

The Rise of Industrial AI in America: Microfoundations of the Productivity J-curve(s)*

Kristina McElheran[†] Mu-Jeung Yang[‡] Zachary Kroff[§] Erik Brynjolfsson[¶]

Preliminary

December, 2025

Abstract

This study examines the productivity dynamics of artificial intelligence (AI) in American manufacturing. Working with the U.S. Census Bureau to collect detailed large-scale data for 2017 and 2021, we find significant initial productivity losses preceding medium-term gains from industrial AI deployment. We attribute this to costly adjustment (not just mismeasurement), which we observe directly via increased work-in-progress inventory, investment in industrial robots, and labor shedding. Over time, however, early AI adopters exhibit stronger growth on average, conditional on weathering the initial “dip.” Losses vary considerably across firms and establishments. A key contingency is age, with young firms faring better than older incumbents—particularly startups with growth-oriented business strategies. Management practices and production-process design significantly shape both the uptake and effects of industrial AI use. Firm structure also matters, as we observe significant cross-establishment spillovers inside large, multi-unit firms. Overall, our detailed findings provide novel evidence regarding AI-related *J-curve* effects, unveiling key mechanisms and extending our understanding of this emerging General Purpose Technology.

JEL Classification: D24, O33, M11, L60,

Keywords: Artificial Intelligence, General Purpose Technologies (GPTs), Technology Adoption, Management Practices, Organizational Structure, Manufacturing, Productivity.

*Thanks are due to Tim Bresnahan, Mert Demirer, Shane Greenstein, Rob Seamans, Prasanna Tambe, Anders Humlum, Tim DeStefano, Todd Zenger, Matt Higgins, Dan Gross, Luca Fontanelli, Manav Raj, Tobias Kretschmer, Cristina Tello-Trillo, Martha Stinson, and Cheryl Grim for helpful comments, as well as to participants in the Rotman Strategy brownbag seminar, Wharton’s AI and the Future of Work conference, the Columbia/Wharton Management of Analytics and Data conference, the University of Utah, the 2024 Empirical Management Conference, the NBER Productivity Lunch, the Wharton Technology and Innovation Conference, and the NBER Summer Institute. We are grateful to the Stanford Digital Economy Lab and the Social Sciences and Humanities Research Council (SSHRC) of Canada for partial funding. Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure avoidance protection (Project 7512395: CBDRB-FY24-CED006-0016, CBDRB-FY25-CED006-001). This paper was written while Kroff was an economist at the U.S. Census Bureau.

[†]University of Toronto, Email: k.mcelheran@utoronto.ca

[‡]University of Colorado, Email: mj.yang@colorado.edu

[§]Analysis Group, Email: zachary.kroff@analysisgroup.com

[¶]Stanford University, Email: erikb@stanford.edu

1 Introduction

Advances in artificial intelligence (AI) are poised to reshape firm performance, labor demand, and the organization of production. Productivity gains from AI have been identified in narrow settings such as software development (e.g., Cui et al., 2024), certain types of drug and materials discovery (e.g., Lou and Wu, 2021; Strieth-Kalthoff et al., 2024), and customer support (Brynjolfsson et al., 2025). However, evidence outside of modular, data-rich, and already prediction-focused applications (Bresnahan, 2024) remains scarce—particularly beyond the level of a specific task or “point solution” (Agrawal et al., 2024).

Yet a broader firm- or establishment-level view is essential, given the importance of the organizational context in shaping the returns to technology use (e.g., Bresnahan et al., 2002; Bloom et al., 2012) and recent theory emphasizing the need for a “system-wide” approach to deploying AI (Agrawal et al., 2024). That said, representative data on firm-level AI adoption has been in short supply (Seamans and Raj, 2018), while adoption has also lagged expectations (Bonney et al., 2024) and varied significantly across industry and organizational contexts (Calvino and Fontanelli, 2023; McElheran et al., 2024). This has made it difficult to generate broad-based facts about AI’s prevalence and average productivity effects—let alone pin down the organizational complements and constraints that have historically affected technology adoption and productivity (David, 1990; Bresnahan and Trajtenberg, 1995; Brynjolfsson and Hitt, 2000; Feigenbaum and Gross, 2024).

We address this gap with new large-scale evidence on the use of AI¹ among U.S. manufacturers and performance data across nearly a decade, drawing on two surveys conducted for this purpose in collaboration with the Census Bureau. In addition to providing representative statistics on both early AI use and barriers to adoption in physical-goods production, we leverage a novel instrumental variable (IV) strategy and visibility to within-firm changes to triangulate on the varied impacts of AI on business performance. Our findings are consistent with a *J-curve* pattern at the micro-level: the average effects of AI deployment are initially quite negative, followed by growth along multiple dimensions over time. These early losses, however, do not unfold uniformly; they vary substantially by age, business strategy, and other organizational characteristics and practices.

¹Defined as “a machine-based system that can perceive and learn about its environment and then make relevant predictions, recommendations or decisions.” See <https://www.census.gov/programs-surveys/mops/technical-documentation/questionnaires.html>

Even when voluntary, technology adoption does not guarantee productivity gains, particularly in the short run (Tambe and Hitt, 2012). With respect to AI, we find the initial relationship between adoption and performance to be *negative*, despite self-selection into AI use. Specifically, a one-standard deviation increase in our AI index is associated with a 1.33% lower total-factor productivity (TFP), controlling for a large number of time-varying (and usually unobserved) establishment characteristics. Moreover, when we employ quasi-experimental techniques commonly used to address endogeneity and measurement error, losses grow considerably in magnitude. IV estimates yield productivity losses of roughly 44%, which we interpret as a local average treatment effect (LATE) experienced by marginal businesses whose AI adoption is sensitive to our instrument.² This nuance is important, as we do not interpret such a large productivity hit to be representative for infra-marginal “always adopters,” such as digitally-advanced businesses already poised to reap early returns from AI use in compatible applications. Yet we find negative effects, on average, across a range of specifications—ultimately underscoring the significant short-term costs of AI deployment, in practice. It further underscores the heterogeneity endemic to this phenomenon, reflecting the much greater prevalence in our sample of smaller firms and establishments that lack the advantages of larger, publicly-traded firms central to other studies (e.g., Babina et al., 2024).

Critically, these initial losses do not imply that AI adoption is systematically a “mistake.” With time for adjustment to unfold (and conditional on survival), we find *growth* to be significantly higher among early AI adopters, on average. This is particularly true among younger firms and establishments. For example, a one standard-deviation increase in AI use is associated with roughly 2% higher labor productivity growth among younger plants from 2017 to 2021. Among firms of all ages, greater AI use is associated with 0.8% employment growth over a four-year period that saw a 1.5% sector-wide decline, overall.

Data limitations prevent us from tracing out the precise inflection point where net returns become positive or characterizing true long-term implications.³ Yet, our results provide empirical micro-foundations for *J-curve* patterns previously observed at the macro level (Brynjolfsson et al., 2021b). In particular, we find direct evidence of costly adjustment within firms. Not only do productivity and profits initially suffer, inventory buffers also increase due to AI use—in a sector

²The “compliers” in the Angrist et al. (1996) framework. In this case, it reflects businesses whose adoption of AI was not impeded by the level of AI-specific expertise or skills at the plant.

³We do not observe the exact year of AI adoption and the panel of performance data ends in 2022.

where “Lean” practices are prevalent. Investment in industrial robots also increases, suggesting adjustment to physical processes. And, potentially linked to this, we observe significant labor shedding in the short term, indicating that workers also experience AI adjustment costs.⁴

Perhaps most importantly, we also gain insight into specific mechanisms shaping AI’s impacts. We disentangle how certain business strategies help mitigate the *J-curve* “dip” and also find that much of the productivity losses are driven by older manufacturing establishments. We further show that productivity losses among older businesses are closely tied to a loss of structured (“Lean”-style) management practices (Bloom et al., 2019)—consistent with an important role for key organizational intangibles (Levitt et al., 2013). Specifically, our quantitative analysis shows that roughly half of the TFP losses at older establishments can be explained by de-adoption of previously-productive management practices. We also find significant spillovers to non-AI using establishments in multi-unit firms. IV estimates point to 40% higher labor productivity among non-adopting units in AI-using firms. These details reveal specific channels through which losses and gains are transmitted, pointing out potentially valuable ways to improve overall returns to AI use, and helping to explain conflicting estimates across differently-composed samples and empirical settings.

Our empirical context of manufacturing is a particularly interesting one for understanding AI’s nuanced impacts. It encompasses a highly innovative set of industries, accounting for 60% of all patents (National Science Foundation, 2021) and 7 % of all corporate patents (Autor et al., 2020a). Some estimates (e.g. Deloitte, 2020) suggest that manufacturing generates twice as much digital information as other data-intensive industries such as media and finance/banking. High levels of physical automation also provide potential complements to AI-related technologies (McElheran et al., 2024) and may reinforce labor-displacing effects (Acemoglu and Restrepo, 2018). Yet, manufacturing production processes are often not modular (Bresnahan, 2024), but instead depend on complex interactions of physical and virtual environments, requiring coordination of machinery, material flows, management practices, differently-skilled employees, business strategy, information, and technology (e.g., Ichniowski et al., 1995; Milgrom and Roberts, 1990, 1995; McElheran and Jin, 2020). Heterogeneity in how to accomplish this—much less adjust it all to align with new technological possibilities—abounds. Thus understanding how organizational precursors and more-mutable

⁴We do not observe worker-level outcomes, making it impossible to infer either individual-level or aggregate labor-market impacts.

managerial decisions determine returns to AI use is critical to gauging its broader implications for firms, workers, and the economy.

Our findings contribute to several streams of research. First, they contribute to longstanding questions on the economic impacts of general-purpose technologies, or GPTs (Bresnahan and Trajtenberg, 1995), which AI-related technologies are increasingly argued to be (e.g., Brynjolfsson et al., 2021b; Goldfarb et al., 2023). This literature emphasizes the implementation lags and complementary organizational adaptations typically required for GPTs to achieve measurable gains (e.g., David, 1990; Bresnahan and Greenstein, 1996; Brynjolfsson and Hitt, 2000; Feigenbaum and Gross, 2024). We extend this line of work by documenting *micro-level J-curve* patterns (Brynjolfsson et al., 2021b) in which short-run losses reflect production-process and organizational disruptions (rather than primarily mismeasurement of intangible investments) during early Industrial AI adoption, followed by medium-term performance improvements for most firms.

Within this broader literature, we contribute to a growing subfield pinpointing specific organizational complements to modern technology use. Previous studies have revealed aspects of how earlier digital technologies interacted with organizational characteristics (e.g., Bartel et al., 2007; Bresnahan et al., 2002; Tambe et al., 2012; Aral et al., 2012; Bloom et al., 2012; Brynjolfsson et al., 2021a). Extending arguments that that production systems consist of mutually reinforcing clusters of technologies, processes, and practices (e.g., Ichniowski et al., 1995; Milgrom and Roberts, 1990, 1995; Brynjolfsson and Milgrom, 2013), we provide the first large-scale evidence of technological, process-design, and organizational contingencies affecting AI uptake and outcomes.

Third, we add to the nascent empirical evidence on AI adoption and its performance consequences. A burgeoning approach relies on field studies of specific applications of AI in specific activities—e.g., computer programming (e.g., Peng et al., 2023; Hoffmann et al., 2024; Cui et al., 2024), customer support (Brynjolfsson et al., 2025), consulting (Dell’Acqua et al., 2023), entrepreneurship coaching (Otis et al., 2024), or scientific discovery (Lou and Wu, 2021; Aspuru-Guzik, 2023)—that may be considered “deep but narrow” (Bresnahan, 2024), and thus difficult to generalize. Contrasting with this approach, other recent studies have inferred less-specified AI use from job postings, finding mixed performance effects among more-diverse yet relatively large publicly-traded firms (e.g., Alekseeva et al., 2020; Babina et al., 2024). A complementary stream of work relies on direct survey measurement by administrative agencies, yielding more-representative

statistics on adoption (Hoffreumon et al., 2023; Fontanelli et al., 2024; Bonney et al., 2024) along with contradictory findings regarding short-term productivity effects (Czarnitzki et al., 2023; Acemoglu et al., 2022; Calvino and Fontanelli, 2023; McElheran et al., 2024). By combining two Census Bureau surveys conducted four years apart, along with administrative data covering tens of thousands of manufacturing businesses across nearly a decade, we bridge these literatures to generate insights that are unusually representative yet detailed, and reflecting the passage of time. Our panel data reveal how AI’s negative *and* positive effects evolve over different horizons, while the survey modules allow us to trace heterogeneity by production system, age, strategy, and within-firm spillovers, offering a richly contextualized view of industrial AI use and its impacts.

Taken together, our findings highlight AI’s dual role as a transformative technology and catalyst for initial organizational disruption, echoing patterns familiar to scholars of technological change. They further underscore the importance of complementary practices, structures, and strategies that mitigate adjustment costs and enhance longer-term returns, providing practical guidance to managers and policy-makers on how to flatten the *J-curve* dip and realize AI’s longer-term productivity potential at scale.

2 Literature and Motivation

Rising digitization and the diffusion of advanced tools to extract value from data have had profound impacts on firm organization, management, competitive strategy, and performance (e.g., Bresnahan et al., 2002; Bloom et al., 2012; Goldfarb and Tucker, 2019; Adner et al., 2019; Tambe et al., 2020). In this section, we extend insights from digitization research in economics and management to motivate our analyses of both the adoption and impact of AI use in industrial (i.e., physical goods) production.

2.1 Understanding AI as a Specific Type of Production Input

2.1.1 Scale-Biased Technology

Information and knowledge are well-understood to be non-rival goods whose consumption does not reduce availability to others and whose reproduction occurs at near-zero marginal cost (Arrow, 1962; Romer, 1990; Goldfarb and Tucker, 2019). A recent return to these ideas focuses attention

on the organizational and strategic implications of “digital resources, such as data, software and AI that are essentially scale free” (Giustiziero et al., 2023).⁵ AI algorithms, once trained, share these features. As a result, they are prone to significant economies of scale, creating more value at larger production volumes. These scale economies will be further reinforced by large up-front training costs (Svanberg et al. 2024).⁶

These characteristics work in favor one of the few robust “stylized AI facts” across countries and industries: early AI adoption rates are systematically higher among larger firms (e.g., Calvino and Fontanelli 2023; McElheran et al. 2024; Bonney et al. 2024; Hoffreumon et al. 2024). Note, however that properly identifying size-related effects is challenging due to the high correlation between size and other potential drivers of digital technology adoption and performance such as age (Kueng et al., 2014), organizational structure (e.g., multi-unit status, per Atalay et al., 2014), and production-process design (Brynjolfsson et al., 2021a; McElheran et al., 2019). We control for size in all of our specifications and lean on the richness of our plant- and firm-level data to tease apart these often-confounded relationships.

