AI in High-Tech Manufacturing: Firm-Level Evidence from a Multi-Plant Rollout [*]

Didier Fouarge[†] Marie-Christine Fregin[†] Pelin Özgül[†] Nicholas Rounding[†]

May 4, 2026

Abstract

This paper examines the adoption of AI-based automation and its implications for labor demand using detailed firm-level data from a multinational electronics manufacturer. We study the deployment of an AI visual inspection system across 14 plants over 15 months, tracking decision-level data from quality assurance tasks involving printed circuit board inspection. We document three main findings. First, AI adoption unfolds gradually and incrementally rather than as discrete technological shocks, with plants expanding automation primarily along the extensive margin (deploying AI on new products and component libraries) rather than the intensive margin (increasing AI decision thresholds on existing deployments). Second, we observe substantial heterogeneity in adoption levels and pace across plants, reflecting organizational constraints, data availability, and customer requirements. Third, using two-way fixed effects estimation, we find that a one percentage point increase in the AI automation rate is associated with approximately a 1.2% decline in the growth rate of operator decisions, indicating significant labor displacement within the quality assurance task. We also observe modest improvements in inspection efficiency, with the overall pass rate increasing without compromising quality. Our findings highlight that even when AI systems are technically capable and rigorously validated, deployment proceeds cautiously due to trust, safety, and reliability considerations, suggesting that the labor market effects of AI are shaped not only by technological capability but also by complex organizational dynamics and the gradual building of confidence in algorithmic decision-making.

Keywords: artificial intelligence, technological change, technology adoption, labour demand

JEL Classification: J24, O31, O33

---

[*] This study is part of the research project ai:conomics funded by the German Federal Ministry of Labour and Social Affairs (Bundesministerium für Arbeit und Soziales (BMAS)/Denkfabrik Digitale Arbeitsgesellschaft) by resolution of the German Bundestag.

[†] Research Center for Education and the Labor Market (ROA)

1 Introduction

Recent advances in Artificial Intelligence (AI) have sparked numerous debates on the impact of automation on labour demand and allocation of worker tasks, igniting perennial anxieties about technological unemployment. While AI is predicted to be a General Purpose Technology with profound effects on productivity in labour markets, its employment consequences remain contested (Brynjolfsson et al., 2025a; Hampole et al., n.d.; Humlum and Vestergaard, 2025). Despite a growing body of direct firm-level evidence suggesting that AI tools can improve labour productivity (Fouarge et al., 2025; Brynjolfsson et al., 2025b; Dell'Acqua et al., 2025; Dillon et al., 2025), evidence on the direct labour displacing effects of AI adoption remains limited, leaving critical questions regarding its extent and scope unanswered.

Due to a lack of reliable adoption data, scholars have relied on indirect estimates of the impact of AI through occupational or firm-level exposures. An extensive literature has developed and utilised occupational exposure scores, mapping AI capabilities to occupational task descriptions (Brynjolfsson and Mitchell, 2017; Gmyrek et al., 2025; Webb, 2020; Felten et al., 2021; Frey and Osborne, 2017; Prytkova et al., 2024). While these exposure measures have proven valuable for assessing which occupations and tasks are most susceptible to AI-related changes, they are designed to capture potential exposure rather than actual adoption. By construction, they often assume that technological capability equates to substitution potential, and do not take into account the technological nor economic feasibility of task replacement. As a result, the exposure scores risk overestimating the potential impact of AI adoption on labour displacement (Svanberg et al., 2024). Furthermore other barriers to task displacement, such as trust in AI systems or regulatory constraints, could impede the progress of task displacement. Therefore, while exposure scores provide important insights into which parts of the labour market may be affected by AI, they cannot reveal when, how quickly, or to what extent firms actually adopt these technologies. To understand the true labour displacing effects of AI, we need direct evidence on how firms actually adopt and integrate these technologies into production.

This paper addresses this gap by providing detailed firm-level evidence on AI adoption and its implications for labour displacement in a high-tech manufacturing company. Specifically, we study how an AI-based visual inspection system is adopted and integrated into the quality-assurance (QA) process within a large multinational electronics manufacturer supplying advanced components to the automotive industry. The QA stage is critical for product reliability and safety**, requiring** production workers to inspect printed circuit boards (PCBs) under strict accuracy, consistency, and throughput requirements. High precision requirements, low adjustment costs, and strong organisational incentives for reliability make this setting particularly interesting to study AI adoption and its consequences for task reallocation and labour demand in an industrial environment. Our analysis relies on firm-level data containing information on QA decisions made by either the AI tool or a human operator over a 15-month period from September 2022 to November 2023. This unique dataset allows us to investigate (i) how AI adoption unfolds over time and across plants, (ii) how firms expand AI automation along the extensive and intensive margins, and (iii) how increasing automation relates to changes in labour demand.

We document two findings. First, we find that AI adoption unfolds gradually and continues to increase long after initial deployment. Rather than discrete technological shocks, we observe plants progressively expanding automation along both the extensive margin (deploying AI on additional products and component libraries) and the intensive margin (adjusting decision thresholds and retraining algorithms to increase AI's share of decisions within existing deployments). We split plants into two groups depending on when they began deployment, we refer to those who began 8 months or more prior to the start of the adoption period as early adopters and those who start later as late adopters. Early adopters reach higher automation levels more quickly, while late adopters exhibit faster relative growth once the technology is introduced. Interestingly, expansion occurs primarily along the extensive margin, with plants deploying AI on new products rather than systematically increasing its intensity on already-enabled components.

Second, increasing automation is associated with changes in labour demand within the QA task. As the AI takes on a larger share of the inspection workload, we observe that the volume of operator decisions declines. This shift suggests both task reallocation and a reduction in labour input as automation intensifies. Taken together, these findings demonstrate that despite substantial global investment in AI capabilities, deployment within firms can occur incrementally, shaped by trust, safety, and organizational considerations rather than by discrete technological shocks.

Our study contributes to the growing literature on the labour market effects of automation and AI in several ways. First, while the task-based approach of Acemoglu and Autor (2011) and Acemoglu and Restrepo (2019) has provided the foundation for a rapidly expanding literature on the displacement effects of automation across industries and occupations, there remains limited empirical evidence about how such technologies are actually adopted and integrated into production settings. We build on this empirical task literature by studying the pace and extent of automation under which the AI deployment occurs. In doing so, we highlight how the deployment of AI unfolds over time within a high-stakes industrial setting, reshaping task allocation and labour input at the firm level.

Second, our study relates to the literature on the diffusion of new technologies and advances the empirical measurement of AI adoption at the firm level. Recent work has focused on measuring firm-level AI exposure by identifying AI-related skills in job advertisements (Acemoglu et al., 2022; Peede and Stops, 2024; Gonschor and Storm, 2023; Babina et al., 2024), while other studies have examined early-stage adoption patterns in knowledge work settings (Brynjolfsson et al., 2025a). However, vacancy postings represent hiring intentions rather than actual technology implementation and cannot capture the timing, intensity, or evolution of deployment once adoption occurs. Our study overcomes these limitations by directly observing AI deployment decisions as they unfold, allowing us to document not only whether a technology is adopted, but how extensively it is used and how its dynamics evolves over time. Moreover, a key insight from the technology diffusion literature is that different innovations follow distinct adoption patterns. Prior automation technologies, such as robotics, typically involved lumpy capital-intensive investment spikes (Bessen et al., 2023; Domini et al., 2021, 2022). These discrete investment episodes often coincided with sharp shifts in labour demand, as firms reorganised production and substituted capital for labour in routine tasks (Acemoglu and Restrepo, 2020; Bessen et al., 2023). Such episodes typically entail substantial adjustment costs—in terms of capital and labour—arising from new equipment installation, retraining, and task reallocation. In contrast, we show that the diffusion of AI-based technologies follows a different pattern. Because AI tools are frequently integrated into existing production processes without large fixed investments or major reconfiguration of production lines, adoption can proceed sequentially and at relatively low adjustment cost (Asphjell et al., 2014). In our setting, we observe this gradual expansion of automation along both the extensive margin (deploying AI on new products and component libraries) and the intensive margin (adjusting thresholds on deployed models and retraining algorithms). This sequential nature of adoption generates considerable heterogeneity across plants, reflecting differences in technological readiness, data availability, and organisational practices.

Finally, we highlight the central role of trust, safety, and reliability in the adoption of AI-based systems (Daly et al., 2025). Even when AI tools are technically capable of taking over tasks, firms are often reluctant to delegate decisions without sustained performance validation. In our setting—where quality assurance is mission-critical—we show how adoption proceeds cautiously, with AI initially used in decision support roles before gradually taking on autonomous tasks. This underscores the importance of organisational dynamics and confidence in shaping the trajectory of AI deployment, and suggests that even disruptive technologies may diffuse in an incremental and human-centric manner when deployed in high-stakes production environments.

The rest of paper is structured as follows. In Section 2 we provide the company background and describe the AI adoption process within the firm. Section 3 outlines the data overview and presents descriptive statistics. Section 4 presents our main findings. We first analyze how AI deployment evolves across plants, capturing both the extensive margin (deploying AI on additional products and component libraries) and the intensive margin (adjusting decision thresholds and retraining algorithms to increase AI's share of decisions within existing deployments). We then describe our empirical strategy and present results on the relationship between AI adoption and changes in labour demand. Section 5 provides a range of robustness checks and sensitivity analyses. Section 6 concludes.

2 Empirical Context

2.1 Company Setting

We evaluate the the effects of AI adoption in a large a global international electronics manufacturer who develop, supply and produce technologically advanced products and services for original equipment manufacturers in the automobile industry. The company is one of the largest in its market, with a global reach, around 100,000 employees, and sales figures in the tens of billions of euros. Despite this, competitive pressures are high, both from traditional competitors and new comers into the market. This pressure demands the company invest in continuous innovation and process improvement to maintain their market position. The successful adoption of AI is viewed by the company as crucial step in forward planning.

