

Yilmaz, Erdem Dogukan; Peukert, Christian

Working Paper

Who Benefits from AI? Project-Level Evidence on Labor Demand, Operations and Profitability

CESifo Working Paper, No. 11321

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Yilmaz, Erdem Dogukan; Peukert, Christian (2024) : Who Benefits from AI? Project-Level Evidence on Labor Demand, Operations and Profitability, CESifo Working Paper, No. 11321, CESifo GmbH, Munich

This Version is available at:

<https://hdl.handle.net/10419/305563>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Who Benefits from AI? Project-Level Evidence on Labor Demand, Operations and Profitability

Erdem Dogukan Yilmaz, Christian Peukert

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Who Benefits from AI? Project-Level Evidence on Labor Demand, Operations and Profitability

Abstract

We examine how the adoption of digital automation technology affects labor demand, operations and profitability in the context of the logistics industry. Our data covers 9,300 digital automation projects in a multinational company involving service robots and machine learning-based software from 2019 to 2021, alongside fine-grained labor and operations data. To identify causal effects, we leverage exogenous variation from supply-chain disruptions and travel restrictions during COVID-19 and an import ban on information and communication technologies imposed by the Trump administration. We find that total labor cost increased after the adoption of digital automation technology, attributable to increased labor demand and more reliance on temporary workers. However, managerial hours declined, possibly due to increased efficiency. Furthermore, digital automation technology increased revenue and profit through a reduction in operational cost, improved utilization of warehouse space, and higher profit margins. However, the effects of digital automation technology are not homogeneous. We highlight substantial complementarities between hardware and software technologies. Management units that only use software technology experience only half the increase in revenue and profit.

Keywords: digital automation technology, robots, artificial intelligence, future of work.

Erdem Dogukan Yilmaz
INSEAD
Fontainebleau / France
erdemdogukan.yilmaz@insead.edu

Christian Peukert
Faculty of Business and Economics
University of Lausanne / Switzerland
christian.peukert@unil.ch

August 28, 2024

We are grateful to Lucas Franco and Rachel Moser for their excellent research assistance. We thank Murat Tarakci and conference and seminar participants at the ICEA Wealth of Nations Online Conference and HEC Lausanne for helpful suggestions. We acknowledge support from the Swiss National Science Foundation for the project 100013 197807.

1 Introduction

Artificial Intelligence (AI) is often referred to as a new general-purpose technology (GPT) with far-reaching implications across various sectors of the economy and society. As a GPT, AI is expected to automate certain tasks traditionally performed by labor, complement human efforts in other tasks, and consequently increase productivity over time ([Acemoglu et al., 2022](#); [Autor et al., 2003](#); [Bresnahan et al., 2002](#); [Cockburn et al., 2018](#)). However, whether it takes a significant amount of time for AI-based productivity gains to be observed in the economy remains an open question. On one hand, similar to previous information and communication technologies, AI integration may require system-wide changes affecting various organizational levels and functions. This can lead to a lag between AI adoption and observable productivity gains, a phenomenon known as the “productivity J-curve” ([Brynjolfsson et al., 2019b](#)). On the other hand, AI systems possess unique characteristics that may accelerate productivity gains compared to static technologies. AI’s ability to efficiently learn from accumulated data and provide “point solutions” without requiring complete system overhauls could lead to quicker realizations of productivity benefits ([Wang et al., 2023](#); [Agrawal et al., 2022](#); [Feigenbaum and Gross, 2024](#)).

In this paper, we aim to make progress on this front by leveraging detailed project-level data from one of the largest logistics companies in the world. In 2019, the company launched a program involving the adoption and implementation of machine learning-based algorithmic task optimization (e.g., resource and inventory management), and specialized service robots that also rely on machine learning. This initiative, referred to as the Digital Automation Technology (DAT) program, was implemented across its facilities in 60 countries. We document the effects of this program, scrutinizing its roll-out across various management units (MUs) worldwide.

Identifying a clear causal relationship between DAT adoption and outcomes such as labor demand, productivity, and profitability is challenging, as many internal and external factors influence both technology adoption and firm performance. However, the coincidental overlap of our study’s timeline with the COVID-19 pandemic introduced an opportunity for identifying causal effect of AI on different firm level outcomes. Due to pandemic-induced disruptions, the implementation of DAT became quasi-random across facilities, given significant delays and disruptions in the installation and training for the AI-based software and the arrival of AI-based

robotics. Leveraging this serendipitous randomness, we apply a difference-in-differences approach to estimate the effects of DAT adoption and interpret our results with insights from company managers.

Our findings offer a set of nuanced insights. First, we show that the adoption of DAT led to an increase in revenues and profitability. However, we also observe an increase in labor costs. Upon closer examination, this rise in labor costs appears to be linked to an increase in overtime hours and greater reliance on temporary labor, despite managers experiencing a reduction in working hours. Our interviews with managers suggest that AI adoption automates certain coordination and monitoring tasks, while the increase in overtime hours mostly stems from better utilization of space and capacity due to DAT adoption, which demands additional work from workers. Furthermore, DAT adoption leads to the standardization of tasks, making it easier to hire temporary workers. Supporting this explanation, we show that the positive effects of DAT manifest through a decrease in operational costs and better utilization of capacity, which balances the increased labor costs and leads to higher revenues and profitability. Furthermore, we uncover that the benefits of DAT were most pronounced in establishments that simultaneously adopted both software (i.e., AI-based software) and hardware (i.e., AI-based robotics) technologies.

To confirm the robustness of our findings, we first leverage variation introduced by an executive order issued by President Trump in late 2019, which effectively banned the import of information and communication technology (ICT) from adversarial countries such as China and Russia. Given our access to internal data from a global organization, we can contrast US operations before and after the enactment of the executive order with operations conducted elsewhere in the world. Specifically, we account for whether the MUs had previously procured DAT from vendors located in China and Russia to assess technology categories directly impacted by the import ban. Subsequently, we utilize the variations observed across countries, technologies, and time as instruments to analyze the DAT adoption decisions made by different MUs of the company. Furthermore, we replicate our analyses using an estimation method robust to staggered difference-in-differences designs. The results remain consistent across different identification strategies and model specifications.

Our paper makes several contributions to the literature. First, we add to the sparse work examining AI and firm productivity. While previous studies do not find measurable impacts of

AI-related investment on firm productivity, they also caution that the productivity effects of AI technologies may be challenging to identify at more aggregated levels because AI-adopting firms differ substantially from non-adopters (Acemoglu et al., 2022; Babina et al., 2024). Furthermore, the intangible nature of many AI-driven improvements (e.g., better decision quality, personalization) may introduce unique challenges in measuring productivity gains, potentially leading to underestimation of the impact of AI (Brynjolfsson and Mitchell, 2017). With more granular data, we overcome some of these challenges and find that AI increases productivity through better utilization of capacity and reduced operational costs.

Second, we contribute to the literature on AI and workers (Fügener et al., 2021; Felten et al., 2021; Lebovitz et al., 2021; Acemoglu et al., 2022; Bauer et al., 2023; Chen et al., 2024) by showing that the adoption of technology can create additional tasks for existing workers while reducing the need for managerial attention, requiring strategic reconfiguration of human capital. In particular, our results suggest that the standardization of tasks through AI adoption may lead to decreased transaction costs, such as search and information costs, making it easier for the company to hire a temporary workforce and reducing the need for managerial oversight. These results also inform the literature examining the labor market effects of general-purpose technologies, which show that technologies like IT and electricity favor high-skilled labor but displace medium-skilled workers (Autor et al., 1998, 2003; Acemoglu and Restrepo, 2019). Adding to this conversation, we find that the adoption of AI-based software and hardware can indeed increase the demand for low-skilled workers, through standardization and an increase in the intensity of residual tasks that cannot be automated.

