each factor on a Likert scale with possible responses being (1) Not applicable, (2) Not important, (3) Slightly important, (4) Important, (5) Very important, and (6) Crucial. Here, we consider the factors of “reducing labour costs” and “improving product/service quality” separately for analysis. For our measure of strategic priority of each factor, we redefine values of (2) on the Likert scale to be equal to (1) and reset the scale to be ascending from 1 to 5, as an increase from the original (1) to (2) and vice versa does not clearly capture the changes in strategic priority that we aim to measure.<sup>16</sup>

*Decision authority for training and choice of production technology.* The WES employer survey data contain detailed information regarding decision-making authority for tasks across different layers of the organizational hierarchy, drawn from survey questions similar to those used by Bresnahan et al. (2002) and Bloom et al. (2014) measuring worker autonomy. The survey asks, “who normally makes decisions with respect to the following activities?” Here, we consider the activities of “training” and “choice of production technology” as they are directly relevant to the firm’s investments in human capital and use of robotics for productivity. For the 2003 and 2005 waves of the survey, survey respondents were given the following five possible responses to the question of who makes decisions: 1) non-managerial employees, 2) work supervisors, 3) senior managers, 4) individuals or groups outside the workplace (typically corporate headquarters for multi-establishment firms), and 5) business owners. To create distinct categories that correspond to hierarchical levels within organizations, we create three dummy variables, each equal to one if: 1) non-managerial employees were assigned decision authority over the task, 2) work supervisors or senior managers were assigned authority over the task, to capture managerial employees, or 3) business owners or corporate headquarters were assigned authority over the task.

---

<sup>14</sup> The survey provides a variety of examples of what is included in the definition of managers: “Examples: president of a single location company; retail store manager; plant manager; senior partners in business services firms; production superintendent; as well as vice-presidents, assistant directors, junior partners and assistant administrators whose responsibilities cover more than one specific domain, department heads or managers (engineering, accounting, R&D, personnel, computing, marketing, sales, etc.); heads or managers of specific product lines; junior partners or assistant administrators with responsibilities for a specific domain; and assistant directors in small locations (without an internal department structure).”

<sup>15</sup> Definitions of middle, low, and high-skilled workers in the survey match the definitions stated earlier in the paper.

<sup>16</sup> This modification does not change the sign or significance of our results from using each original variable. Simply dropping all values of (1) also produces results of identical sign and significance level.

*Supervisor span of control.* To capture supervisor span of control, the WES employee survey asks individual respondents whether they “supervise the work of employees on a day-to-day basis,” and if so to report the total number of employees who either directly report to them or who report to their subordinates. Here, we use this total count as our measure of supervisor span of control, and consider only those managers who are not promoted during the two-year period they are followed in the data.<sup>17</sup>

*Unpredictability of work schedule.* To assess the unpredictability of the work schedule of employees, the WES employee survey asks respondents “how far in advance do you know your weekly hours of work?” with possible responses being (1) always known; (2) more than one month (more than 31 days); (3) one month (22 to 31 days); (4) 3 weeks (15 to 21 days); (5) 2 weeks (8 to 14 days); (6) 1 to 7 days; and (7) Less than one day. For our main measure of work schedule unpredictability, we use the numerical value associated with each response, with increasing values denoting a shorter time period where employees know their work schedule in advance.<sup>18</sup>

*Controls.* A number of control variables are also included in our analysis. In all our NALMF and WES employer sample specifications, we include organization fixed effects to address concerns of unobserved heterogeneity across firms and year fixed effects to control for aggregate shocks and trends. In our WES employee sample regressions, we also estimate models including individual employee fixed effects. We control for organization size, measured by logged total assets in our NALMF sample, logged total revenues in our WES employer sample, and logged total employees in our WES employee sample.<sup>19</sup> We also include a dummy variable control for firms that have multiple business units in our NALMF sample, or organizations that are part of a multi-establishment firm in our WES employer sample. In our WES employer sample analysis, we include separate dummy variables to control for business establishments that have an organized union, or implement outsourcing as an organizational change.<sup>20,21</sup>

---

<sup>17</sup> Promotions were recorded as a separate question in the WES employee survey. Including supervisors who were promoted and controlling for promotions with a dummy variable also produces results of identical sign and significance level.

<sup>18</sup> To ensure our results are not driven by scaling differences at values of (1), (2), and (7), we also repeat our estimation including only those values ranging from (3) to (6) and find similar results.

<sup>19</sup> Logged total assets is used in the NALMF sample since logged total revenues is a left hand side variable in our productivity regressions. Logged total revenues is used in our WES employer sample since changes in employment are part of our main left hand side variables being tested. Logged total employees is used in the WES employee sample as a more precise control for firm size that may affect our span of control dependent variable.

<sup>20</sup> These variables were not available in the NALMF data.

<sup>21</sup> Regarding outsourcing, the survey specifically asks each year “Has your workplace experienced any of the following forms of organizational change?” with “Greater reliance on external suppliers of products / services (outsourcing)” as a possible response. Because organizational changes are likely to be quasi-fixed in the short time period of our data, the dummy variable control remains equal to one in subsequent years after outsourcing is reported.

## 4 Patterns and trends in robot adoption

Figure 1 shows aggregate robot capital stock in Canada for each year from 1996-2017.<sup>22</sup> Overall, investment has been steadily increasing since the late 1990s, with a substantial decline in investment growth corresponding roughly to the timeframe of the Great Recession.<sup>23</sup> Since 2014, investment in robotics has again continued to increase. Since 2007, robot adoption has also proliferated more broadly across different sectors of the economy, with the most dramatic growth coming from outside the automotive sector (see Appendix Figures A4 and A5).