Specifically, we evaluate the adoption of AI in a particular task within this company. We focus on a QA task that involves inspecting soldering points on Printed Circuit Boards (PCBs) to detect anomalies. Performed by quality control workers, this task is particularly demanding, requiring sustained visual attention and precision to identify minor defects at high production volumes. In what follows, we provide details on the company's production process, the nature of the QA task before AI, and the subsequent introduction and deployment of AI technologies that automate parts of this process.

2.2 Printed Circuit Board Assembly

The company assembles PCBs using surface-mount technology (SMT), a highly automated and technologically advanced production method with widespread use in PCB assembly. An exemplary representation of a circuit board and the SMT production environment is shown in Figure 1. In SMT production lines automated pick-and-place machines put electronic components—such as capacitors, diodes, resistors and transistors—onto circuit boards that have had solder paste painted on at specific solder points. These boards are then passed through a reflow oven, melting the solder and permanently connecting the components to the board. For a more detailed overview of the process, see Prasad (2012).

Following the reflow soldering process, the assembled PCBs undergo visual inspection to detect manufacturing defects, such as faulty soldering points or misaligned components. Because PCBs serve as the central nervous system of all electronic devices, their precise and flawless performance is essential. Even minor defects, such as micro-cracks in solder joints or insufficient soldering, can lead to functional failures. In many of the company's end products, which include safety- and mission-critical applications (e.g., automotive safety and control systems), such failures can have severe consequences. As the SMT production process allows for smaller component sizes and higher circuit densities than alternative assembly processes (such as through-hole assembly), there is an increased risk of defective units. As a result, the QA process is a critical stage in the production process, crucial for identifying even minor deviations from manufacturing standards (Acciani et al., 2006). This ensures product reliability, operational safety, and compliance with stringent industry regulations and customer demands. In our case, customers impose particularly high accuracy requirements and detailed inspection protocols, ultimately shaping how extensively the AI system can be utilised in the QA task.

2.3 QA before AI

Prior to the introduction of AI, the visual inspection task of PCBs was carried out by human quality control operators and a rule-based Automated Optical Inspection (AOI) system. AOI systems capture high-resolution, multi-angle images of PCBs and compare them to a predefined "golden" template (Abd Al Rahman and Mousavi, 2020). If the captured images match exactly, then the component is marked as non-defective (or OK) and continues in production. If there are any detected deviations, components are flagged as defective. Because the rule-based algorithms operate with limited flexibility, they tend to generate a substantial share of false negatives, flagging components as potentially defective even when they conform to production standards.[1] Since defective components must be either scrapped or manually repaired, this high false negative rate would introduce significant waste without further inspection. Human operators therefore play a vital role in improving the efficiency of the process: they review all components flagged by the AOI as defective, reclassifying OK components while catching any truly defective ones that must be removed from production.

Operators are selected from unskilled production line workers called assemblers, whose primary task is to load materials into the SMT production line. Assemblers can apply to become operators, though the position offers no pay premium and provides only the benefit of being seated during the task. The visual inspection task itself is cognitively demanding, requiring sustained visual attention and precision to identify the few defects within a high volume of non-defective images. Of the images seen by operators, 90-95% are eventually classified as OK. Due to these demand, workers must pass a simulation-based assessment at the beginning of each shift before being allowed to do the inspection task. Furthermore, operators are only assigned to this task for a portion of their shift, typically around 2 hours, after which they return to regular assembly duties.

2.4 AI-Automated QA Decisions

The traditional QA workflow, while effective at catching defects, remains inefficient: due to stringent quality requirements, operators are trained to err on the side of caution. This results in a higher scrap rates than necessary. Developments in machine learning technologies created an opportunity for the company to tackle this inefficiency, while simultaneously reducing labour costs without compromising on quality. Based on an algorithm design taken from a Google white paper, the company's AI and data science division developed a supervised learning classification algorithm for the QA task.

The AI system is designed to reduce the volume of images requiring human review by acting as an intermediate classifier between the AOI and human operators. Positioned after the rule-based AOI, the AI reviews components flagged as defective and autonomously reclassifies those it can confidently predict to be OK, allowing them to continue in production without operator review. These are actual share of AI-automated decisions as previously they would have required an operator to classify them as OK or defective. To ensure ongoing quality control, a sample of AI-classified OK components are also sent to operators for verification. Components that the AI flags as defective, or cannot classify with sufficient confidence, are forwarded to operators for further inspection (see Figure 2 for an overview of this process). This changes the nature of the operator's task: rather than reviewing the full set of AOI-flagged components, operators now handle a filtered subset consisting primarily of ambiguous or defective cases, alongside a small subset of OK decisions for monitoring.

2.5 Training and Deployment of the AI

The company's AI and data science division developed the algorithm, training procedures, and deployment pathways at a global level. Actual training and deployment, however, occur at the plant level due to differing production technologies across plants. Therefore, an AI model is trained for each plant using a locally curated test dataset. Each dataset contains samples of regular production data, as well as

---

[1] We define false negatives as components flagged as defective that are actually OK. True negatives are components flagged as defective that are actually defective. True positives are components flagged as OK that are OK, false positives are components flagged as OK that are actually defective. Note that in quality assurance literature, this terminology is often inverted, with 'positive' referring to defect detection rather than non-defective classification, e.g., Abd Al Rahman and Mousavi (2020).

sufficient numbers of images of each known error class. These images are then labeled by certified PCB experts before being used to train the AI. While operator-generated labels can be incorporated into the training process, they are typically included only after being verified by experts, and the extent to which they are used depends on the maturity of the model at the respective plant.

The images used for AI training are grouped based on the component captured. These groupings are referred to as component libraries, where a component library contains physically similar components sharing common visual features (e.g., all resistors of a specific size). Training is entirely supervised and component-specific, with no ongoing learning after deployment. Before deployment, the AI must achieve a false positive rate under 2%, and a false negative rate under 10%. As all components flagged as defective are eventually reviewed by operators, there is higher tolerance for false negatives. The false positive rate offers a substantial improvement over human operators, who are estimated by the company to have rates as high as 10 or 15%. Once the AI has exceeded these thresholds, plants can begin the deployment process.

Deployment occurs through a multi-step, gradual process contingent on the technical readiness of the AI model, customer approval, plant confidence in the technology, and plant organisational readiness. As a result, even highly capable AI models are not deployed uniformly or at full scale. This process proceeds along two margins: extensive and intensive. At the extensive margin, the AI is deployed on new products and component libraries. Once a model is trained and approved for release by the plant, it cannot be used on any product without customer approval. The company must therefore convince the customer that the AI provides a substantial quality improvement over the traditional workflow.

Once the AI has been deployed on a component, plants can adjust the intensive margin through several mechanisms. They may lower certainty thresholds over time, allowing the AI to handle a larger share of decisions. They may retrain models with updated data, improving the AI's confidence on previously marginal cases. Finally, they may reduce the sampling rate for operator verification of AI-approved decisions. These adjustments vary considerably across plants, resulting in substantial heterogeneity in how intensively the AI is utilized. We explore adjustments at either the extensive or intensive margin further down, in Section 4.1.

2.6 Role of human operators after AI deployment

Following AI deployment, human operators continue to play a crucial role in the QA process. They review all components for which no AI model has been trained, handle cases where the AI cannot decide with sufficient confidence, and inspect all components the AI flags as defective. Additionally, operators verify a sample of OK decisions made by the AI, allowing plants to monitor the AI's accuracy over time.

This changing role has significant potential implications for the work of operators. With a fully deployed AI, their task composition changes as they see a higher proportion of defective decisions, and potentially more challenging components to classify. They also become an important mechanism for ensuring the AI continues to operate acceptably. Further, they are exposed to decreasing work demand in the verification task due to task replacement. This could either free up time for them to spend on the production line as assemblers, or reduce the overall demand for assemblers within the plant, potentially leading to job losses.

In the next section we investigate the relationship between AI deployment on the QA process through the analysis of decision data provided to us by the company. We focus on changes in labour demand within the QA task, total decision throughput, and the outcomes of the QA process to elicit insights into the changing nature of the QA task.

3 Data

3.1 Data Overview

The company provided us with access to customized QA decision data compiled specifically for this research project. The data cover 16 plants operating in 8 countries across Asia, the Americas, and Europe, and include the date, time, and result of all QA decisions made by human operators and the AI. The data also include anonymized identifiers for the components inspected, the products on which they appear, and their associated component libraries.

QA decisions are made on images of individual components. Each component belongs to both a specific product (the PCB being inspected) and a component library (a grouping of similar components, as described in Section 2.5). The relationship between these levels is illustrated in Figure 3. Our primary analyses use a week-plant panel, though we also construct a week-component library panel for robustness checks. The relationship between products, component libraries, and components is illustrated in Figure 3. Our primary analyses are conducted using a week-plant panel, and we also construct a panel at the week-component library level for additional analyses.

Our main variable of interest is the **AI Automation Rate. This is defined as the share of QA decisions actually automated by AI relative to the total number of QA decisions. Automated decisions are those that are independently classified as "OK" by the AI without any operator review.**[2] This measure captures the intensity of AI-driven automation in the QA process across plants and over time, reflecting adjustments at both the extensive and intensive margins, and forms the basis for the analyses presented in the subsequent sections. Our focus on "OK" decisions follows directly from how the AI system operates within the QA workflow. As described in Section 2.4, the AI is only allowed to make an autonomous decision when it classifies an image as non-defective. These cases constitute the part of the inspection task that the system can execute independently, without human verification. By contrast, any image the AI flags as defective, or is uncertain about, is routed to the human operator for review. We do not focus on the AI defect rate as the thresholds for the AI to make a decision are set significantly lower, and thus operators overturn up to 75% of AI defect decisions. Measuring the share of AI-approved OK decisions relative to all QA decisions therefore provides a direct indicator of the portion of the QA task that the AI can execute without human review.

