Scaling the Returns to Automation: Algorithms, Robots and Organizational Learning*

Erdem Dogukan Yilmaz[1] and Christian Peukert[2]

[1]INSEAD
[2]University of Lausanne, Faculty of Business and Economics

May 4, 2026
Preliminary draft. Please do not cite or circulate.

Abstract

Digital automation technologies (DAT), including algorithmic optimization and advanced robotics, are increasingly adopted, yet evidence on how organizations scale their benefits remains limited. We study the operational and labor effects of DAT using unique data from a global logistics firm covering 13,000 management units in 60 countries. Along-side granular performance and operational data, we can follow 4,000 projects and deployment decisions of algorithmic optimization tools and robotic devices. Exploiting quasi-random delays in deployment timing and policy-induced variation in technology availability, we can identify the causal effects of DAT. We find that DAT increases revenue and profit, as well as total labor demand by about 20 percent. These effects are driven by higher capital utilization and task expansion, alongside a reduction in operative managerial hours. We show that technological complementarity matters to surface these benefits. The largest gains arise when facilities deploy algorithmic tools and robotic devices jointly rather than in isolation. Crucially, we find that implementation experience is a key driver of scale. Organization-wide experience with the implementation of DAT explains between 20 and 40 percent of the observed increases in labor demand, revenue, and profit. Together, these findings highlight the organizational challenge of scaling automation. The benefits of robots and algorithmic optimization do not emerge automatically, but increase with accumulated and well-managed implementation experience and the deliberate exploitation of complementarities. Organizations scale the returns to automation by learning how to deploy, adapt, and combine technologies across operations, rather than by expanding technology use alone.

Keywords: Digital Automation Technology, Robots, Artificial Intelligence, Future of Work

*Yilmaz: erdemdogukan.yilmaz@insead.edu, Peukert: christian.peukert@unil.ch. We are grateful for excellent research assistance from Lucas Franco (initial analysis of quantitative data) and Rachel Moser (collection, transcription and interpretation of qualitative data). 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

A new wave of Digital Automation Technology (DAT), especially artificial intelligence (AI), algorithmic optimization and robotics are proliferating across organizations (Allen and Choudhury, 2022; Bailey et al., 2022). While recent attention has centered on generative AI and its effects on knowledge work (Noy and Zhang, 2023; Hui et al., 2024; Brynjolfsson et al., 2025), comparatively less is known about how software (e.g. algorithmic optimization) in combination with physical robots transform organizational hierarchies and managerial roles. Previous literature on technology adoption suggests that achieving the benefits of new technologies requires significant organizational changes and alters the division of labor and coordination patterns (Puranam et al., 2014; Bailey et al., 2022; Joseph and Sengul, 2025). This requirement for organizational changes has historically led to considerable lags between technology adoption and observable performance improvements (Barley, 1986; Leonard-Barton, 1988).

This question is particularly important because the effects of automation may differ substantially depending on the type of technology and the occupational categories affected. Existing research on AI and labor has predominantly examined software-based tools that augment white-collar workers—from machine translation affecting trade (Brynjolfsson et al., 2019) to AI-assisted medical coding (Wang et al., 2023) and customer support (Brynjolfsson et al., 2025). Yet physical robots represent a distinct technology category with different implementation requirements, coordination demands, and potential to restructure organizational hierarchies. Our study addresses this gap by examining the organization-wide deployment of algorithmic optimization tools and robotics across a global logistics firm.

A central premise of our study is that the benefits of DAT do not arise mechanically from adoption, but depend on how technologies are combined and implemented over time. In particular, we argue that two factors are critical for scaling automation benefits within organizations: technological complementarities between robots and algorithmic optimization, and organization-wide implementation experience accumulated across projects. These factors shape whether automation translates into sustained gains in capacity, revenue, and profitability rather than isolated local improvements.

The new wave of DATs may fundamentally alter the established pattern of delayed productivity gains following technology adoption. Unlike earlier technological solutions, DATs can learn and

adapt to existing workflows by mimicking human cognition and action (Von Krogh, 2018; Wu et al., 2020; Bailey et al., 2022), and they often function as targeted "point solutions (Agrawal et al., 2022; Wang et al., 2023; Feigenbaum and Gross, 2024). As a result, they can be implemented with relatively minimal disruption to existing processes. Moreover, by substituting for certain tasks, DATs can free up labor for more value-added activities (Agrawal et al., 2022; Wang et al., 2023; Hui et al., 2024). Importantly, these technologies rarely operate in isolation. Physical robots execute tasks, while algorithmic optimization governs coordination, routing, and resource allocation.

These arguments raise an important question: Can these new, adaptable DATs overcome the historical lag between adoption and performance improvements, and if so, through what organizational mechanisms do firms learn to scale their benefits over time? What is the role of learning from experience — both positive and negative? Historically, studying this question has been difficult due to limitations in available data. Most prior research has relied on aggregate industry-level metrics or cross-sectional surveys that lack the temporal and organizational granularity needed to assess how these technologies affect firm performance and the composition of the workforce. In contrast, our unique data from a global logistics company offers a rare opportunity to examine the impact of digital automation at scale and with fine-grained detail. In 2018, 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 across its facilities in 60 countries. We draw on a rich proprietary dataset covering over 4,000 automation projects, managed by close to 13,000 units, covering the firm's over 3,000 facilities worldwide. We observe the life-cycle of each project, from screening, to planning, to implementation or cancellation. This granular data enables us to document strong positive correlations between DAT adoption and operational outcomes, with additional analyses exploiting quasi-experimental variation to strengthen causal interpretation. This dataset includes structured information on the type of technology used, deployment status, timing, location, and key variables to capture outcomes related to work, operations and performance. A distinctive feature of our setting is that automation projects are implemented across thousands of sites under a centralized program that tracks both successful implementations and canceled projects. This allows us to observe how experience accumulates at the organizational

level, rather than only within individual facilities. Managers learn not only from successful deployments, but also from projects that are reviewed and ultimately abandoned, refining their ability to identify where automation is likely to pay off. To complement the interpretation of results from the observational data, we incorporate qualitative insights from conversations with the global program manager as well as 13 semi-structured interviews with employees at various hierarchical levels involved in individual projects at various locations worldwide.