## 5 Empirical strategy

A primary concern in estimating the effect of robotics is that robot adoption is unlikely to be random, potentially biasing our coefficient estimates. We address this issue in two ways, in addition to our robustness tests. First, for our total employment regression (using the NALMF sample) we instrument for robot investment using the percentage of workers in each 4-digit NAICS code in occupations with high “manual dexterity” and low “verbal ability” in 1995 multiplied by the inverse of the median price per robot in Canada for each year.<sup>24,25</sup> Measures of occupation-level manual dexterity and verbal ability are obtained from the Career Handbook 2003, a dataset created by Employment and Social Development Canada, which contains ratings of the level of manual dexterity and verbal ability associated with over 920 distinct occupations on a four-point scale. We define high and low levels as the top and bottom two points on the scale, respectively. The median price per robot for each year in Canada is calculated from the import data provided by the Canadian Border Services Agency (CBSA). The percentage of workers in each 4-digit NAICS code in occupations with high manual dexterity and low verbal ability in 1995 provides a cross-sectional measure of industries that have a higher proportion of workers that may engage in activities that more closely match the capabilities of robots, which is multiplied by the inverse median robot price in Canada to create a time-varying instrumental variable. As robot prices decrease over time, those industries with a higher percentage of workers doing work similar to the capabilities of robots are presumably more likely to adopt them. In using this as our instrument, we argue that both cross-sectional

---

<sup>22</sup> We note the graphs in this section use all available robot import data to show aggregate distribution and trends, not our NALMF or WES regression samples.

<sup>23</sup> We note that the Great Recession did not begin at the same time in Canada and the United States. The United States entered the Great Recession in December 2007, while Canada did not enter a recession until October 2008, which ended in July 2009.

<sup>24</sup> Our instrument for robot investment is similar to that used by Graetz and Michaels (2018).

<sup>25</sup> Examples of occupations include electronics assemblers, boilermakers, metal mold makers, metal patternmakers, tool and die makers, and machinists.



industry employment composition in 1995 and the national median price of robots serve as plausibly exogenous predictors of firm-level robot adoption.

Second, we implement Coarsened Exact Matching (CEM) (Iacus et al. 2012), where we match robot-adopting organizations with non-robot adopting organizations on key observables, and repeat the estimation of our main regressions on matched samples for comparison. For our NALMF sample, we match firms in our sample that adopt robots to non-robot adopting firms by industry (measured by 4 digit NAICS code), year, province, whether the firm is a multi-unit enterprise, total assets, firm age, average annual earnings of the firm's employees, and capital stock. Matching is done exactly by industry, year, province, and multi-unit status, with coarsening allowed for the other variables. For our WES sample, matching is done exactly by industry, year, and province, with coarsening allowed for total revenues, age of the organization, average annual employee earnings, and capital stock.<sup>26</sup>

## 6 Results

### 5.1 Main findings

Results for our baseline tests of the relationship between robot investments and total employment are presented in Columns 1 and 2 of Table 1, for both our full and matched samples created from Coarsened Exact Matching. As Columns 1 and 2 show, the coefficient for our measure of robot investment is positive and statistically significant, predicting an increase in total employment and supporting Hypothesis H1a. Column 3 shows the results from our instrumental variable estimation, which is directionally consistent with both Columns 1 and 2 and very similar in magnitude to our matched sample results in Column 2. For both the matched sample and IV estimations, a one percent increase in robot investment predicts roughly a 0.015 percent increase in total employment within the firm. Considering robot capital is only 0.05% of factor share, this is a substantial effect and suggests there are complementary firm practices associated with robots. As an additional step, we estimate the same regression shown in Column 1, but now replace our robot investment measure with a series of time-indexed dummy variables for the years before and after robot adoption. We plot the dummy variable coefficients graphically along with 95% confidence intervals, with results shown in Figure 2. Prior to robot adoption, we find no evidence of differences in total employment trends with non-robot adopting firms, but an increase in total employment occurs beginning in the first year of robot adoption. Results examining the relationship between robot investments and non-managerial employment by different skill types are shown in Columns 4 through 9 of Table 1. As Columns 4 and 5 show, we find consistent evidence of a negative and statistically significant relationship with middle-skilled employment, supporting Hypothesis H2. We also find evidence of a positive and statistically

---

<sup>26</sup> Multi-unit status was excluded because of too few available matches for estimation.

significant relationship for both low-skilled (Columns 6 and 7) and high-skilled (Columns 8 and 9) employment, supporting Hypotheses 3 and 4.

Results for our tests examining the relationship between robot investment and managerial and total non-managerial employment are shown in Table 2, again presenting both full and matched sample results. In Columns 1 and 2, we find evidence of a negative and statistically significant relationship between robot adoption and managerial employment. Similar to our exercise in Figure 2, we again estimate the same regression shown in Column 1, but now replace our robot investment measure with a series of time-indexed dummy variables for the years before and after robot adoption and plot the coefficients graphically in Figure 3. Prior to robot adoption, we find no evidence of differences in total managerial employment with non-robot adopting organizations, but a substantial decrease in managerial employment occurs beginning in the first year of robot adoption. In Table 3, we examine how robot investment may predict hiring and departures of managerial and non-managerial employees. Robot adoption predicts less hiring of new managers (Columns 1 and 2), as well as an increase in the number of managerial departures (Columns 3 and 4), suggesting that both contribute to the change in managerial headcount.

As additional confirmation, we test whether total employment increases may be explained by an increase in total non-managerial employment, with results shown in Columns 3 and 4 of Table 2. If our total employment or managerial employment results are due to measurement error of either variable, we are unlikely to observe a corresponding change in non-managerial employment. The coefficient for robot investment is positive and statistically significant, consistent with total employment increases being driven by non-managerial employees. In Table 3, Columns 5, 6, 7, and 8 examine whether these results are explained by changes in hiring or turnover for non-managerial employees. The coefficient for robot investment is positive and significant across all specifications, suggesting that investments in robotics increase both non-managerial hiring (Columns 5 and 6) as well as non-managerial departures (Columns 7 and 8). While both hiring and turnover increase, the net effect of the two (Columns 3 and 4 in Table 2) ultimately predicts a net gain in total employment for non-managerial employees. Increases in hiring and departures for non-managerial employees also suggest a compositional change in the workforce, consistent with our findings in Table 1 that show a decline in middle-skilled workers and increase in low- and high-skilled workers.