Using this automation variable as a starting point, we introduce several measures in the following sections that allow us to (i) characterise how AI deployment unfolds over time and across plants, and (ii) examine how changes in automation intensity relate to labour demand within the QA task.

3.2 Sample Selection and Summary Statistics

We implement several sample restrictions to ensure data quality and consistency. First, we exclude two plants that contained no information on operator decisions, leaving 14 plants for analysis. Second, we address a data reporting issue that emerged toward the end of our observation period: across multiple plants, operator decision data were not recorded despite ongoing operator activity. To mitigate potential measurement error from this reporting gap, we truncate the sample, removing all observations from the affected period. This restriction eliminates 27.7% of plant-week observations, corresponding to 22.4% of total decisions. Our final estimation sample comprises 14 plants observed over 60 weeks, from September 2022 through November 2023. This approach helps us to account for a systematic data collection problem across plants, however, there remain persistent issues with the data throughout our observation period relating to missing product and decision information.

---

[2] Section 2.4 provides a detailed description of the AI system and the automated decision making process.

Within our final sample, per week on average 9% of decisions are missing product and component identifiers. Thus we cannot assign them into component libraries or identify the products they were made on. Missing product information occurs across all plants, throughout the observation period. Decisions with products missing have a 21% lower AI Automation Rate than non-missing, and are 5% less likely to be passing decisions. For our primary analysis, we remove those decisions with missing product information. This is likely to overestimate the AI Automation Rate, therefore we run robustness checks to test the sensitivity of our results to missing product information by including those decisions in our final total. We also run sensitivity analyses using our variable construction, but excluding weeks with high levels of missing product information. A further data issue arises from lines in the data missing decisions information, but containing product identifiers. On average, around 6% of data rows per week contain no decision information but have an assigned product identifier. To account for this, we identify weeks with a high level of missing product information and in a sensitivity analysis remove them from our sample.

Table 1 summarizes key characteristics of these variables across 14 plants and 60 weeks during the observation period in the plant-week panel. In Panel A, we observe that most plants have the full set of 60 week observations. However two plants have missing data resulting in only 53 and 57 week observations, thus averaging 59 week observations for plants. Plants make decisions across an average portfolio of 902 different products and 1,118 component libraries during the observation period. Plant size varies substantially: the smallest plant accounts for only under 0.2% of total QA decisions whereas the largest contributes to 42%. This heterogeneity in plant sizes presents challenges for our empirically strategy. On the one hand, we want to avoid results being driven disproportionately by the largest plants; on the other hand, we want to prevent very small plants having an outsized influence. To mitigate these concerns, we express outcomes in logs where appropriate and report both unweighted and weighted regression results.

Table 1 Panel B reports on the average AI Automation Rates and QA Outcomes. We find that the lowest plant has an average weekly AI Automation Rate of only 1.8%, whilst the highest plant has 67.9%. The average AI Automation Rate across plants is 23.1%, thus on across plants roughly 75% of decisions are taken by human operators. For all QA decisions, irregardless of whether they have been taken by the AI or a human operator, we find a high average OK rate of 95.3%, with some variation between plants ranging from 86.7% to 99.8%. The Operator OK Rate is slightly lower, at 93.8%. This discrepancy suggests that the AI influences the OK rate at the plant level. Operators could be more likely than the AI to flag a component as defective,alternatively this could reflect changes in task reallocation with Operators handling a higher share of defective images as the AI automates only OK images, we explore the this in Section 4.4.

To assess the impact of AI Automation, we rely on weekly first differences, therefore in Table 1 Panel C, we report on week-to-week changes in QA decisions. We first report on our explanatory variable, AI Automation Rate. We see an average weekly growth in the AI Automation Rate of 0.164 percentage points across plants, across our observation window. Turning to Total Throughput, defined as the total number of decisions taken by operators and AI, in general we observe small, positive average weekly changes of about 0.8%, suggesting that the plants increase there throughput during our observation window. Turning to Total Operator Decisions, we similarly see a slight increase of 0.6%. Whilst similar in magnitude, albeit slightly slower, providing an indication that there is substitution between labour and capital within this task. Measuring the increase in AI decisions, we see that the Total AI Decisions increases by about 3% per week, a higher growth rate than either Total QA Throughput or Total Operator Decisions. The above descriptive evidence suggests that the increase in decisions being assigned to AI is outstripping increases in total throughput, leading to a widening gap between productivity increases and labour demand.

4 Results

In this section, we first document how AI deployment unfolds across plants during our observation period, highlighting differences between early and late adopters and the relative contributions of the extensive and intensive margins. We then turn to labour demand, where we assess how week-to-week changes in the intensity of AI automation relate to changes in operator activity at the plant level and the empirical strategy guiding our analysis. Finally, we examine two additional margins of adjustment, overall QA volume and QA accuracy, to understand whether AI adoption affects productivity or the quality of the inspection process.

4.1 Mapping AI Deployment

We start by exploring the progression of AI deployment across plants during our observation period using the AI Automation Rate. We have precise start dates for all plants that began deployment within 15 months before the start of our observation period, with the remainder deploying earlier. Of the fourteen plants in our sample, nine had already begun AI deployment prior to the start of the observation period, while five deployed during our observation window. Among the nine that deployed before our sample, seven began more than 8 months prior, while two began 1-4 months before. The five plants that deployed during the observation period began at various points in time: three within the first 3 months and two between 8-11 months after the sample start. Table 2 provides the full distribution of deployment timing. To capture heterogeneity in deployment phases, we classify plants as "early adopters" if they began deployment more than 8 months before our sample period (N=7), and as "late adopters" if they began deployment 8 months or less before the sample start or during the observation window (N=7).

Figure 4 plots the AI Automation Rate averaged for all plants and split between early and late adopters. Figure 4a shows that the overall AI Automation Rate across plants remains relatively low, starting at around 16% at the beginning of the observation period and reaching 27% by the end. Between early and late adopters, we observe a persistent difference in levels, with early adopters maintaining higher automation rates throughout our observation period. Figure 4b compares the trajectories of early and late adopters by indexing to Week 1, revealing that both groups exhibit similar patterns of increase over time. This suggests that long after initial deployment, plants continue to gradually increase automation rather than reaching a stable steady state. This indicates that AI implementation is an ongoing process rather than a discrete event. Figure A.1a in the appendix shows individual plant trajectories, confirming that the group means are representative despite variation across plants within each adopter group.

What drives this sustained automation growth? In Section 2.5 we distinguished two mechanisms: at the extensive margin, automation can increase as AI is progressively enabled on additional products or a newly trained algorithm is released on an additional component library. At the intensive margin, automation can rise as the thresholds at which the AI can make an autonomous decision are adjusted, or a retrained algorithm with additional images is deployed.

To examine the extensive margin of automation, Figure 5 plots the share of total QA decisions made on components where the AI has been deployed, which we refer to as Potential AI Automation Rate. Unlike the AI Automation Rate, this measure contains both automated decisions and those taken by operators. Thus, it captures the expansion of AI deployment onto new products and new components but not necessarily the expansion of AI Automation. Figure 5a shows that the share rises from 30% to 70% across plants, through our observation period. Again, we observe a levels difference between early and late adopters, with a higher extensive margin for early adopters throughout. However, moving to Figure 5b, we begin to see differing patterns across the two groups. When indexed to the first week's value, we see that late adopters increased the extensive margin by 40 percentage points compared to just 20 percentage points for early adopters. Comparing the automation rate to expansions at the extensive margin, we see that the increase at the extensive margin is higher than the overall increase in the automation rate. This suggests that plants deploy the AI at a relatively low intensity.

To understand the intensity of AI use, we investigate the AI Automation Rate on components where the AI has been deployed (i.e., the intensive margin). The average AI Automation Rate once the AI has been deployed on a component is around 36% for the entire period, showing that the AI only automates around one third of decisions leaving a substantial share for operators to review even after deployment Splitting into early and late adopters, we find a rate of 44% for early adopters and 28% for late adopters suggesting that plants who began using the AI earlier use it more intensely at both the extensive and intensive margins. However, even after reaching maturity, plants are still cautious about how much to automate requiring human oversight on over half of the AI's decisions.

This pattern is illustrated in Figure 6, which plots the AI Automation Rate, removing all decisions made on components where the AI is not deployed. Figure 6a reflects the averages described above; we find that on average there is a lower *intensive* usage of the AI by early adopters compared to the late adopters. If we move to the indexed change, in Figure 6b we see that for both early and late adopters the adjustments at the intensive margin appear to be more random and show no systematic growth or decline over our observation period.

Taken together, these figures help us to understand the average progression of AI deployment. In plants, the AI Automation Rate increases, both across early and late adopters. These changes appear to be largely driven by the extensive margin, that is the deployment of AI on new products or component libraries as we are unable to see a systematic rise in the intensive margin over our observation period.

4.2 Labour Demand

4.2.1 Empirical Approach

We now turn to examining the relationship between AI adoption and labour demand within the QA task. Two factors impede attempts at causal analysis: (1) deployment timing of the AI at product or component level is endogenous, plants are able to anticipate deployment and adjust their labour demand occurring; (2) products contain both treated and untreated components and operators work across lines with both treated and untreated components, creating spillovers that violate the Stable Unit Treatment Value Assumption (SUTVA) (see Appendix 2 for challenges to causal identification).