Third, our findings on the superior outcomes when both hardware and software components of DAT are adopted together extend the literature on complementarities in technology adoption (Milgrom and Roberts, 1995; Bloom et al., 2012; Sirmon et al., 2011) by providing new evidence in the context of AI.

2 Background

2.1 AI and Workers

The impact of AI on workers has been an important focus of research for several years, attracting significant scholarly attention. Central to this discourse is the notion of skill-biased technological change (SBTC), widely employed as a theoretical framework to understand the

interaction between technology and workers (Autor et al., 2003). SBTC suggests that routine tasks, characterized by repetition and following of explicit procedural rules, are more susceptible to automation than non-routine tasks that demand cognitive agility and problem-solving skills (Autor et al., 2003). This gives rise to the concept of routine-biased technological change, where technology substitutes routine tasks (Acemoglu and Restrepo, 2020). For workers primarily engaged in routine tasks, this often means replacement, whereas for those involved in non-routine cognitive tasks, it often means enhanced productivity as they will have more time to focus on higher-level, non-routine tasks.

Traditionally, workers who conduct non-routine tasks have been assumed to be immune to automation. However, AI technologies are likely to challenge this status quo and may impact labor that involves non-routine tasks (Brynjolfsson et al., 2018; Frank et al., 2019; Agrawal et al., 2019; Krakowski et al., 2023). In fact, AI technologies have superior prediction skills compared with those of humans when a sufficient amount of data is available (Peukert et al., 2023), thus enabling more accurate forecasting and resource planning. For example, using AI in medical diagnoses reduces errors (Aron et al., 2011; Jussupow et al., 2021; Ribers and Ullrich, 2024), AI-based translation software delivers faster translation services (Brynjolfsson et al., 2019a), and AI adoption supports innovation activity in drug discovery (Lou and Wu, 2021).

Oftentimes, substitution and complementation co-occur when firms adopt new technologies (Stadler et al., 2022). To determine whether AI substitution or complementary effect dominates, Acemoglu et al. (2022) analyzed online job advertisements from as far back as 2010. They categorized these ads to discern if they were related to AI-based employment. To estimate employers' exposure to AI, they evaluated whether the tasks described in the advertisements were consistent with the prevailing capabilities of AI. In their analysis, they utilized three metrics to gauge current exposure to AI: Felten et al. (2021)'s "artificial intelligence occupational impact measure"; Brynjolfsson et al. (2018)'s "suitability for machine learning index"; and Webb (2019)'s "artificial intelligence exposure score". Their results suggest that firms with exposure to AI tend to recruit more AI-specialized workers but decrease their overall hiring. This points towards a trend where AI is being adopted primarily for labor substitution, with the displacement effects seemingly overshadowing the complementary effects within organizations.

2.2 AI and Productivity

The adoption of AI represents a significant strategic decision for firms seeking to enhance their competitive position. AI has the potential to make firms more efficient through intelligent automation and by helping workers solve novel problems, which may lead to value creation through the design of new products and the improvement of existing ones (Brynjolfsson and McAfee, 2014; Brynjolfsson and Mitchell, 2017; Dell’Acqua et al., 2023). For instance, AI’s ability to learn from data may lead to faster decision-making for operational tasks and increase efficiency (Raj and Seamans, 2019). Similarly, AI may help developers pinpointing software bugs (Peng et al., 2023).

Previous literature suggests that technology adoption often triggers a process of organizational reconfiguration (Teece et al., 1997; Helfat and Peteraf, 2015). This process requires firms to develop new capabilities and routines, encompassing not only technical learning about the new technologies but also organizational learning about how to restructure work processes and re-define roles (Edmondson et al., 2001). In fact, when a firm adopts new production methods without making wider changes to its production system and investing in complementary assets, changes that seem value-enhancing can be value-destroying (Milgrom and Roberts, 1990; Henderson and Clark, 1990).

Similar to other technologies, the integration of AI into operations may require a system-wide change, affecting various organizational levels and functions. Therefore, the realization of productivity gains from AI adoption may follow a similar pattern to previous ICTs (Brynjolfsson et al., 2019b), meaning that there may be a lag between AI adoption and observable productivity gains until firms make investments in both tangible and intangible capital, often referred to as the “productivity J-curve.” (Brynjolfsson et al., 2021). Supporting these insights, recent studies do not find any impact of AI-related investment on firm productivity at the macro level (Acemoglu et al., 2022; Babina et al., 2024).

However, AI systems are different from other static technologies, which may lead to quicker, accelerated productivity gains. If a company has already accumulated data, AI systems can learn about the company’s inner workings efficiently. Because workers no longer need to explicitly code their knowledge into rules for AI, this eliminates what used to be an effort-intensive and expensive task of knowledge acquisition (Wang et al., 2023). Instead, machine learning algorithms

detect patterns from data by iteratively and recursively exploring complex connections between inputs and outputs to establish the relationship (Jordan and Mitchell, 2015). Furthermore, AI applications can target individual tasks and may provide “point solutions” without requiring the design of entirely new systems (Agrawal et al., 2022). This means that AI may only require adjustment of a few interrelated tasks, leaving the majority of existing organizational structures and routines intact (Feigenbaum and Gross, 2024).

However, we still lack solid empirical evidence on how AI adoption is likely to affect aggregate firm productivity and, eventually, firm performance. In this paper, we aim to address this shortcoming in the literature by using unique project-level data from one of the largest logistics companies in the world.

3 Data and Methods

3.1 Global Logistics Company

For our research, we have collaborated with a leading global logistics company. The company offers warehousing, transport, and value-added services in over 3,000 facilities in 60 countries worldwide. It serves a broad customer base of industry verticals including automotive, consumer, retail, and healthcare. In response to the rising demand and labor constraints, the company started the DAT program in April 2018, which included the adoption of new software and robotics technologies.

The company chose 12 warehouse technologies based on their expected impact (i.e., estimated labor hours reduction in certain tasks) and scalability across sites (i.e., expected applicability across the company’s whole operational footprint). These technologies are expected to substitute some of the tasks previously executed by labor, but also increase the productivity rather than taking over all human labor tasks. Hence, human labor can be freed up to focus more on the value-adding activities.

A range of ML-based hardware technologies are employed for tasks such as replenishment, picking, and the point-to-point movement of items. These systems exemplify key characteristics of artificial intelligence, including perception, decision-making, and learning capabilities. For instance, computer vision systems, often using deep learning models like Convolutional Neural Networks, enable the recognition and classification of items. Concurrently, AI algorithms optimize picking routes and replenishment schedules based on real-time demand and inventory

levels. These systems demonstrate learning capabilities, improving their performance over time through algorithms such as reinforcement learning. In addition to these core operational robots, supporting robots, collectively referred to as service robots, are adopted for tasks such as cleaning, security, and dimensioning. They often incorporate Natural Language Processing (NLP) for understanding and responding to voice commands, and unsupervised learning algorithms for anomaly detection. In contrast to traditional industrial robots, these robots use machine learning algorithms and are aware of their surroundings and can quickly locate humans and other objects.