We next examine how robot investments may be related to changes in the strategic priorities of organizations, with results displayed in Table 4. The pattern of employment changes from robot adoption, especially the decrease in managerial employment, may be related to firms' need to reduce labor costs. If true, our results may reflect a reverse causality where firms that focus on reducing expensive managers choose to adopt robots. As Columns 1 and 2 show, the coefficient for robot investment is not statistically significant, providing no evidence that purchases of robots by firms are motivated by a desire to reduce

labor costs. In Columns 3 and 4, we find a positive and significant coefficient for robot investment with respect to the strategic importance of improving product/service quality. Overall, the results suggest that robot investments are more likely to be motivated by improving the quality of production output, as opposed to efficiency gains from labor cost reductions. This suggests that the possibility of reverse causality where firms may choose to reduce managers and subsequently adopt robots is less likely. These results also corroborate evidence in the field, especially in manufacturing, that suggests robots are often used to improve consistency and reduce production variance.<sup>27</sup>

## **5.2 Changes in organizational practices and the nature of work**

Here, we explore whether the allocation of decision authority to managers within the organization has changed after robot adoption. If firms are simply downsizing managers to reduce slack, we would not necessarily expect to observe a change in decision authority for managers remaining within the firm; downsizing may instead suggest that the remaining managers are doing more than before, and experience an increase in decision authority granted. To explore this possibility, we examine how robot investments predict the allocation of decision authority over training activities and the choice of production technology, with results shown in Tables 5 and 6. These two decisions are particularly relevant as they pertain to human capital management within the firm. Table 5 shows results for the allocation of authority for training decisions, with the coefficient for robot investment being positive for non-managerial employees (Columns 1 and 2) and negative for managerial employees (Columns 3 and 4), with no significant relationship found for business owners/corporate headquarters (Columns 5 and 6). The results provide evidence of decentralization of responsibilities for training from managerial to non-managerial employees within the firm as a response to robot adoption. Table 6 shows results for the allocation of decision authority over the choice of production technology, with no significant relationship found for non-managerial employees (Columns 1 and 2), a negative and significant relationship for managerial employees (Columns 3 and 4), and a positive and significant relationship for business owners/corporate headquarters (Columns 5 and 6). In contrast with training activities, the results suggest the choice of production technology becomes centralized upwards from managerial employees to business owners/corporate headquarters. Although we cannot measure the allocation of decision authority for all managerial tasks, these results suggest that the type of work managers are doing is changing with robot adoption. The downsizing of managers is not just a reduction in headcount, but also a change in their decision authority and the nature of tasks they perform. These results also suggest that robot adoption is also associated with fundamental changes in organizational design.

To further confirm our results at the organization level and consider how the nature of work may be

---

<sup>27</sup> <https://www.robots.com/faq/why-should-my-company-use-industrial-robots>

changing with robot adoption at the individual employee level, we begin by testing whether robot adoption at the organization level predicts changes in the span of control for managerial employees, with results shown in Column 1 of Table 7. The coefficient for robot investment is positive and statistically significant, suggesting that robot adoption predicts increases in the span of control for managers remaining within the organization. An increase in the span of control at the individual manager level is consistent with our earlier organization-level findings of a reduction in managerial headcount and an increase in non-managerial employees.

As an additional test, we also examine how robot investment may change the routine nature of work for individual employees. Here, we consider a specific definition of routine: the degree to which workers can predict their schedule in advance, corresponding with the measure we use.<sup>28</sup> As shown in Column 2 of Table 7, we find a positive relationship between robot investment and the unpredictability of work in advance.<sup>29</sup> The results are consistent with the notion that, as robots automate a larger proportion of tasks within the organization and reduce variance in the production process, human workers are left to focus on work that is less predictable in nature.

### **5.3 Robots and performance measurement mechanism checks**

Here, we conduct two separate tests exploiting available measures in the WES employer survey to investigate whether robot investments affect the ability of firms to measure performance, as proposed in our theoretical arguments. For our first test, we examine whether robot investments increase the likelihood of improvements in performance measurement when organizational change occurs within the workplace. The WES employer survey asks whether any organizational changes occurred during the year, defined as a “change in the way in which work is organized within your workplace or between your workplace and others.”<sup>30</sup> If any organizational changes occurred, the survey subsequently asks respondents whether the impact of the organizational change that affected the most employees increased the “ability to measure performance” for the workplace. Here, we create a dummy variable equal to one if the workplace reports having implemented an organizational change that increased the firm’s ability to measure performance. To address sample selection concerns, we estimate a first stage probit regression predicting the occurrence of organizational change, using the strategic priority of “reorganizing the work process” to the firm as an exogenous predictor, and include the inverse Mills ratio from this regression as an additional control

---

<sup>28</sup> Hage and Aiken (1969) use a similar definition in their study of routine work.

<sup>29</sup> As an additional robustness check, we repeated our analysis only on individuals that reported 3, 4, 5, or 6 (equidistant time gaps), and found similar and statistically significant results.

<sup>30</sup> The survey then provides a list of possible organizational changes for respondents to choose from, ranging from “greater integration among different functional areas” to an “Other” category. Because none of the choices obviously improve performance measurement in all work contexts, we focused on the reported impact of the organizational change on performance measurement rather than on the choice of organizational change itself.

variable.<sup>313233</sup> As shown in Column 1 of Table 8, the coefficient for robot investment is positive and significant, suggesting that robots contribute to improved performance measurement when organizational changes are implemented.

For our second test, we examine whether robot investments are positively related to the strategic priority of improving performance measurement to the firm. For our measure of strategic priority, we again exploit the section of the WES employer survey which asks respondents to “please rate the following factors with respect to their relative importance in your workplace general business strategy,” but now consider the factor of “improving measures of performance.”<sup>34</sup> As the results in Columns 2 and 3 of Table 8 show, the coefficient for robot investment is positive and significant, suggesting that robot adoption and the strategic importance of improving measures of performance are positively related.

### 5.3 Robustness checks

Here, we conduct a series of additional robustness tests for our results. We examine the relationship between robot investment and total employment across different industries (Appendix Tables A9-A11); control for IT investment as a possible omitted variable (Tables A12-A16); investigate whether unobserved purchases from wholesalers and resellers within Canada (instead of direct import purchases) may be affecting our results (Table A17); control for general improvements in firm performance, which may explain our finding of increases in total employment (Table A18), implement an applied Heckman-style correction for the choice to adopt robots (Tables A19-A21); and control for imports from the US and China (Tables A22-A26). Ultimately, we find similar results across our different tests.