We therefore aggregate to the plant-level and rely on a Two-Ways Fixed Effect (TWFE) estimation in a first differenced framework to control for unobserved, time-invariant plant characteristics and estimate how week-to-week changes in the intensity of automation relate to changes in the demand for human labour. This allows us to capture the relationship between AI and labour demand throughout the plant, including any potential spillover effects on products or components where the AI had not previously been enabled.

Figure 7 plots the weekly change in AI Automation Rate against the weekly change in the total operator decisions at the plant level, indicating a negative relationship. To quantify this, we estimate the following model:

$$Y_{i,t} = \alpha + \beta \cdot \Delta \text{AI Automation Rate}_{i,t} + \gamma_t + \theta_i + \varepsilon_{i,t} \tag{1}$$

Here, $Y_{i,t}$ denotes the weekly change in the outcome of interest in plant i and week t e.g Δ Total Operator Decisions (Log). Our main explanatory variable, ΔAI Automation$_{i,t}$, reflects the weekly change in the AI Automation Rate at the plant level. We include week fixed effects γ_t to control for common time shocks affecting all plants, such as the output growth. We also include plant fixed effects, θ_i, to account for plant specific growth trends. We cluster standard errors at the plant level. To account for the heterogeneity in plant size, we run both unweighted and weighted regressions, weighting for the total number of decisions recorded in the plant in that week.

4.2.2 Labour Demand Outcomes

We first explore the relationship between an increase in the AI Automation rate and labour demand, as proxied by the growth rate of total operator decisions (i.e. Δ Total Operator Decisions (Log)). A limitation with this approach is that we rely on the number of decisions taken by operators as a proxy for total labour demand due to a lack of data on specific operators, including total employment numbers, working hours, and time spent per decision. It could be that the total demand for operators stays the same as they are able to spend more time on the decisions that remain after automation. However, discussions with the company suggest that such adjustments are unlikely to occur at scale and that the technology is labour-displacing. Therefore, we believe that our proxy is a good approximation of labour demand, while acknowledging that it may represent a theoretical upper bound on the true displacement effect.

Moving to the results, as reported in Table 3 we find a statistically significant and economically meaningful negative relationship between weekly changes in the AI Automation Rate and the growth rate of operator decisions. The estimated coefficients, ranging from -1.20 in the unweighted specification to -1.16 in the weighted specification, imply that a 1 percentage point increase in the Automation Rate is associated with a 1.17%–1.2% lower growth rate in operator decisions. These results are consistent with the hypothesis that, as the AI is granted increasing autonomy in making QA decisions, the role of human operators diminishes accordingly.

A potential concern is that plants may anticipate the introduction of AI systems and adjust their production numbers in the weeks prior to a deployment. As demonstrated in Appendix 2, there is endogeneity in deployment timing at the component library level. This could be replicated at the plant level, with reduced productivity in the weeks leading up to a large AI deployment. Thus, we may not be capturing the entirety of task reallocation. A related concern is that plants could use medium- and long-term forecast to determine AI deployment patterns, thus increases in the Automation Rate may be correlated with anticipated shifts in customer demand.

We are unable to test for medium and long term endogeneity, however we can account for the possibility of short-run endogeneity. To achieve this, we re-estimate Equation 1 by adding one- and two- week leads of Δ AI Automation Rate. In other words, we test whether the growth in operator decisions in the current week is correlated with future growth in AI Automation Rate. The results in Table 3 show no evidence of such anticipation effects: future AI deployment is not systematically related to current operator decisions. Instead, the observed decline in operator decisions coincides with contemporaneous increases in AI decision shares, suggesting that the detected relationship maps to increases in AI deployment.

4.3 Total QA Volume (i.e. Productivity)

We next explore whether changes in AI adoption affect overall plant productivity. We proxy productivity with the growth rate of total QA decisions, estimating Equation (1) using Δ Total QA Decisions (Log) as the outcome variable. A limitation with this approach is that total QA decisions may not directly reflect productivity as the volume of total decisions can reflect unobserved characteristics of the AOI System, limiting the number of images seen by the AI or the human operator.

Panel B of Table 3 reports on the estimation results. The baseline specification shows no statistically significant relationship between changes in AI Automation Rate and total throughput in either the unweighted or weighted specification. This suggests that increasing AI automation neither expands nor contracts the overall volume of QA activity at the week- plant level.

We again test for anticipation of AI deployment by including one- and two-week leads of Δ AI Automation Rate. The results provide no consistent evidence of anticipation effects. While the two-week lead shows marginal significance in the unweighted specification, this relationship disappears in the weighted specification and the one-week lead is insignificant in both. The weighted specification also yields a negative coefficient on the current Δ AI Automation Rate, though this pattern is not robust to the unweighted regression. Overall, we find no systematic relationship between AI deployment and total QA volume, suggesting that automation is associated with reallocation from operators to AI, rather than leading to increase in the number of QA decisions taken.

4.4 Total and Operator OK Rate

Next, we turn to the relationship between AI Automation and the outcomes of the QA process. While we are unable to detect a relationship between AI automation and total QA volume, productivity could be affected by making the QA process more efficient. To test this relationship, we estimate 1 using Δ OK Rate as the outcome variable. Panel C of 3 reports on the estimation results for both the Overall OK Rate and Operator OK Rate. We are able to detect a small but significant increase in the Total OK Rate of 0.08 percentage points in both the weighted and unweighted specification whilst we are unable to detect any change in the Operator OK Rate.

Interpreting these results, the effect of the increase in the Total OK Rate is somewhat ambiguous. Our data does not allow us to identify a ground truth for the components, therefore an increasing OK Rate could either reflect an increase in true positives or false positives. However, several pieces of evidence help us infer the most likely interpretation. First, operators very rarely overturn OK decisions made by the AI. Using operators as the ground truth, we detect a false positive rate of under 1%, indicating that the AI consistently and accurately detects OK components. Second, interviews with the company revealed that in their own internal assessments AI has a better false positive rate than operators. Thus, this leads us to conclude that the detected increase in the Total OK Rate does not represent an increase in defective components being flagged OK (i.e. false positives), but rather a decrease in OK components flagged as defective (i.e. false negatives). A decrease in false negatives increases the efficiency of the QA process, as fewer products are scrapped due to incorrect flagging.

5 Robustness Checks/ Sensitivity Analysis

Although our dataset is rich, allowing us to assess the relationship between AI Automation and our outcomes, it also contains features that motivate a number of robustness checks (as described in Section 3.1). First, certain plants contribute disproportionately to the sample. In the weighted regressions, the effects may be driven by the plants with many decisions, whilst the unweighted regressions may be effected by plants with low numbers of decisions. Compounding this, some plant-week observations record extreme swings in the AI Automation Rate, as can be seen in Figure 7, highlighting both swings of ± 30 percentage points. Second, there are data missing product and decision information throughout. Third, our main estimation approach is conducted at the plant level. This allows us to capture any potential spillover effects, however, detected relationships may be driven by compositional effects rather than changes in AI Automation. To address these concerns, we conduct a series of robustness checks to assess the validity of the results presented in the previous subsection.

5.1 Leaving-One-Plant-Out

First, to ensure that our findings are not driven by any particular plants, we re-estimate Equation (1) leaving out one plant at a time. This allows us to assess whether the estimated relationship between AI adoption and outcomes is not driven by one particular plant. Figure 8 displays the results from this leave-one-plant-out exercise. The red marker corresponds to the baseline estimate using the full sample, while the grey markers show the coefficients obtained when excluding each plant in turn. We find limited variation in the estimated coefficients across plants. Importantly, no single plant seem to drive the overall relationship between changes in AI Automation and the growth rates of our outcomes. This reinforces the interpretation that the estimated effect reflects a general pattern across plants rather than being specific to any individual production unit.

5.2 Sample Restriction

Second, we turn to the effects of our sample selection and issues relating to missing data. Table 4 reports on the results from various sample selections, using the weighted specification from Equation (1). First, in Column 1, we present the results from our preferred specification as shown in Table 3. In Column 2, we remove the largest positive and negative weekly swings of AI Automation. While the relationship with labour demand remains consistent,the effects on OK rate shift. We observe that the relationship with the overall pass rate is no longer statistically significant and the coefficient of the Operator OK Rate becomes negative, this is consistent with findings at the Component Library level (see Section 5.3). In Column 3, we remove weeks where the number of lines missing decisions is at least 25% as large as the total number of decisions. In Column 4, we remove weeks where the number of lines missing product information is at least 25% as large as the total number of decisions. Column 5 excludes weeks meeting both criteria. The last three columns demonstrate that the identified relationships between our outcomes and AI Automation are robust to weeks with missing information.

5.3 Component Library Analysis

To address the concerns relating to compositional effects driving our detected relationships we re-estimate Equation (1) at the component library level. Just under 80% of component libraries record under 200 decisions per week, yet this only corresponds to 8% of decisions with component library information. To reduce the effect of small component libraries at the week level, we construct a monthly panel at the component library level. To reduce the effect of shifts in our AI Automation Rate being driven by small library sizes, we remove all component library-month observations where we detect less than 200 decisions corresponding to 2.5% of total decisions.

Table 5 reports the results using our final estimation sample, applying the library size restrictions listed above. The results at the component library level are broadly consistent with our plant-level findings. In Panel A, we observe a statistically significant negative relationship between changes in AI Automation Rate and operator decisions, with coefficients of approximately -2.07 in both weighted and unweighted specifications. This is somewhat larger in magnitude than the plant-level estimates, suggesting that the negative effects of automation at the component level are offset by libraries with lower levels of AI Automation. When we include anticipation leads, the primary estimation remains negative and significant. However, we detect a relationship with future AI deployment predicting current labour demand further representing the endogeneity of AI deployment timing at the component library level (see Appendix 2.