As part of the global program, the company has also adopted various ML-based software solutions. First, algorithmic optimization, often based on neural network-based reinforcement learning, is employed to optimize diverse aspects of warehouse management. These algorithms utilize predictive models to forecast demand and adjust inventory levels accordingly. Second, AI-based asset tracking and monitoring systems pair with digital maps to provide accurate visualization of assets while predicting potential failures and detecting anomalies. These systems typically combine Internet of Things sensors with machine learning models for predictive maintenance. Intelligent process automation represents another key area of AI adoption. This involves the use of AI technologies like computer vision and NLP to automate tasks that traditionally required human perceptual skills. For instance, automated quality control systems employing deep learning models can inspect products with high accuracy and consistency. Lastly, smart operations software integrates existing software and labor with robotics to automate process management and provide predictive analytics. These systems often utilize ensemble models or deep learning architectures to analyze large amounts of data and make predictions about various aspects of warehouse operations.

The company's operations are organized around management units (MUs). There is typically at least one MU assigned to each customer in each country. A warehouse can host a couple of MUs and in rare cases, the same MU may be spread throughout a couple of warehouses. The DAT program is closely aligned with the needs of individual customers. Adoption decisions, hereafter referred to as DAT projects, are taken locally, sometimes with co-investment of the customer, at the MU – warehouse level.

Table 1: Variables and Summary Statistics

Variables	Uniques	Observed period
<i>DAT Projects</i>	9,299	Q1/2019–Q3/2021
...of which		
Deployment in progress	6%	
Implemented	77%	
Hardware (Robots/Devices)	36%	
Software	64%	
Vendors	147	
<i>Management Units</i>	12,894	
...of which		
implement DAT projects	18%	
in US	11%	
in US and affected by import ban	10%	
<i>Labor</i>		Q1/2020–Q3/2021
Total Cost		
Cost/Hour		
Regular Hours		
Overtime Hours		
Management Share		
Temporary Share		
<i>Operations</i>		Q1/2019–Q3/2021
Revenue		
Profit		
Operational Cost		
Warehouse Utilization		
Profit Margin		

3.2 Variables

We have access to project-level information for almost 3 years, observe different deployment statuses of projects, the type of DAT used, and – if not self-developed – the vendor from which the company purchases the technology. This data maps to MUs, across the globe. Importantly, only for 18% of the MUs in our dataset, we observe the use of DAT. The company started the initiative just shortly before our observation period begins, so it is unlikely that we miss a substantial amount of DAT projects that have been implemented before we can observe it.

We further observe information about labor and operations of the MUs. We have complete information regarding labor cost, volume handled, revenue and profit. For the remaining variables, our observation period only starts in January 2020. Table 1 provides an overview of the information we observe in our dataset.

3.3 Identification Strategy

It is useful to consider the ideal experiment to cleanly estimate the effect of DAT adoption. We would, (1) randomly select a sample of MUs from the population of MUs, (2) randomly select from the set of available technologies, (3) purchase and deploy a random DAT at the selected MUs. Then, over time, the effect of the adoption could be compared between the treatment and control group MUs by observing changes in productivity and financial performance. Given the global scale of our industry partner, this set-up is practically infeasible. However, we can exploit different sources of quasi-experimental variation that come close to the ideal.

First, we make use of the fact that our observation period significantly overlapped with the COVID-19 pandemic — a time characterized by notable logistical challenges (Meier et al., 2020) and stringent travel restrictions (Deb et al., 2022). The global pandemic provides us with a unique opportunity: it inadvertently introduced substantial delays in the adoption and implementation of DAT. For example, interviews with senior management at our partner company revealed that it was often difficult to stick to the implementation timeline of particular DAT projects because technicians needed for the ground installation (such as the installation of IoT sensors on the equipment to collect real-time data) could not travel internationally or were quarantined. Consequently, these disruptions grant us exogenous variation in adoption timelines, which, in our analysis, closely resembles an almost ideal randomized experiment scenario. Thus, after accounting for the time-invariant characteristics of the MUs and any overarching temporal effects, our identification assumption relies on the notion that the actual timing of adoption is as good as random. This is confirmed by a simple test, where we split our sample into two periods – before COVID-19 and during the COVID-19 pandemic – regress MU characteristics on an indicator of DAT adoption. Table 2 shows that quarterly variation in labor costs and profits is correlated to adoption decisions before COVID-19 (column 1), but not during the COVID-19 pandemic.

With this, we can attempt a difference-in-differences comparison of MUs that chose to adopt DAT versus those that do not, before and after adoption takes place. Despite the exogenous variation in the timing of implementation caused by COVID-19, there might be unobserved reasons why some MUs choose to adopt DATs and other do not. For example, particularly large and historically productive MUs could be more likely to adopt DAT. To capture that, we use

Table 2: DAT adoption is not explained by MU characteristics during COVID-19

	(1) Before COVID-19	(2) During COVID-19
Labor Cost	-0.0274* (0.01043)	-0.0002 (0.00860)
Revenue	-0.0116 (0.00759)	-0.0034 (0.00337)
Profit	-0.0198* (0.00714)	0.0073 (0.00494)
MU FE	Yes	Yes
Quarter FE	Yes	Yes
Observations	12,654	13,464

Note: The dependent variable is an indicator of whether a MU has adopted any digital automation technology during the observation period. Standard errors clustered at the quarter level in parentheses. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

MU fixed effects to account for all time-invariant unobserved heterogeneity across MUs. We also add time fixed effects to accommodate temporal variations that affect each MU equally.

