Accepted Manuscript

Title: Lucy and the Chocolate Factory: Warehouse Robotics and Worker Safety

Authors: Gordon Burtch, Brad Greenwood, and Kiron Ravindran

DOI: https://doi.org/10.1177/00197939251333754

Publication: Industrial and Labor Relations Review

Citation: Burtch, G., Greenwood, B., & Ravindran, K. (2025). Lucy and the chocolate factory: warehouse robotics and worker safety. ILR Review, 78(4), 587-613.

Lucy and the Chocolate Factory: Warehouse Robotics and Worker Safety

Gordon Burtch
Boston University – gburtch@bu.edu

Brad N. Greenwood
George Mason University – bgreenwo@gmu.edu

Kiron Ravindran
IE University – kiron.ravindran@ie.edu

ABSTRACT

We examine the implications of robotics for warehouse worker safety. While warehouse automation has the potential to reduce injuries by eliminating high-risk tasks, it may also increase injuries among remaining non-automated tasks, due to reduced task variety and an accelerated pace of work. Our findings provide evidence of both effects: warehouse robotics are associated with a 40% decrease in severe injuries but a 77% increase in non-severe injuries. We provide subsequent evidence that the rise in non-severe injuries is at least partially attributable to the increased pace of work at robotics facilities. The implications of our findings for regulators, policymakers, workers, and firms are discussed within.

Keywords: Robotics, Automation, Worker Safety, Logistics, Fulfillment, Warehousing

Advancements in computer vision and machine learning have led to the deployment of new forms of robotics, which are mobile and capable of operating alongside human laborers in a variety of novel working environments. This has raised questions about the implications of robotics for worker safety. As with more established forms of automation, it has been argued that these new robots could significantly improve worker safety by eliminating the need for repetitive or hazardous work (Amazon.com 2022). However, those assertions are at odds with recent media accounts that warehouse robotics are associated with an *increase* in the volume of worker injuries (Greene and Alcantra 2021). Which claim is correct? We address that question herein.

Somewhat counterintuitively, we posit that both may be true, and that the answer depends on the unit of analysis one considers. At the task level, robotics may mitigate injuries by reducing human involvement in hazardous activities. However, at the process level, the incorporation of robotics leads to at least two negative externalities for the laborers who continue to perform manual tasks. First, because robotics reduce the variety of tasks available to human workers (Parker and Grote 2022), those workers become less attentive (Gutelius and Theodore 2019). Second, because robots are capable of operating at faster pace, managers' productivity expectations may rise. Workers may therefore take less care in their work, as they seek to keep pace (ibid). Further, a reduction in tax variety combined with an increased pace of work may lead human workers to experience greater repetitive motion and increased ergonomic stress. Accordingly, at the process level, robotics may lead to a rise in worker injuries.

We explore this tension in the context of warehousing, with a specific focus on Amazon Fulfillment Centers. This focus is deliberate and motivated by several considerations. First, the warehousing sector is a large and economically important. More than 1.7 million people were employed in warehouses across the United States as of 2023 (BLS 2023b). Second, the warehousing sector is characterized by high rates of workplace injury (BLS 2023a). Finally, the warehousing sector is characterized by factors that are likely to exacerbate the relationships discussed above, including rapid automation, high rates of turnover, a heavy reliance on temporary workers, and limited worker protections.

We focus our analysis on the Amazon Fulfillment Center (FC) network, which began incorporating robotics in 2012, following Amazon's purchase of Kiva Robotics. We first compare warehouse level annual injury reports obtained from the Occupational Health and Safety Administration (OSHA), contrasting Robotic FCs with Legacy FCs. We consider injuries in general, as well as injuries stratified by severity. Importantly, this data is a particularly unique aspect of our study. Reported at the establishment level, our data is more nuanced and less likely to suffer from various sources of unobservable heterogeneity that may influence industry-level estimates (Bekhtiar et al. 2021).

To better understand the underlying mechanisms behind observed differences in Robotic FCs, we also examine injury-level data for 26 Amazon FCs in 2018. This analysis enables us to look at specific categories of injuries (e.g., sprains and strains versus cuts, bruises, or fractures), employing a difference-in-difference-in-differences design (Olden and Møen 2022) that relates periods of high operational load (i.e. pace) to differences in injuries occurring at Robotic and Legacy FCs. Finally, to probe the robustness of our main findings, we conduct analyses at the state-level, leveraging BLS annual reports of injury rates in states' warehousing sectors.

Our findings have implications for both practice and theory. Theoretically, our work establishes a new empirical beachhead in the academic understanding of the implications of workplace automation. While scholarly work in management and economics has examined robotics for decades, the focus has typically been on the antecedents of adoption (Pillai et al.

2022) or the downstream implications for labor and productivity (Bessen 2015, Bessen 2016, Dixon et al. 2021). We pivot on this work by examining the changing nature of collaborative co-creation between humans and robots and its impact on worker safety. We hope this work serves as a call for greater attention in academic scholarship not only on the quality and productivity implications of automation, but the impact on workers.

RELATED LITERATURE

In what follows, we briefly review two broad but related literatures. First, we consider prior research that addresses how robotics can affect worker health, and the associated implications of the changing nature of robots. Second, we discuss prior research on task displacement in the wake of robotics adoption, which highlights that certain tasks are more or less likely to be displaced. We close by arguing that the certain aspects of the warehousing sector imply that prior scholarship may not generalize well to this setting, before motivating the tension between increased and decreased worker safety and possible shifts in injury types.

Robotics Evolution and Worker Health

Canonical work in robotics and automation argues that, beyond productivity gains, the central reason for robotization is the expectation of improvements in worker safety. The intuition is simple. As robots can be used in place of humans when performing dangerous or repetitive work the worker can be spared from potential injury (Murashov et al. 2016). Support for this logic is easy to find. Gihleb et al. (2022), for example, estimate that a one-standard deviation increase in industrial robotic adoption within a US commuting zone is associated with 12 fewer injuries per 1,000 workers, across industries. However, the broad, macro relationship between robotics and injury is likely to mask heterogeneity across contexts. Indeed, other examples in the literature have come to the opposite conclusion. For example, Karwowski et al. (1988) report that injuries rose substantially in a Kentucky manufacturing plant following the adoption of robotics. These mixed findings emphasize the importance of context when studying the impacts of robotics.

Bearing this in mind, robotics has witnessed tremendous advancement in recent years, taking on novel forms and enabling their deployment in new work environments. Changes in robotics are perhaps most evident when one considers shifts in the language that has appeared within regulatory guidance. In the 1970s and 1980s, OSHA guidance on robotic safety emphasized that robots were "reprogrammable, multifunctional, mechanical manipulators." This is in stark contrast to contemporary categorizations, which are more nuanced and emphasize the collaborative capabilities of robots (Peshkin and Colgate 1999). What is critical to recognize is that, unlike industrial robots, which typically operate in isolation from human workers, collaborative robots operate alongside workers, taking on different tasks, in new contexts. This has led to renewed interest in the effects of collaborative robots on safety (Gualtieri et al. 2021).

Two key points of consideration that have subsequently been raised are the implications of these collaborative robotics for i) job satisfaction and ii) the pace at which work is conducted. Regarding job satisfaction, introducing collaborative robots into the work environment will necessarily change the nature of tasks conducted by human workers and the environment in which those tasks are conducted. Recent work, for example, notes that robots may lead to a decline in mental health among retained workers (Gihleb et al. 2022). This arises from concerns about job displacement as well as an increased sense of monotony, as robots reduce worker task variety, autonomy, and discretion (Berg et al. 2023). Ultimately, these shifts can give rise to general sense of job dissatisfaction (Welfare et al. 2019), leading in turn to boredom and inattentiveness, conditions that set the stage for accidents and injuries to take place.

Regarding the pace of work, because robots can perform standardized tasks more rapidly than a human, their presence may enable heightened expectations of productivity for retained human workers (Antón et al. 2023) when firms attempt to maximize the productivity gains gleaned from robotics. Reflecting this logic, Borenstein (2010) posits: “if robots can help surgical procedures to be completed more rapidly, […] demands on surgeons [may] increase so they will have to perform more procedures per day.” Thus, while robotics might be adopted with the initial intent of improving worker safety, by reducing human involvement in dangerous tasks, unintended negative spillovers may arise for the safety of retained workers, who continue to perform non-automated tasks, ranging from increased rates of boredom (Antón et al. 2023) to selection out of the labor force (Lerch 2020).

Robotics and Task Displacement

Given that robots can displace highly repetitive and standardized tasks (Autor 2015), it should come as no surprise that much of the literature dealing with the employment implications of robots has focused on job displacement. Still, this literature is careful to recognize that jobs and tasks are not equivalent; a job is best viewed as a bundle of complementary tasks (Frank et al. 2019, Görlich 2010). In many cases, only some tasks within a job can be automated, namely those tasks that are highly standardized and repeatable. Tasks that cannot be automated must continue to be performed by a human worker. As the original bundle of tasks are generally complementary to begin with (Görlich 2010), the new division of labor that results is one in which the human laborer’s tasks remain complementary to the tasks undertaken by the robot.

Depending on whether a job’s entire task bundle can be automated, versus only a portion, we might expect the consequences for workers to differ drastically. For example, Barrett et al. (2012) examined the effects of introducing robotics into a pharmacy setting. Those authors observed highly heterogeneous effects. Pharmacists benefited from the packing ability of the robots, and they were able to focus a greater portion of their time on higher-value, difficult-to-automate tasks. Pharmacy technicians also benefitted from the opportunity to upskill, much as bank tellers did in the 1990s when banking underwent a period of technological change (Hunter et al. 2001). In contrast, pharmacy assistants found themselves largely replaced by robots, because most of their tasks were easier to automate. Pharmacy assistants were thus relegated to a secondary role, with little meaningful work, undermining both their morale and morale in the organization. Work by Beane (2019) highlights a similar dynamic, that the introduction of robotics leads to a division or separation between tasks that can and cannot be automated within a job. Studying trainee surgeons learning to work with robotic tools, Beane (2019) observed that surgeons began to focus more on specialist skills and lost generalist skills, as robots were increasingly relied upon to perform basic tasks. These examples highlight that job and process redesign (i.e., system level change) is likely to be an important consideration for retained workers when it comes to the integration of robotics into work processes, and the workplace more generally (Agrawal et al. 2021, Bresnahan 2019).

A notable aspect of these prior studies is the focus on knowledge workers. Whereas pharmacists or surgeons have ample opportunity to specialize or upskill, permitting them to offload portions of their tasks to the robot, this is less likely to be true of a warehouse worker, where tasks are frequently commodified. Further, whereas a pharmacist or a surgeon might shift their attention to more value-added work, warehouse workers collaborate with robots under a relatively more equal division of labor, as the tasks performed by either agent are of similar value. As such, while the human worker may guide the pace of work in a knowledge intensive setting, this is less true for physical labor. Accordingly, the pace of work that results may be driven instead by robots, to the worker's detriment. It should therefore come as no surprise that prior scholarship has emphasized nuanced implications when robotics are implemented. Although hazards may decline in general, new and complicated risks can emerge due to process- or system-level changes that come with automation (Vautrin and Dei Svaldi 1986). We next draw on the above observations from the literature to formulate a pair of testable hypotheses.

HYPOTHESES

There is ample reason to expect that robotics will lead to a reduction in the prevalence of traumatic or severe injury in warehousing. This is because, most simply, robots can replace workers in performing dangerous or hazardous tasks (Brosque and Fischer 2022), e.g., carrying heavy loads or handling sharp objects. As discussed, this expectation is the primary motivation for their adoption into many work settings (Murashov et al. 2016). Although the potential exists that traumatic injury may rise as workers operate in proximity to robots (Vasic and Billard 2013), recent advancements in robotic spatial awareness and computer vision (Heo et al. 2019) suggest that those risks are likely to be mitigated. Indeed, Amazon itself often highlights its efforts to institute safety measures in the design of its collaborative warehouse robots (Gantenbein 2023).