In Panel B, examining total QA decisions, we find a small positive relationship at the component library level, contrasting with the null result at the plant level. This suggests that at the component level, AI deployment may be associated with modest throughput increases, though these effects do not aggregate to significant plant-level productivity gains. Again, this may reflect the fact that plants are aware of AI deployments and adjust production of a particular product centered around an AI deployment. Turning to the anticipation effect, we see positive correlations between current throughput and concurrent changes in the AI Automation Rate, alongside the 1 and 2 month leads. Turning to the weighted specification, this relationship disappears. As with Table 3 and **????????**

Lastly, in Panel C, we observe that the Total OK Rate increases by approximately 0.05 percentage points per 1 percentage point increase in AI Automation Rate, consistent with our plant-level findings. However, we now observe a small negative relationship with the Operator OK Rate, in contrast to our results at the plant level. This is potentially an intuitive result, overall we are unable to detect an effect on the composition at the plant level, however, the AI takes a share of the OK decisions leaving operators with a pool of decisions with a higher defect rate. This indicates that as AI Automation progresses, the operator task may change. However, this association also must be viewed with caution as we are unable to test whether operators change their behaviour following an AI deployment, nor whether this is affected by spillovers across products. Taken together, these component library-level results reinforce our main findings while highlighting that the plant-level estimates capture both direct task substitution and compositional shifts across products and components.

6 Conclusion

This paper investigates the adoption of AI and its implications for labour demand within a high-tech manufacturing setting. Using detailed decision-level data from a multinational electronics manufacturer, we investigate the deployment of an AI system in a visual inspection task across 14 plants over 15 months. Our analysis reveals some key findings that provide insight into how firms adopt and integrate AI, raising questions about conventional narratives around rapid AI-driven displacement.

First, we document that AI adoption unfolds gradually and incrementally, even long after initial deployment. Rather than one-time technological shocks, plants expand automation progressively along both the extensive margin, by deploying AI on new products and component libraries, and the intensive margin, by adjusting decision thresholds and retraining the algorithm. Notably, growth occurs primarily through the extensive margin, with plants adding AI to new components rather than systematically intensifying its use on already-enabled products. This pattern holds across both early and late adopters, though we observe substantial heterogeneity in adoption levels.

Second, we document a significant negative relationship between increases in AI automation and labour demand within the quality assurance task. A one percentage point increase in the AI Automation Rate is associated with approximately a 1.2% decline in the growth rate of operator decisions. Importantly, we find no evidence that plants anticipate AI deployment and adjust labour demand accordingly; rather, substitution occurs once the AI is deployed into the decision-making process. We also observe modest improvements in inspection efficiency, with the overall pass rate increasing slightly as AI automation expands, suggesting fewer false negatives without compromising quality.

What do these findings imply? Even in a setting where the technology is rigorously tested, meets strict performance thresholds and is proven capable of autonomous decision-making, adoption does not immediately scale to full potential. Despite the AI system's technical readiness and demonstrated reliability, plants deploy AI cautiously. While our quantitative data cannot fully capture all underlying mechanisms that drive adoption decisions, qualitative insights from the firm highlight several key constraints. Responsibility for product quality is decentralized at the plant level, and plant managers are accountable for quality outcomes. As a result, adopting or expanding AI is not perceived as a default technological upgrade, but as part of a managerial decision that must be justified. Plants differ in how they evaluate whether AI deployment improves efficiency, reduces waste, or enhances performance relative to existing processes. In settings where managers perceive current quality outcomes as satisfactory, the incentive to expand AI adoption may be weaker, even when the technology is readily available. These governance structures, together with trust in algorithmic decision-making, regulatory and liability concerns, and customer-specific quality requirements act as institutional guardrails. These organizational constraints shape not only whether AI is adopted, but also the pace, scope, and extent of its deployment within production units.

These dynamics have broader implications for understanding AI diffusion in manufacturing environments where reliability, safety, and quality are paramount. Our findings suggest that employers do not simply deploy automating AI systems and displace labour. Instead, our results suggest that such systems requires ongoing human oversight, trust building, and sustained organizational commitment. In high-stakes production settings the path from initial deployment to full-scale automation is potentially neither automatic nor immediate. It is mediated by human judgment, institutional safeguards, and the gradual building of confidence in machine performance.

Ultimately, this study underscores that the labour market effects of AI are not determined solely by technological capability but by the complex interplay of technical performance, organizational practices, and trust. While AI has the potential to reshape task allocation and reduce labour demand, the timeline and extent of these changes depend both on the technical capability of technologies and how firms manage the transition to algorithmic decision-making in high-stakes settings. As AI continues to diffuse across industries, understanding these adoption dynamics will be essential for anticipating its broader economic and social consequences.

References

Acciani, G., G. Brunetti, and G. Fornarelli, "Application of neural networks in optical inspection and classification of solder joints in surface mount technology," *IEEE Transactions on Industrial Informatics*, August 2006, *2* (3), 200–209.

Acemoglu, Daron and David Autor, "Skills, Tasks and Technologies: Implications for Employment and Earnings," *Handbook of Labor Economics*, January 2011, *4* (PART B), 1043–1171. Publisher: Elsevier.

_ and Pascual Restrepo, "Automation and new tasks: How technology displaces and reinstates labor," *Journal of economic perspectives*, 2019, *33* (2), 3–30.

_ and _ , "Robots and jobs: Evidence from US labor markets," *Journal of political economy*, 2020, *128* (6), 2188–2244.

_ , David Autor, Jonathon Hazell, and Pascual Restrepo, "Artificial intelligence and jobs: Evidence from online vacancies," *Journal of Labor Economics*, 2022, *40* (S1), S293–S340.

Asphjell, Magne K, Wilko Letterie, Øivind A Nilsen, and Gerard A Pfann, "Sequentiality versus simultaneity: Interrelated factor demand," *Review of Economics and Statistics*, 2014, *96* (5), 986–998.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, "Artificial intelligence, firm growth, and product innovation," *Journal of Financial Economics*, 2024, *151*, 103745.

Bessen, James, Maarten Goos, Anna Salomons, and Wiljan Van den Berge, "Automatic reaction-what happens to workers at firms that automate?," *The Review of Economics and Statistics*, 2023, (Feb. 6, 2023).

Brynjolfsson, Erik and Mitchell, "What can machine learning do? Workforce implications," *Science*, 2017, *358* (6370).

_ , Bharat Chandar, and Ruyu Chen, "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence," 2025.

_ , Danielle Li, and Lindsey Raymond, "Generative AI at work," *The Quarterly Journal of Economics*, 2025, p. qjae044.

Daly, Sarah J, Anna Wiewiora, and Greg Hearn, "Shifting attitudes and trust in AI: Influences on organizational AI adoption," *Technological Forecasting and Social Change*, 2025, *215*, 124108.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, September 2020, *110* (9), 2964–2996.

Dell'Acqua, Fabrizio, Charles Ayoubi, Hila Lifshitz-Assaf, Raffaella Sadun, Ethan R. Mollick, Lilach Mollick, Yi Han, Jeff Goldman, Hari Nair, Stew Taub, and Karim R. Lakhani, "The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise," March 2025.

Dillon, Eleanor W, Sonia Jaffe, Nicole Immorlica, and Christopher T Stanton, "Shifting Work Patterns with Generative AI," *National Bureau of Economic Research Working Paper Series*, 2025, *No. 33795.*

Domini, Giacomo, Marco Grazzi, Daniele Moschella, and Tania Treibich, "Threats and opportunities in the digital era: Automation spikes and employment dynamics," *Research Policy*, 2021, *50* (7), 104137.

_ , _ , _ , and _ , "For whom the bell tolls: the firm-level effects of automation on wage and gender inequality," *Research Policy*, 2022, *51* (7), 104533.

Felten, Edward, Manav Raj, and Robert Seamans, "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses," *Strategic Management Journal*, December 2021, *42* (12), 2195–2217. Publisher: John Wiley and Sons Ltd.

Fouarge, Didier, Marie-Christine Fregin, Simon Janssen, Mark Levels, Raymond Montizaan, Pelin Özgül, Nicholas Rounding, and Michael Stops, "How AI-Augmented Training Improves Worker Productivity," 2025.

Frey, Carl Benedikt and Michael A. Osborne, "The future of employment: How susceptible are jobs to computerisation?," *Technological Forecasting and Social Change*, January 2017, *114*, 254–280. Publisher: North-Holland.

Gmyrek, Pawel, Janine Berg, Karol Kamiński, Filip Konopczyński, Agnieszka Ładna, Balint Nafradi, Konrad Rosłaniec, Marek Troszyński, and International Labour Organization. Research Department, *Generative AI and jobs: a refined global index of occupational exposure*, Geneva: ILO, 2025.

Gonschor, Myrielle and Eduard Storm, "The Diffusion of Artificial Intelligence: New evidence from German Online Job Vacancy data," 2023.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D W Schmidt, and Bryan Seegmiller, "NBER WORKING PAPER SERIES."

Humlum, Anders and Emilie Vestergaard, "Large Language Models, Small Labor Market Effects," May 2025.

Peede, Lennert and Michael Stops, "Artificial intelligence technologies, skills demand and employment: evidence from linked job ads data," 2024.

Prasad, Ray P., *Surface Mount Technology: Principles and Practice*, Springer Science & Business Media, December 2012. Google-Books-ID: hqDdBgAAQBAJ.

Prytkova, Ekaterina, Fabien Petit, Tommaso Ciarli, Deyu Li, and Sugat Chaturvedi, "The Employment Impact of Emerging Digital Technologies | Publications | CESifo," 2024.

Rahman, M Abd Al and Alireza Mousavi, "A review and analysis of automatic optical inspection and quality monitoring methods in electronics industry," *Ieee Access*, 2020, *8*, 183192–183271.