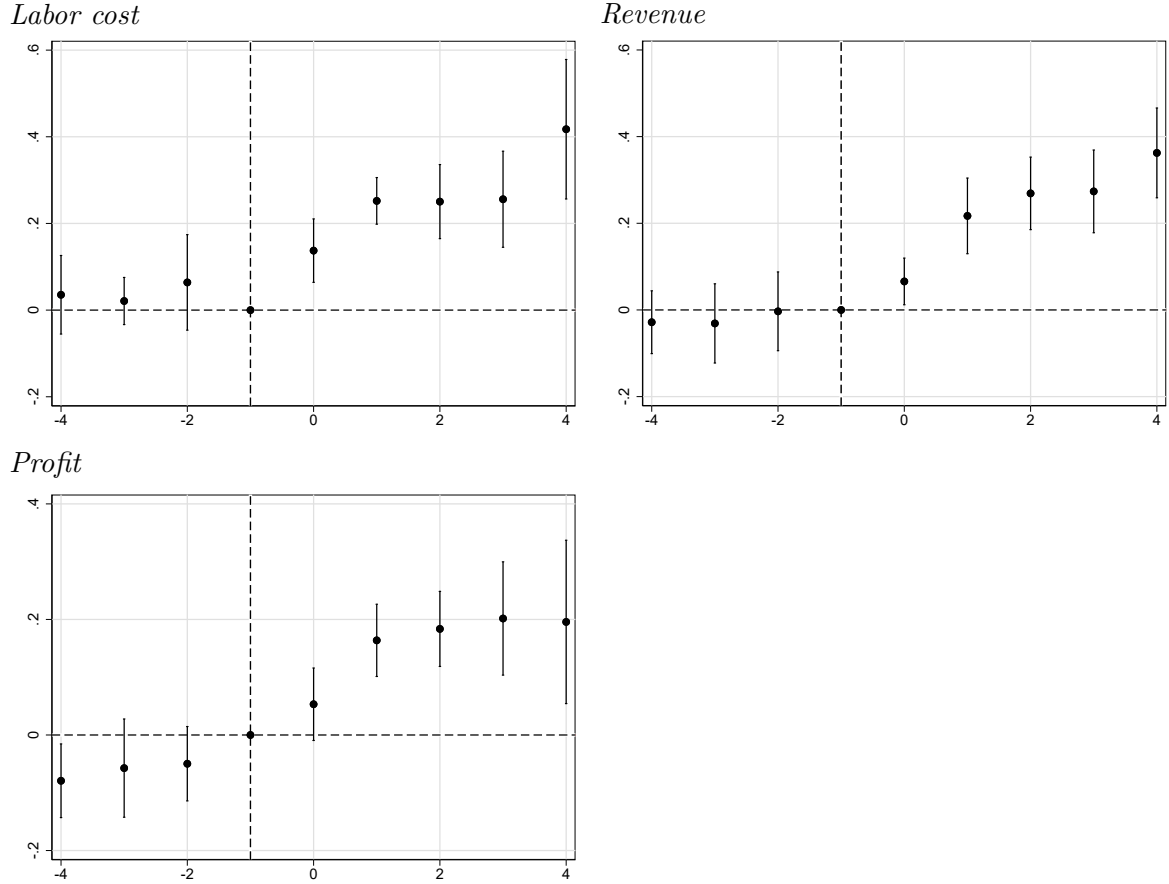
Our baseline specification links labor and operations outcomes to DAT adoption:

$$Outcome_{it} = \alpha + \delta DAT_{it} + \mu_i + \gamma_t + \varepsilon_{it}, \quad (1)$$

where the unit of observation is MU i in quarter t . We use standardize values for each of our different outcome variables to (1) help comparison across different outcomes, (2) ease interpretation of effect sizes, and (3) preserve sensible information from our industry partner. Our preferred model specification includes a MU fixed effect μ_i and a time fixed effect γ_t . We cluster robust standard errors at the MU level. We measure adoption using a dummy variable DAT_{it} , which turns one in month t where MU i implements the first DAT project of any type (“go live”).

Next, we exploit exogenous variation that affects the company’s ability to implement AI endogenously. In particular, we use the timing of an Executive Order of the Trump administration that put restrictions on the import of ICTs and related services by US corporations or US subsidiaries of non-US corporations. Effective since October 12, 2019, the order on “Securing the Information and Communications Technology and Service Supply Chain” (Executive Order 13873) bans imports from adversary countries – most importantly from China and Russia – and applies to

Figure 1: Leads and Lags, main outcomes



Note: Leads and lags plot similar to the difference-in-difference identification strategy shown in table 3. The vertical axis gives the regression coefficient, the horizontal axis indicates time relative to the actual implementation date of DAT in quarters. Bars indicate 90% confidence bands based on standard errors clustered at the MU level.

technologies “integral to artificial intelligence and machine learning, [...], drones, autonomous systems, or advanced robotics.”¹ Thanks to the company’s global operations, we can use this event to effectively compare affected MUs to not affected MUs before and after the Executive Order comes into force. We do so in an instrumental variable approach. The instrumental variable flags MUs in the US after October 2019 that have reported a DAT project before October 2019 (any status prior to “go live”) in technology categories in which the company has sourced DAT from China or Russia.

Table 3: Results: Main Outcomes

	Entire Period			Only During Covid		
	(1) Labor Cost	(2) Revenue	(3) Profit	(4) Labor Cost	(5) Revenue	(6) Profit
Post \times DAT	0.2414*** (0.03696)	0.2725*** (0.03618)	0.2038*** (0.03095)	0.2370*** (0.03860)	0.2222*** (0.03803)	0.1316*** (0.03687)
MU FE	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	35,761	32,816	36,929	20,488	18,857	21,135

Note: The dependent variables are standardized and (log+1)-transformed. *DAT* indicates whether a MU has adopted any digital automation technology during the observation period. Standard errors clustered at the MU level in parentheses. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

4 Results

4.1 Baseline results

We first discuss the results of our baseline difference-in-differences identification strategy. In Figure 1 we perform a leads and lags analysis, comparing adopters to non-adopters quarter-by-quarter. For all outcome variables (labor cost, revenue and profit), we see that pre-trends are not statistically different for adopters vs. non-adopters. This makes us confident that the necessary condition for identifying a causal effect in a difference-in-differences model is not violated. Further, we see that the effects appear to be immediate and lasting, at least during our study period of 2 years before and after the first DAT implementation.

The estimated average effects are reported in Table 3. In columns (1)–(3) we look at the entire observation period, whereas in columns (4)–(6), we exclude the period before the COVID-19 pandemic. The estimated coefficients in columns (1)–(3) are statistically not different from the estimated coefficients in columns (4)–(6) in the sense that 90% confidence intervals overlap. We find that, relative to non-adopters, DAT adopters see an increase in labor cost, revenue, and profit after the first DAT project is implemented. Our specification uses standardized outcome variables, making the interpretation of effect sizes straightforward. The point estimates suggest that the adoption of at least one DAT increases labor cost by about 24%, revenues by 22–27%, and profits by 13–20%.

¹See <https://www.federalregister.gov/documents/2019/05/17/2019-10538/securing-the-information-and-communications-technology-and-services-supply-chain>.

Table 4: Results: Main Outcomes, Staggered DiD

	(1) Labor Cost	(2) Revenue	(3) Profit
ATT	0.1976** (0.07868)	0.2755*** (0.07440)	0.1428** (0.05842)
Observations	24,879	22,549	26,089

Note: The dependent variables are standardized and (log+1)-transformed. *DAT* indicates whether a MU has adopted any digital automation technology during the observation period. Standard errors clustered at the MU level in parentheses. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

4.2 Results from alternative approaches

4.2.1 Staggered difference-in-differences

In Table 4, we explicitly take into account that different MUs are treated with DAT adoption at different points in time, which can obscure the comparisons that a simple difference-in-differences approach makes when estimating the average treatment effect (Baker et al., 2022). Applying a staggered difference-in-differences model to the matched sample and using the estimator of Callaway and Sant’Anna (2021), we continue to find positive and significant effects of DAT on labor cost, revenue and profit. The point estimates are similar and statistically not distinguishable from those reported in Table 3 in the sense that 90% confidence bands overlap. Note that the number of observations is smaller than in our baseline specification because the staggered difference-in-differences approach requires a balanced panel setup, whereas the two-period difference-in-differences approach provides a smooth average across MUs with different numbers of observed time periods.