In contrast, the presence of robots is known to result in worker boredom, disillusionment, reduced job satisfaction and reduced mental health (Welfare et al. 2019); all of which can lead to behaviors like drug or alcohol abuse (Berg et al. 2023, Gihleb et al. 2022). Although these factors set the stage for inattention on the job and thus serious injury, such effects would be second order in nature. Accordingly, we expect that the presence of robotics will in general be associated with reductions in the rate of severe or traumatic injury in warehouses. Thus:

Hypothesis 1: *Amazon Robotic Fulfillment Centers will have a lower rate of severe injury, as defined by injuries which require days of missed work, as compared with Legacy Fulfillment Centers.*

Our second contention relates to less severe injuries (viz., non-traumatic, typically stress-based injuries). Superficially, employing the logic above, we might expect that workers would again experience a significant decline in these sorts of injuries. Beyond simply replacing workers in hazardous tasks, such as heavy lifting, robots are also well-suited to performing highly repetitive tasks (Murashov et al. 2016). As a result, we might expect that robots would replace humans in performing highly repetitive tasks, reducing ergonomic strain, and in turn the prevalence of less severe injuries, e.g., injuries associated with repetitive motion.

However, several of the observations made above suggest that the opposite may be true (Vautrin and Dei Svaldi 1986). First, many of the un-automated tasks in a warehouse environment are highly repetitive. While robots can handle many basic tasks in this work environment, others remain difficult (e.g. manipulating or picking objects that lack flat surfaces). Second, as numerous scholars have pointed out, robotics eliminate some tasks from the bundle a worker would previously have performed (Frank et al. 2019, Görlich 2010). In a warehousing context, the potential for upskilling or a shift toward more specialized work is limited, hence we may expect a simple reduction in the variety of tasks that workers undertake (Barrett et al. 2012, Parker and Grote 2022), implying greater repetition. Finally, it has been observed that robots may lead to increases in the pace of work, with potentially undesirable consequences for the human workers who continue to perform complementary, non-automated tasks (Antón et al. 2023, Borenstein 2010). Scholars have long argued that automation increases the pace of work as firms seek to capture productivity gains (Gutelius and Theodore 2019). As robots operate quickly, precisely, and without pause, managers are likely to impose greater expectations on throughput (Gutelius and Theodore 2019). As the pace of work rises, ergonomic risks are also likely to increase. Taken together, these observations suggest that the introduction of robotics is likely to lead to a counterintuitive rise in the prevalence of less severe, e.g., stress-based, injuries.

Hypothesis 2: Amazon Robotic Fulfillment Centers will have a higher rate of non-severe injury, as defined by injuries which do not require missed work, as compared with Legacy Fulfillment Centers.

METHODS

Study Context & Data

As discussed, the focus of our analysis is the warehousing industry, namely Amazon's Fulfillment Center network. We draw our data from a variety of sources. We begin by drawing on textual data obtained from a large online community of past and present Amazon Fulfillment Center workers. This data includes worker forum postings which speak anecdotally to workers' perceptions of, and experience with, robotics in Amazon's Robotic Fulfillment Centers.

Second, we employ regression analyses to examine an annual warehouse-level panel of worker injuries. This panel is constructed by combining data from two sources. First, annual establishment-level injury reports from OSHA (2016 – 2020). Establishments meeting industry and size criteria are required to provide OSHA with annual reports of all work-related injuries. From the list of establishments that report injuries to OSHA, we identify Amazon Fulfillment Centers. We then parse the Amazon warehouse identifier code from the reporting establishment name. It should be noted that the resulting sample yields an unbalanced panel of injury reports. This is because some centers were constructed after the start of the sample.

After constructing the set of warehouses, we next identify each warehouse type. To do so we leverage investigative reports from *Reveal.* This journalism examined the largest Amazon warehouses in the United States and leveraged data from a warehousing industry consultancy, MWPVL (Al-Elew and Oh 2020). A deeper discussion of data sources is available in Appendix A. Four warehouse types exist: Sortable warehouses, Large-item warehouses, Receiving Centers, and Other (which are comprised of warehouses with more specialized activities, such as customer return processing centers). Our sample includes 45 Large item centers (L), 12 return centers (R), 58 sortation centers (S), and 26 ‘other’ (O) centers.

This dataset also identifies ‘Robotic Fulfillment Centers’ (as opposed to Legacy Fulfillment Centers). After acquiring Kiva robotics in 2012 Amazon began to introduce robotic-focused warehouse facilities that leverage autonomous robots to complete large portions of day-to-day work. These Robotic FCs employ human workers who operate alongside robots to fulfill tasks that are not yet easy to automate. This indicator of whether a facility is a Robotic or Legacy Fulfillment Center is our key independent variable of interest. Considering descriptive differences between Robotic and Legacy FCs, we observe that Robotic centers tend to have more employees, and more total hours worked (see Figure E1). Combining these data, we are left with a panel comprised of 141 distinct Amazon warehouses.

Next, using each warehouse’s address, we map it to local census records. From this we obtain the population of the warehouse’s census tract as well as measures of per capita wealth (median household income) and education (fraction of residents that has attended at least some college). Table 1 contains the main variables that enter this panel, including variable definitions and summary statistics. Here, it should be noted that a given establishment does not necessarily appear in the OSHA sample in each year following its construction, as a given warehouse may move in and out of reporting requirements. Our panel thus includes 89 observations in 2017, 105 observations in 2018, 139 observations in 2019, and 135 observations in 2020. However, repeating our estimations (described below) while omitting warehouses for which the panel is ‘interrupted’ (i.e., where there are missing values for establishment injury reports following a warehouse’s first report in the data) yields consistent results. Figure 1 depicts a map of the locations of the Amazon warehouses that comprise the panel. A correlation matrix for the variables in this estimation sample is presented in Figure 2.

-- INSERT TABLE 1 HERE --

-- INSERT FIGURE 1 HERE –

-- INSERT FIGURE 2 HERE --

Third, we construct another panel where each observation reflects a triple of warehouse, month, and injury type (sprain/strain vs. other), based on injury level log data obtained from 26 Amazon Fulfillment Centers for the year 2018. These logs were sourced from the same investigative journalism conducted by *Reveal* news. This second panel includes an indicator of whether a warehouse is a Robotic or Legacy facility. Descriptive statistics associated with this second panel are in Table 2. As can be seen, approximately 60% of these warehouses are Robotic Fulfillment Centers. Of note is that sprains are by far the most common injury type among these warehouses. The frequency of sprains and strains is 3-4 times all other injury types combined.

-- INSERT TABLE 2 HERE --

Finally, to address the concern that there is no change in warehousing type (i.e., all FCs are built as either Legacy or Robotic Centers with no retrofitting), we construct a panel at the state-year level. This panel captures injury rates collectively reported among all privately owned establishments operating in a given state's Warehousing and Storage sector (NAICS 493), *i,* in a given year, *t*. The injury rate measures, obtained from the BLS's Survey of Occupational Injury and Illness (SOII), capture the count of injuries per 100 full time workers, both for injuries involving job transfer and injuries involving missed work. This panel also includes a binary indicator of whether the state, *i*, has received any Robotic FCs as of year, *t.* Descriptive statistics for the state-year panel are reported in Table 3, below.

-- INSERT TABLE 3 HERE --

Empirical Approach

Textual Data: We begin our analysis by exploring the textual data supplied by past and present employees of Amazon Fulfillment Centers who participate in a large community on Reddit. These workers' posts offer a natural source of insight for understanding worker perceptions of robotics in Fulfillment Centers. We first consider an explicit poll that was conducted among community members, soliciting workers' opinions on the relative safety and 'interestingness' of stock picking work conducted at Amazon's Legacy versus Robotic FCs. We also delve into the text of posts and comments.

We collected all posts and comments from the sub-reddit made between 2022 and 2023. We then filtered the posts based on keywords, identifying the subset that make explicit mention of the terms workers use to describe FCs that make use of Robotics and those that do not. They key identifying terms are 'robot' and 'legacy' or 'traditional'. This process yielded roughly 1,800 posts that referenced Robotic FCs, and another 500 posts that referenced traditional or legacy FCs. After pre-processing the text, i.e., converting to lower case, removing URLs, punctuation, and stop words, we calculate the frequency of all pairs of words appearing alongside one another in a post or comment (i.e., bigrams). Based on their relative frequency, we then constructed word clouds based on the top 50 most frequent bigrams appearing in posts that referenced Robotic FCs, and the top 50 most frequent bigrams appearing in posts that referenced Legacy FC. Further details of this process can be found in Appendix B of the online supplement.

Quantitative Data: To execute our quantitative analysis of establishment-level injury reports published by OSHA we estimate the relative difference in the frequency of various injuries across Robotic and Legacy FCs. Our initial estimations are based on Equation (1), below, where i indexes warehouses and t indexes years. We consider several outcome measures,

Y, beginning with general volumes of injury. OSHA defines a workplace injury as "any wound or damage to the body resulting from an event in the work environment.'' It is important to note that these data are self-reported by each establishment.

We then consider injuries by level of severity. We distinguish injuries that required employees to miss work altogether and those requiring temporary job transfer or reassignment. Our logic is simple. Injuries that necessitate worker absence are more likely to be serious because traumatic injuries require rest or on-going medical care. This is consistent with prior work. Indeed, prior research conducted by the BLS has reported, for example, that head injuries are more likely to result in missed work, whereas hand injuries are more likely to result in temporary job transfer (BLS 2015). Note that these are not mutually exclusive classifications, because a single injury may require both missed work and subsequent job transfer. Further, we consider measures of total employee-days of work missed due to injury, and total employee-days of work reassignment due to injury. In addition to simple linear regressions, as expressed in Equation (1), we also estimate a model that explicitly addresses the non-negative count nature of the various outcomes, namely a bivariate Poisson regression (Xu and Hardin 2016). Standard errors are robust and clustered on the warehouse. Formally:

$$Y_{i,t} = Robotics_i + X_{i,t} + L_i + WHType_i + Year_t + State_i + \epsilon_{i,t} \qquad (1)$$

Our key independent variable of interest in these regressions is *Robotics*, a binary indicator of whether the warehouse is a Robotic Fulfillment Center. In these models, we control for time-varying levels of workload and labor supply at each warehouse (reflected by X in the equation), measures of annual average employee volumes, and total hours worked. We also control for population characteristics, reflected by L in the equation, based on the American Community Survey. These measures include census tract population, median household income, and the fraction of the local population having at least some college education. Finally, we include fixed effects for warehouse type (which may be Sortable warehouses, Large-item warehouses, Receiving Centers, or Other), year, and state. Year fixed effects account for country-wide unobservable shocks to working conditions (e.g., COVID-19), while state fixed effects account for local unobservable factors at the state level, e.g., local regulation or policies regarding workers' rights. Note that, in deference to multicollinearity concerns (i.e., poison controls), we also estimate these models with subsets of the controls. Results remain consistent.

We then explore the relationship between robotics and injury in a more granular manner by evaluating the effect of an exogenous increase in the pace of work, driven by the occurrence of specific sales events (Black Friday / Cyber Monday and Amazon Prime Day) on the volume and mix of injuries that occur at a given warehouse. We consider the monthly volume of injuries arising at warehouses, by injury type, leveraging the injury-level log data for 26 Amazon warehouses in 2018. These more detailed injury logs allow us to operationalize injury volumes seasonally, and to distinguish injuries that are plausibly due to repetitive stress, (e.g., sprains and strains), from those that are due to trauma, (e.g., cuts, bruises, fractures, concussions). To exploit within-year seasonality in operational load, we examine shifts in the volumes of different injuries, across Robotics and Legacy Fulfillment Centers, during periods of higher load, namely Black Friday, Cyber Monday, and Amazon Prime Day in November and June. We then estimate a triple difference regression (Olden and Møen 2022) using Equation (2), where i indexes warehouses, j indexes injury types, and t indexes months. Here, τ captures our vector of month fixed effects, and δ captures our vector of warehouse fixed effects. Note that the main effect of our *Robotics* Fulfillment Center indicator is not identified in this regression as it is subsumed by our warehouse fixed effects, δ.

$$InjuryVolume_{i,j,t} = \tau_t + Sprain_j + \tau_t \cdot Robotics_i + \tau_t \cdot Sprain_j + Robotics_i \cdot Sprain_j + \tau_t \cdot Robotics_i \cdot Sprain_j + \delta_i + \epsilon_{i,j,t} \quad (2)$$

Finally, to ensure the robustness of our main results to the lack of within warehouse change in robotics, and better-establish the causal nature of the impact of robotic facilities on injuries, we change the unit of analysis to the state's warehousing sector. In doing so, we use a state-year panel instead of an establishment-year panel. Such an approach is appealing because Amazon FCs are typically the largest warehousing operation in a geography and their entry has detectable effects on local labor markets (Pathania and Netessine 2022, Rudolph et al. 2023). With this analysis, we treat the first construction of a Robotic FC as a shock to the injury rates collectively reported among all privately-owned establishments operating in a given state's Warehousing and Storage sector (NAICS 493), *i,* in a given year, *t*. Our state-sector injury rate measures are obtained from the BLS's Survey of Occupational Injury and Illness (SOII).