Identifying a clear causal relationship between DAT adoption and operational as well as workforce-related outcomes is challenging, as many internal and external factors influence both technology adoption and performance. First, we address the endogeneity of adoption timing, e.g., that management units (MUs) deploy DAT in times when faced with a shortage of labor. For this, the overlap of our study's timeline with the COVID-19 pandemic provided a unique opportunity: pandemic-induced disruptions rendered the timing of DAT implementation quasi-random across facilities, delaying installation and training for both software and robotics. Leveraging this serendipitous randomness, we employ a difference-in-differences approach to estimate the effects of DAT adoption. Second, we address identification issues stemming from the choices of MUs to deploy DAT in the first place. We 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.

Our analysis reveals four core findings. First, warehouses implementing digital automation experience immediate gains in operational performance, reflected in about 20% higher revenue and profit, with effects strongest when robots and algorithmic optimization are deployed jointly. Second, automation is associated with an expansion of total labor demand and labor costs (about 20%), driven by longer operating hours, higher capacity utilization, and new execution and sup-

port tasks. Third, we document a decline in managerial labor intensity, as algorithmic systems substitute for traditional supervisory coordination. Fourth, and central to scaling, we show that organization-wide implementation experience is a major driver of these effects. Accumulated experience from prior automation projects explains between 20% and 40% of the observed increases in labor demand, revenue, and profit. Importantly, learning occurs not only from successful deployments but also from canceled projects, which sharpen opportunity selection and amplify subsequent returns.

With these results, our study contributes to ongoing debates in the productivity, labor, and organizational design literatures. First, while recent research has examined how generative AI affects white-collar and creative work (Noy and Zhang, 2023; Hui et al., 2024; Brynjolfsson et al., 2025), we show that physical robots combined with AI software also transform white-collar roles—specifically managerial coordination—challenging the assumption that robotics primarily automates blue-collar tasks. Second, we demonstrate that DATs can deliver performance gains without the extended productivity lags historically associated with technology adoption (Brynjolfsson and Hitt, 1996; Brynjolfsson et al., 2019). We demonstrate that DATs not only improve operational efficiency and profit margins, but also reshape workforce dynamics: increasing labor costs through extended operations, reducing managerial intensity via algorithmic coordination, and enabling more flexible deployment of frontline labor. These findings contribute to the growing literature on AI and labor (Choudhury et al., 2020; Yilmaz et al., 2023; Hui et al., 2024). More broadly, our results show that digital automation functions as a capacity-expanding technology whose returns scale with organizational learning. Firms that combine robots with algorithmic optimization and accumulate implementation experience are able to expand operations, reallocate labor away from managerial oversight, and translate automation into sustained revenue and profit growth. With these new insights we contribute to the literature on technology and organizational design (Puranam et al., 2014; Joseph and Sengul, 2025).

2 Related literature

Research on technology adoption and organizational performance has a long history, highlighting both the potential benefits of technology and the challenges and time lags in reaping those benefits (Barley, 1986; Leonard-Barton, 1988; Pisano, 1996; Edmondson et al., 2001). Early studies highlighted how the adoption of new technologies often requires complementary invest-

ments in worker skills, process redesign, and new routines (Barley, 1990; Mitchell and Zmud, 1999). These changes frequently generate considerable time lags between the initial technology deployment and the realization of performance improvements, as organizations gradually develop the required coordination mechanisms and structures (Puranam et al., 2014; Bailey et al., 2022; Joseph and Sengul, 2025). A consistent theme is that technologies, in isolation, rarely increase productivity unless accompanied by organizational realignments. This perspective suggests technology adoption follows a J-curve pattern, where initial implementation disrupts existing routines before eventually yielding performance improvements (Brynjolfsson et al., 2021). By contrast, technologies such as artificial intelligence and robotics differ from earlier enterprise systems in their ability to learn from large data streams and adapt to changing conditions (Allen and Choudhury, 2022; Bailey et al., 2022). Unlike legacy ERP and SCM systems—which relied on extensive custom coding, centralized control, and rigid rule sets—digital automation technologies (DATs) exhibit two distinctive attributes that compress or eliminate traditional implementation lags. First, DATs are modular, meaning they can be introduced as "point solutions" with relatively few dependencies on other business processes. Second, they rely on self-learning algorithms that adapt in real time to operational data rather than static programming. Together, these features reduce the need for top-down organizational reconfigurations and allow faster capture of productivity gains compared to earlier, more monolithic enterprise systems. This modularity, which helps isolate changes to localized subsystems (Sanchez and Mahoney, 1996; Schilling, 2000), reduces the scope of disruption for individual business units, enabling incremental integration of new tools and faster capture of productivity gains than would be possible in a more centralized, monolithic adoption.