2.1.2 Dependent on Prior Digital Transformation

Beyond firm size, another critical determinant of AI adoption is the availability of appropriate digital infrastructure and data inputs (e.g., Tambe et al., 2020; Goldfarb et al., 2023; Babina et al., 2024). To date, AI adopters have tended to report significant levels of digitization and reliance on cloud-based IT infrastructure (McElheran et al., 2024), implying that AI adoption is embedded in a broader, complementary process of digital transformation that may easily be mis-attributed to AI, alone.⁷ Omitting key features of the technological context may obscure drivers of adoption and also bias estimates of AI’s marginal contribution to firm performance (e.g., Calvino and Fontanelli, 2023). A contribution of our approach is the ability to control for other digital technology deployed

⁵See Helfat et al. (2023) for a review of related work.

⁶Consider the benefits of machine vision technology to analyze digital pictures of finished products to diagnose any quality defects that might often be too subtle to quickly notice for human quality control employees. The initial training of this technology with training data will be expensive but once a sufficient number of training instances is used, the supervised learning algorithm will eventually outperform human vision. In terms of the unit cost economics, training the machine vision algorithm is a large fixed cost, which will be spread out across many units, leading to a decline in marginal cost. Additionally, ongoing training of the machine vision technology will increase accuracy, thereby increasing value created from the technology by continued quality improvements of products

⁷The idea that technology adoption is a multi-stage process dependent on prior investments is an emerging theme beyond the study of AI-related technologies (e.g., Comin et al., 2025).

alongside AI. Specifically, we disentangle both cloud-based and on-premises IT infrastructure, as well as prior investments associated with “big data” initiatives focused on digitizing information and applying descriptive and predictive analytics to increasingly-important data resources.

2.1.3 Prediction-Biased

The need to account for prior technology investments and their complementary organizational adjustments (e.g., Bresnahan et al., 2002; Bresnahan and Greenstein, 1996) goes deeper than potentially misattributing the effects of digitization too-narrowly to AI use. Specifically, factors driving both the use of AI—usefully conceptualized as a dramatic advance in “prediction technology” (Agrawal et al., 2018)—and adoption of pre-AI digital tools may have their origins in fundamental differences in production system design.

Certain aspects of production design reflect the resolution of deep economic and organizational tradeoffs and are essentially fixed once a production facility is established (McElheran et al., 2019). At least two distinct approaches, based on differing solutions to essential uncertainty (Galbraith, 1973) have emerged in U.S. manufacturing: “make-to-order” versus “make-to-stock.” Milgrom and Roberts (1988) demonstrate mathematically that profits are convex in the market share supplied from inventory, driving firms to specialize in one of these production modes or the other. They further argue that each approach tends to be complementary to different organizational and technological solutions (Milgrom and Roberts, 1990, 1995), leading firms to adopt clusters of practices, rather than individual innovations in isolation.⁸

“Make-to-order” production, exemplified by Lean Manufacturing (Womack et al., 2007; Holweg, 2007) and related structured management practices (Bloom and Van Reenen, 2007; Bloom et al., 2019), prioritizes responsiveness to realized demand over prediction. This approach to addressing uncertainty depends heavily on visibility, minimal inventory buffers, and empowering workers to manage exceptions in real time. It has successfully diffused throughout U.S. manufacturing based on “low-tech” elements drawn from the Toyota Production System such as *kanban* cards and *andon* techniques (Ohno, 1988) that support information-sharing with relatively little IT.

At the other extreme, traditional “make-to-stock” manufacturing has evolved in U.S. manufacturing via substantial investments in both physical and IT capital—and relies heavily on prediction.

⁸See also Ichniowski et al. (1995); Brynjolfsson and Milgrom (2013); Aral et al. (2012); Tambe et al. (2012).

More-common in settings where lead times are long and/or production flexibility is limited, this approach relies heavily on forecasting and inventory buffers to address uncertainty in supply and demand conditions (Galbraith, 1973; Toktay and Wein, 2001).

More recently, as data availability has grown and IT costs have fallen, manufacturers of all types have adopted increasingly sophisticated predictive analytics (Brynjolfsson et al., 2021a). This has not been universally productivity-enhancing, however. In particular, these systems appear to be complementary to an organizational focus on external information from customers and supply chains (Tambe et al., 2012), as well as being contingent on a low-product-mix, high-volume production system design (Brynjolfsson et al., 2021a).

Such path dependencies and contingencies pose challenges for sharp predictions about interdependencies among AI and certain other technologies. Take basic digital tools such as descriptive analytics (Berman and Israeli, 2022; Galdon-Sanchez et al., 2022). On the one hand, if we consider digital capability development to be an evolving and cumulative process (e.g., Comin et al., 2025), the presence of descriptive analytics tools and techniques may represent greater readiness to adopt advanced tools like AI. Yet, given the organizational and technological interdependencies at work, they could also proxy for deep organizational misalignment with prediction-focused management. This could work against successful deployment of AI. Thus, the relationship between older digital technologies and AI is ultimately a nuanced empirical question.

Regarding more-advanced digital tools such as predictive analytics, there is arguably less ambiguity. Production systems already-organized to leverage forecasting, exemplified by prior reliance on predictive analytics (Brynjolfsson et al., 2021a), should more easily adapt to AI use. To the extent that

2.2 Industrial Production, Adjustment, & the J-Curve

Since our study takes place in the manufacturing sector, it is useful to bridge insights from the study of “digital firms” in the information sector with longstanding intuitions about industrial production, which has become increasingly automated and digitized in the United States. The strategic implications of economies of scale in physical-goods production have been understood

since at least the early 1960s (Arrow, 1962; Henderson, 1968).⁹ One is early-mover advantage. Economies of scale can be an important isolating mechanism, protecting early movers (Rumelt, 2005). First-movers are by definition able to accumulate production experience earlier. Under economies of scale, this increases value creation (by reducing unit costs, increasing quality, or both), and feeds back into yet more demand and higher production volumes (Cabral and Riordan, 1994). While industrial firms do not have the same demand-side economies as firms whose products and services are also digital (Giustiziero et al., 2023), we anticipate that “traditional” economies of scale combined with the non-rival attributes of AI technologies will mutually reinforce early-mover benefits from AI adoption.

A key limiting factor, however, arises from the general-purpose nature of the technologies in question. General-purpose technologies (GPTs), such as the steam engine, electricity or digital computers, are a class of economically and socially transformative technologies that share characteristics of being widely used, constantly improving, and innovation-spawning (Bresnahan and Trajtenberg, 1995). When we look at the family of technologies increasingly referred to as “AI” (e.g., machine learning, machine vision, speech recognition, etc.) they are already observed in every sector of the economy, in far-ranging use cases, advancing at unprecedented rates, and closely linked to innovation (Cockburn et al., 2018; Iansiti and Lakhani, 2020; Felten et al., 2021; Miric et al., 2023; McElheran et al., 2024). While there is some debate at the moment as to whether AI-related technologies will definitively emerge as the next influential GPT (Goldfarb et al., 2023; Eloundou et al., 2023; Bresnahan, 2024), the GPT lens is nevertheless useful for focusing attention on early deployment challenges.

According to work in this vein, technologies with such broad potential, ongoing improvement, and often-uncertain trajectories typically require significant investments in co-invention (Bresnahan and Greenstein, 1996) to align technological capabilities with core business activities, processes, products, and resources—and vice versa. The pervasiveness and magnitude of this adjustment challenge is credited with widespread and lengthy delays in the early diffusion and productivity impacts of important technologies (e.g. David, 1990; Brynjolfsson and Hitt, 1996). Recent

⁹Indeed, Kiechel (2010) documents that the first popular tool sold by the Boston Consulting Group (BCG) in the 1960s was the “experience curve” or the empirical claim of a “consistent” decline of 20-30 percent of costs for each doubling of output (Henderson, 1968). In turn, this tool was used in turn by companies such as Black & Decker to undercut its competitors in terms of price, anticipating lower future unit costs.

GPT research examines this process among public U.S. firms, predicting and finding evidence for a “productivity J-curve” of initially declining measured productivity followed by a sharp rise in returns for investments such as software and R&D that arguably require significant—often intangible—complementary investment (Brynjolfsson et al., 2021b).

The implications of costly co-invention for scale-free AI-related technologies is that short-term returns from deploying AI may be minimal (or even negative) on average, and any early-mover advantages will emerge with a delay. Yet firms that can weather initial performance losses stand to enjoy significant returns to AI use, as observed in early narrow applications among digital giants like Amazon, Netflix, Google, and Facebook (Bresnahan, 2024).

How might this play out in our manufacturing context? Consider that a new AI-enabled production line might use a reinforcement-learning algorithm to optimally adapt production parameters, starting with a set of initial parameter values that are locally but not globally optimal.¹⁰ However, this exploration might prove costly in the short run, if trial parameter values lead to foregone output that would otherwise have been produced if the AI system had stayed at the initial, locally optimal parameter values. At the same time, however, exploration of new parameter values as production volume increases enables reinforcement learning algorithms to try out new parameter values for each newly produced unit, ultimately leading to rapidly improving performance at higher volumes of production.

This example describes a *J-curve* at the level of a given algorithm. However, similar mechanisms will arguably be at work beyond any individual AI deployment. Increasingly, the system-wide challenges of deploying AI beyond a “point solution” are receiving attention (Agrawal et al., 2024). Reconfiguring production processes can lead to short term coordination failures and operational inefficiencies that require additional buffers in the system (e.g., work-in-progress, or WIP). If responsiveness is improved at one point in the production process, increased automation (e.g., robots or additional AI technologies) may be required elsewhere. Uncertainty about the reliability of AI-enabled activities could lead to running parallel “backup” production lines to prevent stoppages or stockouts. Factor inputs may shift, particularly if this increased automation substitutes for human labor (Acemoglu and Restrepo, 2018). Workers may then directly experience AI adjustment costs.

¹⁰An example for a similar technology from practice can be found here: <https://www.pwc.de/en/digitale-transformation/the-perfect-match-digital-twins-and-reinforcement-learning.html>.

Working through these adjustments requires time and investment, with associated initial losses before returns fully materialize. As this plays out at the producer level, we anticipate that returns to AI use will exhibit a “J-curve” pattern, with performance declining upon initial adoption and eventually improving over time.

Furthermore, we should observe organizational adjustments along other dimensions (WIP, physical automation, and labor demand). We test these additional implications, empirically, recognizing that many margins of adjustment will remain intangible (Brynjolfsson et al., 2021b).

2.3 Strategic Responses

Thus, far, our predictions have concerned average tendencies and are of intense interest in economics, management, and public policy. Yet the managerially-relevant question quickly becomes how strategic decisions might help “flatten” the adjustment dip for a given producer.

Traditional learning-by-doing in manufacturing is well known to eventually run into diminishing returns, based in physical or cognitive limits of workers (e.g., Thompson, 2012; Levitt et al., 2013). In contrast, scale economies of digital resources such as AI are unlikely to face similar limits.¹¹ As a result, acceleration of benefits (the upward part of the “J-curve) from AI will tend to be supported by business strategies that emphasize growth. That said, we argue that the nature of a given growth-oriented strategy will matter.

To understand this, consider the BCG consultants that sold insights from the experience curve to manufacturing firms in the 1960s. This typical “cost-leadership” strategy (Porter, 1980) was designed to increase sales volume. Yet cost-leadership strategies quickly run into decreasing returns to scale on the demand side (Giustiziero et al., 2023). In contrast, growth through market expansion—via novel offerings and tapping new customer segments, including new domestic and international markets (Yang et al., 2015, 2021)—avoids these demand-side diseconomies¹². Indeed, co-occurrence of growth-oriented strategies and AI use has recently been documented among U.S. startups (McElheran et al., 2024); however, the performance implications in a broader population remain unknown. We anticipate that business strategies focused on adding scale via market

¹¹This does, of course, not mean that we believe that AI resources are literally scale-free. Indeed Bajari et al. (2019) argue that dataset size runs into diminishing returns at a “square-root N” rate for forecasting tasks.

¹²One way to reduce such demand-side diseconomies is to use product-differentiation, as argued by Porter, 1980. See Babina et al., 2024 for empirical evidence of predictive analytics leading to higher product differentiation

expansion or innovation will attenuate any initial performance declines due to AI use, while cost-leadership strategies will worsen initial performance declines.

2.4 AI vs. the Experience Curve

Our examination of *J-curve* effects at the micro level further has important implications for how mature, more-experienced firms compete with young startups. This “creative destruction” debate about whether larger incumbents or smaller startups are best equipped and/or incentivized to undertake this adjustment dates back at least to Schumpeter (1934, 1950),¹³ and is too large to summarize, here. Our contribution to this long-standing conversation, however, is to highlight that, if *J-curve* effects are both strong and uniform across incumbents and entrants, incumbents will naturally benefit from scale economies earlier, leading to entry barriers for startups (Cabral and Riordan, 1994) and eventually higher industry concentration, as incumbents continue to benefit disproportionately from AI use. This dynamic seems to be emerging with regard to “superstar” firms in settings where intangible capital is very important (Autor et al., 2020b; Tambe et al., 2020). Yet, if short-run performance losses are instead worse for mature incumbents, then AI adoption might conversely lead to less industry concentration, in the long run.

Anchoring again in our industrial production settings, we argue that old firms or establishments are more likely to derive productivity from accumulated experience (Thompson, 2012; Levitt et al., 2013), the accumulation of firm-specific resources such as customized IT capital (Jin and McElheran, 2024), and an established set of operational capabilities (Helfat and Peteraf, 2003). Corresponding management practices (Bloom et al., 2019), vintage-specific human capital (Chari and Hopenhayn, 1991; Barth et al., 2023), and other core competencies (Henderson, 1993) will also tend to sustain incumbent performance derived from a familiar set of technologies. When old firms intensively adopt a novel, AI-based technology, experience with prior technologies may not be transferable to the new system, leading to relatively larger productivity losses for more-established firms. In contrast, young startups by definition have none of this accumulation and will thus have lower opportunity costs of innovating (Arrow, 1962 *inter alia*). Thus, we anticipate that older producers will experience greater performance losses from AI adoption than younger ones.