Using the resulting state-year panel, we estimate a difference-in-differences regression, as per Equation (3). The variable *AnyRoboticFCs* is a binary treatment indicator of whether a state*, i,* has received any Robotics FCs as of year, *t*. Given the staggered nature of the treatments, we perform this estimation using the two-stage estimator proposed by Gardner (2022).

$$Log(InjuryRate_{i,t}) = AnyRoboticFCs_{i,t} + \tau_t + \delta_i + \epsilon_{i,t} \qquad (3)$$

RESULTS

Anecdotal and Textual Evidence from Reddit

We first report results from the subreddits focused on Amazon Fulfillment Center. While not statistical in nature, this evidence is valuable as it provides a check on the face validity of our hypotheses. We begin with the descriptive poll, wherein workers with experience at Robotic and Legacy Fulfillment Centers were asked about their perceptions of the safety and 'interestingness' of stock-picking work that is supported by robots versus not. The poll received 148 responses.

Most workers (N = 93/148, or 63%) reported a belief that stock picking is safer when performed at a Robotic Fulfillment Center. However, workers reported mixed opinions regarding whether the work was more interesting when conducted at a Robotic Fulfillment Center.

More nuanced details were also available in workers' unstructured posts. Several anecdotal posts discussed the differences in the nature of the work conducted with and without robotic support. Many workers observed that the work was less physically stressful at Robotic FCs, but also more repetitive and mentally challenging. As noted by one worker: "the robot-enabled sites are typically less walking [sic], but still physical because picking/stowing is much faster. Many people burn out from the monotonous nature of it. Some people like it and just zone out." A second worker stated that stock picking at a Robotic FC is "not physically exhausting. It's mentally [exhausting]." A third stated that picking at a Robotic FC is "better than working at a legacy FC … [where] we easily clock 20k steps a day." Finally, numerous posts indicated that productivity expectations are higher at Robotic FCs. Posts indicated that management expected pick rates 2-3 times higher at Robotic versus Legacy FCs (300-450 items per hour versus 100 items per hour). These descriptive data are consistent with our hypotheses, as they suggest workers in Robotic FCs perceive their work environment to be: i) less physically demanding, ii) safer, iii) more prone to boredom and distraction, and iv) faster paced.

To increase the rigor of this analysis, we turn to the isolated bigrams, and we calculate their frequencies to create word clouds. The resulting word clouds for Robotic FC and Legacy FC posts are in Figures B1 and B2 of Supplementary Appendix B. Several interesting contrasts emerge from these visuals. First, posts related to Legacy FCs are much more likely to make mention of walking. That is, associates discussing Legacy FCs frequently mention that their job involves "walking around" for several "miles [a] day." Second, Legacy FC posts more frequently mention safety equipment, specifically “safety shoes,” which presumably serve as protection for workers feet, e.g., in the face of potentially dropping “large items.”

In contrast, posts associated with Robotic FCs make no mention of walking, nor any mention of safety equipment. Instead, these posts highlight concerns about “mental health,” standing in “one spot,” while nonetheless performing “hard work,” under the threat of receiving “write ups” from management. Amazon’s write-up policy flags workers for various behaviors, most notably failure to meet stock picking quotas. Media coverage of Amazon has explicitly discussed Amazon’s write-up policy, and its practice of penalizing workers for operating too slowly (Dastin 2022, Pottenger 2020). Taken in sum, these anecdotal reports open the door to the fact that injury rates and types might vary drastically across the two types of facilities.

Annual Injury Panel

Our main regression results, employing ordinary least squares (OLS) regression, are in Table 4. In Model 1, we initially see that robotics are not significantly associated with the volume of workplace injuries occurring at a warehouse in an average year. The only significant associations we observe are, intuitively, that injuries are more prevalent when a warehouse has more employees, and more employee hours worked. This is consistent with the scope of work overall. However, Models 2 through 5 indicate that these baseline results mask important heterogeneity. Models 2 and 3 indicate that robotics are associated with a reduction in injuries that require employees to miss work, whereas Models 4 and 5 indicate that robotics are associated with an increase in injuries that require employees to be temporarily assigned to different roles within the facility. This pattern of results is consistent with the anecdotal evidence reported earlier and underscores the notion that while the incorporation of robotics is associated with a reduced frequency of severe, traumatic, or life-threatening injuries, it is concomitantly associated with a rise in less-severe injuries.

-- INSERT TABLE 4 HERE --

The effects we observe are practically meaningful. The average number of annual injuries involving days of missed work is approximately 58, whereas the average reported volume of injuries involving days of job transfer is approximately 40. The estimated effect of robotics on days of missed work thus translates to decline equal to roughly 27% of the average, whereas the effect on days of job transfer translates to a nearly 70% rise over the average.

We next repeat the analysis employing a bivariate Poisson regression. This has two main benefits. First, this estimator can accommodate the non-negative count nature of the outcomes we study. Second, it enables us to explore the joint effects of robotics on each outcome, in tandem. Results are in Table 5 and remain consistent. Estimates indicate that robotics are simultaneously associated with a reduction in injuries requiring work absences and a rise in injuries requiring job transfer or job restriction. The estimates in Table 5 also have the benefit of being interpretable as elasticities, enabling a comparison of relative magnitudes. The coefficient associated with injuries involving work absences translates to an approximate 40% reduction ($\exp(-0.499) - 1 = -0.393$), while the coefficient associated with injuries involving job transfer or restriction translates to an approximate 70% increase.

-- INSERT TABLE 5 HERE --

Monthly Injury Panel by Injury Type

Before discussing results from the difference-in-difference-in-differences (DDD) estimates, we consider some descriptive details. In Appendix C we present a plot depicting the proportion of 2018 injuries that are of a particular type, e.g., sprain/strain, bruise, cut, fracture, by warehouse type (Robotic vs. Legacy). Sprains and strains are the most common injury type. More importantly, such injuries are systematically more common in Robotic facilities.

Further, we present a descriptive plot of the monthly average trends in injury volumes, contrasting Robotic and Legacy centers, as well as injury types (Figure 3). The lines of best fit reflect Locally Weighted Exponential Scatterplot Smoothing (LOESS) curves, fit to the four injury samples (Robotics vs. Not * Sprains vs. Other), relating total injury volumes to month of year. The initial pattern is consistent with expectations. In general, sprains and strains become systematically more prevalent in Robotic Fulfillment Centers, particularly during periods of high operational load, i.e., June and July, in proximity to Amazon Prime Day, and November, in proximity to Black Friday, Cyber Monday, and the holiday shopping season.

We depict the triple difference estimates graphically in Figure 4, based on Equation (2). Because our panel includes just 26 warehouses, we employ wild cluster bootstrap standard errors when constructing our 95% confidence intervals, to avoid bias that can arise with a small number of clusters (MacKinnon and Webb 2018). March is the reference period, as it appears to be a relative low point in injury volumes. Consistent with the plot, we see that, relative to that month, Robotic FCs exhibit statistically significant increases in the volume of sprain injuries, relative to Legacy FCs and non-sprain injuries, in June and July, as well as November and December. As such, our results are consistent with the notion that Robotic Fulfillment Centers exhibit spikes in injuries associated with repetitive motion and stress under periods of high operational load.

-- INSERT FIGURE 3 HERE --

-- INSERT FIGURE 4 HERE --

State-Year Analysis

We conclude with our analysis of Robotic FC construction and the subsequent impact on injury rates reported in each state's warehousing and storage sector. Details of the panel are in

Appendix D. Employing the state-sector as our unit of analysis enables us to assess our main research question in a different way. As Amazon FCs are generally large and are known to have substantial impacts on local labor markets, it is plausible that their construction will have a material impact on state-level injuries in the warehousing sector. We investigate this question using a difference in difference approach. In doing so, we cast Robotic FC construction as a shock to state injury rates.

Our estimator is the two-stage difference-in-differences approach proposed by Gardner (2022), which addresses potential bias that can arise when treatment is staggered across units and time (Roth et al. 2023). Results are in Figure 5. As can be seen, estimates indicate that the initial construction of a Robotic FC leads to a significant rise in injuries that involve job transfer, with a weaker insignificant effect on the rate of injuries involving days of missed work. Moreover, pre-treatment trends across treated and untreated locations are not significantly different from each other, suggesting that the parallel trends assumption is satisfied. These results indicate a significant effect of Robotic FC construction on a state's rate of injury involving job transfer, in contrast to a weaker relationship with more traumatic injuries, consistent with our prior findings. Note that, repeating these estimations taking the entry of the first Legacy FC as treatment, we observe no clear pattern of effects; that is, we observed no statistically significant increase in injuries of either type of the 3 years following first FC construction.

-- INSERT FIGURE 5 HERE –

ROBUSTNESS CHECKS

Placebo Tests

We first estimate a series of models to examine the correlation between the presence of robotics and the volume of illnesses arising at an establishment. Intuitively, these conditions are not plausibly related to robotics. In doing so, we estimate the relationship between robotics and the volume of poisonings, respiratory conditions, skin conditions, and hearing loss. We estimate each model employing the same set of controls as earlier. We estimate all models employing OLS regression, clustering standard errors by warehouse. Results are in Table E1 of the Appendix, where we observe no significant relationships between robotics and any illness type.

Inverse Propensity for Treatment Weighting

Next, we report the results of our primary regressions employing a reweighted sample. We implement the reweighting based on covariate-balancing propensity scores (Imai and Ratkovic 2014), employing all available measures, including annual average employees, total hours worked, warehouse type, the size of the local population, the median household income, the fraction of residents with some college education, and the year. Our reweighting procedure reduces imbalance across the set of observable covariates, as shown in Figure E1 of the Appendix. Importantly, all standardized mean differences are less than or equal to 0.2, typically considered a lower threshold for a 'small' effect size with standardized mean differences, i.e., Cohen's *d* (Cohen 2013).

Replicating our main estimations with the reweighted sample, we observe consistent results (Table E2 in the Supplement). We see an insignificant relationship between the presence of robotics and the average annual number of injuries reported by an establishment. However, once we separate injuries into those that require work absence and those that require job transfer, we observe the same set of countervailing outcomes. That is, the presence of robotics is significantly and negatively associated with the volume of reported injuries that involve work absences, yet significantly and positively associated with the volume of reported injuries involving job transfer. We observe the same pattern of effects around the total number of days of injury work absences and injury job transfers.

Alternative Regression Specifications

As another robustness check, we consider alternative regression specifications. Specifically, rather than merely condition on hours worked, we consider an estimation wherein we normalize injury volumes by the total employee hours worked at a facility. We accomplish this by re-estimating our injury count models using Poisson regression, and incorporating the total hours worked (in 1000s) as an exposure term. As can be seen in Table E3, results remain consistent. Further, in Table E4, we repeat our estimation employing a log-linear regression with a rate outcome, i.e., explicitly normalizing the injury outcomes with respect to annual average employees. Results remain consistent.