Such "simplification"—the reduction of interdependencies through targeted coordination mechanisms—can accelerate organizational adaptation (Puranam et al., 2014). In logistics setting, for example, DATs may embed algorithmic decision rules for tasks such as order picking or inventory management. From an organizational design perspective, this form of 'algorithmic coordination' reshapes traditional decision rights and may manifest in a measurable decrease in managerial work hours, while simultaneously allowing the organization to scale operations more efficiently. Additionally, the self-learning features of DAT help it adapt to local operational conditions with minimal need for large-scale restructuring (Agrawal et al., 2019; Brynjolfsson et al., 2019). While

prior enterprise technologies relied on static rules and constant oversight, these AI-driven point solutions refine their routines in near-real time using continuous data inputs. This learning-by-doing aspect is especially critical in mitigating the typical implementation dip: rather than waiting for extensive corporate reconfigurations or staff retraining, companies can "plug in" algorithms that immediately adjust different tasks. As a result, organizations can avoid the protracted productivity trough that often accompanies major technology overhauls.

For example, tasks such as inventory management, route optimization, or item recognition can be partially automated by algorithms that continually refine themselves via real-time feedback without requiring radical workflow changes. Similarly, in office and finance operations, robotic process automation (RPA) executes digital tasks through the same interfaces employees use, thereby allowing firms to integrate RPA into legacy IT systems without substantial modifications[1]. Together, these modular and self-learning features accelerate the path to productivity gains, minimize redesign requirements for complementary processes, and streamline integration with existing systems.

Still, even modular DATs can reverberate throughout broader organizational processes (Puranam et al., 2014; Joseph and Sengul, 2025). Substituting traditional managerial oversight with algorithmic coordination may reconfigure decision rights, reshape the managerial workforce, or require new supervisory skill sets (Barley, 1990). Moreover, algorithms may lack the domain-specific expertise to frame complex problems, potentially requiring human intervention in exception handling (Allen and Choudhury, 2022; Acemoglu et al., 2022).

Importantly, recent empirical work on AI and labor has predominantly focused on software-based and generative AI technologies. Studies have examined how machine translation affects international trade (Brynjolfsson et al., 2019), how AI feedback influences employee performance (Tong et al., 2021), and how generative AI tools affect freelancers (Hui et al., 2024), customer support agents (Brynjolfsson et al., 2025), and professional writing (Noy and Zhang, 2023). While this literature has illuminated the effects of AI on knowledge work and individual productivity, it has paid less attention to how physical robotics—combined with AI software—might reconfigure organizational hierarchies. Our study complements this work by examining a context where robots and ML-based software jointly automate not only manual tasks but also coordination

[1] https://www.uipath.com/rpa/robotic-process-automation

functions traditionally performed by managers.

Against this backdrop, our study examines First, can modern DATs overcome the historical lags between technology adoption and productivity gains? Second, and distinct from the prevailing focus on software-based AI, do physical robots combined with AI software affect managerial work—not just frontline operations? By documenting a significant decline in managerial labor share following DAT adoption, we provide evidence that robots and AI can automate white-collar coordination tasks, extending the scope of automation's organizational impact beyond blue-collar labor.

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. The firm's global footprint is organized into MUs, each typically aligned with a specific customer in a given country. A single warehouse may host multiple MUs, and occasionally, a single MU may span multiple warehouses. MUs oversee various logistics tasks, including inventory management, picking and packing, and shipment scheduling. At the local level, each MU complies with host-country regulations and market norms, which influence how costs like labor and IT investments are categorized. Despite this localized reporting, the company employs a standardized global reporting system. We use this standardized data in our analyses to ensure robustness and comparability across different units and regions.

This system consolidates financial and operational data from all MUs using uniform accounting principles and cost-allocation standards. Labor costs, recorded at the MU-level, include expenditures on regular and overtime hours, as well as temporary staff. Operational costs, tracked systematically according to company rules, encompass materials, utilities, and maintenance. MU-level profits are calculated by subtracting labor and operational costs from revenues generated by the services provided. This approach ensures a consistent and comparable measure of financial performance across the company's global network.

In April 2018, the company initiated what it refers to as the Digital Automation Technology (DAT) program. Specifically, DAT refers to integrated systems that combine machine learning-

based software, robotics, and related automation tools to learn from operational data, adjust to changing conditions, and execute repetitive tasks with minimal human intervention. In contrast to conventional enterprise technologies such as Enterprise Resource Planning (ERP), Supply Chain Management (SCM), or Customer Relationship management (CRM) systems—which rely on static rule sets and continuous human input—DAT can autonomously refine its own processes, offering a more adaptive and data-driven approach to optimizing operations.

Prior to adoption, the company's operations relied on a combination of human labor and older-generation IT systems. While human workers carried out tasks such as item recognition, order picking, and inventory management, these processes were supported—albeit in a limited capacity—by legacy software and hardware solutions. These systems, however, were largely rule-based and lacked the adaptability and predictive capabilities of contemporary machine learning-driven applications. For instance, consider the point-to-point movement of items within warehouses. Before DAT, workers often followed static, predetermined routes to transport goods from one location to another. Although some rudimentary IT tools were in place, they could not dynamically adjust pathways in response to real-time conditions such as congestion, equipment failures, or sudden changes in order priority. This meant that even incremental adjustments to routing or scheduling had to be planned and executed manually, limiting the firm's ability to respond quickly to shifting operational demands.