Svanberg, Maja, Wensu Li, Martin Fleming, Brian Goehring, and Neil Thompson, "Beyond AI Exposure: Which Tasks are Cost-Effective to Automate with Computer Vision?," January 2024.

Webb, Michael, "The Impact of Artificial Intelligence on the Labor Market," Technical Report 2020.

Figures:

(a) Printed circuit boards (PCBs) (b) Surface Mount Technology (SMT) Line

Figure 1: Exemplary representations of the production environment

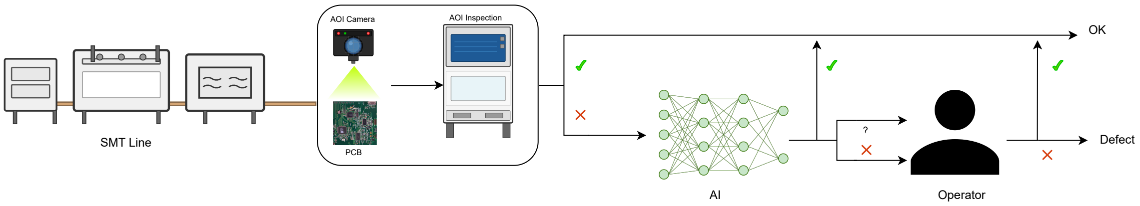

Figure 2: Quality Assurance Process Flow: AOI, AI, and Operator Decision Logic

Notes: This diagram demonstrates the Quality Assurance process flow once the AI has been enabled. The Automated Optical Inspection (AOI) system classifies components as OK or defective, OK decisions are forwarded directly to the back-end. Components classified as defective by AOI are forwarded to the AI system, the AI makes autonomous decisions where it has the required level of confidence (OK components proceed to back-end production). If an image is classified as defective, or the AI is unable to make a decisions, it is forwarded to human operators for final verification. Operator decisions result in either acceptance to back-end production or component rejection.

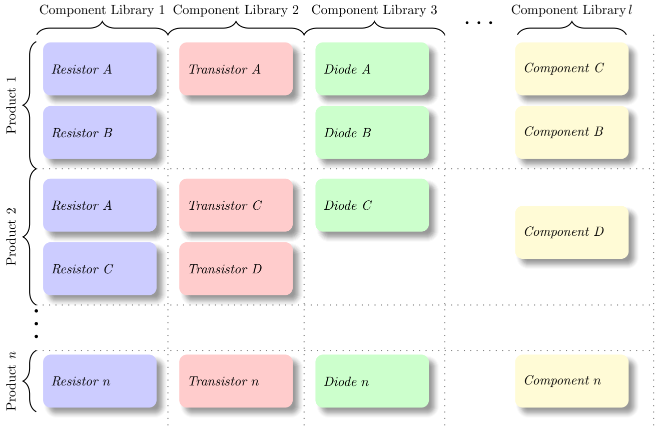

Figure 3: Relationship between Products, Component Libraries, and Components

Notes: This figure displays the representation of our data structure. Components exist at the intersection of Products and Libraries: each product contains multiple components from multiple libraries, and each library contains components appearing across multiple products. This representation shows a fictionalised distribution of components across products and libraries. Component libraries are groupings of similar types of components, here represented as Resistors, Transistors and Diodes. In reality, component libraries will be more granular, representing for example Resistors of a particular size and shape.

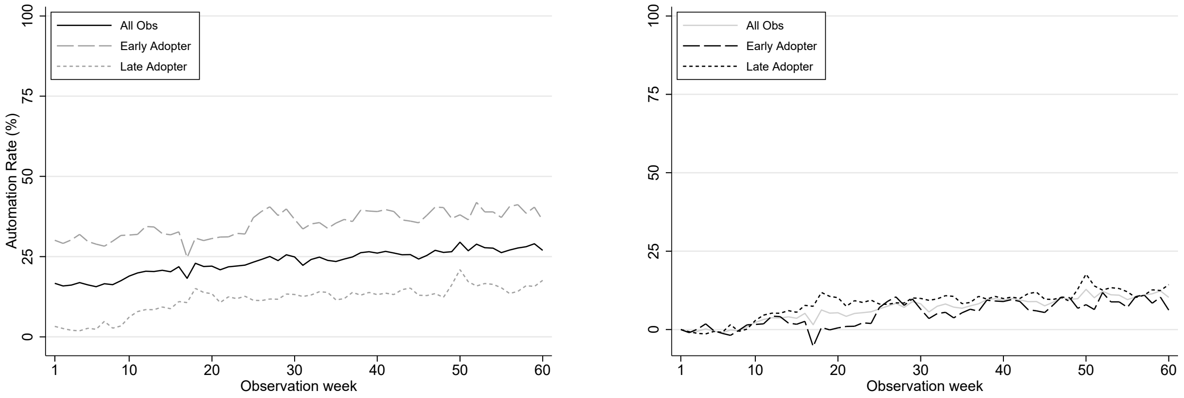

(a) Automation Rate - Levels

(b) Automation Rate -Indexed Growth

Figure 4: AI Automation Rate

Notes: Panel (a) shows the share of QA decisions automated by AI over the 60-week observation period for all plants, early adopters (deployed >8 months before sample start), and late adopters (deployed ≤8 months before sample start). Panel (b) shows the same series indexed to Observation Week 1 to highlight relative growth trajectories.

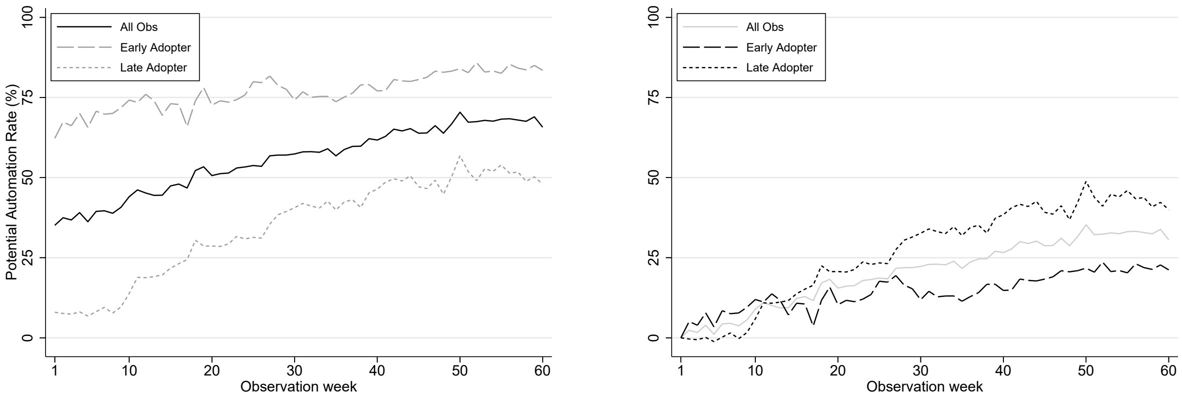

(a) Potential Automation Rate - Levels

(b) Potential Automation Rate Indexed Growth

Figure 5: AI Potential Automation Rate

Notes: Panel (a) shows the share of QA decisions made that could be automated by AI over the 60-week observation period for all plants, early adopters (deployed >8 months before sample start), and late adopters (deployed ≤8 months before sample start). Panel (b) shows the same series indexed to Observation Week 1 to highlight relative growth trajectories.

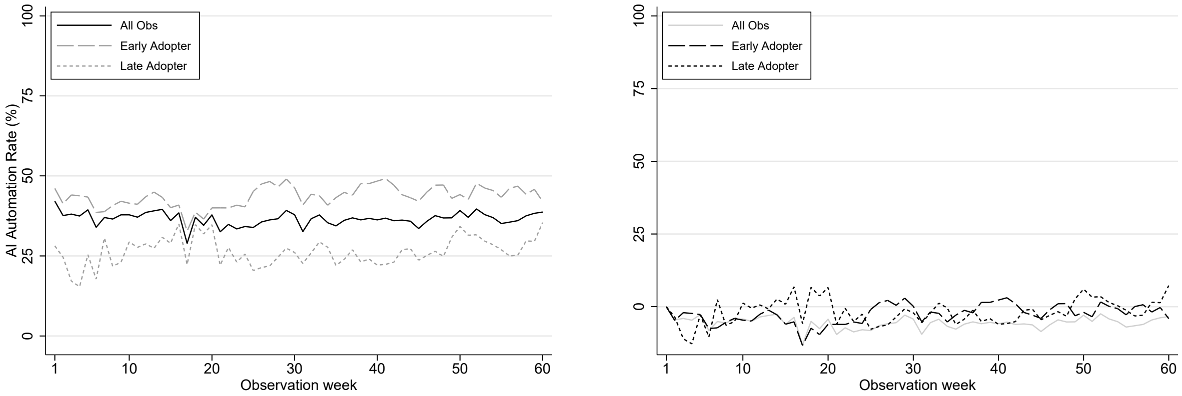

(a) AI Automation Rate - Levels

(b) AI Automation Rate Indexed Growth

Figure 6: Progression of AI Automation Rate on components where the AI has been deployed

Notes: Panel (a) shows the share of QA decisions automated by AI on components where the AI has been deployed, early adopters (deployed >8 months before sample start), and late adopters (deployed ≤8 months before sample start). Panel (b) shows the same series indexed to Observation Week 1 to highlight relative growth trajectories.