¹³Schumpeter (1934) initially suggested that technological change and innovation would be led by small entrepreneurial firms yet subsequently (Schumpeter 1950) argued that large incumbents possessed superior incentives and resources to innovate and appropriate the returns to innovation.

2.5 Mechanisms Rooted in Organizational Routines and Practices

To better understand the specific channels by which older manufacturing firms might exhibit larger initial performance losses in response to AI adoption, we build on recent work showing that certain structured management practices—often drawing from the Lean tradition—work to capture, retain, and exploit organizational knowledge and capabilities (Bloom et al., 2012; Scur et al., 2021; Scur and Wolfolds, 2024).

By way of example, Levitt et al. (2013) analyze a single car manufacturing plant in the U.S. and show the critical role of management practices in realizing learning-by-doing gains from production. Specifically, they document the use of real-time reports and quality audits by engineers, combined with whiteboard system that allowed frontline workers to report individually-experienced production problems on an ongoing basis, thereby leading to generation and sharing of valuable knowledge,¹⁴ which Levitt et al. (2013) argue can be understood as a form of “organizational capital.” Structured management practices as measured by Bloom et al. (2013; 2019) similarly scaffold the generation and aggregation of employee knowledge by defining standard operating practices, promoting proactive investigation of root causes of problems, routinizing the review of key performance indicators (KPIs) by managers and employees, and enabling widespread awareness of production targets among all factory workers to make production exceptions quickly salient and less costly to address (Womack et al., 2007; Holweg, 2007). But such structured practices are fundamentally still ‘analog’ in that they do not rely on digitalized information or much technology at all, but rather on the interaction of standardized operating procedures and employee training and practices. Older factories will tend have more accumulated organizational capital of this type. The implication, then, is that AI adoption could interfere with reliance on the old knowledge generation system, yielding observable declines in structured management practices and, potentially, the employees in which they are embodied.

¹⁴Levitt et al. (2013) describe the interaction of one structured practices, namely the use of whiteboard for problem tracking and frontline employees in knowledge creation as follows: “(...) large amounts of information about production still originated from the workers on the line. (...) Aggregation and diffusion of this knowledge were the purposes of the whiteboard system. Workers were encouraged to note problems on the boards, (...) The system therefore quickly pulled information from individual line workers and allowed management to manipulate the production process in ways that benefited any worker at a similar position (...). The system therefore acted as a conduit that gathered worker knowledge and, through the complementary efforts of management, transformed it into plant knowledge that became embodied in the plant’s physical and organizational capital.”

2.6 Spillovers within Multi-Unit Firms

The non-rival or scale-free nature of AI as a digital resource opens up additional considerations for large multi-unit manufacturing firms. E.g., Pratt (2015) has argued that robotics or AI systems that are remotely connected via cloud computing could lead to increased knowledge spillovers across these systems, as newly valuable information of each system is immediately accessible and valuable to all connected systems. A similar argument should apply to establishments or factories as unit of analysis: AI systems in one factory will generate productivity spillovers in other connected factors, even enabling “parallelization” of learning if production problems are modular across factories. Thus AI adoption at other establishments within the same firm should tend to increase productivity at non-adopting manufacturing establishments.

3 Data and Measurement

3.1 Two Datasets and Samples

To provide robust estimates and study dynamics, we use two datasets on AI use by U.S. manufacturers at different time periods. Our main dataset is the 2021 Management and Organizational Practices Survey (MOPS), supplement to the Annual Survey of Manufactures (ASM). The ASM is one of the oldest large-scale Census Bureau data collections, sampling 10% of the roughly 300,000 manufacturing establishments across the country. Stratified by size and industry to provide representative statistics, the ASM nevertheless oversamples large establishments, covering over 70% of the sector’s total value added in the certainty sample. A government-mandated survey, the MOPS had a response rate of 68% in 2021.¹⁵ Panel A of Table 1 reports that the average establishment in our main sample has roughly 170 employees, is 29 years old, ships more than 60% of its products via e-commerce, and has approximately \$11 million in profits.

Linking the MOPS to the broader ASM yields a useful breadth of organizational measures. In addition to fine-grained data on the use of AI-related technologies, it provides information on the use of other digital technologies (predictive analytics, descriptive analytics, cloud computing,

¹⁵For more details go to <https://www.census.gov/programs-surveys/asm.html>. We also link to the Longitudinal Business Database (LBD), a dataset of tax records covering the entire non-farm employer economy of the United States, to acquire data on establishment and firm age. Establishment age is measured as the number of years since the establishment first reported having March 12 employment on their tax records. Firm age is determined by the age of the oldest establishment.

digitalization of information, IT capital) and structured management practices (Bloom et al., 2013, 2019). Dimensions of organizational and production design are captured in the ASM-MOPS linked data, as well (see below).

Our second main dataset is a panel of approximately 55,000 manufacturing firms combining data from the 2018 Annual Business Survey (ABS)¹⁶ with the Economic Census of Manufacturing (CMF) from 2012 to 2017. This panel dataset allows us to estimate the effects of AI adoption on performance while controlling for firm fixed effects. The ABS-CMF data are more representative of the large population of single-unit firms in the U.S. economy (Zolas et al., 2020). This different size coverage adds robustness to our findings, while also supporting an identification strategy focused on within-firm changes over time.

Panel B of Table 1 describes manufacturing firms in the ABS-CMF sample. The average firm in this sample is by no means small or young, with mean employment of 344 workers and age of 25. However, when comparing the MOPS-ASM data to the ABS-CMF sample, it is useful to keep in mind that the latter is a firm-level dataset whereas the former is at the establishment level. Moreover the majority of establishments in MOPS-ASM belong to larger, multi-unit firms (hence our ability to identify within-firm spillovers).

3.2 Measurement

3.2.1 New Measures of AI Use in the MOPS

Members of the author team worked with the Census Bureau to create new, designed-for-purpose measures of AI-related technologies and their applications for the 2021 MOPS. Measuring fast-emerging technologies systematically and at scale is a perennial challenge. Given the recency and some lack of consensus about what “AI” encompasses, our measurement approach is multifaceted. First, we clearly define what we mean by “AI” in the survey, as well as “predictive analytics,” which we separately measure to avoid confusion (and to distinguish from prior work, e.g., Brynjolfsson et al., 2021a). These definitions underwent systematic cognitive testing involving members of the research team and dedicated Census experts, as well as respondents pulled from the official ASM sampling frame, to ensure validity and reliability of responses.¹⁷ According to the survey, “Artificial

¹⁶Note that 2017 is the reference year for the 2018 ABS. See Zolas et al. (2020) for details.

¹⁷MOPS respondents are typically plant managers (Bloom et al., 2013).

Intelligence is a machine-based system that can perceive and learn about its environment and then make relevant predictions, recommendations or decisions for an objective that is determined by humans.” This definition is close to the legal definition of AI in the recent EU AI Act.¹⁸ Two points are worth emphasizing about this definition of AI. First, it shares with the EU AI Act the emphasis on systems that learn from their environment and adapt. Second, the last part was added because respondents were resistant to the idea that adaptive technologies might actually be choosing the objectives (more in line with “artificial general intelligence”). Including this caveat increased confidence and reduced cognitive burden for respondents, as required by the U.S. Office of Management and Budget to approve the survey.

A second feature of our approach to measuring AI is that we rely on two question “blocks” to triangulate AI use. In one, the survey asks about AI use in six business functions: production scheduling/monitoring, quality control, environmental and safety compliance, equipment maintenance, logistics, and sales forecasting. In the other, we ask about specific AI technologies, such as machine vision, speech recognition, automated guided vehicles, and AI-enabled industrial robots. This combined measurement allows us to capture two related but distinct ways in which respondents tend to be aware of AI use in their establishments, mitigate concerns that respondents might overlook AI technology “under the hood” of factory equipment or software. For instance, users of tools with AI capabilities tend to recall “we do have AI in our sales forecasting module,” or “our predictive maintenance software relies on AI.” At the same time, managers of establishments that are training their own machine learning or machine vision algorithms readily respond to that specific, more-technical terminology.

Contrary to a number of unofficial accounts, baseline adoption of “predictive” AI (i.e., not the LLMs introduced after 2021) in the U.S. has been low for some time, including in manufacturing (see McElheran et al. 2024; Bonney et al. 2024). Combining all of our AI-related questions to capture a measure of *at least some* use of AI, we estimate only about 23 percent of plants used any form of AI by 2021. Intensity of use is even lower, at around 8 percent (Panel C of Table 1).

For our regression analysis, we construct a continuous measure of AI use. First, we construct

¹⁸Specifically, the EU definition is: “a machine-based system that is designed to operate with varying levels of autonomy that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments.”

dummy indicators of whether the establishment reports using AI in specific business functions¹⁹ or in terms of adopting specific AI technologies, across the two question blocks.²⁰ We normalize these, i.e. subtract the sample mean and divide by the sample standard deviation, and then sum them into an overall index, which in turn is normalized again. This facilitates interpretation in terms of a standard-deviation change and is a longstanding practice in the related literature (e.g., Bloom and Van Reenen, 2007).

3.2.2 AI Measurement in the ABS

Our AI measure in the MOPS is even more expansive than that captured in the ABS, which focuses on a range of AI-related technologies such as machine vision, voice recognition etc. (McElheran et al., 2024), but not applications. Replicating this narrower measure in the MOPS data, we estimate AI use in 2021 at only 13% with intensity even lower at 2.3% as shown in Panel C of Table 1. Since our assessment of AI use between 2012 and 2017 relies on the ABS, we follow the narrower McElheran et al. (2024) definition for that analysis.

3.2.3 Other Variables in the MOPS

Barriers to AI use. A critical set of variables for our identification strategy includes questions on what respondents report to be the main barriers to AI use. Additionally, AI adoption barriers help us understand why AI adoption is low despite rising interest and attention.

As Panel D of Table 1 shows, the leading consideration preventing or delaying AI use was cost, which approximately 43% of plant managers consider prohibitive. In descending order of prevalence, other barriers include difficulty of identifying business use cases for AI (28%), the level of AI expertise at the establishment (12%), and employee attitudes towards AI (9%). Only about 1% of respondents see uncertainty about government regulations or industry standards as a barrier, while nearly 6% of plant managers reported no barriers to AI use and stated that AI is widely used in the establishment.

¹⁹Each dummy for business functions is one if AI is relied upon “Up to 50%”, “More than 50% but not all” and “All or nearly all” the time. The dummies are otherwise zero.

²⁰Each dummy for AI technologies is one if AI is reported to cover “Up to 50% of direct production”, “More than 50% of direct production but not all” and “All direct production”. It is otherwise zero, including if the AI technology is reported to be used in “Testing or piloting only”.

Non-AI Digitization. We directly measure other digital technologies in order to empirically separate them from AI. Closest to AI is predictive analytics (PA), defined as “statistical or algorithm-based models that analyze historical and current data to make predictions about future or unknown events.” Cognitive testing indicated that respondents found this description helped them separate out algorithms that can learn on their own from predictions based on statistical models (e.g., multivariate regression) or older rule-based algorithms. Our question on PA ask respondents to report its use across the same six business functions discussed above for AI usage. In contrast to AI, PA use is high at nearly 65% of respondents reporting at least some use, with 30% reliance on this technology. This high adoption rate of PA reinforces the notion that respondents are unlikely to overlook “AI under the hood”, as a majority of plant managers clearly recognizes PA technologies and is likely to similarly be able to recognize AI technologies.

Another potentially related technology is the use of descriptive analytics and statistics (DS) which is defined as “descriptive analyses of data . . . typically used to support making key decisions,” such as summary statistics, time trends, and real-time dashboards. Panel E of Table 1 shows that DS adoption is also very high, with 73% of establishments using at least one DS technology. Intensity is also high at 52 % of production relying on DS.

Underlying these data-intensive technologies is the background digitalization of information, defined on the survey as “information stored in digital format” and measured across the previously-described six business functions. The vast majority of plants report some level of data availability: 91% report at least some information being digitized, along with high intensities at 64%. The MOPS and ASM also allow us to measure IT infrastructure, either measured by IT capital—cumulative computer equipment expenses over the last 3 years—or cloud computing expenses.

Management Practices. In addition to digital technologies, the MOPS also includes a comprehensive list of variables to measure structured management practices. These measures have been extensively discussed by Bloom et al. (2013; 2019) in the Census data and more-generally by Scur et al. (2021; 2024), so we refer the interested reader to this work.

3.2.4 Business Strategy Measures in the ABS

A particularly useful set of variables from the 2018 ABS are questions about the business strategies that firms pursue. Prior work has shown systematic relationships between business strategies and

AI adoption (McElheran et al., 2024), technology use (Wu et al., 2020) firm organization and innovation outcomes (Yang et al., 2015), and firm performance responses to competitive shocks (Yang et al., 2021). The ABS asks respondents about the importance of 14 different business strategies, from a focus on “introducing new products or services” and “expand into new markets” (domestic and international) to “low price” and “reduction costs.” We use principal component analysis to extract three factors. The first factor captures the intensity of pursuing growth through new markets, including expansions into new (domestic or international) existing markets or creation of new markets through innovation, and we call it “new markets strategy.” The second factor is about pursuit of growth via cost leadership, close in spirit to Porter’s cost leadership strategy (Porter, 1980). The third factor captures the intensity of pursuing satisfaction of current main customers via excellent product and service quality. We call this “quality strategy,” and it is close in spirit to Porter’s “differentiation strategy.” These three generic business strategies are similar to ones found by Yang et al. (2015) in the Workplace and Employee Survey (WES) in Canada.

3.2.5 The Implications of “Shadow AI”

Although the timing of our survey precedes the introduction of generative AI tools in late 2022, the issue of “shadow AI” is an increasing concern when it comes to measuring AI diffusion and its impacts. The worry is that there could be AI use taking place in the firm, either undetected in machinery or software— or, as seems to be common with the rise of LLLMs, in use by individual employees— without managerial knowledge. We think the risk of this is lower with the Industrial AI applications studied here (compared to newer generative AI technologies). That said, it is worth considering the implications for our analyses. Ultimately, widespread “shadow AI” might bias downwards our adoption statistics. However, if AI use is prevalent and undetected in our sample, this should work against finding any significant productivity differences between firms based on observed AI usage.