Similarly, item recognition tasks largely depended on human judgment. Workers visually inspected products to ensure they matched order specifications, a time-consuming and error-prone approach. While barcode scanning and other basic identification systems existed, these technologies did not incorporate advanced image recognition algorithms or self-learning capabilities. As a result, misclassification of items or delays in identification could occur, reducing efficiency and occasionally leading to incorrect shipments or inventory mismatches. Order picking and inventory management also leaned heavily on a combination of human effort and simple IT solutions. Inventory records were updated manually or via basic database systems that lacked predictive modeling features. Workers relied on static lists to pick items, and any forecast of future demand or stockouts typically stemmed from aggregated historical data, not from real-time analytics. In short, while older IT systems provided some structure and record-keeping functions, they did not empower workers or managers to optimize their activities based on immediate feedback or

advanced data-driven insights.

The company chose 12 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). In all cases, the baseline scenario is the use of older-generation IT systems that are less adaptive and more reliant on static rules, complemented by substantial human input. This means that the shift we examine—moving from these older systems and manual procedures to new, machine learning-based DAT—is a clear counterfactual, allowing us to measure the incremental effect of DAT adoption on labor demand, operations, and profitability. 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 robotic 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 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, 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 MU. 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.

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

Table 1: Variable overview

Variables	Uniques	Observed period
DAT Projects	4,132	Q1/2019–Q3/2021
...of which		
Implemented	62%	
Canceled	28%	
In planning	20%	
Robotics	36%	
Algorithmic Optimization	64%	
Vendors	147	
Management Units	12,894	
...of which		
implement DAT projects	18%	
in US	11%	
in US and affected by import ban	10%	
Labor		Q1/2020–Q3/2021
Total Cost		
Cost/Hour		
Regular Hours		
Overtime Hours		
Management Share		
Temporary Share		
Operations		Q1/2019–Q3/2021
Revenue		
Profit		
Operational Cost		
Warehouse Utilization		
Profit Margin		

Note: For the baseline specifications, we focus on the effects of implemented projects. For studying heterogeneity, we also take experience from cancelled projects into consideration.

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

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*	-0.0002
	(0.01043)	(0.00860)
Revenue	-0.0116	-0.0034
	(0.00759)	(0.00337)
Profit	-0.0198*	0.0073
	(0.00714)	(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. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

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 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.

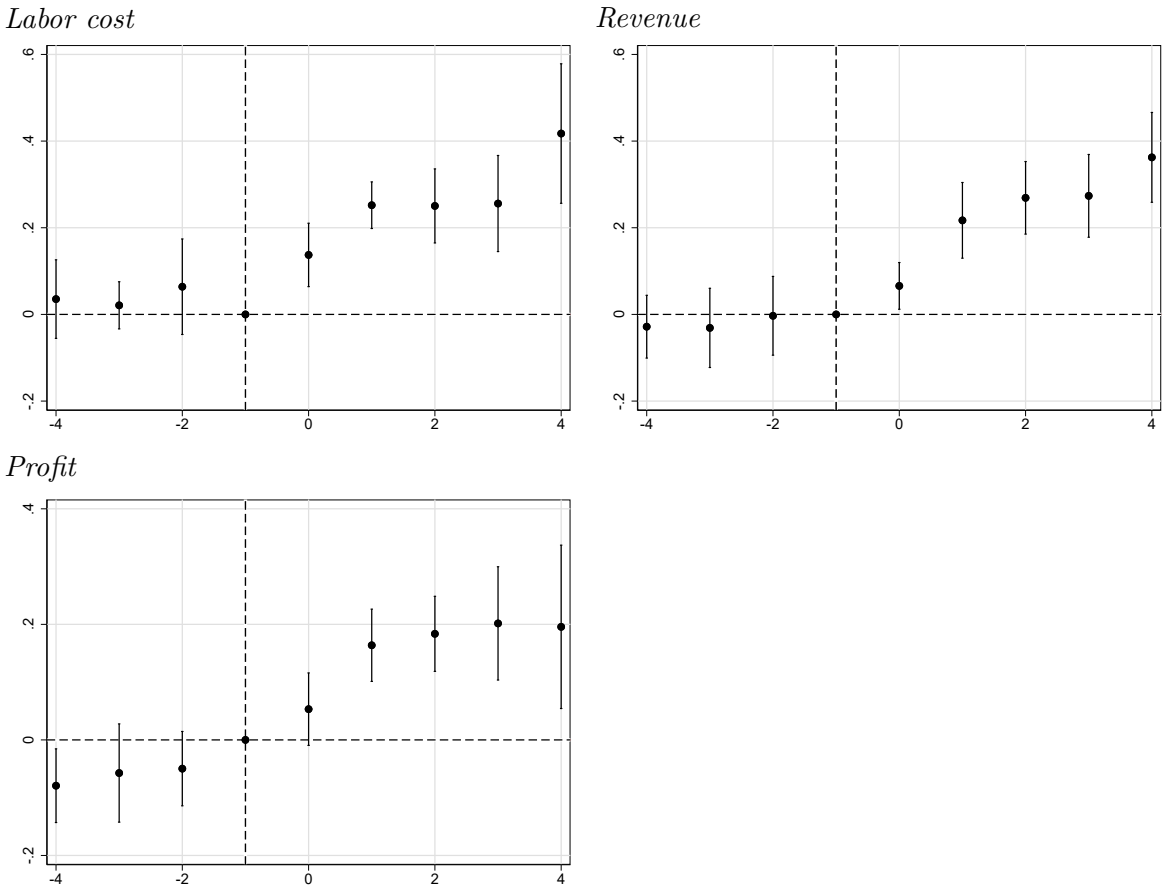
Our baseline specification links labor and operations outcomes to DAT adoption:

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

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 procure and deploy certain DATs. 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 technologies "integral to artifical intelligence and machine learning, [...], drones, autonomous systems, or advanced robotics."[2] 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.[3]

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.