Survival Model

There is also the natural concern that Robotic Fulfillment Centers are not assigned at random, and that pre-existing injury activities might even drive Amazon to construct a Robotic rather than Legacy FC in a location. To assess this, we estimate a Cox Proportional Hazard model using the state-year panel, wherein a binary indicator of the first Robotic FC construction in a state serves as the failure event. We examine whether states' contemporaneous and lagged rates of worker injury involving job transfer and missed work significantly affects the hazard of robotics use. Results of are in Appendix Table D1. We observe no significant relationship between either the contemporaneous or lagged injury rates and the hazard of Robotic FC construction. This null result provides no evidence that Amazon is selective about its use of Robotics in a way that would associate with the pre-existing nature of injury activity in the warehousing sector of a given geography.

DISCUSSION

In this work we examine the relationship between robotic automation and worker health in the warehousing industry, focusing specifically on Amazon Fulfillment Centers. We provide evidence that while the adoption of robotics is associated with a reduction in severe injuries, it is also associated with a significant increase in less-severe injuries. The increase in less-severe injuries can likely be attributed to a simultaneous decline in task variety and increase in the pace of work imposed on workers.

As discussed, these findings have practical and theoretical implications. On the practical side, our work has implications for organizations and regulators. Adopting organizations might address the rise in non-severe injuries in several ways. Note that non-traumatic (less severe) injuries are not accidents; rather, they arise from the normal performance of a task, albeit at a potentially higher pace. Taking this into consideration, it seems apparent that addressing robotics in the workplace will require careful attention to task and job design, to account for increases in the pace of work. This may involve more frequent task rotation or workforce cross-training, as we note elsewhere.

In a warehousing context, cross-training could involve training workers to operate different types of machinery, such as forklifts, conveyors, or automated storage and retrieval systems. This would enable workers to move between different roles within the warehouse, should the need arise, e.g., due to injury. Cross-training could also involve training workers in different aspects of the fulfillment process, such as packing or shipping. This would enable workers to perform different tasks and to have greater task variety, which may help to prevent injuries from repetitive motion, overexertion, or monotony.

There are further implications for regulators. To the extent that workplace injuries are a reality of the warehousing setting, one interpretation of our findings might be that investment in robotics simply displaces one type of injury for another, thereby maintaining the status quo. Yet, to the extent that our findings also suggest that non-traumatic injuries spike during periods of higher load, it appears that firms like Amazon are balancing the pace of work (to capture the efficiencies of robotics) and worker safety. And while it is beyond the scope of this scholarship to determine what the optimal balance between economic output and worker safety is, should regulators determine that worker welfare is being degraded, our findings suggest that worker protections may be warranted. These could come in numerous forms, from federal intervention via OSHA's administrative rulemaking process, to legislation under Congress' ability to regulate interstate commerce, to state level statutory responses. Such regulations could take the form of simple fines / penalties, or more comprehensive regulation of the interaction between workers and robots, to shift the managerial calculus in favor of worker health. Regulation could further include requirements for training, safety inspections, and penalties for unsafe working conditions.

Concomitantly, our findings provide empirical evidence to support the recent trend in collective bargaining at warehousing facilities. Should workers believe the tradeoff between safety and productivity is exploitive, unionization is the traditional means by which workers can increase their ability to sculpt firm policy. Not only would this facilitate the ability of workers to address grievances through the means available to organized labor (e.g., general strikes, worker slowdowns, arbitration) it might also benefit warehousing firms in the long term by reducing turnover. This has the potential to lower administrative overhead for firms (by limiting the scope of onboarding and training for new workers) and mitigate any tortious liability that worker safety exposes the firm to. And while further research is needed to comprehensively understand this calculus for firms, the case for collective bargaining is evident.

Finally, there seems to be no shortage of workers flocking to warehouse jobs, despite the ready availability of worker job evaluations, e.g., on GlassDoor, or in ad hoc venues like Reddit. This suggests a lack of visibility to the problems we document here, or a lack of appreciation for the gravity of the problem. In terms of academic contribution, our work provides a novel consideration of the implications of new forms of robotics for worker safety, in entirely new working environments. While canonical work in this space has considered the effect of automation on worker safety in contexts like assembly lines (Acemoglu and Restrepo 2020, Graetz and Michaels 2018), or examined downstream productivity (Bessen 2015, Bessen 2016, Dixon et al. 2021), little attention has been paid to emerging forms of robotics which work alongside human collaborators. We are thus able to extend not only the scope (context) of consideration, but the style of robots investigated, and the nature of injuries which accrue.

As robotic technology continues to advance, understanding how these new forms of robotics influence worker safety will be critical, especially as they begin to spill over into other industries. Moreover, as the impacts on worker health are highly contextual, it will be important for academics, firms, workers, and regulators to take stock of the impact on worker well-being in each setting, and adapt accordingly. As mentioned, we hope this work serves as a call for increased attention on physical robotic automation and worker safety.

Our work also provides a novel perspective on the relationship between automation, robotics, and worker safety, considering spillovers between tasks at the level of a broader organizational process. Whereas prior work considering the implications of robotics and automation for worker safety has frequently focused on the individual worker (Gihleb et al. 2022, Vasic and Billard 2013), our study highlights the importance of considering the design of the working environment and highlights the importance of coordinating activities between various agents, whether robot or human, i.e., system-level change (Agrawal et al. 2021, Bresnahan 2019). The key insight is that the introduction of robotics may present firms with an incentive to sacrifice worker health to reap the productivity benefits of robotic automation. To the extent possible, future work should explore approaches to mitigate that incentive, whether through modified job design, advancements such as health monitoring technologies, or alternative enforcement mechanisms.

Our work is of course subject to some limitations, which provide rich opportunities for future work. First, our dataset allows us to observe a wide range of injuries in the workplace. That said, we have limited visibility into detailed task performance of a particular worker, and how that task connected with robotics to drive injury. As such, we are unable to offer concrete guidelines or insights about task design to best accommodate the introduction of robotics into the work environment. Future work might look to address this gap via alternative methods, e.g., ethnographic observation (Burtch et al. 2010).

Second, we observe, in aggregate (at the warehouse level) the relationship between injury types and robotic automation. However, our analysis lacks insight into the characteristics of the worker population at a given warehouse, and lacks insight into the health progression of workers over time, as they encounter robotics. As the individual worker is the focus, we hope future work will examine how workers' entry or exit from a roboticized work environment affects worker health over time. Relatedly, although we have confirmed verbally with Amazon representatives that the nature of workflows and processes at Robotic and Legacy FCs are similar, aside from the presence of robotics and conditional on the Fulfillment Center type, we lack any direct means of verifying the extent to which this is true in practice. Accordingly, it is possible that Robotic FCs may differ in unobserved ways beyond the presence of robotics, which may contribute to the effects we observe.

Third, several of our datasets are not comprehensive. Our establishment-level OSHA data is not comprehensive, as it applies only to establishments that meet reporting requirements. Similarly, our injury-level OSHA 300 data pertains to 26 specific warehouses, for which Reveal News was able to obtain the logs via current or past employees. And, our state-level warehouse sector injury panels are subject to missingness, as not all states report injury rate metrics in all time periods for NAICS 493. Given these issues, our results should be interpreted with caution, as it is possible that inclusion of unobserved establishments and states in all years might influence our estimates and conclusions.

Finally, given the data limitation, we cannot speak to why one warehouse in a particular location is designed as a Legacy Fulfillment Center, while another is designed as Robotic. Our triple difference approach, as well as matching, limit the effects of this limitation to some extent, but it nevertheless remains a concern. Future work might consider not only why organizational bodies designate one warehouse as one and one warehouse as another, but how transitioning between the two affects workers.

CONCLUSION

It is critical that firms be transparent about the risks and benefits associated with the integration of robotics into workplace processes, as well as the measures they are taking to manage these risks. This transparency can build trust among workers and prevent both conflict and resistance to the adoption of new technologies. To manage risks and foster trust, collaboration will be key. As with any significant change in the workplace, the adoption of robotics requires a cultural shift. And firms must foster a culture of safety and continuous improvement to ensure that worker safety remains a top priority in the face of rapid technological change. Firms should work closely with workers, unions, regulators, and other stakeholders to ensure that the adoption of new technologies is managed in a safe and responsible manner that prioritizes worker health and safety. This study highlights the need for further research to better understand the complex relationship between robotic automation and worker health and safety. This may involve conducting detailed studies of workplaces or longitudinal studies that track the effects of automation over time. It is our hope that this work can serve as a stepping-stone to future work in the nascent space of robotics.

REFERENCES

Acemoglu, D., Restrepo, P. 2020. Robots and Jobs: Evidence from Us Labor Markets. *Journal of political economy*. 128(6) 2188-2244.

Agrawal, A.K., Gans, J.S., Goldfarb, A. 2021. *Ai Adoption and System-Wide Change*. National Bureau of Economic Research.

Al-Elew, M., Oh, S. 2020. *What Are Injury Rates Like at Amazon Warehouses?* Reveal News, https://revealnews.org/article/amazon-injury-rates/.

Amazon.com. 2022. *10 Years of Amazon Robotics*https://www.aboutamazon.com/news/operations/10-years-of-amazon-robotics-how-robots-help-sort-packages-move-product-and-improve-safety.

Antón, J.I., Fernández-Macías, E., Winter-Ebmer, R. 2023. Does Robotization Affect Job Quality? Evidence from European Regional Labor Markets. *Industrial Relations: A Journal of Economy and Society*. 62(3) 233-256.

Autor, D.H. 2015. Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *Journal of economic perspectives*. 29(3) 3-30.

Barrett, M., Oborn, E., Orlikowski, W.J., Yates, J. 2012. Reconfiguring Boundary Relations: Robotic Innovations in Pharmacy Work. *Organization science*. 23(5) 1448-1466.

Beane, M. 2019. Shadow Learning: Building Robotic Surgical Skill When Approved Means Fail. *Administrative Science Quarterly*. 64(1) 87-123.

Bekhtiar, K., Bittschi, B., Sellner, R. 2021. *Robots at Work? Pitfalls of Industry Level Data*. WIFO Working Papers.

Berg, J., Green, F., Nurski, L., Spencer, D.A. 2023. Risks to Job Quality from Digital Technologies: Are Industrial Relations in Europe Ready for the Challenge? *European Journal of Industrial Relations* 09596801231178904.

Bessen, J. 2015. Toil and Technology: Innovative Technology Is Displacing Workers to New Jobs Rather Than Replacing Them Entirely. *Finance & Development*. 52(001) 16.

Bessen, J.E. 2016. How Computer Automation Affects Occupations: Technology, Jobs, and Skills. *Boston Univ. school of law, law and economics research paper*(15-49).