4 Endogeneity and Identification Strategies

We are interested in regressions of the form

$$y_i = \beta AI_i + C_i\gamma + \epsilon_i \quad (1)$$

where i indexes different establishments, and the dependent variable y_i is typically (log) value added or profits, defined as value of shipments minus salaries/wages, employee benefits, cost of work-in-progress, raw materials, intermediate inputs, fuels, purchased energy, services, and resales. C_i denotes a matrix of control variables typical in productivity regressions, including input controls such as number of employees and non-IT capital stock, as well as other controls described below. The key endogeneity problem for regressions such as (1) is selection bias (e.g. Angrist and Pischke, 2009), where better performing or better managed businesses will tend to use AI more intensively. This will bias upwards any effects estimated via ordinary least squares (OLS). We pursue three identification strategies to address this selection bias.

Our first identification strategy is the use of an unusually rich set of proxy variables to control for typically-unobservable organizational factors (e.g., managerial skill or digital capabilities) and can therefore be understood as a matching, or “selection-on-observables,” identification strategy. As mentioned, the particular strength of the MOPS–ASM data is the availability of four sets of variables that are rarely available at scale. The first set includes controls for data-intensive technologies that are not AI. Similar to the AI index, we construct continuous indices for PA and DS, which we use as controls throughout. The second set captures IT infrastructure, both on-premises IT capital and cloud computing expenses, both of which have been closely linked to the use of prediction technologies and AI (Brynjolfsson et al., 2021a; McElheran et al., 2024). Our third set includes indices for structured management as in Bloom et al. (2019), as well as controls for the percent of college-educated employees and union membership. These directly proxy for the quality of management as well as for worker skill and organization. The fourth set includes indicators for the establishment being part of a multi-unit firm (MU), being co-located with the firm headquarters (HQ), or having one of three different production-process designs: job shops, batch production, or continuous flow manufacturing. These quasi-fixed organizational design characteristics have been

associated with business performance and technology adoption in prior work (e.g., Atalay et al., 2014; Bloom et al., 2013; McElheran et al., 2019; Brynjolfsson et al., 2021a).

A shortcoming of the selection-on-observables approach is that it cannot account for unobserved factors that are correlated with performance but uncorrelated with our controls. We therefore leverage a natural experiment available in the MOPS–ASM data to address unobservable confounders. Specifically, we use the reported perception of “no lack of AI expertise” as an instrument for more intensive AI use. This IV is relevant for AI adoption and plausibly exogenous: lack of AI expertise does not necessarily imply lack of manufacturing operations expertise. The logic underpinning our approach proceeds in steps.

First, we build on recent work on AI use by firms that has argued extensively that AI-related skills are distinct, technology-specific, and sought by firms that seek to adopt AI (Acemoglu et al., 2018; Babina et al., 2024). A specific question on the 2021 MOPS asks about the relationship between AI adoption and AI-specific skills. We document the strong correlation between lack of a reported AI skills gap and AI adoption, below.

The exclusion restriction, however, will be violated if AI-specific skills are correlated with other human capital at the plant that also affects productivity. Thus, a few nuances of our setting and our approach are worth keeping in mind. First, the relative recency of AI applications (as opposed to inventions) suggests that many competently-run organizations will not have recruited plant-managers for AI skills, specifically—particularly if other barriers (such as lack of obvious use cases) are predominant. In this context, the vast majority of our respondents are plant managers with tenure exceeding 5 years (Bloom et al., 2019), which makes it very unlikely that AI skills important consideration when they were hired. For example, managers hired in 2016 (5 years prior to the sample year of 2021), we unlikely to be hired for their AI expertise as machine vision had just begun to outperform human vision in 2015²¹.

Second, lack of AI skills by plant managers may be offset by access to AI skills by other employees, such as engineers. We address this concern to a large extent by directly controlling for the percent of workers at the plant with a formal Bachelor’s degree. Use of structured management practices is also controlled for (Bloom et al., 2019). For these reasons, we assume that an in-house, AI-specific skills gap will increase the costs of AI adoption while being conditionally uncorrelated

²¹With ImageNet achieving an error rate of 5% narrowly beating the average human error rate of 5.1%

with the overall performance of the plant. In addition, we follow Angrist and Pischke (2009) and evaluate the plausibility of the exclusion restriction by showing, for a subset of establishments with “zero first-stage” effects, that reduced-form estimates exhibit zero effects as expected.

Since exclusion restrictions in IV estimation can never be directly tested, we leverage a third identification strategy based on the ABS-CMF panel data. Specifically, as documented by McElheran et al. (2024), AI adoption by U.S. firms by 2017 was very low. Maintaining the assumption that AI adoption as of 2012 (the closest year for the full Census of Manufactures, which maximizes our balanced panel sample) was essentially nonexistent²² allows us to estimate regressions of the form

$$y_{f,t} = \beta \cdot AI_{f,t} + C_{f,t}\gamma + D_f + \epsilon_{f,t} \quad (2)$$

where f is a manufacturing firm, t indicates different years (2012 and 2017), and D_f denotes a full set of firm fixed effects. We thus use within-firm changes to control for permanent or slowly changing unobservable confounders. As before, we use firm productivity (labor productivity and TFP) to measure value creation and profits to measure value captured.

It is useful to remember that first difference approaches like (2) will exacerbate attenuation bias of coefficients towards zero as a result of “over-differencing” (Angrist and Pischke, 2009), so, setting aside potential upward bias due to selection (which is unaddressed in this approach), we expect the estimates for this specification to be smaller in magnitude compared to the IV estimates, as IV automatically corrects for classical measurement error.

5 Results

5.1 Correlates of AI Use

In this section we examine the variables that help predict AI adoption in our MOPS-ASM sample. Controlling for size, which is a well-established correlate of AI uptake (McElheran et al., 2024; Calvino and Fontanelli, 2023), industry (at the 3-digit NAICS level), multi-unit and headquarters status of the plant, production-process design, and the prevalence of bachelor’s degrees among employees, there remains a great deal of heterogeneity in AI adoption among U.S. manufacturers.

²²This identifying assumption is similar in spirit to that used by Forman et al. (2012) and Forman and McElheran (2025) to study the impact of the commercial internet in the 1990s.

5.1.1 Other Technologies

The first two rows in Table 2 demonstrate that AI is distinct from and co-varies with the presence of other digital technologies. The descriptive analytics index, DS, is negatively correlated with AI use, consistent with substitution between AI and less-sophisticated uses of digital information such as summary statistics and descriptive dashboards. In contrast, the index of predictive analytics, PA, is positive and significantly related to AI use. However, the overlap between AI and PA is far from perfect. Quantitatively, a standard-deviation increase in the PA index is only associated with a 0.39 standard deviation increase in the AI index. These results underscore the importance of disentangling different digital technologies from each other and taking a nuanced approach to understanding their interactions with each other and the broader organizational context.

The next two rows in Table 2 document the correlation between IT infrastructure and AI use. The data allow us to separate IT infrastructure into on-premises IT equipment, such as computers and servers, and cloud computing, which have been shown to differently affect the survival and productivity of firms of different ages (Jin and McElheran, 2024). The results across specifications indicate that AI use is systematically correlated with cloud computing but not with on-premises IT capital, consistent with potential benefits of cloud-connected AI systems hypothesized by Pratt (2015) and the technological interdependencies conjectured in McElheran et al. (2024).

5.2 Organizational Complements

A key test for complementarity between technology and organizational characteristics (Milgrom and Roberts, 1990) is the “correlation test” for co-adoption of technology and key organizational features (Brynjolfsson and Milgrom, 2013; Brynjolfsson et al., 2021a). The next rows in Table 2 provide insights along these lines.

Rows 5 and 6 provide nuanced insight into how structured management practices may affect AI adoption. We find a positive relationship only with *process*-focused Lean management practices. Plants with greater managerial attention to key performance indicators, production targets, and other practices related to monitoring and managing production activities report a higher AI index. In contrast, structured *people*-focused management practices are negatively associated with AI use. Tying compensation to production targets, promoting employees primarily based on performance

(rather than tenure) and quickly firing or re-assigning under-performing workers is associated with less AI uptake. This is controlling for unionization, which correlates positively with AI use (see row 7 of Table 2).

Other organizational interactions are less surprising (and thus exact estimates are not disclosed as of this writing). Age is negatively correlated with AI use, although the coefficient is only statistically significant in the richer specification in column 3 (row 8). AI use is also more prevalent alongside greater (non-IT) capital investment (column 3, row 9). This is worth noting as capital stocks are often missing or mismeasured in standard data sets (e.g., Bryzgalova et al., 2025). Statistically, this will cause unobserved returns to capital equipment to load onto correlated observables—which we show here includes AI—potentially biasing upward the estimated returns to AI use (more on this, below).

5.3 Potential for Reverse Causality

While examining potential drivers of AI use, we can probe how earlier plant performance is conditionally associated with 2021 AI usage. Columns 4 and 5 of Table 2 document that prior growth (change in value-added from 2019 to 2021) and 2019 labor productivity are either insignificantly or negatively correlated with AI use. This runs counter to standard concerns in productivity studies (i.e., that reverse causality would lead to a positive correlation between AI use and performance; see, e.g., Brynjolfsson and Hitt, 2000).

5.4 IV Relevance

A key question for this analysis is the relevance of the instrumental variable. Based on the importance of “lack AI expertise” as a reported barrier to AI adoption in panel D of Table 1, we reverse-code this question to construct an indicator equal to one if “lack of AI expertise” is *not* reported to be a barrier. As the final row of Table 2 shows, this proposed instrument significantly and robustly predicts more intensive AI adoption. Plants where a lack of AI expertise at their establishments is not flagged have slightly over 8% higher AI usage, on average, conditioning on a wide range of other factors.

5.5 AI Use is Associated with Initial Performance Losses

5.5.1 OLS Estimates in the 2021 MOPS

We turn next to systematically exploring the relationship between AI use and performance. Panel A of Table 3 provides a cumulative regression build of OLS regressions of value added on AI and increasingly rich sets of controls to examine near-term plant performance in the 2021 MOPS data. The first column indicates that, absent other controls, AI adoption appears significantly correlated with (log) value added, even controlling for the presence of other data-intensive analytics tools (DS and PA). This continues to be true when we control for IT infrastructure in column (2). However, once we control for size and industry in column (3), along with multi-unit status and energy expenditure, the precision of the estimate improves while its magnitude falls to point of statistical insignificance. Columns (4) and (5) add controls for structured management practices and other organizational characteristics including plant age, resulting in a negative yet statistically insignificant correlation between AI use and labor productivity. Finally, adding capital stock controls in column (5) allows us to analyze whether more AI-intensive plants tend to have higher or lower total factor productivity (TFP). Here, the negative coefficient is much larger and statistically significant at the five-percent level., indicating that plants with a standard-deviation higher AI index tend to be 1.33% *less* productive. The conditional correlation between AI use and profits is small and noisy in a similar specification (column 7).

A common concern regarding OLS performance estimates is the potential for bias due to selection bias or plant-level unobservables. Typically, the concern is that larger, better-managed firms with richer organizational complements will be both more likely to adopt a new technology and more likely to gain from its use. This is often referred to as “selection on gains” in related technology and economics studies (e.g.,(Bloom and Van Reenen, 2007)). It should be noted that in our context this concern will bias the results against finding a negative effect of Industrial AI adoption. One might therefore expect that well-identified causal effects may exhibit much more negative productivity effects of AI adoption.

5.6 IV Results

Next we employ quasi-experimental methods to address common endogeneity concerns regarding OLS estimation. Before turning to these results, it is important to underscore that a hypothetical natural experiment “treating” firms randomly with AI technologies would not necessarily yield a realistic estimate of the average treatment effect (ATE) of AI on performance, due to the importance of adjustment and co-invention in realizing gains from transformative technologies (Bresnahan and Greenstein, 1996; Bresnahan et al., 2002), leading to large treatment effect heterogeneity (Angrist et al., 1996).

Our aim in developing a novel IV strategy is not to derive definitive “causal” estimates, but rather to add to a collage of evidence depicting the heterogeneous impacts of early AI use in industrial production. These results are best interpreted as the local average treatment effect (LATE) of AI use among firms that had a greater propensity to adopt AI early for reason conditionally unrelated to anticipated gains. This section discusses these estimates along with a number of tests of the identifying assumptions underlying the approach and exploration of key drivers of heterogeneity.

The first column of Panel B in Table 3 reprises findings in Table 2 that the instrument of “no lack of AI expertise” is indeed relevant. Columns (2) and (3) report second-stage IV results, with column (2) estimating the impact of a one standard deviation higher AI index on labor productivity while column (3) estimates the impact on TFP (i.e., revenues controlling for production inputs including capital stock). The estimated effects are quantitatively very large. Focusing on TFP, column (3) indicates that a one standard deviation increase in AI reduces TFP by roughly 44% ($-0.44 = \exp\{-0.587\} - 1$). Although these effects are substantially larger than the OLS coefficient reported in column (6) of Panel A, these effects should be interpreted with a few things in mind.

First, selection on gains would predict that firms with a high correlation between productivity and AI adoption will not be sensitive to the instrument (“Always Takers” in the treatment effects terminology, e.g., Angrist et al., 1996; Angrist and Pischke, 2009), so this approach plausibly strips out the firms with the highest returns to AI use. Second, OLS is well known to suffer from attenuation bias when classical measurement error is sizable, which is both likely in the context of fast-changing AI technologies (i.e., they are inherently hard to define and measure) and addressed by IV estimation. Third, the quantitative implications of our IV results on TFP may be

usefully situated in prior research on productivity drivers in firms. One reference for well-estimated causal effects on productivity concerns the adoption of structured management on establishment TFP from a field experiment in Indian manufacturing from Bloom et al. (2013). They suggest that a one standard deviation higher structured management score increases plant TFP by 60% ($0.59 = \exp\{0.49\} - 1$). These productivity effects are of the same order of magnitude as the negative TFP effects estimated here. This sheds important light on the experience of an important subset of firms that were early adopters of AI for conditionally uncorrelated reasons.

Column (7) of Panel A and column (4) of Panel B in Table 1 report the relationship of AI adoption and profits. Although positive and quantitatively small in OLS, our IV estimates suggest that a one standard deviation more AI use causes a loss of about \$11 million for the average manufacturing establishment in our sample. Given very similar average profits in the estimation sample (see Table 1, Panel A), this is a quantitatively important magnitude and underscores the short-term risks of industrial AI adoption.

5.6.1 IV Robustness

All econometric identification strategies embed tradeoffs, and the central concern of any IV analysis is whether failure of the exclusion restriction could bias the estimates. For example, the early timing of adoption notwithstanding, one could worry that lack of AI expertise within the plant leading up to adoption by 2021 could indicate of a broader lack of operational capability (again, despite our direct controls for management practices and employee education), which could directly reduce establishment productivity. In this case, establishments with no lack of AI expertise should directly demonstrate higher productivity than those reporting such a barrier. However, as the first column of Panel A in Table 4 shows, the opposite appears (and is statistically significant) when we regress labor productivity directly on the instrument.