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

[2]See https://www.federalregister.gov/documents/2019/05/17/2019-10538/securing-the-information-and-communications-technology-and-services-supply-chain.

[3]The import ban specifically targeted equipment categories related to advanced ICT associated with DAT projects. We find no evidence that it affected other unrelated types of equipment procurement.

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 × DAT	0.2414***	0.2725***	0.2038***	0.2370***	0.2222***	0.1316***
	(0.03696)	(0.03618)	(0.03095)	(0.03860)	(0.03803)	(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. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

Table 4: Results: Main Outcomes, Staggered DiD

	(1) Labor Cost	(2) Revenue	(3) Profit
ATT	0.1976**	0.2755***	0.1428**
	(0.07868)	(0.07440)	(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. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

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%.

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, 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,

Table 5: Results: Import Ban and Project Completion

	(1) Go-Live	(2) Phase
Post × Affected	0.1068***	0.2575***
	(0.01238)	(0.04441)
US	0.0936	0.2263
	(0.09261)	(0.33207)
Post × US	0.0761	0.5768*
	(0.08436)	(0.30250)
Affected × US	0.0804	0.1371
	(0.08669)	(0.31084)
Post × Affected × US	-0.2189**	-0.7994**
	(0.08688)	(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. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

Table 6: Results: ICT Important Ban as IV

	(1) Revenue	(2) Profit
Post × DAT	0.3392***	0.2717***
	(0.05253)	(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. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

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. 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.

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

Table 7: Mechanism results: Labor

	Labor Mechanisms				
	(1) Cost/h	(2) Regular hs	(3) Overtime hs	(4) Mgmt Share	(5) Temp Share
Post × DAT	0.0037	0.0089	0.1276***	-0.1246**	0.1025*
	(0.05405)	(0.04951)	(0.03448)	(0.05143)	(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. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

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.

Table 8: Mechanisms results: Operations

	Operations Mechanisms		
	(1) Operational Cost	(2) Utilization	(3) Profit Margin
Post × DAT	-0.1293**	0.0511***	0.0816***
	(0.05511)	(0.01379)	(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. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

4.4 Complementarities between hardware and software

We now examine potential 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 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 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

Table 9: Heterogeneity results: Technological complementarities

	(1) Labor Cost	(2) Revenue	(3) Profit
Post DAT	0.2721***	0.3759***	0.3107***
	(0.07366)	(0.07249)	(0.06013)
Post DAT × Only Software	-0.0226	-0.1509*	-0.1546**
	(0.08513)	(0.08281)	(0.06943)
Post DAT × Only Hardware	-0.1085	-0.0449	-0.0530
	(0.11092)	(0.11654)	(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. *Only Software* and *Only Hardware* indicates that a MU has only implemented software or hardware technology, respectively. Standard errors clustered at the MU level in parentheses. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

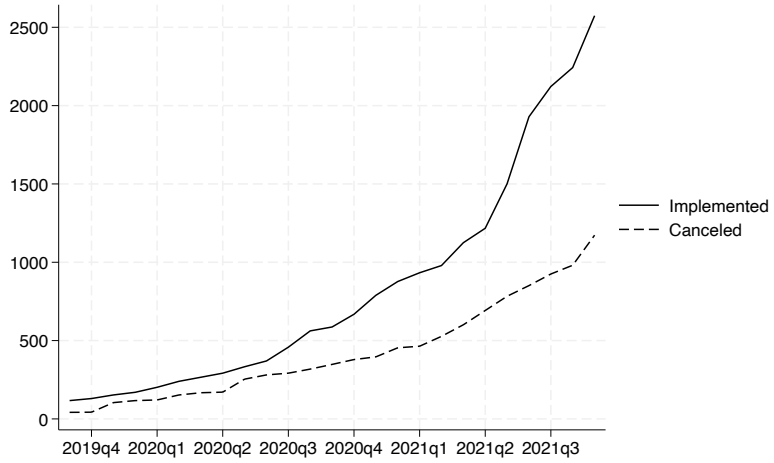
revenues and profits.

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.

4.5 Implementation and cancellation experience

Figure 2 shows the rollout of projects in the organization, differentiating between the number of projects per time period that are canceled after reviewing the opportunities, and projects that are implemented. Both show growth over time, but the implementation rate increases faster than the cancellation rate, suggesting that the management learns to identify opportunities better. Both are important sources of organizational learning: experience improves opportunity identification (screening), reduces implementation friction (rollout speed, integration), and can increases complementarity realization (better pairing of robots and algorithms). As we discuss in detail in the next section, the global logistics company we partner with implemented

Figure 2: Accumulated experience



Note: Number of projects that are implemented or cancelled organized-wide per time period.

Table 10: Heterogeneity results: Implementation experience

	(1) Labor Cost	(2) Revenue	(3) Profit
Post × DAT	0.1046***	0.1565***	0.1599***
	(0.03804)	(0.03703)	(0.03392)
Post × DAT × Implementation Experience	0.0648***	0.0668***	0.0465**
	(0.02288)	(0.02257)	(0.02022)
MU FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Observations	29,407	26,992	30,309

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. *Implementation Experience* measures the lagged (one quarter) cumulative number of implemented DAT projects over the entire organization (standardized). Standard errors clustered at the MU level in parentheses. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

a transparent dashboard where facility managers can at any time learn about positive cases (implementations) as well as negative cases (cancellations of projects). Hence, experience accumulated over the large number pf projects can enable learning — from success as well as from failures.

Regression analysis shows that experience of either kind play a big role in explaining the baseline effects of Table 3. We interact the lagged organized-wide number of implemented projects with the DiD coefficient. Hence, the uncovered relationship is not driven by single MU's outcomes. We further employ time-fixed effects and MU-fixed effects.