4.2.2 ICT Import Ban as Instrument

Finally, we turn to the results of our instrumental variable approach. We first check whether we indeed see an effect of the import ban on the adoption of technology by the company. Descriptively, we do not see any (intended) purchase from Chinese or Russian vendors at all by any US operation of the firm. In Table 5, we estimate a difference-in-differences model where we compare operations in the US to operations in other countries, that are affected by the ban (because they have projects for which the company has purchased technology from Chinese or Russian vendors somewhere in their global operations), before and after the ban comes into force in October 2019. The results in column (1) show that projects are less likely to be completed,

Table 5: Results: Import Ban and Project Completion

	(1) Go-Live	(2) Phase
Post \times Affected	0.1068*** (0.01238)	0.2575*** (0.04441)
US	0.0936 (0.09261)	0.2263 (0.33207)
Post \times US	0.0761 (0.08436)	0.5768* (0.30250)
Affected \times US	0.0804 (0.08669)	0.1371 (0.31084)
Post \times Affected \times US	-0.2189** (0.08688)	-0.7994** (0.31152)
MU FE	Yes	Yes
Category FE	Yes	Yes
Month FE	Yes	Yes
Observations	73,487	73,480

Note: The dependent variable in column (1) indicates whether a project is completed. In column (2) the dependent variable is the deployment phase which can vary from -1 (canceled) to 5 (go-live). *Post* indicates the period after the import ban went into force (October 19, 2019). *US* indicates whether a MU is located in the United States and *Affected* indicates whether a project deals with technology category in which the company has somewhere in its global operations purchased from a vendor in China or Russia. Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

and according to column (2) remain in earlier deployment phases (such as canceled, planning, seeking approval, etc.). These results strongly suggest that the US import ban on technology from China and Russia has broader consequences than just a substitution of vendors. Hence, we are confident that the US import ban provides a suitable source of exogenous variation to try to identify the causal effect of DAT.

More formally, we see that the instrumental variable, an indicator for a MU in the US that at some point has projects in affected technology categories is highly correlated to the coefficient of interest $Post \times DAT$. In a first-stage regression with MU fixed effects and quarter fixed effects, the coefficient estimate is 0.790 with a t-statistic of 158.18. The value of the F test of excluded instruments is 25020.96, such that we can reject the null hypothesis that the first stage equation is under-identified with high confidence.

A key caveat to this identification approach is that we cannot observe labor data from before 2020. Hence, in Table 6, we can only report the second-stage estimates for revenue and profit.

Table 6: Results: ICT Important Ban as IV

	(1) Revenue	(2) Profit
Post \times DAT	0.3392*** (0.05253)	0.2717*** (0.05197)
MU FE	Yes	Yes
Quarter FE	Yes	Yes
Observations	32,816	36,929

Note: The dependent variables are log-transformed in columns (1), (2) and (4), and an indicator variable in column (3). *DAT* indicates whether a MU has adopted any digital automation technology during the observation period. Standard errors clustered at the MU level in parentheses. In the first-stage regression with MU fixed effects and quarter fixed effects, the coefficient estimate is 0.790 with a t-statistic of 158.18. The value of the F test of excluded instruments is 25020.96. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

We find that MUs see an increase in revenue as well as profit. The point estimates are remarkably similar to our baseline specification in Table 3 and statistically not distinguishable in the sense that 90% confidence bands overlap.

4.3 Mechanisms

To dig deeper into the underlying potential mechanisms, we study a variety of labor-specific outcomes that can explain the positive effect of DAT on labor cost. In column (1) of Table 7, we do not see an increase in labor cost per hour, suggesting that the hourly wage rate remains unaffected by DAT, at least in the relatively short time of our study period. However, columns (2) and (3) indicate that the overall increase in labor cost is predominantly attributable to an uptick in overtime work rather than an expansion in regular work hours. This suggests that employees are working additional hours beyond their standard schedules, leading to higher total labor expenditures.

Furthermore, a closer examination reveals that DAT has precipitated a shift in the composition of the workforce. Specifically, there is a noticeable reduction in the proportion of management hours, as shown in column (4). This could imply a restructuring or a reallocation of managerial duties, potentially to streamline operations. Concurrently, column (5) indicates a significant rise in the share of hours worked by temporary workers hired through agencies. This trend suggests that, in response to standardization brought about by DAT, the firm is relying more on flexible, temporary staffing solutions to meet labor needs and possibly to better manage fluctuations in customer demand.

Table 7: Mechanism results: Labor

	Labor Mechanisms				
	(1) Cost/h	(2) Regular hs	(3) Overtime hs	(4) Mgmt Share	(5) Temp Share
Post \times DAT	0.0037 (0.05405)	0.0089 (0.04951)	0.1276*** (0.03448)	-0.1246** (0.05143)	0.1025* (0.05712)
MU FE	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes
Observations	8,190	12,068	12,306	10,038	7,246

Note: The dependent variables are standardized and (log+1)-transformed, except in columns (5) and (6). *DAT* indicates whether a MU has adopted any digital automation technology during the observation period. Standard errors clustered at the user level in parentheses. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

In Table 8, we investigate the potential mechanisms underlying the observed positive effects on revenue and profit following the implementation of DAT. Column (1) shows that the adoption of DAT resulted in a significant reduction in operational costs, suggesting that DAT has enhanced operational efficiency and cost-effectiveness. This cost reduction likely contributes to the overall improvement in the firm’s financial performance.

Furthermore, column (2) provides evidence of increased warehouse utilization post-DAT implementation. Utilization is defined as the ratio of the volume processed in a given quarter to the maximum volume observed for each MU prior to the first introduction of DAT. The observed increase in utilization implies that firms are leveraging their warehouse capacity more effectively, possibly leading to better resource allocation and operational scalability.

In column (3), we observe that profit margins expanded after the adoption of DAT. This increase in profit margins indicates that the cost savings and improved resource utilization associated with DAT have translated into greater profitability. Collectively, the findings in Table 8 suggest that DAT not only reduces operational costs but also enhances the efficiency of resource utilization, thereby contributing to improved financial outcomes.

4.4 Complementarities between hardware and software

As the final step in our analysis, we examine the complementarities in the effects of DAT by distinguishing between MUs that implemented both software- and hardware-based DAT, those that implemented only software-based DAT, and those that implemented only hardware-based DAT. This distinction is crucial for understanding scale effects, as machine-learning-based

Table 8: Mechanisms results: Operations

	Operations Mechanisms		
	(1) Operational Cost	(2) Utilization	(3) Profit Margin
Post \times DAT	-0.1293** (0.05511)	0.0511*** (0.01379)	0.0816*** (0.03033)
MU FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Observations	8,215	12,076	30,057

Note: The dependent variables are standardized and (log+1)-transformed in column (1) and percentages in columns (2) and (3). *DAT* indicates whether a MU has adopted any digital automation technology during the observation period. Standard errors clustered at the MU level in parentheses. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

software can potentially be deployed more widely and cost-effectively across multiple facilities compared to service robots, which may involve higher costs and complexities in implementation. In column (1) of Table 9, the results suggest that the increase in labor costs is statistically similar across MUs that adopted both hardware and software, those that adopted only software, and those that adopted only hardware. This implies that the choice between software and hardware does not significantly affect labor cost outcomes.

However, columns (2) and (3) present a different picture regarding revenue and profit effects. MUs that implemented only software-based DAT experienced approximately half the positive impact on revenues and profits compared to those that implemented both software and hardware. The results suggest that complementarities between software- and hardware-based DATs are crucial for maximizing revenue and profit, yet they appear to have little effect on labor costs. This divergence in impact can be speculated upon by considering the different ways in which these technologies influence operational dynamics.