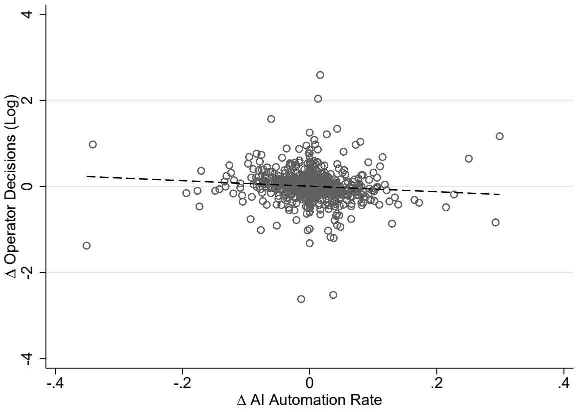

Figure 7: Change in operator decision and AI Automation Rate

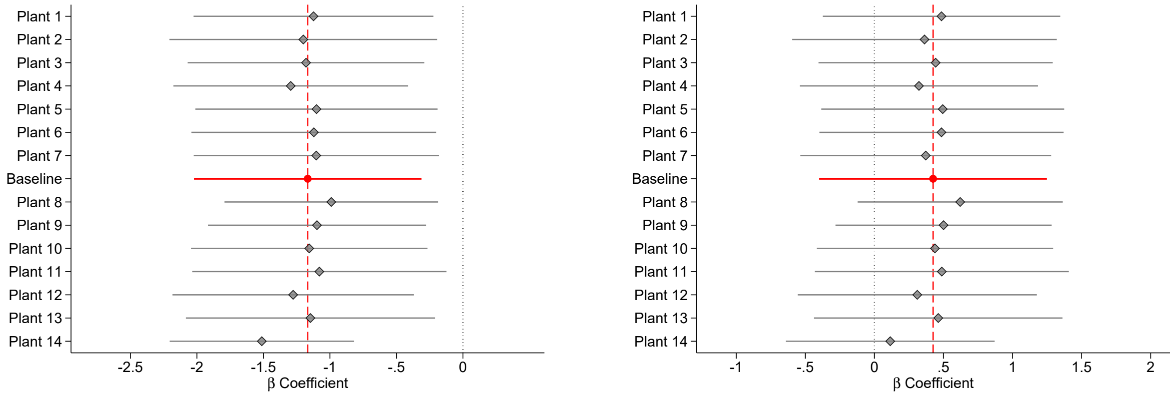

(a) Δ Total Operator Decisions (Log)

(b) Δ Total QA Decisions (Log)

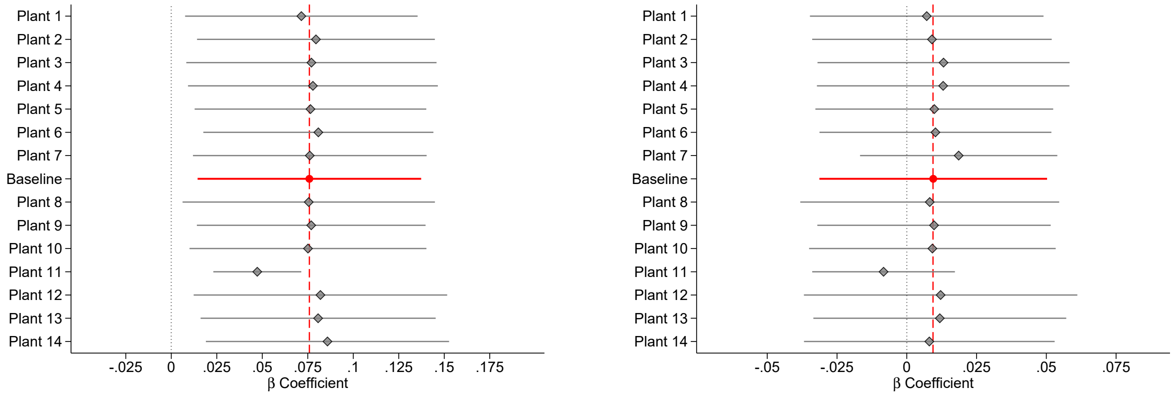

(c) Δ Total OK Rate

(d) Δ Operator OK Rate

Figure 8: Leave-One-Plant-Out Robustness Checks

Notes: Each panel shows coefficient estimates when excluding one plant at a time for each model run including 95% confidence intervals, the Y-axis indicates which plant has been excluded. The red dashed line indicates the baseline estimate using all plants, the black dotted line indicates zero effect. All regressions include plant and week fixed effects with standard errors clustered at the plant level.

Tables:

	Mean	SD	Min	Max
Panel A: Plant Characteristics				
Weeks Observed	59	2	53	60
Unique Products	902	811	130	3303
Unique Component Libraries	1118	671	102	2404
Share of total QA Decisions (%)	7.1	10.6	0.2	42.0
Panel B: AI Automation & QA Outcomes				
AI Automation Rate (%)	23.1	20.6	1.8	67.9
Overall OK Rate (%)	95.3	4.2	86.7	99.8
Operator OK Rate (%)	93.8	5.2	83.4	99.7
Panel C: Week-to-Week Changes				
Δ AI Automation Rate (pp)	0.164	0.109	-0.038	0.374
Δ Total QA Throughput (Log)	0.008	0.018	-0.021	0.046
Δ Total Operator Decisions (Log)	0.006	0.018	-0.021	0.044
Δ Total AI Decisions (Log)	0.030	0.048	-0.024	0.167

Table 1: Summary Statistics

Deployment timing	Number of Plants
Early Adopters (>8 months before)	
>15 months before	4
8-15 months before	3
Late Adopters (<=8 months before/during)	
1-4 months before	2
0-3 months after start	3
8-11 months after start	2
Total number of plants	**14**

Table 2: Plant distribution by AI deployment timing

Table 3: Relationship betwen changes in AI Automation Rate and Main Outcomes

	(1) Unweighted	(2) Weighted
Panel A – Δ Total Operator Decisions (Log)		
Δ AI Automation Rate	-1.201**	-1.167**
	(0.413)	(0.396)
Observations	810	810
Anticipation Effect		
Δ AI Automation Rate (Current)	-1.066**	-1.053**
	(0.410)	(0.392)
Δ AI Automation Rate (lead 1 week)	0.420	0.398
	(0.261)	(0.246)
Δ AI Automation Rate (lead 2 weeks)	0.442	0.371
	(0.250)	(0.216)
Observations	800	800
Panel B – Δ Total QA Decisions (Log)		
Δ AI Automation Rate	0.374	0.425
	(0.397)	(0.382)
Observations	810	810
Anticipation Effect		
Δ AI Automation Rate (Current)	0.507	-0.356*
	(0.397)	(0.174)
Δ AI Automation Rate (lead 1 week)	0.384	0.304
	(0.265)	(0.385)
Δ AI Automation Rate (lead 2 weeks)	0.486*	0.124
	(0.256)	(0.226)
Observations	800	800
Panel C – Δ OK Rate		
Total OK Rate		
Δ AI Automation Rate	0.080**	0.076**
	(0.028)	(0.028)
Observations	810	810
Operator OK Rate		
Δ AI Automation Rate	0.013	0.009
	(0.020)	(0.019)
Observations	810	810

Notes: This table presents two-way fixed effects regressions of weekly changes in labor demand (Panel A), total QA throughput (Panel B), and OK Rates (Panel C) on changes in the AI Automation Rate at the week-plant level. The *Anticipation Effect* specification includes one- and two-week leads to test whether future AI deployment correlates with current outcomes. All regressions include plant and week fixed effects. Column (1) reports unweighted estimates; Column (2) weights by total weekly decisions. Standard errors are clustered at the plant level. *** ($p < 0.01$); ** ($p < 0.05$); * ($p < 0.1$).

Table 4: Sensitivity Analysis on Main Outcomes

	(1) Preferred Specification	(2) Outliers	(3) Missing Decisions	(4) Missing Products	(5) Missing Products & Decisions
Panel A – Δ Total Operator Decisions (Log)					
Δ AI Automation Rate	-1.167**	-1.558***	-0.876**	-1.415***	-1.139***
	(0.396)	(0.329)	(0.389)	(0.318)	(0.294)
N	810	728	758	740	702
Panel B – Δ Total QA Decisions (Log)					
Δ AI Automation Rate	0.425	0.085	0.751*	0.188	0.501*
	(0.382)	(0.396)	(0.351)	(0.320)	(0.275)
N	810	728	758	740	702
Panel C – Δ OK Rate					
			Total OK Rate		
Δ AI Automation Rate	0.076**	0.018	0.075**	0.079**	0.078**
	(0.028)	(0.013)	(0.032)	(0.030)	(0.035)
N	810	728	758	740	702
			Operator OK Rate		
Δ AI Automation Rate	0.009	-0.052**	0.007	0.009	0.005
	(0.019)	(0.021)	(0.022)	(0.019)	(0.022)
N	810	728	758	740	702

Notes: This table presents weighted two-way fixed effects regressions of weekly changes in labor demand (Panel A), total QA throughput (Panel B), and OK Rates (Panel C) on changes in the AI Automation Rate at the week-plant level with various sample restrictions applied. Column (1) reports the weighted estimates reported in Table 3; Column (2) removes the top and bottom 5% of changes in AI Automation Rate; Column (3) removes weeks with a high proportion of lines missing decisions; Column (4) removes weeks with a high proportion of missing products; Column (5) removes both. Standard errors are clustered at the plant level. All regressions include plant and week fixed effects. *** ($p < 0.01$); ** ($p < 0.05$); * ($p < 0.1$).