Table 11: Heterogeneity results: Cancellation experience

	(1) Labor Cost	(2) Revenue	(3) Profit
Post × DAT	0.0971**	0.1474***	0.1533***
	(0.03906)	(0.03790)	(0.03477)
Post × DAT × Cancellation Experience	0.0724***	0.0762***	0.0532**
	(0.02535)	(0.02492)	(0.02224)
MU FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Observations	29,407	26,992	30,309

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. *Cancellation Experience* measures the lagged (one quarter) cumulative number of canceled DAT projects over the entire organization (standardized). Standard errors clustered at the MU level in parentheses. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

In Table 10 we find that with the total accumulated implementation experience of the organization, the effect of DAT on labor cost increases by 62%, the effect of DAT on revenue increases by 43% and the effect of DAT on profits increases by 29%. Put differently, 38% (labor cost), 29% (revenue) and 22% (profit) of the total effect with full implementation experience (Post × DAT + Post × DAT × Experience) is accounted for by experience.

To estimate the effect of learning from canceled projects, we interact the lagged organized-wide number of canceled projects with the DiD coefficient. Local management can cancel a project that is in screening or planning status rather than implementing it and deploying the planned DAT.

Interestingly, there is even more to learn from implementation failure. With the total accumulated project cancellation experience, Table 11 suggests that the effect of DAT on labor cost increases by 75%, the effect of DAT on revenue increases by 52% and the effect of DAT on profits increases by 35%. Put differently, 43% (labor cost), 34% (revenue) and 26% (profit) of the total effect with full cancellation experience (Post × DAT + Post × DAT × Experience) is accounted for by experience.

4.6 Qualitative evidence

To support the interpretation of quantitative results, we draw on data from 13 semi-structured interviews with employees at various hierarchical levels of a large global logistics company. Interviewees were selected using a two-stage process. First, a random sample of 10 management

units was drawn for each project status category—canceled, no-adoption, and go-live—resulting in an initial list of 30 units. This list was then refined with input from the company's program managers, who helped identify individuals with relevant insights and decision-making authority. The final set of interviewees included members of the global project team, supra-regional and regional leads, and on-site managers, ensuring broad representation across geographies, organizational levels, and functional roles. Interview partners had varying tenures with the company, ranging from under a year to over 20 years, and most had been involved with the automation initiative from its early phases. The interviews were conducted between May and August 2022 via video conferencing, each lasting 30 to 60 minutes. A semi-structured guide with open-ended questions was used to explore experiences with the deployment of automation technologies, perceived challenges, required adjustments, and the organizational changes involved. Interviews were recorded, transcribed, and systematically coded into categories of complementary investments to inform the analysis.

The interview data provides rich context for interpreting the four core findings from the quantitative analysis of DAT adoption. First, the observed gains in operational efficiency, particularly when both robotics and software technologies are deployed together, are consistent with interview data emphasizing that automation cannot be effective in isolation. Managers highlighted that processes must be carefully re-engineered and systems aligned across functions to realize these benefits. One explained that "isolated processes work very good by themselves, however having those processes work together is a challenge", underscoring the need for integrated adoption. These insights support the finding that hardware-software complementarities are critical to performance improvements such as higher revenue and profit margins. Second, while automation is often expected to reduce labor costs, the quantitative data show the opposite—labor costs increase following DAT adoption. Our interviews help explain this counterintuitive result. Deployments frequently led to overtime and greater staffing needs during the transition period, driven by delays and the learning curve associated with tuning new systems. As one manager put it, "it can take weeks or even months until a process has reached the desired productivity levels". In addition to longer hours, automation often introduced new tasks requiring human oversight, recalibration, and exception handling, which further contributed to increased labor demand. Third, the decline in the share of managerial labor documented in the data aligns

with qualitative accounts of how automation reshapes organizational structure. With systems increasingly managing scheduling, routing, and coordination, one interviewee observed that "the vertical structure gained a horizontal component", reflecting a shift toward flatter hierarchies and cross-functional integration. This redistribution of coordination roles away from traditional management helps explain the drop in managerial labor even as overall labor demand rises. Finally, our finding that automation does not displace labor but instead redistributes roles, particularly toward execution and support under more flexible arrangements, is directly supported by interviews. "Automation changes some jobs and creates others", noted one manager, referencing the emergence of roles like robot operators, data specialists, and local key users. The transition often involved reconfiguring existing jobs rather than eliminating them, and the use of temporary or rotating staff became more feasible as processes became more standardized. Importantly, interviewees emphasized that communication and involvement are crucial: "people will support what they understand", reinforcing the idea that organizational buy-in and adaptive human capital practices are key to successful adoption. Together, these insights reveal that the impacts of DAT are not merely technological, but deeply shaped by organizational adaptation, learning processes, and strategic change management.

In our setting, organization-wide learning is actively enabled through a centralized digital infrastructure that makes implementation experiences visible well beyond the local project team. The company operates a intranet platform designed as a single source of truth for the global automation pipeline, covering all focus technologies and regions. The platform tracks how projects move through a standardized set of deployment phases (from potential opportunity and site assessment to planning, deployment, and go-live), is refreshed frequently, and is accessible broadly within the organization. Critically, project information is maintained "as close to the action as possible" by country and cluster leads, which improves timeliness and accuracy while enabling cross-site comparison, rapid escalation, and systematic best-practice sharing. The platform is complemented by an internal communications hub that curates implementation collateral such as assessment tools, videos and case studies from recent deployments, and contact lists that help sites replicate or adapt proven solutions. Regular progress reporting and recurring steering calls at country, regional, and global levels further institutionalize review and reflection. Together, these tools create a mechanism for vicarious learning: sites can learn not only from successful

go-lives but also from projects that stall or are deprioritized, sharpening opportunity selection and accelerating subsequent implementations across the network.