BLS. 2015. *Occupational Injuries and Illnesses: A Pilot Study of Job-Transfer or Workrestriction Cases, 2011–2013*. Bureau of Labor Statistics, https://www.bls.gov/opub/reports/job-transfer-or-work-restriction/archive/a-pilot-study-of-job-transfer-or-work-restriction-cases-2011-2013.pdf.
BLS. 2023a. *Employer-Reported Workplace Injuries and Illnesses – 2021-2022*. Bureau of Labor Statistics, https://www.bls.gov/news.release/pdf/osh.pdf.
BLS. 2023b. *Warehousing and Storage: Naics 493*. US Bureau of Labor Statistics, https://www.bls.gov/iag/tgs/iag493.htm.
Borenstein, J. 2010. Work Life in the Robotic Age. *Communications of the ACM*. 53(7) 30-31.
Bresnahan, T. 2019. Artificial Intelligence Technologies and Aggregate Growth Prospects. *Prospects for Economic Growth in the United States*.
Brosque, C., Fischer, M. 2022. Safety, Quality, Schedule, and Cost Impacts of Ten Construction Robots. *Construction robotics*. 6(2) 163-186.
Burtch, G., Yoo, Y., Weiss, A. 2010. Digital Innovation and Craftsmanship: The Case of Cf Martin & Company.
Cohen, J. 2013. *Statistical Power Analysis for the Behavioral Sciences*. Academic press.
Dastin, J. 2022. *Insight: Amazon Issued 13,000 Disciplinary Notices at Single U.S. Warehouse*. Reuters, https://www.reuters.com/technology/amazon-issued-13000-disciplinary-notices-single-us-warehouse-2022-07-12/.
Dixon, J., Hong, B., Wu, L. 2021. The Robot Revolution: Managerial and Employment Consequences for Firms. *Management Science*. 67(9) 5586-5605.
Frank, M.R., Autor, D., Bessen, J.E., Brynjolfsson, E., Cebrian, M., Deming, D.J., 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*. 116(14) 6531-6539.
Gantenbein, D. 2023. *Making Robots Better Co-Workers*. Amazon Science, https://www.amazon.science/latest-news/making-robots-better-co-workers.
Gardner, J. 2022. Two-Stage Differences in Differences. *arXiv preprint arXiv:2207.05943*.
Gihleb, R., Giuntella, O., Stella, L., Wang, T. 2022. Industrial Robots, Workers' Safety, and Health. *Labour Economics*. 78 102205.
Görlich, D. 2010. *Complementary Tasks and the Limits to the Division of Labour*. Kiel Working Paper.
Graetz, G., Michaels, G. 2018. Robots at Work. *Review of Economics and statistics*. 100(5) 753-768.
Greene, J., Alcantra, C. 2021. *Amazon Warehouse Workers Suffer Serious Injuries at Higher Rates Than Other Firms*. The Washington Post, https://www.washingtonpost.com/technology/2021/06/01/amazon-osha-injury-rate/.
Gualtieri, L., Rauch, E., Vidoni, R. 2021. Emerging Research Fields in Safety and Ergonomics in Industrial Collaborative Robotics: A Systematic Literature Review. *Robotics and Computer-Integrated Manufacturing*. 67 101998.
Gutelius, B., Theodore, N. 2019. The Future of Warehouse Work: Technological Change in the Us Logistics Industry. *UC Berkeley Center for Labor Research*.
Heo, Y.J., Kim, D., Lee, W., Kim, H., Park, J., Chung, W.K. 2019. Collision Detection for Industrial Collaborative Robots: A Deep Learning Approach. *IEEE Robotics and Automation Letters*. 4(2) 740-746.

Hunter, L.W., Bernhardt, A., Hughes, K.L., Skuratowicz, E. 2001. It's Not Just the Atms: Technology, Firm Strategies, Jobs, and Earnings in Retail Banking. *ILR Review*. 54(2A) 402-424.
Imai, K., Ratkovic, M. 2014. Covariate Balancing Propensity Score. *Journal of the Royal Statistical Society Series B: Statistical Methodology*. 76(1) 243-263.
Karwowski, W., Rahimi, M., Mihaly, T. 1988. Effects of Computerized Automation and Robotics on Safety Performance of a Manufacturing Plant. *Journal of Occupational Accidents*. 10(3) 217-233.
Lerch, B. 2020. Robots and Non-Participation in the Us: Where Have All the Workers Gone? *Available at SSRN 3650905*.
MacKinnon, J.G., Webb, M.D. 2018. The Wild Bootstrap for Few (Treated) Clusters. *The Econometrics Journal*. 21(2) 114-135.
Murashov, V., Hearl, F., Howard, J. 2016. Working Safely with Robot Workers: Recommendations for the New Workplace. *Journal of occupational and environmental hygiene*. 13(3) D61-D71.
Olden, A., Møen, J. 2022. The Triple Difference Estimator. *The Econometrics Journal*. 25(3) 531-553.
Parker, S.K., Grote, G. 2022. Automation, Algorithms, and Beyond: Why Work Design Matters More Than Ever in a Digital World. *Applied Psychology*. 71(4) 1171-1204.
Pathania, V., Netessine, S. 2022. The Impact of Amazon Facilities on Local Economies. *Available at SSRN 4116645*.
Peshkin, M., Colgate, J.E. 1999. Cobots. *Industrial Robot: An International Journal*. 26(5) 335-341.
Pillai, R., Sivathanu, B., Mariani, M., Rana, N.P., Yang, B., Dwivedi, Y.K. 2022. Adoption of Ai-Empowered Industrial Robots in Auto Component Manufacturing Companies. *Production Planning & Control*. 33(16) 1517-1533.
Pottenger, K. 2020. *Speed-Ups, Write-Ups, and Profits up at Amazon*. New Labor Forum, https://newlaborforum.cuny.edu/2020/07/16/speed-ups-write-ups-and-profits-up-at-amazon/.
Roth, J., Sant'Anna, P.H., Bilinski, A., Poe, J. 2023. What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature. *Journal of Econometrics*.
Rudolph, K.E., Williams, N.T., Milazzo, F., Venkataramani, A., O'Brien, R. 2023. Has the Opening of Amazon Fulfillment Centers Affected Demand for Disability Insurance? *Plos one*. 18(11) e0294453.
Vasic, M., Billard, A. 2013. *Safety Issues in Human-Robot Interactions*. IEEE.
Vautrin, J., Dei Svaldi, D. 1986. Manipulating Industrial Robots in France—Effects on Health, Safety and Working Conditions: Results of the Inrs—Cram Survey. *Journal of Occupational Accidents*. 8(1-2) 1-12.
Welfare, K.S., Hallowell, M.R., Shah, J.A., Riek, L.D. 2019. *Consider the Human Work Experience When Integrating Robotics in the Workplace*. IEEE.
Xu, X., Hardin, J.W. 2016. Regression Models for Bivariate Count Outcomes. *The Stata Journal*. 16(2) 301-315.

TABLES

Table 1. Annual Panel: Variables and Descriptive Statistics

N	Variable	Variable Description	Mean	S.D.	Min.	Max.
468	Injuries	Total number of work-related injuries, by facility and year.	106	83.8	0	566
468	InjuriesMissedWork	Total number of work-related injuries where at least one day of work was missed, by facility and year.	57.9	57.3	0	384
468	InjuriesJobTransfer	Total number of work-related injuries where at least one day of job transfer was required, by facility and year.	40.2	52.6	0	338
468	DaysMissedWork	Total days of missed work due to workplace injuries, by facility and year.	2,981	3,531	0	28,164
468	DaysJobTransfer	Total days of job transfer due to workplace injuries, by facility and year.	2,960	4,112	0	28,004
468	Robotics	A binary indicator of whether this is a ‘robotic fulfillment center’.	0.342	0.475	0	1
468	Year	Integer value indicating the year of the annual report.	2019	1.08	2017	2020
468	HoursWorked	Total number of employee hours worked, by facility and year (in 1000s)	2,989.23	2,153.48	43.07	9,511.09
468	AvgEmployees	Average number of individuals employed, by facility and year	1,953	1,388	21	6,263
468	Poison	Total number of job-related poisonings, by facility and year.	0.013	0.113	0	1
468	Repiratory	Total number of respiratory conditions arising from work, by facility and year.	0.041	0.228	0	2
468	Skin	Total number of skin disorders arising from work, by facility and year.	0.004	0.065	0	1
468	Hearing	Total number of cases involving hearing loss, arising from work, by facility and year.	0.021	0.159	0	2
468	Population	Total population in the surrounding census tract	4,497	1,953	0	14,513
446	MedianHHI	The median household income in the surrounding county (in $s).	80,063	29,525	24,688	178,750
450	CollegeEducated	The proportion of residents who have some college education in the surrounding county.	0.408	0.115	0.107	0.767

Table 2. Descriptive Statistics and Variable Definitions for Monthly Stacked Panel Based on Injury-level Log Data for 26 Fulfillment Centers in 2018

N	Variable	Description	Mean	S.D.	Min.	Max.
262	Sprain	A binary indicator of whether the observation in question pertains to sprain injuries or not.	0.50	0.50	0.00	1.00
262	InjuryVolume	Total number of injuries of type j in month t for warehouse i.	7.41	9.35	0.00	95.00
262	Month	The month of the year.	6.80	3.44	1.00	12.00
262	Robotics	A binary indicator of whether this is a 'robotics fulfillment center'.	0.59	0.49	0.00	1.00

Table 3. Descriptive Statistics and Variable Definitions for State-Year Warehousing and Storage Sector (NAICS 493) Panel

N	Variable	Description	Mean	S.D.	Min.	Max.
292	InjuryRateMissedWork	Total number of work-related injuries where at least one day of work was missed, per 100 warehousing and storage workers, in a state-year.	0.584	0.739	0.100	6.200
294	InjuryRateJobTransfer	Total number of work-related injuries where at least one day of job transfer was required, per 100 warehousing and storage workers, in a state-year.	0.585	0.791	0.100	7.000
294	AnyRoboticFC	A binary indicator of whether the state has received any Robotic FCs.	0.350	0.478	0.000	1.000

Table 4. Ordinary Least Squares (OLS) Regression Results Based on Annual Establishment-level Injury Data from the Occupational Safety and Health Administration (OSHA)

	Dependent Variable (DV)				
Explanatory Variable	**All Injuries**	**Injuries Missed Work**	**Days Missed Work**	**Injuries Job Transfer**	**Days Job Transfer**
Robotics	10.450 (10.472)	**-16.470**** **(7.391)**	**-895.975*** **(483.213)**	**27.515***** **(8.145)**	**1,579.967***** **(593.071)**
HoursWorked (1,000s)	0.024*** (0.003)	0.014*** (0.003)	0.686*** (0.148)	0.009*** (0.003)	1.042*** (0.228)
AvgEmployees (100s)	0.972** (0.461)	0.562 (0.393)	48.816+ (25.913)	0.534 (0.335)	14.132 (29.318)
Population (1,000s)	-1.218 (3.096)	-0.379 (1.816)	70.006 (138.388)	0.038 (2.123)	-124.043 (187.531)
Median HHI ($1,000s)	0.099 (0.192)	-0.001 (0.085)	5.306 (5.312)	0.057 (0.148)	5.080 (11.826)
CollegeEducated (%)	-17.233 (39.563)	-6.157 (22.619)	-595.193 (1,323.607)	-12.814 (30.500)	-1,762.907 (2,299.219)
WH Type FEs	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes
State FEs	Yes	Yes	Yes	Yes	Yes
Observations	446	446	446	446	446
Adj. R^2	0.679	0.545	0.504	0.602	0.619

*Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$; Robust standard errors in parentheses, clustered by warehouse.*

Table 5. Bivariate Poisson Regression Results Based on Annual Establishment-level Injury Data from the Occupational Safety and Health Administration (OSHA)

Explanatory Variable	**DV = Injuries Job Transfer**	**DV = Injuries Missed Work**
Robotics	**0.373*** (0.094)**	**-0.499*** (0.087)**
HoursWorked (1,000s)	0.00005 (0.00003)	0.0002*** (0.00002)
AvgEmployees (100s)	0.041*** (0.006)	0.016*** (0.004)
Population (1,000s)	0.026 (0.017)	-0.037** (0.013)
Median HHI ($1,000s)	-0.003** (0.001)	0.003*** (0.001)
CollegeEducated (%)	-0.846*** (0.323)	-0.724** (0.0.194)
WH Type FEs	Yes	
Year FEs	Yes	
State FEs	Yes	
Observations	446	
Wald Chi^2	2.5e+10 (63)	

*Notes: *** p < 0.01, ** p < 0.05, * p < 0.10; Robust standard errors in parentheses, clustered by warehouse; Frank copula distribution.*