First, the lack of a differential effect on labor costs across the various implementations might be due to the fact that both software- and hardware-based DATs, whether used individually or together, streamline operations in ways that do not necessarily increase or decrease the total labor required. For instance, the introduction of software-based DAT, such as machine-learning algorithms, might optimize scheduling, inventory management, or customer service, while hardware-based DATs, like service robots, could automate specific physical tasks. However, both types of DATs may primarily reallocate labor rather than reduce it, leading to similar

Table 9: Complementarity results: Hardware vs software

	(1) Labor Cost	(2) Revenue	(3) Profit
Post DAT	0.2721*** (0.07366)	0.3759*** (0.07249)	0.3107*** (0.06013)
Post DAT \times Only Software	-0.0226 (0.08513)	-0.1509* (0.08281)	-0.1546** (0.06943)
Post DAT \times Only Hardware	-0.1085 (0.11092)	-0.0449 (0.11654)	-0.0530 (0.09305)
MU FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Observations	35,761	32,816	36,929

Note: The dependent variables are standardized and (log+1)-transformed. *DAT* indicates whether a MU has adopted any digital automation technology during the observation period. Standard errors clustered at the MU level in parentheses. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

labor cost outcomes regardless of the combination used.

On the other hand, the significant complementarities observed in revenue and profit could stem from the synergistic effects that arise when both types of DAT are implemented together. Software-based DATs may enhance decision-making processes, predict demand more accurately, or improve customer interactions, leading to more efficient operations. When combined with hardware-based DATs, these software-driven insights could be executed more effectively on the ground, such as through faster order fulfillment, more precise production processes, or improved service delivery. This synergy likely amplifies the overall operational effectiveness, driving higher revenues and profits.

Moreover, the increased warehouse utilization observed in previous analyses (column 2 of Table 8) might be more fully realized when both technologies are present. Software alone might identify opportunities for greater efficiency, but without the physical capabilities provided by hardware-based DATs, the full extent of these opportunities might not be capitalized upon. Similarly, the profit margin increases (column 3 of Table 8) could be a result of these enhanced efficiencies and the ability to scale operations more effectively when both software and hardware are working in tandem.

5 Discussion

5.1 Theoretical implications

By using data from one of the largest global logistics companies, we examine the effect of ML-based software and hardware on labor, productivity, and financial performance. Our analysis reveals that DAT adoption leads to increased labor demand, particularly among temporary workers, alongside improvements in financial performance. Specifically, we observe an increase in total hours worked and labor costs, coupled with higher revenues and profitability. These effects are most pronounced in establishments that adopted both hardware (robotics) and software (AI) components of DAT.

In contrast to popular discourse that AI will be predominantly adopted for substitution and recent macro-level findings ([Acemoglu et al., 2022](#)), we observe an increase in total hours worked. Our interviews with managers reveal three concurrent explanations for these findings. First, DAT adoption is boosting the efficiency and capacity of individual Management Units (MUs), enabling them to handle higher volumes of work. This local productivity enhancement is evidenced by our finding of a significant increase in warehouse utilization (Table 8, column 2) following DAT implementation. This improved capacity utilization is a key driver of the observed increases in both labor demand and financial performance. Second, and in line with Acemoglu and Restrepo (2019)’s task-creation effect, DAT implementation is generating new tasks and intensifying existing non-automated tasks. This phenomenon is directly supported by our empirical results showing a significant increase in overtime hours (Table 7, column 3) after DAT adoption. The creation of new tasks and intensification of existing ones necessitates more labor hours to keep pace with the enhanced operational capacity, explaining the observed increase in overall labor demand despite productivity improvements. Third, DAT adoption is standardizing certain processes within the organization. This standardization reduces the firm-specific human capital required for some roles, thereby lowering the barriers to using temporary labor. Our data provide direct evidence for this mechanism, showing a significant increase in the share of temporary workers (Table 7, column 5) following DAT implementation. This shift towards temporary workers allows the firm to more easily scale its workforce in response to fluctuating demand without incurring the costs associated with training and retaining permanent employees for standardized tasks.

Importantly, we find a decrease in managerial hours following DAT adoption. This suggests that AI may be particularly effective at automating or streamlining certain managerial tasks, such as monitoring and coordination. In line with our interpretation, our interviews with managers suggests that DAT improves the collection, processing, and dissemination of operational data. This increased transparency and accessibility of information often reduces the need for middle management to serve as information conduits, allowing for more decentralized decision-making. Our results contribute to SBTC theory ([Katz and Murphy, 1992](#); [Autor et al., 1998, 2003](#); [Autor, 2015](#)), which suggests that new technologies tend to favor high-skilled workers over low-skilled workers. However, our findings present a more complex picture. The reduction in managerial hours challenges the simple SBTC narrative, indicating that AI and robotics can substitute for some high-skill tasks, particularly in management. Furthermore, the increase in overall labor demand, especially for temporary workers, suggests that DAT may create opportunities for workers across skill levels, not just high-skilled workers.

We contribute to the growing literature on AI and workers ([Fügener et al., 2021](#); [Felten et al., 2021](#); [Jain et al., 2021](#); [Lebovitz et al., 2021](#); [Acemoglu et al., 2022](#); [Chen et al., 2024](#)) by showing how AI adoption reshapes labor demand and task allocation within organizations. Our findings reveal a complex interplay between technology adoption, task creation, and human capital reconfiguration. Specifically, we show that the adoption of AI-based technologies can simultaneously create additional tasks for existing workers while reducing the need for managerial attention. This dual effect necessitates a strategic reconfiguration of human capital within the organization. Our results suggest that AI adoption leads to standardization of certain tasks, which in turn decreases transaction costs such as search and information costs. This standardization makes it easier for companies to hire temporary workers and reduces the need for constant managerial oversight. These insights align with the predictions of Transaction Cost Economics theory ([Williamson, 2016](#); [Poppo and Zenger, 2002](#)), suggesting that AI can significantly reduce coordination costs, leading to more outsourcing and increased use of market mechanisms for coordination.

Our results also contributes to the ongoing debate about the relationship between technology and productivity ([Bresnahan et al., 2002](#); [Bloom et al., 2012](#)). The original productivity paradox, observed in the 1980s and early 1990s, referred to the apparent contradiction between significant

investments in information technology and relatively slow productivity growth (Brynjolfsson and Hitt, 1996). Brynjolfsson et al. (2019b) maintains that a similar productivity paradox might be observed with AI technology. However, our results suggest that, it may be easier to achieve productivity gains via AI given it’s adaptive nature.

Finally, the superior outcomes observed when both hardware and software components of DAT are adopted together extend the literature on complementarities in technology adoption (Milgrom and Roberts, 1995; Bloom et al., 2012). The synergistic effects of combined hardware and software adoption suggest that the full potential of AI technologies may only be realized when they are integrated with physical systems. This finding has important implications for firms’ technology investment strategies and highlights the potential limitations of purely software-based AI solutions.

5.2 Managerial implications

In light of the insights generated from our study, several managerial implications emerge. First, given that establishments benefit most from adopting both software and hardware components of DAT, managers should be inclined towards a comprehensive approach rather than partial adoption. Integrative strategies might offer the dual advantage of operational efficiency.

Second, contrary to conventional wisdom, our findings highlight that DAT adoption can lead to increased working hours, particularly among temporary workers. Managers should proactively prepare for such shifts in work patterns. This includes forecasting potential increases in overtime costs and adapting scheduling practices to accommodate extended work hours. Additionally, a potential dependence on temporary labor might arise. However, with effective recruitment and training, this may allow firms to more easily scale their workforce up or down, reducing the costs associated with long-term labor commitments.