While we find consistent evidence of increases in total employment, we also find contrasting declines in employment for both managers and middle-skilled workers in our WES. For these two findings, we formally examine the sensitivity of our results to possible bias from omitted variables using the bounding method developed by Oster (2019). The procedure exploits observable control variables that are correlated with unobservable controls, and examines how the coefficient of interest and regression R-squared change when observable controls are included in the specification. If the coefficient of interest remains relatively stable and the R-squared increases substantially, this increases confidence in the

---

<sup>31</sup> We note that the organizational change and strategic priority measures are recorded in separate section of the WES survey.

<sup>32</sup> This is because improvements in the firm’s ability to measure performance are only recorded for firms that implement an organizational change.

<sup>33</sup> Similar to the measure we use for the strategic priority of improving measures of performance, the WES questionnaire also records the strategic priority of “reorganizing the work process” on a Likert scale. We adjust and rescale this variable in the same manner as our measure of strategic priority of improving measures of performance.

<sup>34</sup> Similar to our other strategic priority measures, we redefine on the Likert scale values of (2) to be equal to (1) and reset the scale to be ascending from 1 to 5. This does not change the sign or significance of our results from using the original variable, and simply dropping all values of (1) also produces results of identical sign and significance level.

direction of our coefficient estimates. To assess the level of confidence, Oster (2019) derives a simple parameter  $\delta$  that represents how strong selection on unobservables would have to be relative to selection on observables to diminish the coefficient of interest to zero. Oster (2019) suggests that an absolute value of  $\delta$  greater than 1 implies a sufficient degree of confidence in the direction of the coefficient estimate. Results of our analysis are shown in Tables A27 in the Appendix. For both our managerial and middle-skilled employment specifications, the absolute values of  $\delta$  are above 1, suggesting our results are relatively robust to concerns of selection on unobservables.

## **7 Discussion and conclusion**

Utilizing novel data capturing investments in robotics for a population of businesses in a developed economy, we provide the first firm-level evidence of the effect of robot adoption on employment and management and the associated changes in organizational practices. The results suggest that robots do not affect employment within the firm uniformly, leading to net increases in the headcount of non-managerial employees, but also decreases in the headcount of managerial employees. This is consistent with the notion that by taking on a subset of responsibilities and activities in the production process of the firm, robots affect the demand for workers engaged in other activities within the firm. Employees whose skills have greater complementarity to robot investments are more likely to experience net gains in employment, depending on the degree to which their skills are complementary. In our study, we find evidence of skill polarization of the non-managerial workforce, with decreases in middle-skilled employment and increases in low and high-skilled employment, consistent with prior findings on automation (Autor and Salomons 2018, Autor et al. 2003). Surprisingly, we find evidence of displacement of specific higher cognitive-skilled jobs such as managers that were previously less vulnerable to skill-biased technical change from earlier waves of technology. We find that this reduction may be the consequence of both a decrease in the need for certain types of supervisory work from robot adoption as well as an indirect effect from the changing composition of non-managerial employees. Consistent with a decline in managerial employment and increase in total employment, we find that the span of control for managers has also increased after robot adoption. We also find evidence that managerial work has fundamentally changed after robot adoption, as their decision authority has been reduced. However, we find no evidence that job losses are caused by firms desiring to cut labor costs, and instead find evidence that firms primarily adopt robots to improve product and service quality.

In addition to changes in employment, we also observe that organizational practices change with robot adoption as the allocation of decision authority for certain tasks shifts to different layers of the hierarchy, away from managers. Human resource-related decisions with respect to training are decentralized from managers to non-managerial employees, while the choice of production technology is centralized from

managers to business owners and corporate headquarters. This differs from effects of earlier generations of IT that tended to decentralize decision authority (Acemoglu et al. 2007). However, with robot adoption rapidly increasing in prevalence and capability, we expect that the allocation of decision authority and other complementary work practices are likely to continue to evolve. Those firms that can best match their capabilities and work practices to productive opportunities can benefit substantially from robot investments and develop potential competitive advantages, highlighting the need to understand the different types of complements to robots as a new technology.

Overall, our findings using organization-level data suggest the effect of robots on labor is more nuanced than earlier work has predicted and requires a deeper examination beyond the level of industry or region to understand how they are used to complement and substitute labor and how organizational practices need to evolve around the changing nature of work. While our analysis suggests that robot adoption is associated with using different types of labor, the associated implication for wages is also an important question. The extent to which wages may change depends on the type of jobs that are created and eliminated. Our initial evidence suggests that although labor cost reduction is not the primary reason for why firms adopt robots, the reduction in managerial and middle-skilled employment and increase in low and high-skilled employment ultimately predicts an ambiguous result for average wages. However, complementing our finding of a decline in demand for middle-skilled employment, Dauth et al. (2018) use industry-level robot investment to examine effects on employee wages, and find that robot adoption leads to substantial wage decreases for middle-skilled workers.

Changes in employee types and skills as a result of robot adoption would also lead firms to implement complementary work practices to accommodate the skill change, similar to earlier generations of skill-biased technical change (Bresnahan et al. 2002, Murnane et al. 1999). To understand these effects, the collection of microdata, especially at the firm level, is crucial. Additionally, better data about robot investment across different contexts are critical to understanding whether the effects we observe on employment and work practices can be generalized to other economies (Buffington et al. 2018, Frank et al. 2019). While we provide detailed firm-level evidence on robotics and show that work practices have already evolved in response to robot technologies, future research should continue to examine how robotics technologies in general affect different firms, occupations, industries, and geographical regions (Felten et al. 2019). With rapid advances in robotics capabilities, understanding their implications is critical as investments in robots are likely to have profound effects on both employment and organizations.