Table 5: Relationship between changes in AI Automation Rate and Main Outcomes using Component Library Panel

	(1) Unweighted	(2) Weighted
Panel A – Δ Total Operator Decisions (Log)		
Δ AI Automation Rate	-2.066***	-2.071***
	(0.067)	(0.068)
Observations	24361	24361
Anticipation Effect		
Δ AI Automation Rate (Current)	-1.920***	-1.936***
	(0.075)	(0.075)
Δ AI Automation Rate (lead 1 month)	0.125***	0.142***
	(0.046)	(0.046)
Δ AI Automation Rate (lead 2 months)	0.206***	0.218***
	(0.043)	(0.043)
Observations	18750	18750
Panel B – Δ Total QA Decisions (Log)		
Δ AI Automation Rate	0.282***	0.277***
	(0.062)	(0.063)
Observations	24598	24598
Anticipation Effect		
Δ AI Automation Rate (Current)	0.373***	-0.190***
	(0.068)	(0.052)
Δ AI Automation Rate (lead 1 month)	0.115***	0.009
	(0.043)	(0.060)
Δ AI Automation Rate (lead 2 months)	0.190***	-0.038
	(0.042)	(0.045)
Observations	18909	18909
Panel C – Δ OK Rate		
Total OK Rate		
Δ AI Automation Rate	0.049***	0.048***
	(0.006)	(0.006)
Observations	24598	24598
Operator OK Rate		
Δ AI Automation Rate	-0.028***	-0.028***
	(0.007)	(0.006)
Observations	24361	24361

Notes: This table presents two-way fixed effects regressions of weekly changes in labor demand (Panel A), total QA throughput (Panel B), and OK Rates (Panel C) on changes in the AI Automation Rate at the month-component library level. The *Anticipation Effect* specification includes one- and two-week leads to test whether future AI deployment correlates with current outcomes. All regressions include plant and week fixed effects. Column (1) reports unweighted estimates; Column (2) weights by total weekly decisions. Standard errors are clustered at the plant level. *** ($p < 0.01$); ** ($p < 0.05$); * ($p < 0.1$).

Online Appendix

1 Appendix Figures

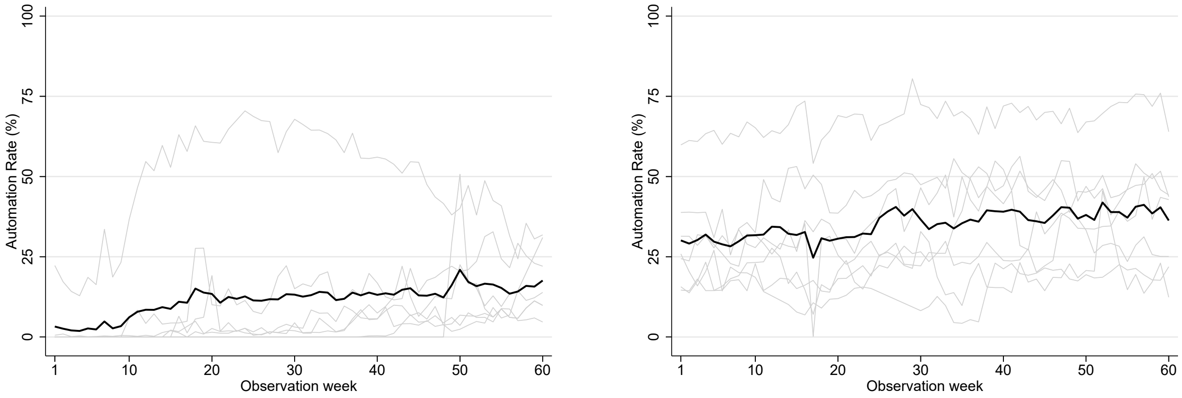

(a) Late Adopters

(b) Early Adopters

Figure A.1: Plant-Level Automation Rate Heterogeneity by Adopter Group

*Note:*Individual plant trajectories (thin gray lines) and group mean (thick black line) for automation rates over the 60-week observation period. Panel (a) shows late adopters (N=7 plants, deployed $\leq$8 months before sample start). Panel (b) shows early adopters (N=7 plants, deployed >8 months before sample start).

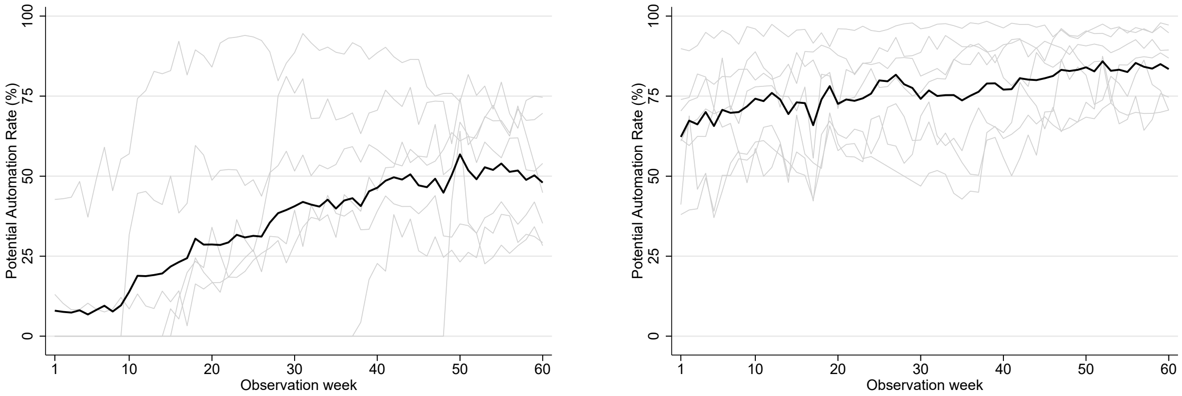

(a) Late Adopters

(b) Early Adopters

Figure A.2: Plant-Level Potential Automation Rate Heterogeneity by Adopter Group

*Note:*Individual plant trajectories (thin gray lines) and group mean (thick black line) for potential automation rates over the 60-week observation period. Panel (a) shows late adopters (N=7 plants, deployed $\leq$8 months before sample start). Panel (b) shows early adopters (N=7 plants, deployed >8 months before sample start).

2 Identification Challenges

There are a number of factors that limit causal identification in our setup. At more granular levels than the plant month we encounter violations of the Stable Unit Treatment Value Assumption (SUTVA). Consider deployment at the component library level, where treatment is defined as whether AI has been approved for any product containing a component from that library. Because individual products contain components from multiple libraries, spillovers are inevitable. For example, if a PCB contains both a resistor (with AI deployed) and a diode (without AI deployed), AI adoption for the resistor affects outcomes for the diode as they are produced on the same production line. Operator workload may decrease due to reduced labour demands, but this could be offset by faster production times. They may also face increased cognitive demand due to a higher proportion of defective images, impacting their decision making on the diode. If we assigned treatment at the product level, we would be able to rule out this possibility. However, operators work across multiple lines. As we are unable to identify the lines on which operators work, we are unable to rule out the fact that decisions on untreated products may be affected by treated ones. Imagine a factory with two lines, each producing separate products. One product has the AI enabled, the other does not. An operator may be making decisions on products from both lines, thus the AI automating a proportion of decisions from one line could affect decisions made on products from another.

Aside from the SUTVA violations, we also encounter problems relating to the endogeneity of treatment timing. Plants do not deploy AI randomly across products or component libraries. As described in Section 2.5, deployment requires customer approval, trained models with sufficient labeled data, and plant-level confidence in the technology's performance. Critically, plants may strategically time AI deployment to coincide with increases in production volume associated with new orders or high-volume production periods. Figure A.1 illustrates this problem using the estimator of de de Chaisemartin and D'Haultfœuille (2020) to trace the evolution of total QA decisions around AI deployment at the component library level. Total decisions increase sharply in the period leading up to and immediately after deployment, rising by approximately 50% in the month of deployment. This spike demonstrates that deployment timing is correlated with production shocks rather than being exogenous. Because we cannot distinguish whether changes in operator decisions reflect genuine task displacement or compositional shifts in the product mix being inspected, we cannot construct a credible counterfactual for labor demand at the product or component library level, even with staggered difference-in-differences estimators designed to handle dynamic treatment effects.

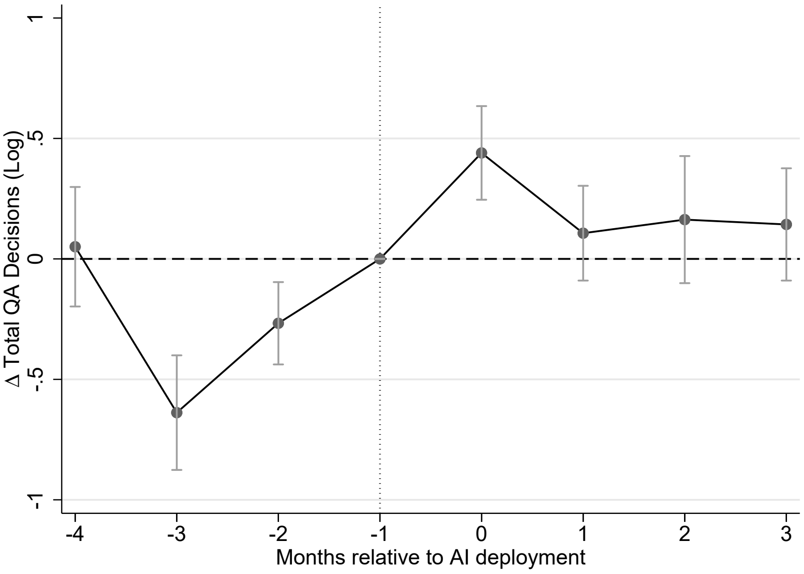

Figure A.1: Endogeneity of AI Deployment Timing

Notes: This figure shows coefficients from the de Chaisemartin and D'Haultfœuille (2020) estimator at the component library level. The vertical line indicates AI deployment (t=1). Total QA decisions increase sharply following deployment, rising by approximately 50% at t=1. This spike demonstrates that deployment timing coincides with production increases, preventing causal identification at this level of aggregation.