A related selection-based bias could arise if the response of “no lack of AI expertise” were indicative of plant managers who expect disproportionate gains from AI. If true, this would imply that respondents that instead reported “AI is widely used” in the same question block should exhibit systematically higher AI-related productivity. We therefore use the response “AI is widely used” as a placebo instrument, for which we report the first stage in the second column of Panel A in Table 4. As expected, this response is strongly positively correlated with AI use in the first stage.

However, as column (3) shows, the second stage effect is statistically indistinguishable from zero, albeit positive. This is consistent with plants led by managers who expect higher returns to AI being more likely to adopt in the first place. However, they cannot rule out that the realized costs exceed anticipated gains, at least in the short term. Stories of big AI and automation implementations needing to be reversed due to unanticipated performance problems are an increasing feature of recent news cycles²³. Note, further, that the sign and significance of the placebo effects differ from our main IV results, ruling out a mechanical relationship driven by this “barriers” question block.

To further probe the plausibility of the exclusion restriction, we follow Angrist and Pischke (2009) in using a “zero first stage” approach. For this, we require a subsample of plants for which no lack of AI expertise nevertheless fails to lead to more AI adoption. In other words, in such a subsample, the instrument is irrelevant to being treated. If the exclusion restriction holds, one should then observe no correlation between the eventual performance outcome and the instrument, since the exclusion restriction requires that the instrument *only* impact performance via the treatment (in our case, AI use). There are good theoretical reasons to believe that the IV is irrelevant for at least two particular subsamples of manufacturing establishments.

The first is “job shops,” which flexibly make a variety of made-to-order products in small batches or even one-off prototypes. Products and production parameters will typically vary widely from order to order. In this high-mix, low-volume context (Hayes and Wheelwright, 1979), uncertainty is high by design. An AI technology such as machine vision will typically not have enough data or repeated use cases to train and leverage prediction algorithms. Previous work on the use of predictive analytics found a similar constraint on prediction’s benefits in these production environments (Brynjolfsson et al., 2021a).

The second subsample is single-unit (SU) establishments. Systematically smaller than multi-unit plants, they will be less likely to adopt AI systems due to lack of scale (Svanberg et al., 2024).

For both of these subsamples, we find that the instrument fails as predicted in the first stage, as documented in columns (1) and (3) of panel B of Table 4. Importantly, not only are the IV first stages for the subsamples statistically insignificant, they exhibit opposite signs compared to the first

²³Swedish firm Klarna made headlines for replacing large numbers of employees with AI chatbots, only to have to hire them back because of low quality results (<https://www.techradar.com/pro/over-half-of-uk-businesses-who-replaced-workers-with-ai-regret-their-decision>). Tesla also famously confirmed that challenges associated with automation hindered vital Model 3 production, echoing earlier challenges with automation in car manufacturing (<https://www.iqsdirectory.com/resources/teslas-big-problem-excessive-automation.html>).

stage in column (1) of Panel B in Table 3. Further, in each of these cases, the reduced-form labor productivity regressions (columns (2) and (4) of panel B) show a quantitatively small, statistically insignificant, and negative relationship between AI use and. This is exactly the result one would expect if the IV exclusion restriction holds.

5.6.2 Understanding LATE

Given the large magnitude of the IV estimate, we are interested in better understanding treatment heterogeneity and how broadly we should extrapolate from the local average treatment effect, or LATE (Angrist and Pischke, 2009). First, we explore heterogeneity based on exogenous organizational characteristics. Consistent with prior work linking IT performance to product mix and volume (McElheran and Jin, 2020), we find that the IV estimates vary across production designs. Specifically, AI-related productivity is higher in continuous-flow plants, where production mix focuses on few products in service of higher volume. Plants with this type of stable (i.e., predictable) production exhibit less of a *J-curve* dip (Column 2, Panel A of Table 5). This lends indirect support to the exclusion restriction, as these plants tend to be relatively capital-intensive, more-intensively managed, and relatively higher-productivity than other plant types (McElheran et al., 2019). It further underscores the fit of AI for prediction-friendly production environments, as discussed in section 2.1.3.

In contrast, older establishments show a more-pronounced negative LATE (Column 3, panel A, Table 5), consistent with our discussion in section 2.4. We return to the challenges of older businesses, below.

To understand what types of plants are sensitive to the instrument and therefore driving the IV results (the “Compliers” if we had a binary treatment, per Angrist et al., 1996), in panel B of 5 we report the odds ratio for various firm characteristics among the population of plants for which “no skills gap” is associated with being above the median AI-index. These plants, which represent about 3% of the analysis sample (recalling the low average adoption rates) should be thought of as the “marginal” adopters: they would not adopt Industrial AI, “but-for” plant managers’ view of no AI-related skills gap. To better understand marginal AI adopters conceptually and interpret our results, it is useful to contrast Compliers with two other groups of establishments, following the logic of Angrist et al. (1996). On the one hand, “Always Takers” are infra-marginal establishments

whose perceived net benefits of AI adoption are so high that they would adopt this technology, irrespective of whether they happen to have sufficient AI expertise right now. In other words, they are willing to incur the additional costs of acquiring the necessary skills if needed. In terms of observable characteristics, we would expect Always Takers to have the digital inputs and cloud computing infrastructure previously found to be associated with AI use (Goldfarb et al., 2023; Calvino and Fontanelli, 2023; McElheran et al., 2024). We further would expect them to benefit from the scale-effects of digital technologies (Giustiziero et al., 2023), exemplified by being part of large multi-unit firms, by using e-commerce to ship their goods nation-wide, or by deploying AI technologies at their headquarters (HQ) to coordinate other production units. The adjustment costs might further be lower at HQ, all else equal, if proximity to firm leadership and managerial capabilities is greater in these locations.

On the other hand, “Never Takers” are least likely to adopt Industrial AI—even if they happen to have sufficient AI expertise—due to insufficient anticipated net benefits. Empirically, we would expect them to be the opposite of Always Takers: not digitalized, not reliant on cloud computing, not using e-commerce, single-unit in structure or, if not, located away from HQ within multi-establishment firms.

As a result of this characterization of Never Takers and Always Takers, Compliers should be between these extremes: they will demonstrate some characteristics of Always Takers while resembling Never Takers on other dimensions. Panel B of 5 shows exactly this. More like Always Takers, Compliers are 33% more likely than the average plant to ship all their goods through e-commerce, 23% more likely to exhibit above-median cloud use, and 42% more likely to be part of a multi-unit firm. At the same time, Compliers are similar to Never Takers on the dimensions of digitalization and headquarters status. They are 60% less likely to be highly digitalized than the average plant and 48% less likely to be co-located with HQ.

Our empirical profiling of Complier plants can help to shed light on why the LATE from our IV analysis is so negative. Specifically, despite these plants exhibiting characteristics that make them likely to eventually benefit from Industrial AI adoption (such as multi-unit status), being less-advanced in their digitization journey likely diminishes the quantity and/or quality of training data to hand, leading to worse short-term performance. The lower likelihood of being co-located with HQ may further represent uncertainty about the system-wide benefits or lack of organizational

support, again negatively impacting short-term returns to AI use above and beyond the average adjustment costs observed in the broader population.

5.7 Within-Business Changes

Despite the plausibility of our IV exclusion restriction, exclusion restrictions can never be conclusively tested unless the instrument is generated by an actual randomized controlled trial (RCT). We therefore pursue yet another identification strategy to triangulate on the causal effects of AI use on business performance. Leveraging the panel structure of the ABS-CMF data, we estimate equation (2) with either plant or firm fixed effects to control for time-invariant organizational confounders.

The first three columns of Panel A in Table 6 show that more-intensive AI adoption between 2012 and 2017 is associated with declining sales, lower TFP, and lower profits. The magnitudes of these effects are all an order of magnitude smaller than effects estimated via IV. This is unsurprising, since any type of classical measurement error (discussed above) will be magnified in simple first-difference estimators, attenuating their magnitude compared to the IV results. In addition, this approach “over-controls” for slow-moving organizational characteristics such as those discussed in the complier analysis, above. Finally, the effect reflects more time for adjustment among any plants that adopted early in the five-year window. While we do not directly estimate how long it takes for the *J-curve* to reverse direction or net out, overall, some improvement amongst survivors in this more medium-term analysis is to be expected.

Those caveats aside, the sign of the effects is consistent with both the OLS estimates and the IV estimates from the MOPS–ASM data. AI adoption is associated, again, with short-term within-firm productivity declines. In these analyses, the magnitude of a one standard deviation higher AI use comes with a performance loss of around 2% ($-0.019 = \exp\{-0.02\} - 1$).

The panel data also allow us to investigate the dynamic patterns predicted by the *J-curve* theory, expressed in section 2.2. Specifically, we can analyze the longer-term growth and performance outcomes of manufacturing firms as a function of their AI use in 2017. For this, we track growth in terms of employment, sales, and labor productivity in the Longitudinal Business Database (LBD) from 2017 to 2021.

Columns (1) to (3) in Panel B of Table 6 show growth in employment, sales, and labor productivity to be significantly higher for manufacturing firms that deployed AI in production (the

ABS definition of adoption) by 2017 and persisted in our sample through 2021. These positive growth effects are consistent with section 2.2 (and *J-curve* arguments more generally), that initial performance losses should be considered investments in co-investment and intangible capital accumulation that yield returns in the longer run.

Columns (5)–(7) of Panel B in Table 3 report additional IV estimates that reinforce this adjustment explanation and deepen our understanding of the causal effect. Motivating the specification in column (5) is recognition of how important “Lean” production has become in U.S. manufacturing (e.g., Womack et al. 2007), as discussed in Section 2. Key characteristics of Lean include purposefully driving down inventory in the production process so as to make problems and defects more visible, to force responsiveness to process exceptions, and to improve (shorten) the lead times and inventory carrying costs of the entire system (e.g., Holweg 2007). It also tends to be a very “pull-based” and often low-tech approach to operating a production process, focused more on responsiveness to demand and preventing defects than buffering for them (Milgrom and Roberts, 1988). As such it is arguably not a straightforward fit for more-digitized methods associated with sophisticated prediction, make-to-forecast, and ultimately AI. Disrupting established “analog” processes rooted in the Lean tradition should therefore become visible in other metrics such as inventory levels and carrying costs. We find empirical evidence consistent with this in column (5) of Panel B in Table 3, which shows that WIP systematically increases in response to AI adoption.

Another consideration is that AI systems might be replacing manual, human-led activities with physical automation. If this is the case, then early adjustment should also manifest as increased investment in production capital—especially industrial robots. Robot use and reliance on AI in production have been correlated in prior work (McElheran et al., 2024), but the timing of adoption and causality remain poorly understood. Column (6) of Panel B in Table 3 shows that the number of active and purchased industrial robots increases as a result of AI adoption. Again, this is consistent with significant co-invention in production processes due to moving to reliance on AI.

Finally, AI use causes the number of workers to decline, as shown in column (7) of Panel B in Table 3. While this is, at first glance, consistent with substitution of human labor with automation (Autor and Salomons, 2018), it is a short-term result that must be interpreted in the context of employment growth over time. We do not observe which workers are shed, nor which ones contribute to employment growth over time. We note, however, that this pattern is consistent with a need to

shuffle labor inputs as businesses shed prior practices and reconfigure their operating systems.

5.8 The Role of Business Strategy

In this section we investigate what firm characteristics might impact how AI shapes competition. We first begin with generic business strategies. Following section 2.3, we expect that the strategic decisions of how firms pursue growth matter for the initial performance decline of their *J-curve*.

Column (4) of Panel A in Table 6 yields results consistent with this hypothesis. Specifically, firms with more strategic emphasis on growth through market expansion and innovation (“new markets strategy”) tend to exhibit significantly lower initial productivity losses as seen in the positive and significant interaction term. This result is consistent with early benefits from scale economies helping to reduce initial foregone output from adopting AI. At the same time, pursuing scale through cost leadership does not work as well, as shown by corresponding interaction term in column (4) of Panel A in Table 6. Although we do not find statistically significant results, the negative sign of this effect is consistent with the view that growth through cost leadership still implies demand-side dis-economies of scale, which limit how much firms can avoid initial *J-curve* losses from foregone output.

We now move to the tests of our predictions from section 2.2, which predicts that old firms will exhibit greater performance losses from AI adoption because they disproportionately lose productivity benefits from prior accumulated experience.

Results in Panel A of Table 6 confirm this intuition. Column (5) of Panel A shows that after controlling for the interaction of firm age and AI adoption, the main effect of AI adoption on productivity is actually positive and statistically significant. The last two columns of Panel B in Table 6 build on this analysis and estimate growth outcomes from 2017 to 2021 in the LBD as a function of AI adoption and firm age. These estimates show that older firms systematically benefit less in the wake of AI adoption and that once the interaction of AI adoption and firm age is controlled for, other firms also gain substantially more in terms of sales growth and labor productivity.

5.9 Mechanisms

In this section we dig into the specific channels underlying the age effects documented in the last section as well as the implications of AI as a scale-free resource within large multi-unit firms.

5.9.1 What is Driving Age Effects of AI Adoption?

We begin by documenting that age effects documented for the ASM-CMF data are indeed also present in the MOPS-ASM data. For purposes of this analysis, we define old establishments as establishments that are at least 25 years old, with the average plant age being almost 29 as shown in Table 1.

The first two columns of Panel A in Table 7 show that the negative effects of AI adoption on labor productivity (in column 1) and TFP (in column 2) are either substantially larger in old establishments, or even completely driven by old establishments.

We then move to exploit a key strength of the MOPS data to investigate whether AI adoption causes a change of structured management practices. Column (1) of Panel B in Table 7 shows that more intensive AI adoption causes a de-adoption of structured management practices, concentrated in old establishments. Quantitative magnitudes of the de-adoption of structured management are large: estimates from column (1) of Panel B in Table 7 imply that a one standard deviation in AI intensity reduces structured management by 0.83 standard deviations. This result is meaningful in at least two ways. First, de-adoption of structured management practices is likely to have a direct effect by reducing the quality of management. Second, de-adoption of structured management practices is indicative of a broader loss of production experience/know-how and “organizational capital” (Levitt et al. 2012) as argued in section 2.5.