FIGURES

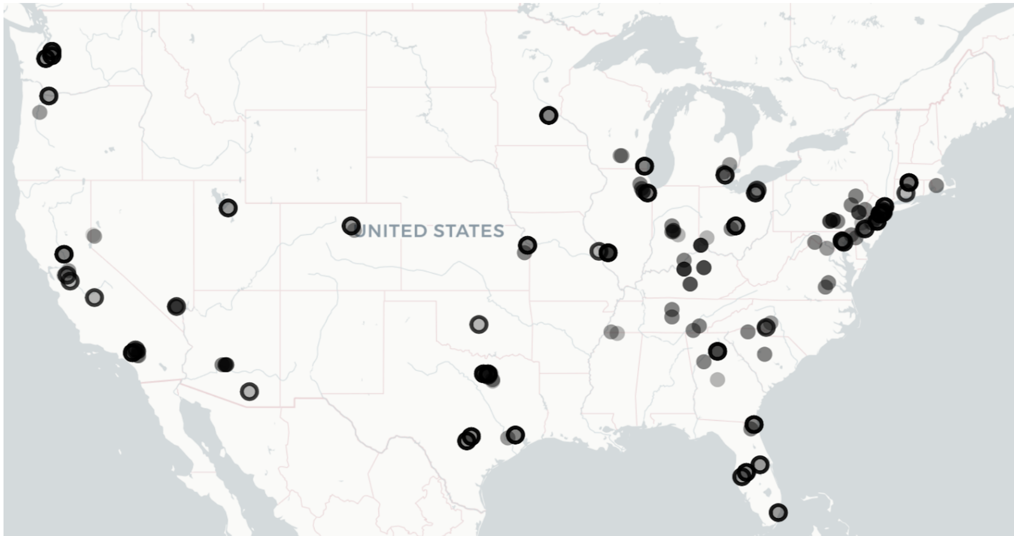

Figure 1. Plot of Amazon Warehouse Locations in our Sample (Bolded outlines indicate Robotic facilities)

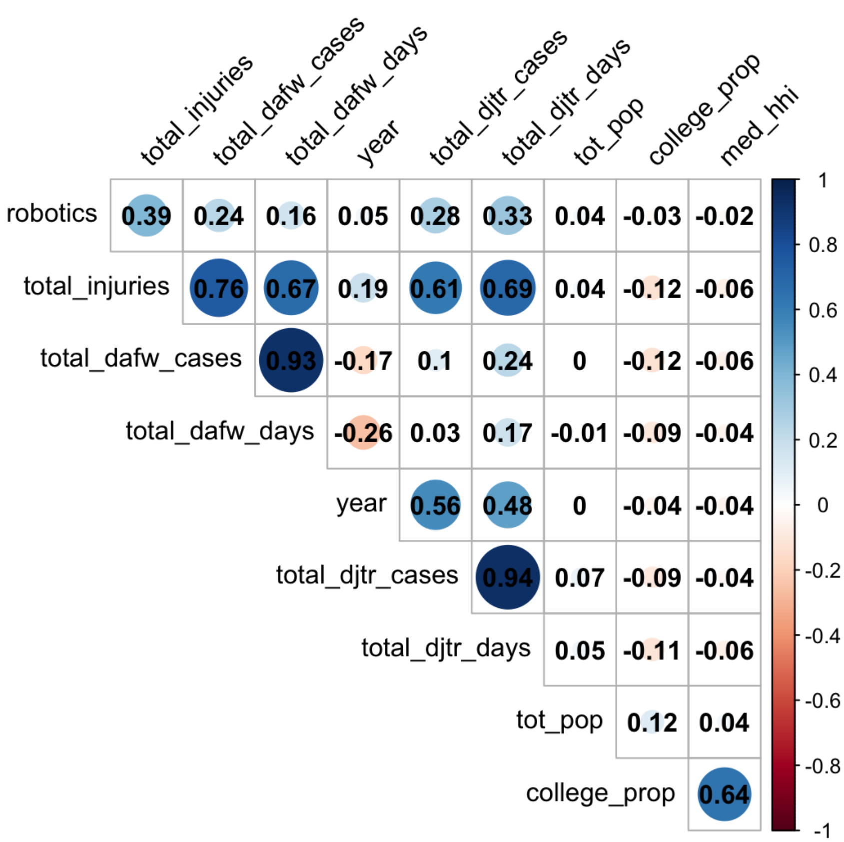

Figure 2: Correlation Matrix

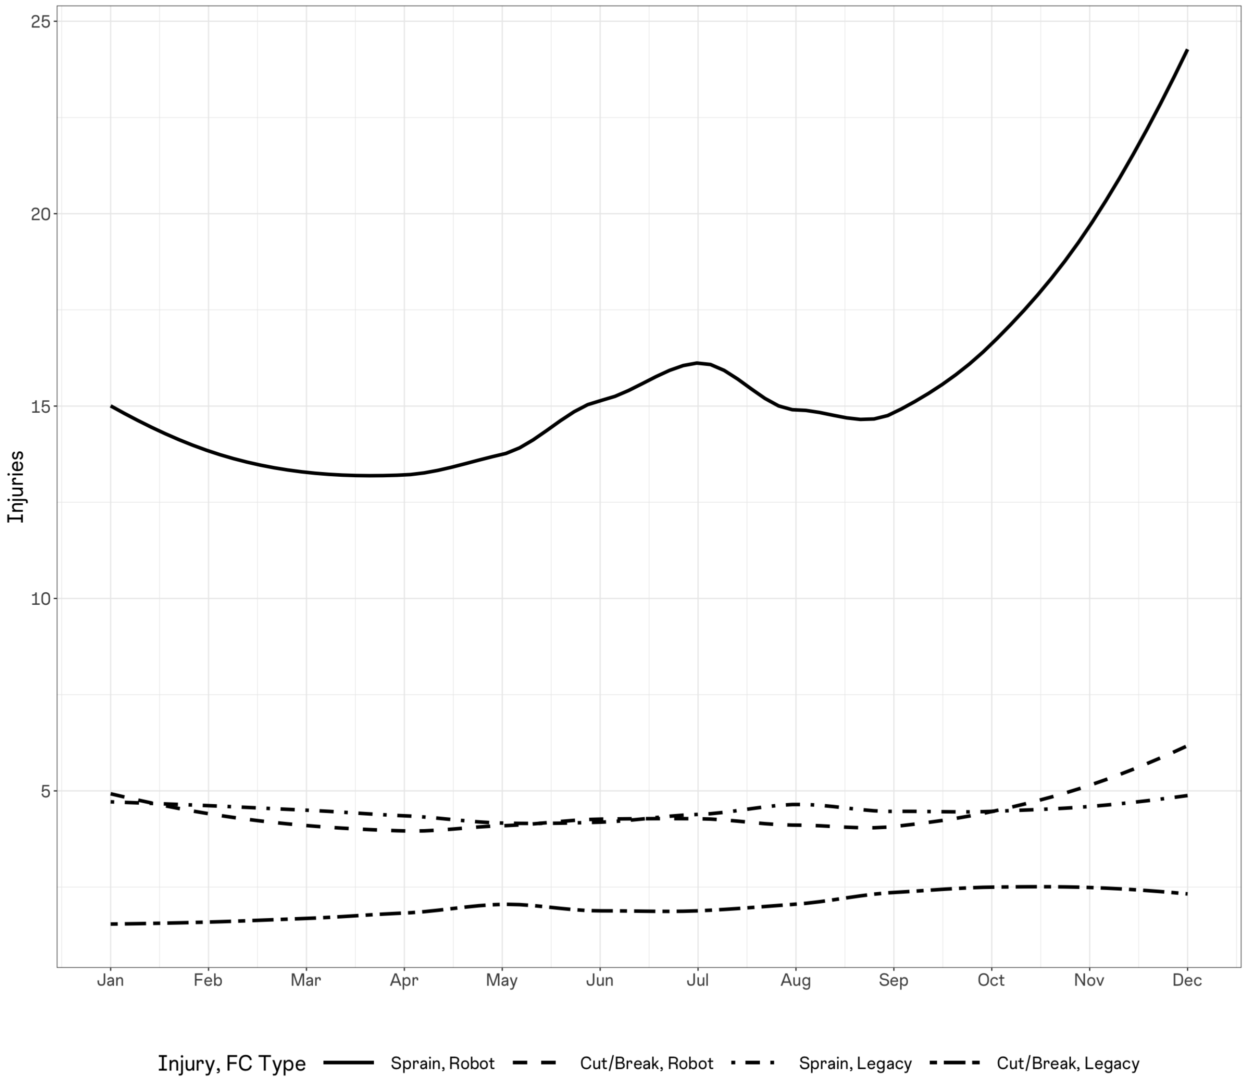

Figure 3. Best Fit Lines of Injuries per Warehouse by Calendar Month by Combination of Injury Type (Strain/Sprain vs. Other) and Fulfillment Center Type (Robotic vs. Legacy) for 2018. Lines Reflect Local Exponential Scatterplot Smoothing (LOESS).

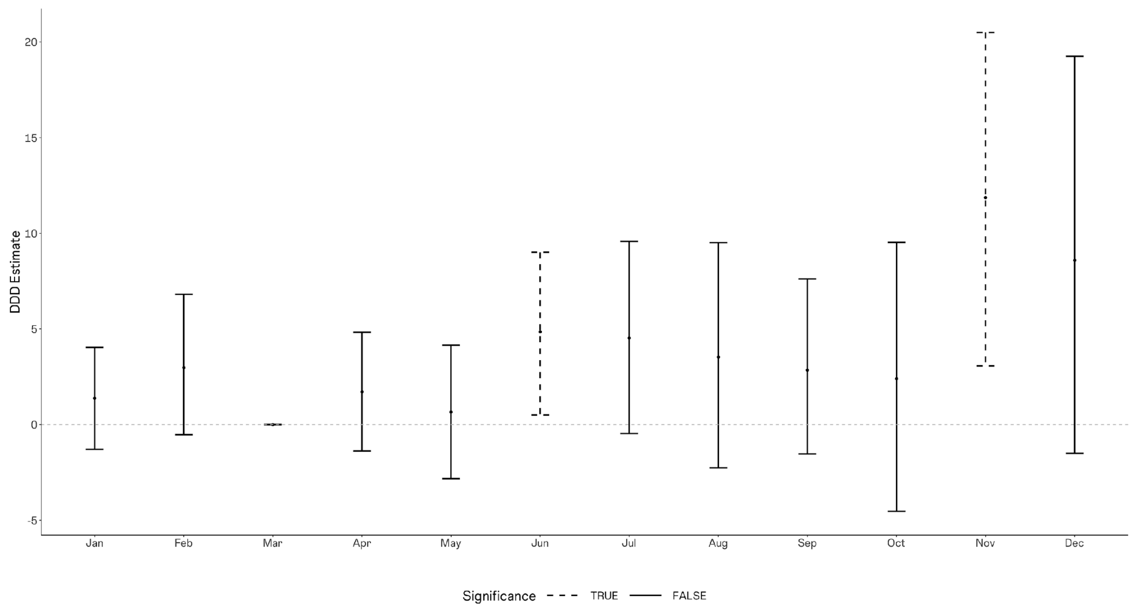

Figure 4. Difference-in-Difference-in-Differences Estimates with 95% Confidence Intervals Based on Wild Cluster Bootstrap Standard Errors (1,000 Bootstrap Iterations)

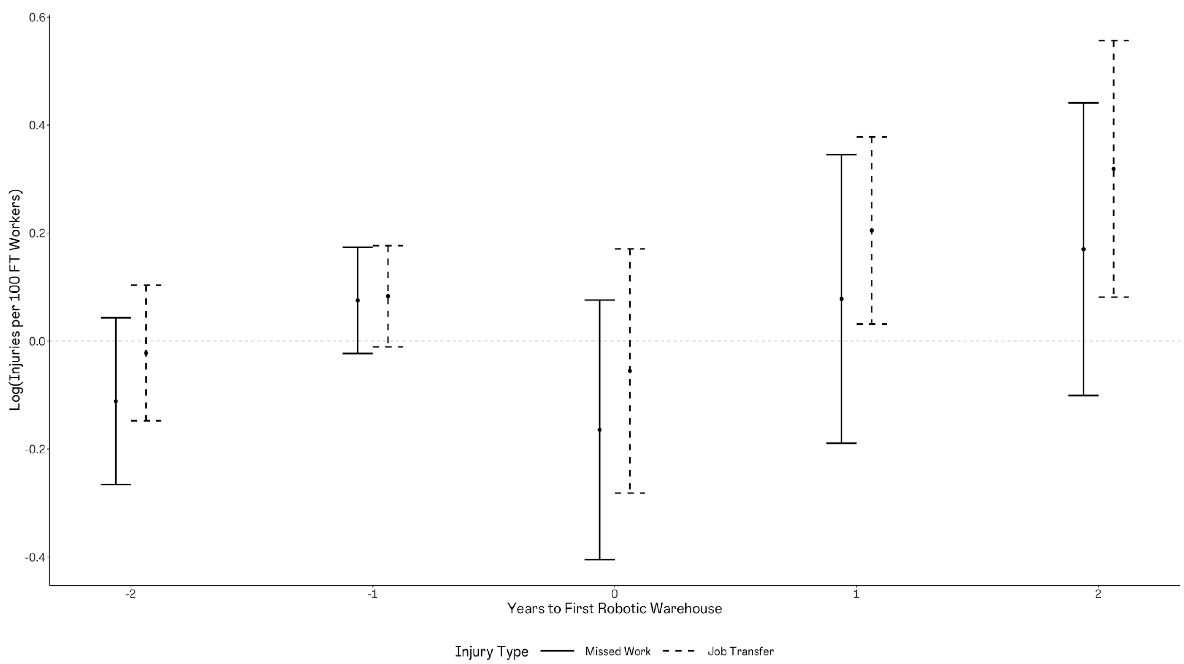

Figure 5. DID-2S Estimate of the Effect of Robotic FCs on Rates of Injury Involving Job Transfer and Effect on Rates of Injury Involving Missed Work (Error Bars Reflect 95% CI's)

SUPPLEMENTARY APPENDICES

APPENDIX A: ACQUISITION OF DATA BY REVEAL NEWS

The data on Amazon Fulfillment Centers was obtained by Reveal News from a third-party consulting firm that specializes in transportation and logistics, MVPVL. The data on annual establishment-level injuries is publicly available and posted by OSHA. The data on state-level injury rates in states' warehousing sectors is publicly available and posted by the BLS. Most importantly, the 2018 injury-level log data for 26 warehouses was obtained by Reveal directly from current and former employees at the warehouses in question. We discuss that process below.