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 productivity, labor, 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 components of DAT.

One key insight from our analysis is that DAT adoption does not uniformly displace human labor. Different from previous technological innovations that required comprehensive system adjustments, we find that DAT adoption often generates new tasks and intensifies existing ones while enabling more flexible deployment of human capital. Our interviews with managers reveal three concurrent explanations for these findings. First, DAT adoption is boosting the efficiency and capacity of individual 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 DAT may be particularly effective at automating or streamlining certain managerial tasks, such as monitoring and coordination. This finding aligns with broader patterns documented in recent research on automation technologies. For instance, Dixon et al. (2021) found similar reductions in middle management following robotics adoption in manufacturing firms, suggesting a consistent pattern across industries where digital technologies streamline managerial oversight functions. Our interviews confirm that DAT improves the collection, processing, and dissemination of operational data, reducing the need for management to serve as information conduits. Similarly, our findings on improved warehouse utilization parallel recent work on predictive analytics adoption (Wu et al., 2020), where enhanced operational efficiency emerged as a key benefit of data-driven decision making. These parallels with existing literature on both robotics and analytics adoption suggest that our findings reflect broader patterns in how digital technologies reshape organizational structure and performance, while extending these insights to show how the integration of both hardware and software components can amplify these benefits.

Our findings make several contributions to the literature. First, we advance the productivity literature by providing evidence on how DATs affect firm performance. While prior work on the productivity paradox suggests significant lags between technology adoption and performance improvements (Brynjolfsson and Hitt, 1996; Brynjolfsson et al., 2019), our findings demonstrate that modern DATs can yield immediate gains in operational efficiency and profitability. Unlike earlier generations of enterprise technologies that caused prolonged implementation dips (Bresnahan et al., 2002; Bloom et al., 2012), we find that automation software and robotics can quickly improve warehouse utilization and expand profit margins. This challenges the conventional J-curve model of technology adoption and suggests that learning algorithms' ability to adapt to existing workflows may accelerate the realization of productivity benefits. Yet, we highlight that the performance and labor outcomes vary substantially with organization-wide experience, both in terms of the number of positive (implementation) as well as the number of negative (cancellation) decisions with respect to prior projects.

Second, our findings contribute to the growing literature on AI and labor (e.g., Choudhury et al., 2020; Yilmaz et al., 2023; Hui et al., 2024; Kim et al., 2024). While recent work has focused on how generative AI affects knowledge workers—including freelancers (Hui et al., 2024), customer support agents (Brynjolfsson et al., 2025), and creative professionals (Noy and Zhang, 2023), we examine a complementary domain: physical robotics combined with AI software. Critically, we find that these technologies do not merely automate manual tasks but also substitute for managerial coordination, as evidenced by the decline in management labor share. This challenges the conventional assumption that robots primarily displace blue-collar work, suggesting instead that algorithmic coordination can replace white-collar supervisory functions. Furthermore, rather than viewing DAT as a cost-reducing tool, our results reframe robotics and AI as levers for capacity expansion.

Third, our results contribute to the literature on organizational design and human capital management in the context of technological change. While previous research has emphasized how new technologies often require substantial organizational restructuring (Barley, 1990; Mata et al., 1995; Mitchell and Zmud, 1999; Puranam et al., 2014; Joseph and Sengul, 2025), our findings suggest that DAT adoption triggers a relatively smaller strategic reconfiguration of human capital, where standardization of operational processes actually increases demand for frontline workers while reducing the need for traditional managerial oversight. Specifically, we show that digital automation-enabled standardization reduces transaction costs in temporary worker deployment while simultaneously shifting managerial attention to more strategic tasks, suggesting new patterns of organizational adaptation to technological change. Finally, a set key theoretical implications arising from our findings include: (1) automation benefits scale as an organizational capability, (2) project cancellations function like disciplined experimentation that improves screening, (3) technological complementarities are important determinants of performance and labor outcomes of DATs.

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.

Finally, our results point to a clear managerial lesson: scaling the returns to automation is an implementation challenge, not just a procurement decision. Managers should invest deliberately in implementation routines that make deployments repeatable, such as playbooks, training, and empowered site champions, and they should build processes that systematically capture and diffuse what teams learn during rollout. Crucially, learning should not be limited to successful go-lives. Canceled projects and stalled deployments contain valuable information about feasibility constraints, hidden complementarities, and local bottlenecks, and organizations can improve opportunity selection by treating these "negative cases" as inputs into future design and screening. Vendor relationships are another asset for scaling: stable partnerships, shared standards, and accumulated interface knowledge can reduce coordination costs and shorten time-to-value across sites. Taken together, our evidence suggests that early automation projects are best managed as a learning portfolio whose purpose is to build organization-wide capability. Sustained gains in revenue and profit arise when firms combine robots with algorithmic optimization and pair these complementarities with accumulated implementation experience, rather than relying on adoption alone.

5.3 Limitations

While this study draws on data from a single multinational logistics firm, the DAT technologies we analyze—such as automated guided vehicles and machine learning-driven scheduling tools—are widely adopted across industries, including manufacturing, retail, and healthcare logistics. This focused approach offers several advantages: our granular data enables particularly convincing empirical tests of DAT's effects on labor demand and profitability, while the homogeneity of operations within the company eliminates many sources of unmeasured heterogeneity that often confound productivity comparisons in broader samples.