To quantify the direct productivity effect of the de-adoption of structured management, we can combine the estimate of column (3) with estimates of the causal effect of structured management on TFP from Bloom et al. 2013. Doing so results in an estimated TFP loss of 32% ($-0.32 = \exp\{-0.83 \times 0.477\} - 1$) at old establishments. Alternatively, one can use the OLS regression estimates of TFP on structured management by Bloom et al. (2019) of 0.209 in combination with the finding in that paper that classical measurement error attenuates the OLS coefficient by 1/2 to obtain a measurement error corrected coefficient of 0.418 ($= 0.209 \times 2$), which then implies a

productivity loss of 29% ($0.29 = \exp\{-0.83 \times 0.418\} - 1$). These TFP losses should be compared to the TFP effects at old establishments shown in column (2) of Table 7, which indicate a 68% TFP loss ($0.68 = \exp\{-1.133\} - 1$). In both cases, this mechanism alone can explain roughly half (32%/68%) of the TFP loss from AI adoption in column (2) of Panel A in Table 7 for old establishments. Another way to restate this result is that the residual loss of experience/organizational capital could potentially explain up to half of the measured TFP loss due to AI adoption.

Additionally, the MOPS data allows us to more fully understand which of the structured management practices are affected in particular. Indeed, columns (2) –(4) of Panel B in Table 7 show that the de-adoption is driven by production monitoring and targeting practices and especially by de-adoption of “KPI reviews by non-managerial employees” and the degree of “Target awareness across employees,” which is higher the more non-managerial employees are typically aware of production targets. These results are consistent with old establishments removing operations practices that previously helped them to utilize frontline employee knowledge (section 2.5). This is also consistent with the results from the last two columns of Panel A in Table 7, which show that more intensive AI adoption causes reduction in employee wages and overall workforce, consistent with the replacement of frontline employees with industrial AI systems such as predictive maintenance and AI-enabled industrial robots.

5.9.2 Within-Firm Spillover Benefits of AI Adoption

Section 2.6 suggests large multi-unit establishments might mitigate some of the direct losses from AI adoption by indirectly benefiting from positive spillovers across establishments within the same firm. We investigate this mechanism by deriving a second instrument from our first for multi-unit firms. Specifically, we construct an AI measure for AI adopted at other establishments of the same firm and call it “elsewhere AI index” (EAI). This EAI index increases with the average AI index for other establishments within the same firm, excluding the focal one. In other words, it is a “leave out mean” AI index with the mean being calculated across all establishments belonging to the same firm as the focal establishment but excluding the focal establishment. We then use the number of plants outside the focal plant but inside the same firm which reported no lack of AI expertise as instrument for EAI.

Columns (1) and (2) of Panel C in Table 7 shows that the system of equations of two IV first

stages confirm that both instruments are relevant. Column (3) of Panel C in Table 7 shows that the causal effect of AI adoption at the focal plant is still negative, but that AI adoption at other plants within the same firm has a positive and sizable causal effect.

6 Conclusion

This paper examined the performance impact of AI on U.S. manufacturing over different time horizons and across diverse organizational settings. Using two distinct but related datasets on tens of thousands of establishments and firms, we document three key findings. First, early industrial AI adoption in U.S. manufacturing robustly causes statistically and economically significant productivity losses in the short run. Second, however, we find evidence for recovery and growth in the longer run. Together these findings are consistent with *J-curve* mechanisms that have not heretofore been established at the micro-level.

Last but not least, we unpack heterogeneity in these effects and their implications for competitive dynamics. Productivity losses are driven by old establishments and old firms. This suggests that industrial AI adoption does not favor incumbents over entrants, at least not in our context of industrial AI in U.S. manufacturing. In addition, we find that growth-oriented startups benefit disproportionately, suggesting that Industrial AI use may tend to promote “creative destruction” in the sector, as AI diffuses. Further, we provide direct evidence on the quantitative importance of a particular causal mechanism leading to the AI-related decline in establishment productivity: the displacement of structured management practices (Bloom et al. 2012; Bloom et al. 2019) and obliteration of knowledge management systems combining worker feedback and structured practices (Grant, 1996; Levitt et al. 2012). The labor shedding we observe is likely linked to this adjustment process, which also manifests in WIP and increased robot investment

Our empirical analysis has several limitations that constitute opportunities for future research. The first is that industrial AI diffusion remains in its early stages. As AI matures as a technology or co-invention know-how also diffuses, productivity losses from initial adoption might be better balanced by productivity gains. Further, our estimates of dynamic effects are not as well-identified as the baseline short-term losses and probably suffer from some remaining (upward) selection bias. So, the long-term effects could be less rosy than portrayed in these estimates. It will be useful to

continue gathering empirical evidence on this question as AI-related technologies continue to evolve and diffuse.

A second limitation of this study is the focus on (U.S.) manufacturing. This focus on a given industry context allows us to better understand and interpret the results. But if AI is really a GPT, then it will be widely adopted across many different industries and sectors of the economy. Future work should therefore continue to explore other business contexts to better understand and evaluate the effects of AI on productivity and workers across the economy.

Finally, more work is needed on the specific uses to which AI-related technologies are being put by firms. The low baseline adoption, combined with Census’ disclosure avoidance rules prevents us from providing more-granular breakdowns of AI use and potentially differing performance implications. Yet a hallmark of GPTs is their broad applicability. Better understanding specific applications of AI in specific contexts is needed to understand other dimensions of heterogeneity, trace dynamics, and “flatten” the curve for more organizations and industries.

Table 1: Summary Statistics

	(1) Mean	(2) Std. Dev.
<u>Panel A: Establishment Characteristics (MOPS–ASM)</u>		
Number of employees	171.5	263.4
Age of establishment	28.54	14.23
Percentage of shipments through e-commerce	62%	42%
Number of active and purchased industrial robots	1.73	6.96
Establishment-level profits (in thousands)	11,000	16,000
Change in work-in-progress inventory (in thousands)	445	6370
<u>Panel B: Firm Characteristics (ABS–CMF)</u>		
Number of employees	344.3	7,678
Age of firm	24.85	12.59
Percentage of shipments through e-commerce	23.5%	37.7%
Establishment-level profits (in thousands)	19,500	42,000
Early adopters	7.5%	9.1%
<u>Panel C: AI Adoption (MOPS–ASM)</u>		
Any AI	22.8%	41.9%
Any technical AI application	12.6%	33.2%
Production using AI	8.0%	20.4%
Production using technical AI applications	2.3%	9.1%
<u>Panel D: Barriers to AI Adoption (MOPS–ASM)</u>		
No applications or business use cases	28.4%	45.1%
Regulation uncertainty	1.0%	10.1%
No expertise	12.3%	32.8%
Cost	43.2%	49.5%
Employee attitudes	9.4%	29.2%
No barriers	5.5%	22.9%
<u>Panel E: Other Digital Technologies (MOPS–ASM)</u>		
Any predictive analytics	64.6%	47.8%
Any descriptive statistics	72.8%	44.5%
Any digitalized data	91.0%	28.6%
Expenditures on cloud computing (in thousands)	7.63	33.21

Notes: Panel A is using the MOPS–ASM sample and is comprised of roughly 28,500 manufacturing establishments. Panel B is using the ABS–CMF data and consists of roughly 55,000 manufacturing firms. Panels C, D, and E report weighted summary statistics based on data from the MOPS–ASM, combined with sample weights that make estimates representative for roughly 300,000 manufacturing establishments in the U.S. The AI, descriptive statistics, and predictive analytics indices are normalized, with zero mean and unit standard deviation, which is why we do not report summary statistics for these indices here.

Table 2: Correlates of AI Use by U.S. Manufacturers
MOPS—ASM: Establishment-level

	(1)	(2)	(3)	(4)	(5)
DS Index	-0.0504*** (0.0177)	-0.0506*** (0.0177)	-0.0526*** (0.0177)	-0.0506***** (0.0177)	-0.0502*** (0.0176)
PA Index	0.386*** (0.0119)	0.385*** (0.0119)	0.385*** (0.0118)	0.385*** (0.0119)	0.385*** (0.0119)
Log IT Capital	-0.0007 (0.0032)	-0.0006 (0.0032)	-0.0047 (0.0033)	-0.0006 (0.0032)	-0.0002 (0.0032)
Log Cloud Expense	0.0116** (0.0049)	0.0117** (0.0049)	0.0117** (0.0049)	0.0117** (0.0049)	0.0119** (0.0049)
Structured Mgmt - Process	0.0433*** (0.0092)	0.0434*** (0.0092)	0.0407*** (0.0091)	0.0434*** (0.0092)	0.0440*** (0.0093)
Structured Mgmt - HR	-0.0283*** (0.0072)	-0.0287*** (0.0072)	-0.0281*** (0.0072)	-0.0288*** (0.0072)	-0.0294*** (0.0071)
Union	0.0771** (0.0363)	0.0813** (0.0363)	0.0755** (0.0361)	0.0812** (0.0363)	0.0826** (0.0362)
Plant Age		(-)	(-)**	(-)	(-)
Log Capital			(+)**		
VA Growth 2019–2021				-0.0013 (0.0077)	-0.0060 (0.0085)
Log Labor Prod. 2019					-1.521* (0.894)
No lack of AI expertise (IV)	0.0827*** (0.0174)	0.0824*** (0.0174)	0.0840*** (0.0175)	0.0824*** (0.0174)	0.0818*** (0.0174)
Additional Controls:	Size, Industry, Skill, Plant Type, MU, Energy, HQ				

Notes: AI Index (the dependent variable) measures adoption of applications of AI in business functions or adoption of specific AI technologies, with the index being normalized to have zero mean and unit standard deviation. DS index measures use of descriptive statistics in decision making with the index being normalized to have zero mean and unit standard deviation. PA index measures use of prediction algorithms across six business functions with index being normalized to have zero mean and unit standard deviation. Other variables include log IT capital (accumulated IT equipment expenses in the past 3 years), logged expenditure on cloud computing (from the MOPS), an indicator that AI skills were not mentioned as a barrier to AI adoption; an index of “structured management-process” practices focused on production monitoring and target setting (Bloom et al., 2019), an index of “structured management – HR” related to incentives and promotion practices (Bloom et al., 2019; Cornwell et al., 2021), extent of unionization, plant age (from the LBD), and logged capital stock calculated using the perpetual inventory method, value added growth from 2019–2021 and log labor productivity in 2019. Unreported controls include Size (logged employment), Industry (3-digit NAICS fixed effects), Skill (percentage of employees with BA degrees), Plant Type (production-process design), MU (multi-unit status), Energy (logged energy expenditure), and HQ (indicator that headquarters for the firm is co-located). Standard errors are clustered at the firm level and reported in parentheses. Statistical significance: *** p<0.01, ** p<0.05, * p<0.1

Table 3: Performance Regressions of AI Adoption
MOPS—ASM: Establishment-level

Panel A: OLS							
	Log Value Added (VA)						Profits
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
AI Index	0.097*** (0.014)	0.078*** (0.013)	0.006 (0.007)	-0.004 (0.007)	-0.003 (0.007)	-0.013** (0.007)	81.63 (130.3)
DS Index	0.559*** (0.014)	0.451*** (0.014)	0.077*** (0.008)	0.031*** (0.009)	0.032*** (0.009)	0.015* (0.009)	7.42 (113.8)
PA Index	0.063*** (0.013)	0.054*** (0.012)	0.008 (0.007)	-0.002 (0.007)	-0.0009 (0.007)	-0.002 (0.007)	-39.46 (107.2)
Controls:							
IT Capital and Cloud	No	Yes	Yes	Yes	Yes	Yes	Yes
Size & Energy & MU Status	No	No	Yes	Yes	Yes	Yes	Yes
Skill, Structured Mgmt, Plant Type, HQ, & Unionization	No	No	No	Yes	Yes	Yes	Yes
Plant Age	No	No	No	No	Yes	Yes	Yes
Capital Stock	No	No	No	No	No	Yes	No
Industry	No	No	Yes	Yes	Yes	Yes	Yes
Panel B: IV							
	AI Index	Log Value Added	Profits	Change in WIP	Log # of	Log	
	(1)	(2) LP	(3) TFP	Inventory	Robots	Employment	(7)
No lack of AI expertise	0.0827*** (0.0174)						
AI Index		-0.775*** (0.271)	-0.587** (0.230)	-11,300*** (4,110)	2,900** (1,408)	0.412** (0.184)	-0.555** (0.243)
DS Index	-0.050*** (0.018)	-0.007 (0.020)	-0.015 (0.017)	-565** (271)	111 (103)	0.039*** (0.0125)	0.087*** (0.0166)
PA Index	0.386*** (0.012)	0.297*** (0.104)	0.219** (0.089)	4339*** (1585)	-1202** (545.8)	-0.155** (0.0710)	0.238** (0.0943)
Additional Controls	See table notes						

Notes: AI Index (the dependent variable) measures adoption of applications of AI in business functions or adoption of specific AI technologies, with the index being normalized to have zero mean and unit standard deviation. DS index measures use of descriptive statistics in decision making with the index being normalized to have zero mean and unit standard deviation. PA index measures use to prediction algorithms across six business functions with index being normalized to have zero mean and unit standard deviation. **Panel A:** As indicated by the with control variable listings, control variables may include: accumulated capitalized IT equipment expenditures over the prior 3 years; expenditures on Cloud computing; size in terms of logged employment; log energy expenses; indicator of multi-unit status; worker skill in terms of the percentage of employee with a BA degree; indexes of structured management practices, for both process and HR separately (see Bloom et al. 2019); a plant type variable capturing production strategy (0 for R&D plants or job shops, 1 for batch production and 2 for continuous flow or cellular manufacturing); percentage of unionized employees; logged plant age from the LBD; logged capital stock; industry at the 3-digit NAICS level. Column 6 includes an indicator for missing data on non-IT capital to stabilize the sample size.

Panel B: All columns include all control variables from column 5 of Panel A. “No lack of AI expertise” is a dummy, reverse-coded from respondents reporting lack of AI expertise as reason not to adopt or to delay AI use (see Table 1, Panel D). “WIP” denotes the value of work-in-progress inventory reported on the ASM. Number of robots includes active and purchased industrial robots from the ASM.

Standard errors for both panels are clustered at the firm level and reported in parentheses. Note that we do not report first stage F-Stats, since according to Angrist and Pischke (2007), a significant first stage and significant second stage are sufficient if the endogenous variable is just-identified (as many endogenous variables as instruments).