The injury-level logs (with employee names redacted) were retrieved from Pro Publica, where they were posted by Reveal. Reveal obtained these data legally, pursuant to 29 CFR 1904.35(b)(2) of the US Federal Register. The administrative rules in question, which carry force of law, were written by OSHA and provide that employees, former employees, their personal representatives, and authorized employee representatives have the right to access the current OSHA 300 Log, as well as any stored OSHA 300 Log(s) for any establishment in which the employee or former employee has worked. The employer is compelled by law to provide the requester one free copy of the OSHA 300 Log(s) by the end of the next business day. Further, according to OSHA spokeswoman Kimberly Darby, employees are free to share the OSHA 300 logs with any third party, as employees are not required to keep the obtained logs in confidence.

Reveal posted these logs as part of their article on the investigation of Amazon's Fulfillment Center network. Downloading the scanned paper logs, we employed graduate research assistants to manually convert the injury records to a digital format that could be used in our analyses.

Reveal collected the OSHA 300 logs organically. It did so by posting requests for Amazon workers to reach out on their website and then to obtain and supply their warehouse's injury logs to reporters. This raises possibility that the sample is perhaps non-representative. To safeguard against this, we assess the covariate imbalance between warehouses in our annual panel dataset those that appear in our OSHA 300 injury-record dataset and those that do not. We depict the resulting standardized mean differences in the Figure below. The dotted lines represent a standardized mean difference (Cohen's d) of 0.2, the typical threshold for a 'small' effect size. Results are in Figure A1 below.

The sample is skewed toward Robotic Fulfillment centers. Because Robotic FCs tend to be sortable item facilities, there is, in turn, over-representation among Sortable Item warehouses. Further, because Robotic FCs tend to employee more workers, the sample is also skewed toward larger warehouses, with more total hours worked, and thus a larger volume of injuries. Beyond the oversampling on Robotic FCs, the sample appears representative in terms of its geographic distribution, and in terms of its representation among other types of Fulfillment Centers.

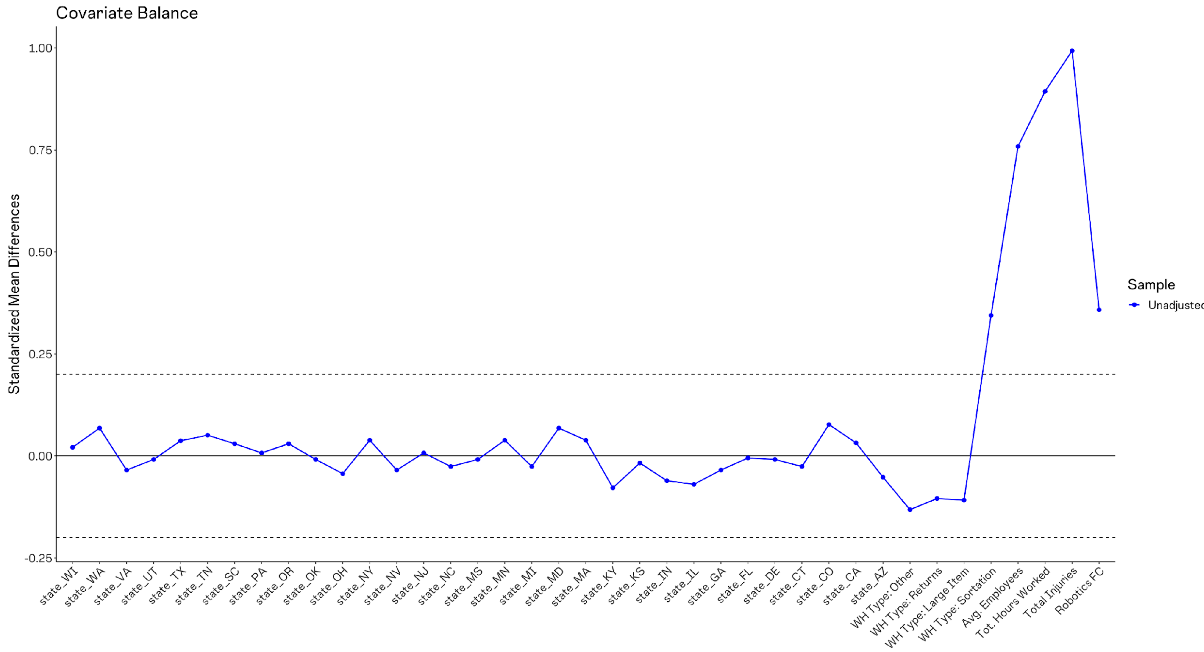

Figure A1. Assessing the Representativeness of the OSHA 300 Warehouse Sample

APPENDIX B. ANALYSIS OF SUB-REDDIT TEXT

To conduct this analysis we collected all posts and comments from the Amazon FC subreddit between January of 2022 and September of 2023. We then filtered those posts and comments to extract all posts that mention the word 'robot', and all posts that mention the words 'traditional' or 'legacy', terms typically used to refer to non-robotic fulfillment centers. This process yielded approximately 1,800+ posts that mention Robotic FCs, and another 500+ that mention traditional or Legacy FCs. After pre-processing the text (i.e., removing URLs, punctuation, English stop words, etc.) we calculated all bigram frequencies, i.e., pairs of words appearing alongside one another in the text. We then built separate word clouds reflecting the top 50 most frequent bigrams associated with Robotic FC posts and the Legacy FC posts, respectively. Results appear in Figures B1 and B2 below, with notable terms circled.

Figure B1. Most frequent bigrams associated with Legacy FC posts.

Figure B2. Most frequent bigrams associated with Robotic FC posts.

APPENDIX C. INJURY LOG DATA DESCRIPTIVES

We present a simple descriptive plot of injury volumes by Fulfillment Center type (Robotic vs. Legacy). As can be seen in Figure C1, sprains / strains and bruises are by far the most common injuries in these warehouses. Of note, however, is that sprains and strains represent a systematically larger proportion of injuries in Robotics facilities than they do in Legacy facilities.

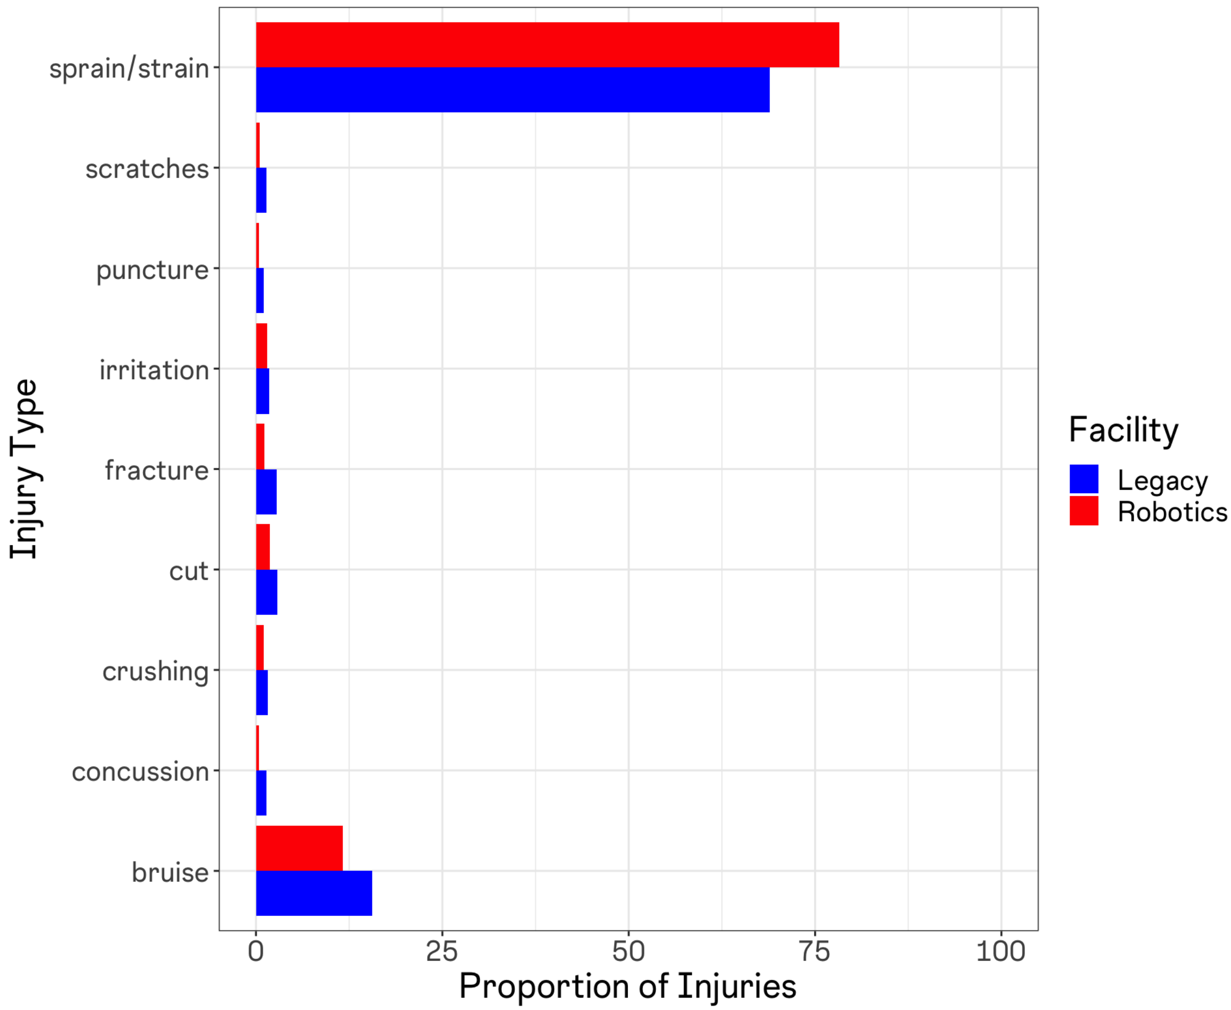

Figure C1. Proportion of 2018 injuries by warehouse type and injury type. The x-axis reflects the total proportion of injuries within a given warehouse that are of the type indicated on the y-axis.