## References


- Acemoglu D, Aghion P, Lelarge C, Van Reenen J, Zilibotti F (2007). Technology, Information, and the Decentralization of the Firm. *The Quarterly Journal of Economics*. 122(4) 1759-1799.
- Acemoglu D, Restrepo P (2017). Robots and jobs: Evidence from US labor markets.
- Acemoglu D, Restrepo P (2018). *Artificial intelligence, automation and work*. National Bureau of Economic Research.
- Agrawal A, Gans J, Goldfarb A. (2018). *Prediction Machines: The simple economics of artificial intelligence*. Harvard Business Press.
- Aral S, Brynjolfsson E, Wu L (2012). Three-way complementarities: Performance pay, human resource analytics, and information technology. *Management Science*. 58(5) 913-931.
- Arntz M, Gregory T, Zierahn U (2016). The risk of automation for jobs in OECD countries.
- Autor D, Salomons A (2017). *Robocalypse Now: Does Productivity Growth Threaten Employment?* National Bureau of Economic Research, Inc.
- Autor D, Salomons A (2018). *Is automation labor-displacing? Productivity growth, employment, and the labor share*. National Bureau of Economic Research.
- Autor DH, Katz LF, Kearney MS (2006). The polarization of the US labor market. *American economic review*. 96(2) 189-194.
- Autor DH, Levy F, Murnane RJ (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*. 118(4) 1279-1333.
- Autor DH, Levy F, Murnane RJ (2003). The Skill Content of Recent Technological Change: An Empirical Exploration\*. *The Quarterly Journal of Economics*. 118(4) 1279-1333.
- Bell, GD (1967). Determinants of span of control. *American Journal of Sociology*, 73(1), 100-109.
- Bepko Jr RJ, Moore JR, Coleman JR (2009). Implementation of a pharmacy automation system (robotics) to ensure medication safety at Norwalk hospital. *Quality Management in Healthcare*. 18(2) 103-114.
- Bessen JE, Goos M, Salomons A, Van den Berge W (2019). Automatic Reaction-What Happens to Workers at Firms that Automate? *Boston Univ. School of Law, Law and Economics Research Paper*.
- Bloom, N, Garicano, L, Sadun, R, Van Reenen, J (2014). The distinct effects of information technology and communication technology on firm organization. *Management Science*, 60(12), 2859-2885.
- Bresnahan TF, Brynjolfsson E, Hitt LM (2002). Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-level Evidence. *The Quarterly Journal of Economics*. 117(1) 339-376.
- Bresnahan TF, Trajtenberg M (1995). General purpose technologies 'Engines of growth'? *Journal of econometrics*. 65(1) 83-108.
- Brynjolfsson E, Hitt L (1996). Paradox lost? Firm-level evidence on the returns to information systems spending. *Management Science*. 42(4) 541-558.
- Brynjolfsson E, Hitt LM, Yang S (2002). Intangible assets: Computers and organizational capital. *Brookings papers on economic activity*. 2002(1) 137-181.
- Brynjolfsson E, McAfee A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.
- Brynjolfsson E, Mitchell T (2017). What can machine learning do? Workforce implications. *Science*. 358(6370) 1530-1534.
- Brynjolfsson E, Rock D, Syverson C (2018). *Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics*. University of Chicago Press.
- Buffington C, Miranda J, Seamans R (2018). *Development of Survey Questions on Robotics Expenditures and Use in US Manufacturing Establishments*.
- Card D, DiNardo JE (2002). Skill-biased technological change and rising wage inequality: Some problems and puzzles. *Journal of labor economics*. 20(4) 733-783.
- Cockburn IM, Henderson R, Stern S (2018). *The Impact of Artificial Intelligence on Innovation*. Working Paper, National Bureau of Economic Research, Cambridge.
- David H, Dorn D (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*. 103(5) 1553-1597.
- Dickens, WT, Katz, LF, Lang, K, Summers, LH (1989). Employee crime and the monitoring puzzle. *Journal of labor economics*, 7(3), 331-347.
- Dickens, WT, Katz, LF, Lang, K, Summers, LH (1990). Why do firms monitor workers?. In *Advances in the Theory and Measurement of Unemployment* (pp. 159-171). Palgrave Macmillan, London.



- Dinlersoz E, Wolf Z (2018). *Automation, Labor Share, and Productivity: Plant-Level Evidence from US Manufacturing*.
- Dustmann C, Ludsteck J, Schönberg U (2009). Revisiting the German wage structure. *The Quarterly Journal of Economics*. 124(2) 843-881.
- Eisenhardt, KM (1989). Agency theory: An assessment and review. *Academy of management review*, 14(1), 57-74.
- Felten E, Raj M, Seamans R (2019). *The Occupational Impact of Artificial Intelligence: Labor, Skills, and Polarization*.
- Ford M. (2015). *Rise of the Robots: Technology and the Threat of a Jobless Future*. Basic Books.
- Frank MR, Autor D, Bessen JE, Brynjolfsson E, Cebrian M, Deming DJ, Feldman M, Groh M, Lobo J, Moro E (2019). Toward understanding the impact of artificial intelligence on labor. *Proceedings of the National Academy of Sciences* 201900949.
- Frey CB, Osborne MA (2017). The future of employment: how susceptible are jobs to computerisation? *Technological forecasting and social change*. 114 254-280.
- Gombolay M, Yang XJ, Hayes B, Seo N, Liu Z, Wadhwanian S, Yu T, Shah N, Golen T, Shah J (2018). Robotic assistance in the coordination of patient care. *The International Journal of Robotics Research*. 37(10) 1300-1316.
- Graetz, G., Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5), 753-768.
- Hage, J, Aiken, M (1969). Routine technology, social structure, and organization goals. *Administrative science quarterly*, 366-376.
- Hales C (1986). What do managers do? A critical review of the evidence. *Journal of Management studies*. 23(1) 88-115.
- Hales, C (1999). Why do managers do what they do? Reconciling evidence and theory in accounts of managerial work. *British journal of management*, 10(4), 335-350.
- Hammer M (1990). Reengineering work: don't automate, obliterate. *Harvard Business Review*. 68(4) 104-112.
- Helper S, Henderson R (2014). Management practices, relational contracts, and the decline of General Motors. *Journal of Economic Perspectives*. 28(1) 49-72.
- Helper S, MacDuffie JP, Sabel C (2000). Pragmatic collaborations: advancing knowledge while controlling opportunism. *Industrial and corporate change*. 9(3) 443-488.
- Hong, B, Kueng, L, Yang, MJ (2019). Complementarity of performance pay and task allocation. *Management Science*, 65(11), 5152-5170.
- Huselid, MA, Becker, BE (1997). The impact high performance work systems, implementation effectiveness, and alignment with strategy on shareholder wealth. In *Academy of Management Proceedings* (Vol. 1997, No. 1, pp. 144-148). Briarcliff Manor, NY 10510: Academy of Management.
- Ichniowski C, Shaw K, Prennushi G (1997). The effects of human resource practices on manufacturing performance: A study of steel finishing lines. *American Economic Review*. 87(3) 291-313.
- Jensen, MC, Meckling, WH (1976). Agency Costs and the Theory of the Firm. *Journal of financial economics*, 3(4), 305-360.
- Kenny M, Florida R (1993). *Beyond mass production: The Japanese system and its transfer to the US* New York: Oxford Univ. Press.
- Koch M, Manuylov I, Smolka M (2019). Robots and firms.
- Kolbjørnsrud V, Amico R, Thomas RJ (2016). *How artificial intelligence will redefine management*.
- MacDuffie JP (1997). The road to “root cause”: Shop-floor problem-solving at three auto assembly plants. *Management Science*. 43(4) 479-502.
- Malone TW (2003). Is empowerment just a fad? Control, decision making, and IT. *Inventing the Organizations of the 21st Century* 49-69.
- Malone TW. (2004). *The future of work*. Audio-Tech Business Book Summaries, Incorporated.
- Mann K, Püttmann L (2017). Benign Effects of Automation: New Evidence From Patent Texts.
- Manyika J (2017). A future that works: AI, automation, employment, and productivity. *McKinsey Global Institute Research, Tech. Rep*.
- McAfee A, Brynjolfsson E. (2017). *Machine, platform, crowd: Harnessing our digital future*. WW Norton & Company.
- Meyer, MW (1968). Expertness and the span of control. *American Sociological Review*, 944-951.
- Mintzberg H (1973). The nature of managerial work.
- Mintzberg, H (1980). Structure in 5's: A Synthesis of the Research on Organization Design. *Management science*, 26(3), 322-341.
- Mintzberg H (2013). *Simply managing: What managers do—and can do better*. Berrett-Koehler Publishers.