Statistical significance: *** p<0.01, ** p<0.05, * p<0.1

Table 4: IV Robustness
MOPS—ASM: Establishment-level

Panel A: Ruling out sample-selected IV

	log VA (1)	AI Index (2)	log VA (3)
No lack of AI expertise	-0.0639*** (0.0176)		
AI is widely used		0.603*** (0.041)	
AI Index			0.065 (0.044)
DS Index	0.0320*** (0.00866)	-0.043** (0.018)	0.035*** (0.009)
PA Index	-0.00175 (0.00673)	0.373*** (0.012)	-0.027 (0.019)
Log IT Capital	(+)** (+)	(-)	(+)** (+)
Log Cloud Expense	(+)** (+)	(+)** (+)	(+)** (+)
Additional Controls	See table notes		

Panel B: Plausibility of Exclusion Restriction

	Job Shops		SU Firms		
	AI Index (1)	Log VA (2)	AI Index (3)	Log VA (4)	
No lack of AI expertise	-0.0204 (0.034)	-0.0213 (0.0308)	-0.0149 (0.0272)	-0.011 (0.0284)	
Additional Controls	See table notes				

Panel A Notes: Placebo instrument is the affirmative response that “Artificial intelligence is widely or increasingly used at this establishment.” Unreported controls include: size in terms of logged employment; log energy expenses; indicator of multi-unit status; worker skill in terms of the percentage of employee with a BA degree; indexes of structured management practices, for both process and HR separately (see Bloom et al. 2019); a plant type variable capturing production strategy (0 for R&D plants or job shops, 1 for batch production and 2 for continuous flow or cellular manufacturing); percentage of unionized employees; logged plant age from the LBD; industry at the 3-digit NAICS level.

Panel B Notes: Instrumenting “no lack of AI expertise” for the AI Index. First stage is reported in columns 1 and 3. Reduced forms are reported in columns 2 and 4. Job Shops are plants reporting high mix, low volume production in a job shop or prototyping production strategy. SU firms are single-unit firms. Controls match the unreported controls of Panel A. Unreported controls in both panels include all controls used in the second column of panel B in Table 3, which are controls used in Panel A of this table, including the DS and PA indexes. Standard errors for both panels are clustered at the firm level and reported in parentheses. Statistical significance: *** p<0.01, ** p<0.05, * p<0.1

Table 5: Heterogeneity in Treatment Effects
MOPS—ASM: Establishment-level

Panel A: Treatment Effects

	Full Sample (1)	Cont. Flow Prod. (2)	Old (25+ years) (3)
First Stage	0.083*** (0.017)	0.136*** (0.031)	0.068*** (0.022)
Second Stage	-0.775*** (0.271)	-0.394 (0.257)	-1.424** (0.556)
Controls	See table notes		

Panel B: Complier Characteristics

Characteristic	Value
Fraction of compliers in overall sample	0.03
<i>(a) Complier characteristics resembling Always Takers</i>	
Odds Ratio for All Sales through E-Commerce	1.33
Odds Ratio for above-median Cloud Computing expenditure	1.23
Odds Ratio for Part of Multi-Unit Firm (MU)	1.42
<i>(b) Complier characteristics resembling Never Takers</i>	
Odds Ratio for Highly Digitalized	0.40
Odds Ratio for HQ Plant	0.52
Controls	See table notes

Notes: Panel A presents IV estimates for different subsamples. Standard errors in parentheses. Statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

Panel B displays complier characteristics, showing the fraction of compliers in the overall sample and odds ratios for various firm characteristics.

Unreported controls in both panels include all controls used in the second column of panel B in Table 3, such as: DS index measuring use of descriptive statistics in decision making with the index being normalized to have zero mean and unit standard deviation; PA index on use of prediction algorithms across six business functions with index being normalized to have zero mean and unit standard deviation; accumulated capitalized IT equipment expenditures over the prior 3 years; expenditures on Cloud computing; size in terms of logged employment; log energy expenses; indicator of multi-unit status; worker skill in terms of the percentage of employee with a BA degree; indexes of structured management practices, for both process and HR separately (see Bloom et al. 2019); a plant type variable capturing production strategy (0 for R&D plants or job shops, 1 for batch production and 2 for continuous flow or cellular manufacturing); percentage of unionized employees; logged plant age from the LBD; industry at the 3-digit NAICS level.

Table 6: Performance of Early AI Adopters
ABS—CMF: Firm-level

Panel A: First Difference 2012–2017					
	Log Sales (1)	Log Value Added (2)	Profits (3)	Log Value Added (4)	Log Value Added (5)
AI Index	-0.0262*** (0.00467)	-0.0187*** (0.00473)	-4073.1* (2241.5)	-0.0247*** (0.00561)	0.0790** (0.0352)
AI Index X Strategy: Quality				-0.00189 (0.00667)	
AI Index X Strategy: New Markets				0.0144*** (0.00491)	
AI Index X Strategy: Cost Leadership				-0.0014 (0.00547)	
AI Index X log firm age					-0.0307*** (0.0103)
Additional Controls			See table notes		
Number of firms			≈ 55,000		
Number of obs			≈ 110,000		

Panel B: Growth during 2017–2021					
	Growth: Employment (1)	Growth: Revenue (2)	Growth: Labor Productivity (3)	Growth: Revenue (4)	Growth: Labor Productivity (5)
AI Index	0.00843*** (0.00225)	0.00473*** (0.00178)	0.00336** (0.00160)	0.0330*** (0.0108)	0.0197** (0.00910)
AI Index X Log Firm Age				-0.00895*** (0.00324)	-0.00514* (0.00274)
Additional Controls			See table notes		
Number of firms			≈ 55,000		
Number of obs			≈ 55,000		

Notes: AI Index measures adoption of applications of AI in business functions or adoption of specific AI technologies, with the index being normalized to have zero mean and unit standard deviation. Strategy measures from the 2018 Annual Business Survey. Firm age per the LBD. **Panel A Notes:** Unreported controls include logged employment, logged capital stock, logged capitalized IT equipment expenses, logged IT outsourcing expenses, logged energy expenses, logged software expenses. Standard errors are clustered at the firm level. Years included are 2012 and 2017 with year fixed effects. Number of firm observations is roughly 55,500. **Panel B Notes:** Growth rates are calculated using Davis, Haltiwanger and Schuh (1996) symmetric growth rates between 2017 and 2021. Additional controls include initial year (2017) logged employment (column 1), initial year logged sales (columns 2 & 4), initial year logged sales per worker (columns 3 & 5). Robust standard errors are reported in parentheses. Panels A and B also include a full set of 3-digit NAICS fixed effects.

Unreported controls in both panels include all controls used in the second column of panel B in Table 3, such as: DS index measuring use of descriptive statistics in decision making with the index being normalized to have zero mean and unit standard deviation; PA index on use of prediction algorithms across six business functions with index being normalized to have zero mean and unit standard deviation; accumulated capitalized IT equipment expenditures over the prior 3 years; expenditures on Cloud computing; size in terms of logged employment; log energy expenses; indicator of multi-unit status; worker skill in terms of the percentage of employee with a BA degree; indexes of structured management practices, for both process and HR separately (see Bloom et al. 2019); a plant type variable capturing production strategy (0 for R&D plants or job shops, 1 for batch production and 2 for continuous flow or cellular manufacturing); percentage of unionized employees; logged plant age from the LBD; industry at the 3-digit NAICS level.

Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 7: Investigating the Causal Mechanism
MOPS—ASM: Establishment-level

Panel A: Causal Effects of AI at Old Establishments

	Log Value Added (LP) (1)	Log Value Added (TFP) (2)	Log Wages (3)	Log Employment (4)
AI Index	-0.1497 (0.2643)	0.0835 (0.2456)	-0.0534 (0.166)	-0.0952 (0.2584)
AI Index X Old	-1.224** (0.5766)	-1.133** (0.4678)	-0.6961** (0.3265)	-0.9205* (0.5199)
Additional Controls	See table notes			

Panel B: Causal Effects of AI on Structured Management at Old Establishments

	Structured Management (1)	Structured Management: Process (2)	KPI Review by Non-Managers (3)	Production Target Awareness (4)
AI Index	0.1189 (0.2051)	0.1795 (0.2088)	0.1734 (0.0865)	0.2534 (0.1109)
AI Index X Old	-0.8322** (0.3991)	-0.8602** (0.397)	-0.286*** (0.1363)	-0.2941** (0.1486)
Additional Controls	See table notes			

Panel C: Within-Firm Spillover Effects of AI

	AI Index (1)	EAI Index (2)	Log Value Added (LP) (3)
No lack of AI expertise	0.0810*** (0.0174)	-0.000746 (0.0172)	-0.822*** (0.278)
No lack of AI expertise elsewhere	0.0922*** (0.0348)	0.790*** (0.0847)	0.343*** (0.0632)
Additional Controls	See table notes		

Notes: Panel A: “Old” is any plant older than 25 years. Additional control variables include: DS index, PA index, accumulated IT equipment expenses in the past 3 years, expenses on cloud computing, size in terms of logged employment, energy expenditures, multi-unit indicator, skill in terms of number of employees with bachelors degrees, a plant type indicator capturing production strategy, and percentage of unionized employees. Columns (1), (2) and (7) and (8) also include an index of structured management practicess (Bloom et al., 2019). Column (2) includes logged capital stock, calculated via perpetual inventory method. **Panel B** includes controls from column 1 of Panel A. Both Panels A and B include 3 digit NAICS industry fixed effects. Number of establishment observations is roughly 28,500. Standard errors are clustered at the firm level and reported in parentheses.

Statistical significance: *** p<0.01, ** p<0.05, * p<0.1

References

- ACEMOGLU, D., U. AKCIGIT, H. ALP, N. BLOOM, AND W. KERR (2018): “Innovation, Reallocation, and Growth,” *American Economic Review*, 108, 3451–3491.
- ACEMOGLU, D., G. W. ANDERSON, D. N. BEEDE, C. BUFFINGTON, E. E. CHILDRESS, E. DILERSOZ, L. S. FOSTER, N. GOLDSCHLAG, J. C. HALTIWANGER, Z. KROFF, ET AL. (2022): “Automation and the workforce: A firm-level view from the 2019 annual business survey,” Tech. rep., National Bureau of Economic Research.
- ACEMOGLU, D. AND P. RESTREPO (2018): “The race between man and machine: Implications of technology for growth, factor shares, and employment,” *American economic review*, 108, 1488–1542.
- ADNER, R., P. PURANAM, AND F. ZHU (2019): “What is different about digital strategy? From quantitative to qualitative change,” *Strategy Science*, 4, 253–261.
- AGRAWAL, A., J. GANS, AND A. GOLDFARB (2018): *Prediction Machines: The Simple Economics of Artificial Intelligence*, Harvard Business Press.
- AGRAWAL, A., J. S. GANS, AND A. GOLDFARB (2024): “Artificial intelligence adoption and system-wide change,” *Journal of Economics & Management Strategy*, 33, 327–337.
- ALEKSEEVA, L., M. GINE, S. SAMILA, AND B. TASKA (2020): “AI Adoption and firm performance: Management versus IT,” *Available at SSRN 3677237*.
- ANGRIST, J. D., G. W. IMBENS, AND D. B. RUBIN (1996): “Identification of causal effects using instrumental variables,” *Journal of the American statistical Association*, 91, 444–455.
- ANGRIST, J. D. AND J.-S. PISCHKE (2009): *Mostly harmless econometrics: An empiricist’s companion*, Princeton university press.
- ARAL, S., E. BRYNJOLFSSON, AND L. WU (2012): “Three-way complementarities: Performance pay, human resource analytics, and information technology,” *Management Science*, 58, 913–931.
- ARROW, K. J. (1962): “The economic implications of learning by doing,” *The Review of Economic Studies*, 29, 155–173.
- ASPURU-GUZIK, A. (2023): “The decade of artificial intelligence in chemistry and materials,” *Digital Discovery*, 2, 10–11.
- ATALAY, E., A. HORTAÇSU, AND C. SYVERSON (2014): “Vertical integration and input flows,” *American Economic Review*, 104, 1120–1148.
- AUTOR, D., D. DORN, G. H. HANSON, G. PISANO, AND P. SHU (2020a): “Foreign competition and domestic innovation: Evidence from US patents,” *American Economic Review: Insights*, 2, 357–374.
- AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. V. REENEN (2020b): “The Fall of the Labor Share and the Rise of Superstar Firms [“Automation and New Tasks: How Technology Displaces and Reinstates Labor”],” *The Quarterly Journal of Economics*, 135, 645–709.
- AUTOR, D. AND A. SALOMONS (2018): “Is Automation Labor Share–Displacing? Productivity Growth, Employment, and the Labor Share,” *Brookings Papers on Economic Activity*, 49, 1–87.

- BABINA, T., A. FEDYK, A. HE, AND J. HODSON (2024): “Artificial intelligence, firm growth, and product innovation,” *Journal of Financial Economics*, 151, 103745.
- BARTEL, A., C. ICHNIOWSKI, AND K. SHAW (2007): “How does information technology affect productivity? Plant-level comparisons of product innovation, process improvement, and worker skills,” *The quarterly journal of Economics*, 122, 1721–1758.
- BARTH, E., J. C. DAVIS, R. B. FREEMAN, AND K. McELHERAN (2023): “Twisting the demand curve: Digitalization and the older workforce,” *Journal of Econometrics*, 233, 443–467.
- BERMAN, R. AND A. ISRAELI (2022): “The value of descriptive analytics: Evidence from online retailers,” *Marketing Science*, 41, 1074–1096.
- BLOOM, N., E. BRYNJOLFSSON, L. FOSTER, R. JARMIN, M. PATNAIK, I. SAPORTA-EKSTEN, AND J. VAN REENEN (2019): “What drives differences in management practices?” *American Economic Review*, 109, 1648–1683.
- BLOOM, N., E. BRYNJOLFSSON, L. FOSTER, R. JARMIN, I. SAPORTA-EKSTEN, AND J. VAN REENEN (2013): “Management in America,” CES Working Paper 13-01, US Census Bureau Center for Economic Studies.
- BLOOM, N., R. SADUN, AND J. VAN REENEN (2012): “Americans do IT better: US multinationals and the productivity miracle,” *American Economic Review*, 102, 167–201.
- BLOOM, N. AND J. VAN REENEN (2007): “Measuring and explaining management practices across firms and countries,” *The quarterly journal of Economics*, 122, 1351–1408.
- BONNEY, K., C. BREAU, C. BUFFINGTON, E. DINLERSOZ, L. S. FOSTER, N. GOLDSCHLAG, J. C. HALTIWANGER, Z. KROFF, AND K. SAVAGE (2024): “Tracking firm use of AI in real time: A snapshot from the Business Trends and Outlook Survey,” NBER Working Paper w32319, National Bureau of Economic Research.
- BRESNAHAN, T. (2024): “What innovation paths for AI to become a GPT?” *Journal of Economics & Management Strategy*, 33, 305–316.
- BRESNAHAN, T. AND S. GREENSTEIN (1996): “Technical progress and co-invention in computing and in the uses of computers,” *Brookings Papers on Economic Activity. Microeconomics*, 1996, 1–83.
- BRESNAHAN, T. F., E. BRYNJOLFSSON, AND L. M. HITT (2002): “Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence,” *Quarterly Journal of Economics*, 117, 339–376.
- BRESNAHAN, T. F. AND M. TRAJTENBERG (1995): “General purpose technologies ‘Engines of growth’?” *Journal of Econometrics*, 65, 83–108.
- BRYNJOLFSSON, E. AND L. HITT (1996): “Paradox Lost? Firm-Level Evidence on the Returns to Information Systems Spending,” *Management Science*, 42, 541–558.
- BRYNJOLFSSON, E. AND L. M. HITT (2000): “Beyond computation: Information technology, organizational transformation and business performance,” *Journal of Economic perspectives*, 14, 23–48.