Nonetheless, several limitations warrant consideration. First, while the technologies we examine are widely used, their impact may vary with local labor market conditions, cultural differences, or availability of complementary assets like skilled technicians. Second, it is possible that observed shifts in labor demand reflect both technological adoption and strategic decisions by individual management units. While our quasi-random timing of DAT implementation helps identify causal effects, the results capture both direct technological impacts and broader strategic responses to these new capabilities. Additional limitations concern our ability to fully explore distributional consequences within the firm and anticipate general equilibrium effects. However, there remains substantial potential for business value creation through expanded DAT deployment, as only 18 percent of MUs had adopted these technologies by September 2021.

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). "The wrong kind of ai? artificial intelligence and the future of labor demand." *Cambridge Journal of Regions, Economy and Society*, *13*(1), 25–35.

Agrawal, A., Gans, J., and Goldfarb, A. (2019). *The economics of artificial intelligence: an agenda*. University of Chicago Press.

Agrawal, A., Gans, J., and Goldfarb, A. (2022). *Power and prediction: The disruptive economics of artificial intelligence*. Harvard Business Press.

Allen, R., and Choudhury, P. (2022). "Algorithm-augmented work and domain experience: The countervailing forces of ability and aversion." *Organization Science*, *33*(1), 149–169.

Bailey, D. E., Faraj, S., Hinds, P. J., Leonardi, P. M., and von Krogh, G. (2022). "We are all theorists of technology now: A relational perspective on emerging technology and organizing." *Organization Science*, *33*(1), 1–18.

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.

Barley, S. R. (1986). "Technology as an occasion for structuring: Evidence from observations of ct scanners and the social order of radiology departments." *Administrative science quarterly*, 78–108.

Barley, S. R. (1990). "The alignment of technology and structure through roles and networks." *Administrative science quarterly*, 61–103.

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., Li, D., and Raymond, L. (2025). "Generative ai at work." *The Quarterly Journal of Economics*, *140*(2), 889–942.

Brynjolfsson, E., Rock, D., and Syverson, C. (2019). "Artificial intelligence and the modern productivity paradox." *The economics of artificial intelligence: An agenda*, *23*, 23–57.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). "The productivity j-curve: How intangibles complement general purpose technologies." *American Economic Journal: Macroeconomics*, *13*(1), 333–372.

Callaway, B., and Sant'Anna, P. H. (2021). "Difference-in-differences with multiple time periods." *Journal of Econometrics*, *225*(2), 200–230.

Choudhury, P., Starr, E., and Agarwal, R. (2020). "Machine learning and human capital complementarities: Experimental evidence on bias mitigation." *Strategic Management Journal*, *41*(8), 1381–1411.

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.

Dixon, J., Hong, B., and Wu, L. (2021). "The robot revolution: Managerial and employment consequences for firms." *Management Science*.

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*.

Hui, X., Reshef, O., and Zhou, L. (2024). "The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market." *Organization Science*, *35*(6), 1977–1989.

Joseph, J., and Sengul, M. (2025). "Organization design: Current insights and future research directions." *Journal of Management*, *51*(1), 249–308.

Kim, H., Glaeser, E. L., Hillis, A., Kominers, S. D., and Luca, M. (2024). "Decision authority and the returns to algorithms." *Strategic Management Journal*, *45*(4), 619–648.

Leonard-Barton, D. (1988). "Implementation as mutual adaptation of technology and organization." *Managing knowledge assets, creativity and innovation*, 431–447.

Mata, F. J., Fuerst, W. L., and Barney, J. B. (1995). "Information technology and sustained competitive advantage: A resource-based analysis." *MIS quarterly*, 487–505.

Meier, M., Pinto, E., et al. (2020). "Covid-19 supply chain disruptions." *Covid Economics*, *48*(1), 139–170.

Mitchell, V. L., and Zmud, R. W. (1999). "The effects of coupling it and work process strategies in redesign projects." *Organization Science*, *10*(4), 424–438.

Noy, S., and Zhang, W. (2023). "Experimental evidence on the productivity effects of generative artificial intelligence." *Science*, *381*(6654), 187–192.

Pisano, G. P. (1996). "Learning-before-doing in the development of new process technology." *Research Policy*, *25*(7), 1097–1119.

Puranam, P., Alexy, O., and Reitzig, M. (2014). "What's "new" about new forms of organizing?" *Academy of management Review*, *39*(2), 162–180.

Sanchez, R., and Mahoney, J. T. (1996). "Modularity, flexibility, and knowledge management in product and organization design." *Strategic management journal*, *17*(S2), 63–76.

Schilling, M. A. (2000). "Toward a general modular systems theory and its application to interfirm product modularity." *Academy of Management Review*, *25*(2), 312–334.

Tong, S., Jia, N., Luo, X., and Fang, Z. (2021). "The janus face of artificial intelligence feedback: Deployment versus disclosure effects on employee performance." *Strategic Management Journal*, *42*(9), 1600–1631.

Von Krogh, G. (2018). "Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing." *Academy of Management Discoveries*, *4*(4), 404–409.

Wang, W., Gao, G., and Agarwal, R. (2023). "Friend or foe? teaming between artificial intelligence and workers with variation in experience." *Management Science*.

Wu, L., Hitt, L., and Lou, B. (2020). "Data analytics, innovation, and firm productivity." *Management Science*, *66*(5), 2017–2039.

Yilmaz, E. D., Naumovska, I., and Aggarwal, V. A. (2023). "Ai-driven labor substitution: Evidence from google translate and chatgpt." *Available at SSRN*.