Third, the observation that software-based and hardware-based DAT can be complementary important managerial implications as well. Managers should exercise caution when introducing automation in a piece-wise fashion. Organizations should strive to identify the ‘sweet spot’ of DAT intervention. This requires careful evaluation of existing processes and targeted implementation of automation where it can deliver the most substantial enhancements. By doing so, managers can avoid potential downsides associated with over-automation and guarantee that technology amplifies performance in a meaningful manner.

5.3 Policy implications

The implications of our research carry significant importance for policymakers as they grapple with the multifaceted challenges posed by the integration of artificial intelligence and robotics into various sectors of the economy. Contrary to the prevailing perception perpetuated by the popular press, our study reveals nuances that call for a recalibration of policy approaches.

One of the predominant fears surrounding the rise of artificial intelligence and robotics is the notion of substitution, wherein technological advancements lead to a reduction in human labor. Our findings offer contrary evidence to this narrative by highlighting that performance gains facilitated by these technologies do not inherently equate to a decrease in labor hours. This nuanced perspective underscores the complex relationship between technology and labor, underscoring that performance enhancements can, in fact, coexist with sustained labor engagement. Policymakers must recognize that the impact of technology on employment is intricate and multifaceted, requiring a more sophisticated approach to workforce planning and policy formulation.

Moreover, it is crucial to recognize that the notion of machines replacing all forms of employment is far from reality. Our findings emphasize that technological advancements can serve as complements to human skills and labor, rather than outright substitutes. Policymakers should capitalize on this potential synergy by promoting policies that facilitate upskilling and reskilling initiatives. By investing in educational and training programs that equip workers with the necessary skills to collaborate with evolving technologies, policymakers can empower the workforce to remain adaptable and competitive in an increasingly automated landscape.

5.4 Limitations

Our analysis is ultimately a case study of one large organization that may or may not be representative of other organizations in the entire economy. While our analysis is ultimately a case study of one large multinational logistics company, which may not be fully representative of other organizations in the broader economy, this focused approach offers several distinct advantages. The unique, granular data we have collected permits particularly convincing empirical tests of the effects of digital automation technologies (DAT), including AI and robotics, on labor demand, operations, and profitability. The homogeneity of operations and processes within the company, combined with our estimation of longitudinal models, eliminates many sources

of unmeasured heterogeneity that often confound productivity comparisons in more aggregated data and broader samples. This allows us to isolate the effects of DAT adoption with a degree of precision that would be challenging to achieve in a multi-industry study.

Moreover, the trends we observe in increased operational efficiency and service customization through DAT adoption are likely to characterize many industries in the current digital era. In manufacturing sectors, DAT can reduce setup times and increase product variety, much like our observed effects in logistics operations. In service industries such as healthcare, retail, or financial services, similar technologies may enhance process efficiency and enable more personalized service delivery.

However, we acknowledge that the magnitude and specific nature of these effects may differ across industries and organizational contexts. Factors such as regulatory environments, labor market conditions, and the nature of the product or service may influence how DAT adoption impacts organizational outcomes. Future research could extend our findings by examining similar dynamics in other industries or by conducting cross-industry comparisons.

Further, it is of course difficult to anticipate general equilibrium effects, but it still seems that there is a lot potential for these effect to aggregate to substantial business value. Especially with increased roll out of DAT through out the organizations operations. As shown in table 1, at the end of our observation period in September 2021 only 18% of MUs in the entire organization have adopted DAT.

Finally, our data do not allow us to fully explore the distributional consequences of DAT adoption within the firm. Understanding how the benefits and costs of AI are distributed across different types of workers and organizational units remains an important area for future research.

6 Conclusion

In this paper, we studied the effects of digital automation technology on labor demand, operations and profitability. Drawing from a rich dataset spanning one of the world’s largest logistics companies, our findings offer fresh perspectives that challenge both conventional wisdom and previously held academic views about digital automation.

Our research underscores the relationship between DAT adoption and its impact on labor. While popular discourse might point to a linear relation where technological enhancements lead to reduced labor demand, our empirical evidence paints a different picture. The increase in

working hours, especially among temporary workers post-DAT adoption, emphasizes that the role of automation extends beyond mere labor substitution of lower-skilled workers, at least in the short run. If at all, our setting shows that digital automation technology reduces demand for managerial tasks.

We contribute to the literature in several ways. We not only bridge the chasm between industry-level studies and firm-specific implications with our granular data but also present a nuanced perspective on the effects of automation on labor demand. Our empirical results suggest the significance of adopting an integrative strategy involving both software and hardware components of DAT. Furthermore, by highlighting the potential diminishing returns of AI when complementarities between software and hardware are ignored, we provide a tempered view, indicating that digital automation technology is not a panacea and its benefits might be constrained in certain contexts.

In sum, as the global economy continues to evolve in the face of rapid technological advancements, it becomes crucial for both managerial decision-makers and policymakers to understand the question of who benefits from robots and AI, and under which circumstances. Our study provides a crucial step in that direction, offering insights that can guide both strategic endeavors and policy initiatives in an era where even fully physical industries go through their digital transformation.

References

- Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). “Artificial intelligence and jobs: Evidence from online vacancies.” *Journal of Labor Economics*, 40(S1), S293–S340.
- Acemoglu, D., and Restrepo, P. (2019). “Automation and new tasks: How technology displaces and reinstates labor.” *Journal of Economic Perspectives*, 33(2), 3–30.
- Acemoglu, D., and Restrepo, P. (2020). “Robots and jobs: Evidence from us labor markets.” *Journal of Political Economy*, 128(6), 2188–2244.
- Agrawal, A., Gans, J., and Goldfarb, A. (2022). *Power and prediction: The disruptive economics of artificial intelligence*. Harvard Business Press.
- Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). “Exploring the impact of artificial intelligence: Prediction versus judgment.” *Information Economics and Policy*, 47, 1–6.
- Aron, R., Dutta, S., Janakiraman, R., and Pathak, P. A. (2011). “The impact of automation of systems on medical errors: evidence from field research.” *Information systems research*, 22(3), 429–446.
- Autor, D. (2015). “Why are there still so many jobs? The history and Future of Workplace Automation.” *Journal of Economic Perspectives*, 29(3), 3–30.
- Autor, D. H., Katz, L. F., and Krueger, A. B. (1998). “Computing inequality: have computers changed the labor market?” *The Quarterly journal of economics*, 113(4), 1169–1213.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). “The skill content of recent technological change: An empirical exploration.” *The Quarterly journal of economics*, 118(4), 1279–1333.
- Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). “Artificial intelligence, firm growth, and product innovation.” *Journal of Financial Economics*, 151, 103745.
- Baker, A. C., Larcker, D. F., and Wang, C. C. (2022). “How much should we trust staggered difference-in-differences estimates?” *Journal of Financial Economics*, 144(2), 370–395.
- Bauer, K., von Zahn, M., and Hinz, O. (2023). “Expl (ai) ned: The impact of explainable artificial intelligence on users’ information processing.” *Information systems research*, 34(4), 1582–1602.
- Bloom, N., Sadun, R., and Reenen, J. V. (2012). “Americans do it better: Us multinationals and the productivity miracle.” *American Economic Review*, 102(1), 167–201.
- Bresnahan, T. F., Brynjolfsson, E., and Hitt, L. M. (2002). “Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence.” *The quarterly journal of economics*, 117(1), 339–376.
- Brynjolfsson, E., and Hitt, L. (1996). “Paradox lost? firm-level evidence on the returns to information systems spending.” *Management science*, 42(4), 541–558.
- Brynjolfsson, E., Hui, X., and Liu, M. (2019a). “Does machine translation affect international trade? evidence from a large digital platform.” *Management Science*, 65(12), 5449–5460.
- Brynjolfsson, E., Jin, W., and McElheran, K. (2021). “The power of prediction: predictive analytics, workplace complements, and business performance.” *Business Economics*, 56, 217–239.