- Murnane RJ, Levy F, Autor D (1999). *Technological change, computers and skill demands: Evidence from the back office operations of a large bank*.
- Oster, E (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business and Economic Statistics*, 37(2), 187-204.
- Parker M, Slaughter J (1988). Management by stress. *Technology Review*. 91(7) 37-44.
- Perrow C (1967). A framework for the comparative analysis of organizations. *American sociological review* 194-208.
- Scott, ED, O'Shaughnessy, KC, Cappelli, P (1994). Management jobs in the insurance industry: Organizational deskilling and rising pay inequity. Wharton Financial Institutions Center, Wharton School of the University of Pennsylvania.
- Scully MA (2000). Manage your own employability: Meritocracy and the legitimation of inequality in internal labor markets and beyond. *Relational wealth: The advantages of stability in a changing economy* 199-214.
- Syverson C (2004). Product substitutability and productivity dispersion. *Review of Economics and Statistics*. 86(2) 534-550.
- Taylor FW. (1977). *Shop management*. Рипол Классик.
- Verl A (2019). *Managing mass customisation with software-defined manufacturing*. IFR, University of Stuttgart.
- Wu L (2013). Social network effects on productivity and job security: Evidence from the adoption of a social networking tool. *Information Systems Research*. 24(1) 30-51.

Figure 1. Aggregate robot stock in Canada, 1996-2017



Note: Robot stock is depreciated using a 12-year useful life assumption, following guidance from the International Federation of Robotics (IFR).

Table 1. Total employment and non-managerial employment by skill type regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	FE	FE	2SLS	FE	FE	FE	FE	FE	FE
				WES	WES	WES	WES	WES	WES
Dataset:	NALMF	NALMF	NALMF	Employer	Employer	Employer	Employer	Employer	Employer
Sample:	Full	Matched	Full	Full	Matched	Full	Matched	Full	Matched
Dependent variable:	ln(Total employees)	ln(Total employees)	ln(Total employees)	ln(Total middle-skilled)	ln(Total middle-skilled)	ln(Total low-skilled production)	ln(Total low-skilled production)	ln(Total high-skilled)	ln(Total high-skilled)
ln(Total assets)	0.191*** (0.013)	0.215*** (0.037)	0.346*** (0.016)						
ln(Total revenues)				0.147 (0.103)	0.106 (0.094)	0.122 (0.086)	0.398** (0.162)	0.040 (0.071)	0.037 (0.074)
Multi-unit enterprise	0.139*** (0.014)	0.144*** (0.022)	0.500*** (0.027)	-0.077 (0.095)	-0.396*** (0.062)	-0.235* (0.132)	-0.049 (0.074)	0.090 (0.063)	0.486** (0.175)
Unionized				0.389*** (0.115)	-0.092 (0.669)	0.200 (0.161)	2.052*** (0.477)	-0.219** (0.095)	-1.041* (0.523)
Outsourcing				-0.001 (0.086)	0.419** (0.187)	0.048 (0.104)	-0.335 (0.322)	0.171** (0.068)	0.162 (0.151)
ln(Robot capital stock)	0.007*** (0.002)	0.015** (0.006)	0.015*** (0.004)	-0.086*** (0.014)	-0.031** (0.012)	0.061*** (0.021)	0.021** (0.009)	0.016** (0.007)	0.018** (0.008)
Industry fixed effects	N	N	Y	N	N	N	N	N	N
Province fixed effects	N	N	Y	N	N	N	N	N	N
Year fixed effects	Y	Y	Y	Y	Y	Y	Y	Y	Y
Organization fixed effects	Y	Y	N	Y	Y	Y	Y	Y	Y
Observations	929,162	41,399	865,759	17,449	1,746	17,449	1,746	17,449	1,746
Adj R-squared	0.92	0.94		0.70	0.74	0.72	0.83	0.59	0.76

Standard errors in parentheses, clustered by industry. All regressions using WES data use sampling weights. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table 2. Managerial and non-managerial employment, hiring, and departure regressions