- BRYNJOLFSSON, E., W. JIN, AND K. MCELHERAN (2021a): “The power of prediction: predictive analytics, workplace complements, and business performance,” *Business Economics*, 56, 217–239.
- BRYNJOLFSSON, E., D. LI, AND L. RAYMOND (2025): “Generative AI at work,” *The Quarterly Journal of Economics*, qjae044.
- BRYNJOLFSSON, E. AND P. MILGROM (2013): “Complementarity in organizations,” *The handbook of organizational economics*, 11–55.
- BRYNJOLFSSON, E., D. ROCK, AND C. SYVERSON (2021b): “The productivity J-curve: How intangibles complement general purpose technologies,” *American Economic Journal: Macroeconomics*, 13, 333–372.
- BRYZGALOVA, S., S. LERNER, M. LETTAU, AND M. PELGER (2025): “Missing financial data,” *The Review of Financial Studies*, 38, 803–882.
- CABRAL, L. M. B. AND M. H. RIORDAN (1994): “The learning curve, market dominance, and predatory pricing,” *Econometrica*, 62, 1115–1140.
- CALVINO, F. AND L. FONTANELLI (2023): “A portrait of AI adopters across countries: Firm characteristics, assets’ complementarities and productivity,” OECD Working Paper 2023/02, OECD.
- CHARI, V. V. AND H. HOPENHAYN (1991): “Vintage human capital, growth, and the diffusion of new technology,” *Journal of Political Economy*, 99, 1142–1165.
- COCKBURN, I. M., R. HENDERSON, AND S. STERN (2018): “The impact of artificial intelligence on innovation,” NBER Working Paper 24449, National Bureau of Economic Research.
- COMIN, D. A., X. CIRERA, AND M. CRUZ (2025): “Technology Sophistication Across Establishments,” Tech. rep., National Bureau of Economic Research.
- CORNWELL, C., I. M. SCHMUTTE, AND D. SCUR (2021): “Building a productive workforce: The role of structured management practices,” *Management Science*, 67, 7308–7321.
- CUI, Z. K., M. DEMIRER, S. JAFFE, L. MUSOLFF, S. PENG, AND T. SALZ (2024): “The effects of generative AI on high skilled work: Evidence from three field experiments with software developers,” *Available at SSRN 4945566*.
- CZARNITZKI, D., G. P. FERNÁNDEZ, AND C. RAMMER (2023): “Artificial intelligence and firm-level productivity,” *Journal of Economic Behavior & Organization*, 211, 188–205.
- DAVID, P. A. (1990): “The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox,” *American Economic Review*, 80, 355–361.
- DELL’ACQUA, F., E. MCFOWLAND III, E. R. MOLICK, H. LIFSHITZ-ASSAF, K. KELLOGG, S. RAJENDRAN, L. KRAYER, F. CANDELON, AND K. R. LAKHANI (2023): “Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality,” *Harvard Business School Technology & Operations Mgt. Unit Working Paper*.

- DELOITTE (2020): “AI Manufacturing Application Survey,” <https://www2.deloitte.com/content/dam/Deloitte/cn/Documents/cip/deloitte-cn-cip-ai-manufacturing-application-survey-en-200116.pdf>, accessed: 2024-03-23.
- ELOUNDOU, T., S. MANNING, P. MISHKIN, AND D. ROCK (2023): “GPTs are GPTs: An early look at the labor market impact potential of large language models,” *arXiv preprint arXiv:2303.10130*.
- FEIGENBAUM, J. AND D. P. GROSS (2024): “Organizational and economic obstacles to automation: A cautionary tale from AT&T in the twentieth century,” *Management Science*, forthcoming.
- FELTEN, E., M. RAJ, AND R. SEAMANS (2021): “Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses,” *Strategic Management Journal*, 42, 2195–2217.
- FONTANELLI, L., F. CALVINO, C. CRISCUOLO, L. NESTA, AND E. VERDOLINI (2024): “The Role of Human Capital for AI Adoption: Evidence from French Firms,” *Available at SSRN 5010602*.
- FORMAN, C., A. GOLDFARB, AND S. GREENSTEIN (2012): “The Internet and local wages: A puzzle,” *American Economic Review*, 102, 556–575.
- FORMAN, C. AND K. MCELHERAN (2025): “Production chain organization in the digital age: information technology use and vertical integration in US manufacturing,” *Management Science*, 71, 1027–1049.
- GALBRAITH, J. (1973): “Designing complex organizations,” *Reading, Mass.*
- GALDON-SANCHEZ, J. E., R. GIL, AND G. URIZ UHARTE (2022): “The value of information in competitive markets: The impact of big data on small and medium enterprises,” *Available at SSRN 4022675*.
- GIUSTIZIERO, G., T. KRETSCHMER, D. SOMAYA, AND B. WU (2023): “Hyperspecialization and hyperscaling: A resource-based theory of the digital firm,” *Strategic Management Journal*, 44, 1391–1424.
- GOLDFARB, A., B. TASKA, AND F. TEODORIDIS (2023): “Could machine learning be a general purpose technology? A comparison of emerging technologies using data from online job postings,” *Research Policy*, 52, 104653.
- GOLDFARB, A. AND C. TUCKER (2019): “Digital Economics,” *Journal of Economic Literature*, 57, 3–43.
- HAYES, R. H. AND S. C. WHEELWRIGHT (1979): “Link manufacturing process and product life cycles,” *June*, 2, 1999.
- HELFAT, C. E., A. KAUL, D. J. KETCHEN JR, J. B. BARNEY, O. CHATAIN, AND H. SINGH (2023): “Renewing the resource-based view: New contexts, new concepts, and new methods,” *Strategic Management Journal*, 44, 1357–1390.
- HELFAT, C. E. AND M. A. PETERAF (2003): “The dynamic resource-based view: Capability lifecycles,” *Strategic Management Journal*, 24, 997–1010.

- HENDERSON, R. (1993): “Underinvestment and incompetence as responses to radical innovation: Evidence from the photolithographic alignment equipment industry,” *The RAND Journal of Economics*, 248–270.
- HOFFMANN, M., S. BOYSEL, F. NAGLE, S. PENG, AND K. XU (2024): “Generative AI and the Nature of Work,” Tech. rep., CESifo Working Paper.
- HOFFREUMON, C., C. FORMAN, AND N. VAN ZEEBROECK (2023): “Make or Buy your Artificial Intelligence? Testing for Complementarities in Technology Sourcing,” *Cornell University Working Paper*.
- HOLWEG, M. (2007): “The genealogy of lean production,” *Journal of Operations Management*, 25, 420–437.
- IANSITI, M. AND K. R. LAKHANI (2020): *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World*, Harvard Business Press.
- ICHNIOWSKI, C., K. L. SHAW, AND G. PRENNUSHI (1995): “The effects of human resource management practices on productivity,” .
- JIN, W. AND K. MCELHERAN (2024): “Economies before Scale: IT Strategy and Performance Dynamics of Young U.S. Businesses,” *Management Science*, forthcoming.
- KUENG, L., M.-J. YANG, AND B. HONG (2014): “Sources of firm life-cycle dynamics: Differentiating size vs. age effects,” NBER Working Paper w20621, National Bureau of Economic Research.
- LEVITT, S. D., J. A. LIST, AND C. SYVERSON (2013): “Toward an understanding of learning by doing: Evidence from an automobile assembly plant,” *Journal of political Economy*, 121, 643–681.
- LOU, B. AND L. WU (2021): “AI on drugs: Can artificial intelligence accelerate drug development? Evidence from a large-scale examination of bio-pharma firms,” *Evidence from a Large-scale Examination of Bio-pharma Firms*, (March 15, 2021). *MISQ Forthcoming*.
- MCELHERAN, K. AND W. JIN (2020): “Digital Strategy in the Age of Cloud Computing,” *University of Toronto working paper*, available at SSRN 4660979.
- MCELHERAN, K., J. F. LI, E. BRYNJOLFSSON, Z. KROFF, E. DINLERSOZ, L. FOSTER, AND N. ZOLAS (2024): “AI Adoption in America: Who, What, and Where,” *Journal of Economics & Management Strategy*, 33, 375–415.
- MCELHERAN, K., S. OHLMACHER, AND M. J. YANG (2019): “Strategy and Structured Management,” Available at SSRN 4660979.
- MILGROM, P. AND J. ROBERTS (1988): “Communication and inventory as substitutes in organizing production,” *The Scandinavian Journal of Economics*, 275–289.
- (1990): “The Economics of Modern Manufacturing: Technology, Strategy, and Organization,” *American Economic Review*, 80, 511–528.
- (1995): “Complementarities and fit strategy, structure, and organizational change in manufacturing,” *Journal of accounting and economics*, 19, 179–208.

- MIRIC, M., N. JIA, AND K. G. HUANG (2023): “Using supervised machine learning for large-scale classification in management research: The case for identifying artificial intelligence patents,” *Strategic Management Journal*, 44, 491–519.
- NATIONAL SCIENCE FOUNDATION (2021): “Data Tables,” <https://ncses.nsf.gov/pubs/nsf23314/assets/data-tables/tables/nsf23314-tab066.pdf>, accessed: 2024-03-23.
- OHNO, T. (1988): *Toyota production system: beyond large-scale production*, Cambridge, MA: Productivity Press.
- OTIS, N., R. CLARKE, S. DELECOURT, D. HOLTZ, AND R. KONING (2024): “The uneven impact of generative AI on entrepreneurial performance,” .
- PENG, S., E. KALLIAMVAKOU, P. CIHON, AND M. DEMIRER (2023): “The impact of AI on developer productivity: Evidence from GitHub Copilot,” *arXiv preprint arXiv:2302.06590*.
- PORTER, M. E. (1980): *Competitive Strategy*, New York: Free Press.
- PRATT, G. A. (2015): “Is a Cambrian Explosion Coming for Robotics?” *Journal of Economic Perspectives*, 29, 51–60.
- ROMER, P. M. (1990): “Endogenous technological change,” *Journal of Political Economy*, 98, S71–S102.
- RUMELT, R. P. (2005): “Theory, strategy, and entrepreneurship,” in *Handbook of Entrepreneurship Research: Interdisciplinary Perspectives*, Boston, MA: Springer US, 11–32.
- SCUR, D., R. SADUN, J. VAN REENEN, R. LEMOS, AND N. BLOOM (2021): “The World Management Survey at 18: lessons and the way forward,” *Oxford Review of Economic Policy*, 37, 231–258.
- SCUR, D. AND S. WOLFOLDS (2024): “Revisiting the world management survey in strategy: Applications to theory and replication,” *Strategy Science*, 9, 58–78.
- SEAMANS, R. AND M. RAJ (2018): “AI, labor, productivity and the need for firm-level data,” Tech. rep., National Bureau of Economic Research.
- STRIETH-KALTHOFF, F., H. HAO, V. RATHORE, J. DERASP, T. GAUDIN, N. H. ANGELLO, M. SEIFRID, E. TRUSHINA, M. GUY, J. LIU, ET AL. (2024): “Delocalized, asynchronous, closed-loop discovery of organic laser emitters,” *Science*, 384, eadk9227.
- SVANBERG, M., W. LI, M. FLEMING, B. GOEHRING, AND N. THOMPSON (2024): “Beyond AI Exposure: Which Tasks are Cost-Effective to Automate with Computer Vision?” Working Paper 4700751, SSRN.
- TAMBE, P., L. HITT, D. ROCK, AND E. BRYNJOLFSSON (2020): “Digital capital and superstar firms,” NBER Working Paper w28285, National Bureau of Economic Research.
- TAMBE, P. AND L. M. HITT (2012): “The productivity of information technology investments: New evidence from IT labor data,” *Information Systems Research*, 23, 599–617.
- TAMBE, P., L. M. HITT, AND E. BRYNJOLFSSON (2012): “The extroverted firm: How external information practices affect innovation and productivity,” *Management Science*, 58, 843–859.

- THOMPSON, P. (2012): “The relationship between unit cost and cumulative quantity and the evidence for organizational learning-by-doing,” *Journal of Economic Perspectives*, 26, 203–224.
- TOKTAY, L. B. AND L. M. WEIN (2001): “Analysis of a forecasting-production-inventory system with stationary demand,” *Management Science*, 47, 1268–1281.
- WOMACK, J. P., D. T. JONES, AND D. ROOS (2007): *The machine that changed the world: The story of lean production—Toyota’s secret weapon in the global car wars that is now revolutionizing world industry*, Simon and Schuster.
- WU, L., L. HITT, AND B. LOU (2020): “Data analytics, innovation, and firm productivity,” *Management Science*, 66, 2017–2039.
- YANG, M.-J., L. KUENG, AND B. HONG (2015): “Business Strategy and the Management of Firms,” NBER Working Paper 20846, National Bureau of Economic Research.
- YANG, M.-J., N. LI, AND L. KUENG (2021): “The Impact of Emerging Market Competition on Innovation and Business Strategy: Evidence from Canada,” *Journal of Economic Behavior and Organization*, 181, 117–134.
- ZOLAS, N., Z. KROFF, E. BRYNJOLFSSON, K. MCELHERAN, D. N. BEEDE, C. BUFFINGTON, N. GOLDSCHLAG, L. FOSTER, AND E. DINLERSOZ (2020): “Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey,” NBER Working Papers 28290, National Bureau of Economic Research, Inc.