- Brynjolfsson, E., and McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.
- Brynjolfsson, E., and Mitchell, T. (2017). “What can machine learning do? workforce implications.” *Science*, 358(6370), 1530–1534.
- Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). “What can machines learn, and what does it mean for occupations and the economy?” In *AEA Papers and Proceedings*, vol. 108, 43–47.
- Brynjolfsson, E., Rock, D., and Syverson, C. (2019b). “Artificial intelligence and the modern productivity paradox.” *The economics of artificial intelligence: An agenda*, 23, 23–57.
- Callaway, B., and Sant’Anna, P. H. (2021). “Difference-in-differences with multiple time periods.” *Journal of Econometrics*, 225(2), 200–230.
- Chen, R., Balasubramanian, N., and Forman, C. (2024). “How does worker mobility affect business adoption of a new technology? the case of machine learning.” *Strategic Management Journal*.
- Cockburn, I. M., Henderson, R., and Stern, S. (2018). *The impact of artificial intelligence on innovation*, vol. 24449. National bureau of economic research Cambridge, MA, USA.
- Deb, P., Furceri, D., Ostry, J. D., and Tawk, N. (2022). “The economic effects of covid-19 containment measures.” *Open Economies Review*, 33(1), 1–32.
- Dell’Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., and Lakhani, K. R. (2023). “Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality.” *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (24-013).
- Edmondson, A. C., Bohmer, R. M., and Pisano, G. P. (2001). “Disrupted routines: Team learning and new technology implementation in hospitals.” *Administrative science quarterly*, 46(4), 685–716.
- Feigenbaum, J., and Gross, D. P. (2024). “Organizational and economic obstacles to automation: A cautionary tale from at&t in the twentieth century.” *Management Science*.
- Felten, E., Raj, M., and Seamans, R. (2021). “Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses.” *Strategic Management Journal*, 42(12), 2195–2217.
- Frank, M. R., Autor, D., Bessen, J. E., Brynjolfsson, E., Cebrian, M., Deming, D. J., Feldman, M., Groh, M., Lobo, J., Moro, E., et al. (2019). “Toward understanding the impact of artificial intelligence on labor.” *Proceedings of the National Academy of Sciences*, 116(14), 6531–6539.
- Fügener, A., Grahl, J., Gupta, A., and Ketter, W. (2021). “Will humans-in-the-loop become borgs? merits and pitfalls of working with ai.” *Management Information Systems Quarterly (MISQ)-Vol*, 45.
- Helfat, C. E., and Peteraf, M. A. (2015). “Managerial cognitive capabilities and the microfoundations of dynamic capabilities.” *Strategic management journal*, 36(6), 831–850.

- Henderson, R. M., and Clark, K. B. (1990). “Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms.” *Administrative science quarterly*, 9–30.
- Jain, H., Padmanabhan, B., Pavlou, P. A., and Raghu, T. (2021). “Editorial for the special section on humans, algorithms, and augmented intelligence: The future of work, organizations, and society.” *Information Systems Research*, 32(3), 675–687.
- Jordan, M. I., and Mitchell, T. M. (2015). “Machine learning: Trends, perspectives, and prospects.” *Science*, 349(6245), 255–260.
- Jussupow, E., Spohrer, K., Heinzl, A., and Gawlitza, J. (2021). “Augmenting medical diagnosis decisions? an investigation into physicians’ decision-making process with artificial intelligence.” *Information Systems Research*, 32(3), 713–735.
- Katz, L. F., and Murphy, K. M. (1992). “Changes in relative wages, 1963–1987: supply and demand factors.” *The quarterly journal of economics*, 107(1), 35–78.
- Krakowski, S., Luger, J., and Raisch, S. (2023). “Artificial intelligence and the changing sources of competitive advantage.” *Strategic Management Journal*, 44(6), 1425–1452.
- Lebovitz, S., Levina, N., and Lifshitz-Assaf, H. (2021). “Is ai ground truth really true? the dangers of training and evaluating ai tools based on experts’ know-what.” *MIS quarterly*, 45(3).
- Lou, B., and Wu, L. (2021). “Ai on drugs: Can artificial intelligence accelerate drug development? evidence from a large-scale examination of bio-pharma firms.” *Evidence from a Large-scale Examination of Bio-pharma Firms*, (March 15, 2021). *MISQ Forthcoming*.
- Meier, M., Pinto, E., et al. (2020). “Covid-19 supply chain disruptions.” *Covid Economics*, 48(1), 139–170.
- Milgrom, P., and Roberts, J. (1990). “Rationalizability, learning, and equilibrium in games with strategic complementarities.” *Econometrica: Journal of the Econometric Society*, 1255–1277.
- Milgrom, P., and Roberts, J. (1995). “Complementarities and fit strategy, structure, and organizational change in manufacturing.” *Journal of accounting and economics*, 19(2-3), 179–208.
- Peng, S., Kalliamvakou, E., Cihon, P., and Demirer, M. (2023). “The impact of ai on developer productivity: Evidence from github copilot.” *arXiv preprint arXiv:2302.06590*.
- Peukert, C., Sen, A., and Claussen, J. (2023). “The editor and the algorithm: Recommendation technology in online news.” *Management Science*, *forthcoming*.
- Poppo, L., and Zenger, T. (2002). “Do formal contracts and relational governance function as substitutes or complements?” *Strategic management journal*, 23(8), 707–725.
- Raj, M., and Seamans, R. (2019). “Primer on artificial intelligence and robotics.” *Journal of Organization Design*, 8, 1–14.
- Ribers, M. A., and Ullrich, H. (2024). “Complementarities between algorithmic and human decision-making: The case of antibiotic prescribing.” *Quantitative Marketing and Economics*, 1–39.

- Sirmon, D. G., Hitt, M. A., Ireland, R. D., and Gilbert, B. A. (2011). “Resource orchestration to create competitive advantage: Breadth, depth, and life cycle effects.” *Journal of management*, 37(5), 1390–1412.
- Stadler, C., Helfat, C. E., and Verona, G. (2022). “Transferring knowledge by transferring individuals: Innovative technology use and organizational performance in multiunit firms.” *Organization Science*, 33(1), 253–274.
- Teece, D. J., Pisano, G., and Shuen, A. (1997). “Dynamic capabilities and strategic management.” *Strategic management journal*, 18(7), 509–533.
- Wang, W., Gao, G., and Agarwal, R. (2023). “Friend or foe? teaming between artificial intelligence and workers with variation in experience.” *Management Science*.
- Webb, M. (2019). “The impact of artificial intelligence on the labor market.” *Available at SSRN 3482150*.
- Williamson, O. E. (2016). “Comparative economic organization: The analysis of discrete structural alternatives.” In *Institutional change*, 75–108, Routledge.