	(1)	(2)	(3)	(4)
	FE	FE	FE	FE
	WES	WES	WES	WES
Dataset:	Employer	Employer	Employer	Employer
Sample:	Full	Matched	Full	Matched
Dependent variable:	ln(Total managers)	ln(Total managers)	ln(Total non-mgr. employees)	ln(Total non-mgr. employees)
ln(Total revenues)	0.084** (0.033)	0.009 (0.168)	0.242*** (0.053)	0.389*** (0.093)
Multi-unit enterprise	0.032 (0.096)	0.307 (0.434)	0.046 (0.049)	1.199 (0.847)
Unionized	0.168 (0.108)	0.594 (0.472)	0.025 (0.033)	-2.309*** (0.488)
Outsourcing	0.001 (0.059)	0.160 (0.152)	0.005 (0.059)	-0.205* (0.115)
ln(Robot capital stock)	-0.080*** (0.011)	-0.073*** (0.014)	0.005** (0.002)	0.016** (0.008)
Year fixed effects	Y	Y	Y	Y
Organization fixed effects	Y	Y	Y	Y
Observations	17,449	1,746	17,449	1,746
Adj R-squared	0.69	0.75	0.88	0.86

Standard errors in parentheses, clustered by industry. All regressions using WES data use sampling weights. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table 3. Managerial and non-managerial hiring, and departure regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	FE	FE	FE	FE	FE	FE	FE	FE
Dataset:	WES Employer	WES Employer	WES Employer	WES Employer	WES Employer	WES Employer	WES Employer	WES Employer
Sample:	Full	Matched	Full	Matched	Full	Matched	Full	Matched
Dependent variable:	ln(Total mgr. hires)	ln(Total mgr. hires)	ln(Total mgr. departures)	ln(Total mgr. departures)	ln(Total non-mgr. hires)	ln(Total non-mgr. hires)	ln(Total non-mgr. departures)	ln(Total non-mgr. departures)
ln(Total revenues)	0.053 (0.070)	0.066 (0.183)	0.023 (0.037)	0.024 (0.044)	0.209** (0.082)	0.045 (0.161)	0.077 (0.084)	0.203* (0.112)
Multi-unit enterprise	0.030 (0.149)	-0.343* (0.185)	-0.050 (0.078)	-0.061 (0.223)	-0.124 (0.208)	-0.584 (0.419)	0.178 (0.118)	0.006 (0.192)
Unionized	0.279** (0.124)	0.514 (0.523)	-0.019 (0.067)	-0.159 (0.106)	-0.149 (0.104)	-1.139*** (0.301)	0.292*** (0.100)	0.119 (0.117)
Outsourcing	0.116 (0.086)	0.109 (0.101)	-0.049 (0.050)	0.089 (0.115)	0.158 (0.119)	-0.343** (0.173)	0.062 (0.079)	0.053 (0.084)
ln(Robot capital stock)	-0.032*** (0.009)	-0.037*** (0.013)	0.025*** (0.004)	0.031** (0.012)	0.044*** (0.007)	0.028** (0.014)	0.037*** (0.009)	0.018*** (0.006)
Year fixed effects	Y	Y	Y	Y	Y	Y	Y	Y
Organization fixed effects	Y	Y	Y	Y	Y	Y	Y	Y
Observations	17,449	1,746	17,449	1,746	17,449	1,746	17,449	1,746
Adj R-squared	0.19	0.22	0.06	0.03	0.60	0.46	0.33	0.15

Standard errors in parentheses, clustered by industry. All regressions using WES data use sampling weights. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table 4. Strategic priority regressions

	(1)	(2)	(3)	(4)
	FE	FE	FE	FE
Dataset:	WES Employer	WES Employer	WES Employer	WES Employer
Sample:	Full	Matched	Full	Matched
Dependent variable (strategic importance):	Reducing labor costs	Reducing labor costs	Improving product/service quality	Improving product/service quality
ln(Total revenues)	-0.014 (0.130)	0.118 (0.322)	0.098 (0.133)	0.180 (0.380)
Multi-unit enterprise	-0.197 (0.121)	0.192 (0.347)	-0.198 (0.173)	0.629* (0.374)
Unionized	-0.144 (0.230)	-0.743*** (0.209)	-0.336* (0.199)	0.093 (0.333)
Outsourcing	0.050 (0.178)	0.488 (0.488)	0.094 (0.169)	0.960 (0.590)
ln(Robot capital stock)	0.027 (0.036)	-0.001 (0.020)	0.108*** (0.013)	0.103*** (0.031)
Year fixed effects	Y	Y	Y	Y
Organization fixed effects	Y	Y	Y	Y
Observations	8,906	889	8,906	889
Adj R-squared	0.32	0.46	0.38	0.21

Standard errors in parentheses, clustered by industry. All regressions using WES data use sampling weights. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table 5. Task allocation regressions, training decisions

	(1)	(2)	(3)	(4)	(5)	(6)
	FE	FE	FE	FE	FE	FE
	WES	WES	WES	WES	WES	WES
Dataset:	Employer	Employer	Employer	Employer	Employer	Employer
Sample:	Full	Matched	Full	Matched	Full	Matched
Training decisions						
Dependent variable:	Non-manual employees	Non-manual employees	Managers	Managers	Business owners or Corp HQ	Business owners or Corp HQ
ln(Total revenues)	-0.003 (0.019)	-0.022 (0.046)	0.003 (0.090)	-0.017 (0.066)	0.027 (0.089)	0.307 (0.230)
Multi-unit enterprise	0.009 (0.013)	-0.094 (0.095)	-0.021 (0.077)	-0.234 (0.640)	0.110 (0.104)	0.755* (0.450)
Unionized	-0.041 (0.139)	0.014 (0.015)	-0.070 (0.212)	-0.027 (0.025)	-0.139 (0.173)	0.018 (0.101)
Outsourcing	0.011 (0.028)	0.121 (0.074)	-0.019 (0.072)	0.066 (0.300)	-0.058 (0.081)	-0.278 (0.195)
ln(Robot capital stock)	0.074*** (0.011)	0.077*** (0.012)	-0.077*** (0.011)	-0.080*** (0.012)	0.003 (0.003)	0.012 (0.009)
Year fixed effects	Y	Y	Y	Y	Y	Y
Organization fixed effects	Y	Y	Y	Y	Y	Y
Observations	6,173	632	6,173	632	6,173	632
Adj R-squared	0.29	0.84	0.33	0.72	0.39	0.75

Standard errors in parentheses, clustered by industry. All regressions use sampling weights. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table 6. Task allocation regressions, choice of production technology