APPENDIX D. STATE-YEAR PANEL

To ensure the robustness of the analysis, we consider a different unit of analysis: the state-year. Specifically, we construct a panel based on annual injury rates published by the BLS regarding private establishments comprising each state's warehousing and storage sector (NAICS 493). Note that state-level industry data is the most granular comprehensive injury data that the BLS publicly provides. We supplement this panel with a binary indicator of whether a given state has received any Robotic FCs as of a given year, based on our warehouse-level dataset. It should be noted that the state-year panel is imperfect in at least two respects. First, our records of Robotic FCs in each state-year are subject to measurement error because our warehouse-level data from OSHA does not include all Amazon FCs. Second, the BLS data is subject to missingness, because some states do not report warehouse-industry specific injury estimates in some years.

As discussed in the main text, there is the natural concern that Robotic Fulfillment Centers are not assigned at random, and that the organic local injury rate might push Amazon in one direction or another (i.e. reverse causality). The logic is that the liability regime which characterizes a state might affect Amazon's decision-making process. To address this, we estimate a hazard model of Robotic FC construction in a state, employing a panel of state-year observations that includes a binary indicator of the first Robotic FC construction as the failure event. The independent variables of interest are the rate of injuries per 100 warehouse workers in the state (both the rate of injuries involving job transfer and the rate involving missed work). The estimator is a Cox Proportional Hazard (PH) regression. Results are in Table D1. As can be seen, we observe no significant relationship between either the contemporaneous injury rates or the lagged injury rates and the hazard of Robotic FC construction.

Table D1. Cox PH Regression of First Robotic FC Construction

Explanatory Variable	Model (1)	Model (2)
Rate of DAFW† Injuries	0.142 (0.259)	--
Rate of DJTR† Injuries	0.079 (1.082)	--
Lag(Rate of DAFW Injuries)	--	0.480 (1.616)
Lag(Rate of DJTR Injuries)	--	0.292 (1.339)
Observations	177	144
Events	19	18
Wald test	0.25 (2)	2.41 (2)
Robust Logrank test	0.24 (2)	2.31 (2)

*Notes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; Failure is defined as a state's first construction of a Robotic FC; Standard errors reported in parentheses clustered by State; †DAFW = days away from work; DJTR = days of job transfer or restriction*

APPENDIX E: ROBUSTNESS CHECKS (TABLES & FIGURES)

We report a series of placebo tests. First, we replace our physical injury outcomes with reports of other workplace ailments that are not plausibly related to the presence of Robotics. We consider poisonings, respiratory issues, skin conditions, and hearing problems. As anticipated, in Table E1, we find no statistically significant relationship between Robotics and the occurrence of these ailments.

Table E1. Annual Panel: Placebo Regression Results (OLS)

Explanatory Variable	**Model (1) Poison**	**Model (2) Respiratory**	**Model (3) Skin**	**Model (4) Hearing**
Robotics	-0.008 (0.017)	-0.047 (0.037)	-0.002 (0.006)	-0.033 (0.040)
HoursWorked (1000s)	-0.004 (0.003)	0.021 (0.010)	0.00001 (0.00001)	0.00001 (0.00001)
AvgEmployees (100s)	0.002* (0.001)	-0.001 (0.002)	-0.001 (0.001)	-0.001 (0.002)
Population (1000s)	-0.002 (0.004)	-0.009 (0.008)	-0.004 (0.003)	0.003 (0.005)
Median HHI (1000s)	0.0001 (0.0003)	0.001 (0.001)	-0.00002 (0.0001)	-0.0004 (0.0004)
CollegeEducated (%)	-0.115 (0.073)	-0.087 (0.117)	0.042 (0.455)	0.192** (0.095)
WH Type FEs	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
State FEs	Yes	Yes	Yes	Yes
Observations	446	446	446	446
Adj. R^2	-0.032	0.098	0.032	0.038

*Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$; Standard errors in parentheses are clustered by warehouse.*

Next, we explore reweighting to deal with imbalance on observable characteristics. We achieve this via inverse propensity for treatment weighting (IPTW), based on covariate-balancing propensity scores. Reweighting in this fashion, rather than relying solely on linear covariate adjustment in our regressions, helps address the concern of model misspecification. In Figure E1, we see that Robotics facilities employee systematically more workers, and have systematically more hours worked, as compared to Legacy fulfillment centers. These facilities are also systematically more likely to be Sortable warehouses (Warehouse Type: S), as opposed to Large-Item facilities (Warehouse Type: L). Re-weighting to achieve balance on these characteristics, we re-estimate our models. Results, reported in Table E2, are consistent with our main findings.

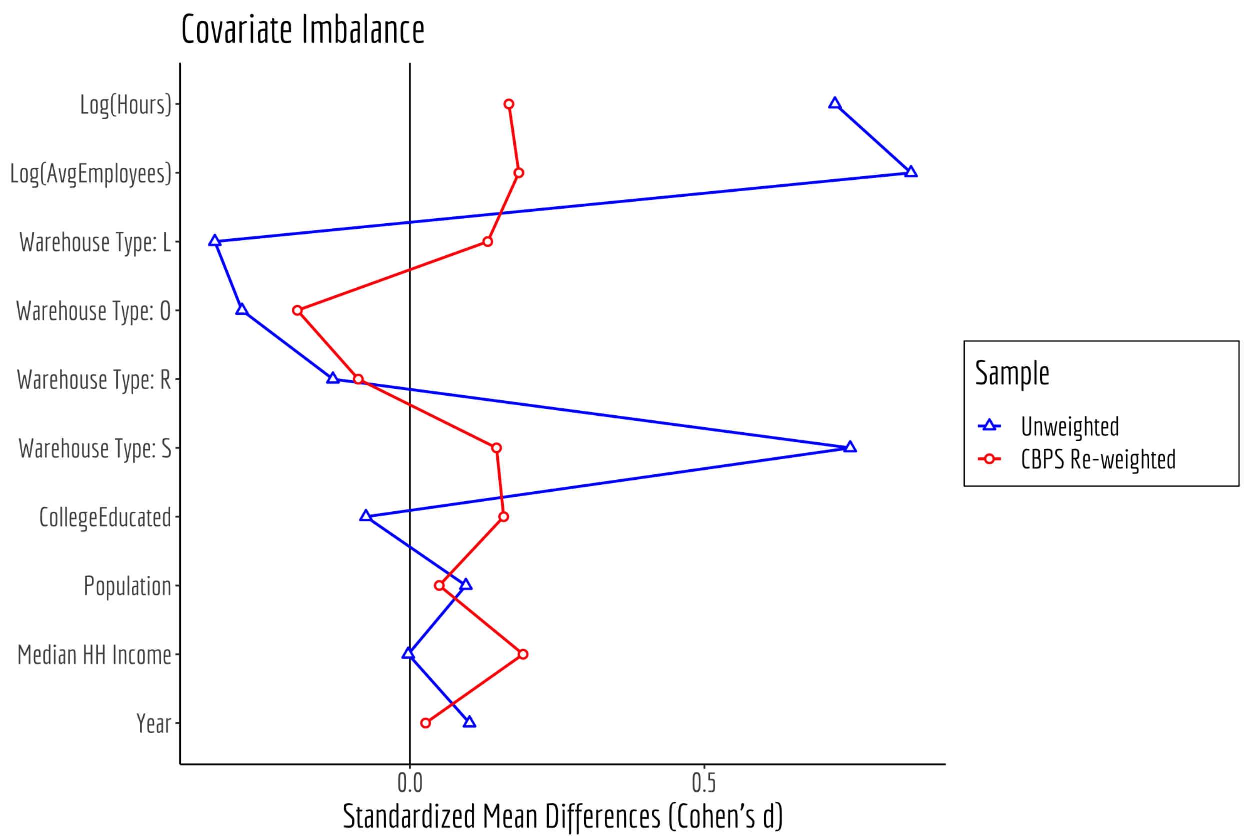

Figure E1. Covariate Imbalance: CBPS Re-weighted vs. Unweighted Sample.

Table E2. Annual Panel /w Re-weighting: Regression Results (OLS)

Explanatory Variable	All Injuries	Injuries Missed Work	Days Missed Work	Injuries Job Transfer	Days Job Transfer
Robotics	6.440 (9.349)	**-15.634**** **(6.008)**	**-877.207**** **(390.708)**	**22.782***** **(8.029)**	**1,319.606**** **(577.250)**
HoursWorked (1,000s)	0.024*** (0.003)	0.014*** (0.003)	0.709*** (0.155)	0.008*** (0.002)	0.835*** (0.209)
AvgEmployees (100s)	1.292** (0.503)	0.725* (0.432)	56.104** (27.680)	0.737** (0.287)	37.552 (27.085)
Population (1,000s)	-2.333 (2.713)	-0.899 (0.001)	33.659 (104.774)	-0.522 (2.017)	-185.908 (161.804)
Median HHI ($1,000s)	-0.001 (0.209)	-0.093 (0.123)	4.232 (6.936)	0.037 (0.139)	4.878 (10.860)
CollegeEducated (%)	4.929 (37.901)	8.700 (24.484)	73.020 (1,561.380)	-0.949 (29.767)	-1,299.912 (2,171.321)
WH Type FEs	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes
State FEs	Yes	Yes	Yes	Yes	Yes
Observations	446	446	446	446	446
Adj. R^2	0.668	0.576	0.533	0.574	0.596

*Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$; Standard errors in parentheses are clustered by warehouse; Sample reweighted using covariate-balancing propensity scores.*

Finally, we consider alternative regression specifications. Specifically, rather than merely condition on hours worked, we consider an alternative estimation wherein we normalize injury volumes by the total employee hours worked at a facility. Put simply, we estimate a rate model. We accomplish this by re-estimating our injury count models using Poisson regression, and incorporating the total hours worked (in 1000s) as an exposure term. As can be seen in Table E3, results remain consistent. Further, repeating our estimations employing log-linear regression, explicitly normalizing the injury measures with respect to average annual employees (i.e., a rate measure), we obtain the estimates in Table E4. Once again, results remain consistent.

Table E3. Poisson Regression Results Taking Total Hours Worked as Exposure, Based on Annual Establishment-level Injury Data from the Occupational Safety and Health Administration (OSHA),

Explanatory Variable	**DV = Injuries Job Transfer**	**DV = Injuries Missed Work**
Robotics	**0.629*****	**-0.280****
	(0.153)	**(0.133)**
Employees	-0.004	-0.005
(100s)	(0.005)	(0.004)
Population	0.005	-0.013
(1,000s)	(0.029)	(0.032)
Median HHI	0.0002	-0.001
($1,000s)	(0.002)	(0.001)
CollegeEducated	0.377	0.552+
(%)	(0.567)	(0.317)
WH Type FEs	Yes	Yes
Year FEs	Yes	Yes
State FEs	Yes	Yes
Exposure	HoursWorked (1000s)	HoursWorked (1000s)
Observations	446	446
Pseudo R^2	0.701	0.663

*Notes: *** p < 0.01, ** p < 0.05, * p < 0.10; Standard errors in parentheses are clustered by warehouse.*

Table E4. Ordinary Least Squares (OLS) Regression Results Based on Annual Establishment-level Injury Data from the Occupational Safety and Health Administration (OSHA)

Explanatory Variable	Log(All Injury per Employee)	Log(Missed Work per Employee)	Log(Days Missed Work per Employee)	Log(Job Transfer per Employee)	Log(Days Job Transfer per Employee)
Robotics	0.095 (0.133)	-0.511* (0.206)	-0.586* (0.282)	0.508** (0.161)	0.529** (0.189)
Population (1,000s)	0.007 (0.019)	0.009 (0.027)	0.005 (0.043)	0.041 (0.030)	0.021 (0.039)
Median HHI ($1,000s)	0.003 (0.002)	0.002 (0.002)	0.005 (0.003)	0.003 (0.003)	0.004 (0.003)
CollegeEducated (%)	-0.780 (0.508)	-0.773 (0.680)	-1.202 (0.828)	-0.620 (0.778)	-1.239 (0.894)
WH Type FEs	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes
State FEs	Yes	Yes	Yes	Yes	Yes
Observations	445	443	443	417	431
Adj. R^2	0.333	0.422	0.336	0.449	0.386

*Notes: *** p < 0.01, ** p < 0.05, * p < 0.10; Robust standard errors in parentheses, clustered by warehouse; some observations are omitted due to log(0) being undefined.*