	(1)	(2)	(3)	(4)	(5)	(6)
	FE	FE	FE	FE	FE	FE
	WES	WES	WES	WES	WES	WES
Dataset:	Employer	Employer	Employer	Employer	Employer	Employer
Sample:	Full	Matched	Full	Matched	Full	Matched
Choice of Production Technology						
Dependent variable:	Non-manual employees	Non-manual employees	Managers	Managers	Business owners or Corp HQ	Business owners or Corp HQ
ln(Total revenues)	0.004 (0.008)	0.006 (0.033)	0.056 (0.072)	-0.131 (0.100)	-0.049 (0.075)	0.365 (0.262)
Multi-unit enterprise	-0.007 (0.012)	-0.010 (0.018)	0.038 (0.066)	-0.498 (0.427)	0.070 (0.096)	0.930*** (0.344)
Unionized	-0.000 (0.004)	0.009 (0.009)	0.231 (0.189)	0.868*** (0.092)	-0.527*** (0.181)	-0.878*** (0.070)
Outsourcing	-0.010 (0.019)	0.024 (0.024)	0.038 (0.075)	0.212 (0.250)	-0.003 (0.077)	-0.324* (0.179)
ln(Robot capital stock)	-0.000 (0.000)	0.002 (0.001)	-0.069*** (0.015)	-0.077*** (0.012)	0.075*** (0.013)	0.082*** (0.017)
Year fixed effects	Y	Y	Y	Y	Y	Y
Organization fixed effects	Y	Y	Y	Y	Y	Y
Observations	6,173	632	6,173	632	6,173	632
Adj R-squared	0.30	0.09	0.31	0.54	0.33	0.54

Standard errors in parentheses, clustered by industry. All regressions use sampling weights. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table 7. Span of control and work predictability regressions

	(1)	(2)
	FE	FE
Dataset:	WES Employee	WES Employee
Dependent variable:	Span of control	Work Unpredictability
ln(Total employees)	22.532* (12.112)	-0.112 (0.317)
Multi-unit enterprise	32.915 (29.069)	0.255 (0.270)
Union member	-6.911 (4.560)	0.067 (0.231)
Outsourcing	-4.066 (5.147)	0.325 (0.229)
ln(Robot capital stock)	0.342** (0.132)	0.158** (0.066)
Year fixed effects	Y	Y
Employee fixed effects	Y	Y
Observations	11,719	10,969
Adj R-squared	0.15	0.59

Standard errors in parentheses, clustered by industry. All regressions using WES data use sampling weights. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Table 8. Performance measurement regressions, WES employer sample

	(1)	(2)	(3)
	FE	FE	FE
Dataset:	WES Employer	WES Employer	WES Employer
Sample:	Full	Full	Matched
Dependent variable:	Increase in ability to measure performance	Strategic priority of improving measures of performance	Strategic priority of improving measures of performance
ln(Total revenues)	0.034 (0.047)	0.090 (0.141)	-0.171 (0.258)
Multi-unit enterprise	0.027 (0.088)	0.167 (0.192)	0.356 (0.251)
Unionized	-0.028 (0.062)	0.039 (0.186)	-0.523*** (0.120)
Outsourcing		-0.011 (0.142)	0.702 (0.582)
ln(Robot capital stock)	0.022** (0.011)	0.076*** (0.014)	0.119*** (0.024)
Inverse Mills ratio	-0.140** (0.068)		
Organization fixed effects	Y	Y	Y
Year fixed effects	Y	Y	Y
Observations	4,947	8,906	889
Adj R-squared	0.42	0.29	0.59

Standard errors in parentheses, clustered by industry. Inverse Mills ratio is from first stage probit regression predicting organizational change. All regressions use sampling weights. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1



Figure 2. Total employment time-indexed dummy regression coefficient plot, NALMF sample

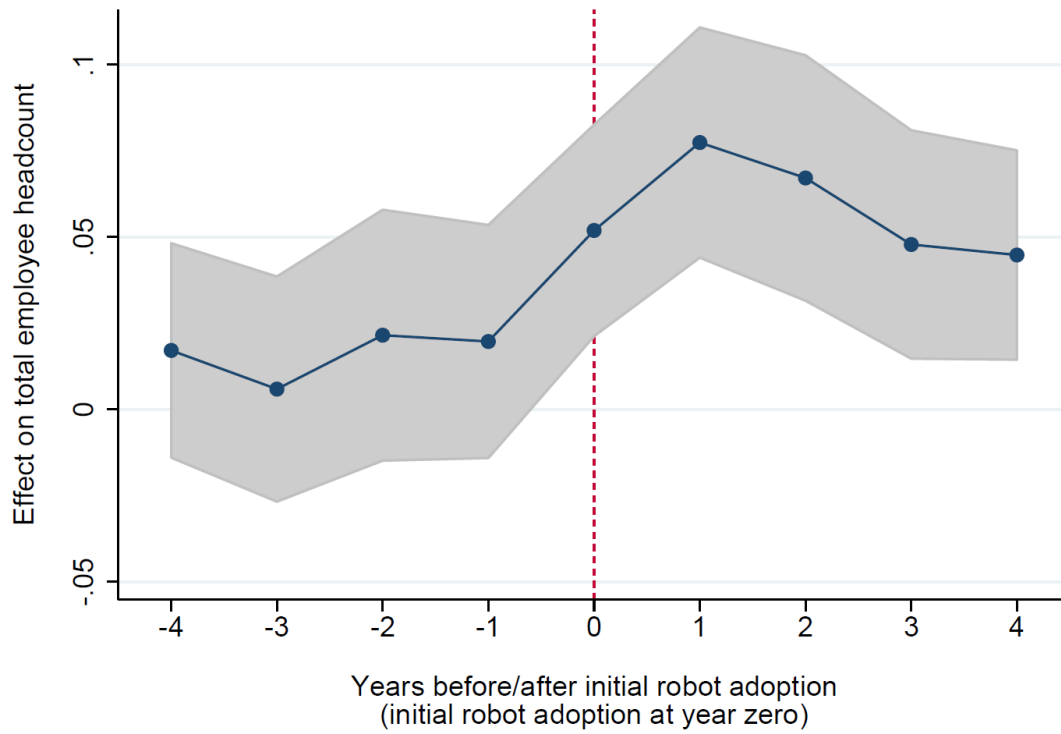


Figure 3. Total manager headcount time-indexed dummy regression coefficient plot, WES employer sample

