Publishing to Coordinate: AI Disclosure and Supply Chain Adoption

Wei Jiao
Rutgers School of Business - Camden
Rutgers University
Wei.jiao@rutgers.edu

Jenny Li Zhang
Sauder School of Business
The University of British Columbia
jenny.zhang@sauder.ubc.ca

Yuxiang Zheng
Rutgers School of Business - Camden
Rutgers University
yuxiang.zheng@rutgers.edu

December 2025

Publishing to Coordinate: AI Disclosure and Supply Chain Adoption

Abstract

Artificial intelligence requires coordinated adoption across supply chain partners to realize its full value. We propose that focal firms' AI-related scientific publications serve as coordination mechanisms that facilitate supplier AI adoption. When AI investments exhibit complementarities, suppliers face uncertainty about others' actions, potentially leading to underinvestment. Scientific publications can resolve this coordination failure by creating common knowledge about a firm's technological commitments—unlike patents, which create exclusionary barriers, or private communication, which cannot coordinate across the supplier network. We find that focal firms' AI publications are positively associated with subsequent supplier AI adoption, with stronger effects when coordination challenges are more severe: shorter relationship duration, more suppliers, and greater supply chain complexity. The publication effect is specific to suppliers and does not extend to customers, partners, or competitors. Only publications—not patents—are associated with supplier AI adoption. Our findings identify a novel coordination function of corporate AI scientific disclosure.

Key words: Corporate scientific publication, Supply chain coordination, AI adoption

JEL codes: M41, D23, L14, O31, O34

1. Introduction

Large corporations have systematically withdrawn from scientific research over the past several decades, increasingly focusing on development rather than basic research (Arora, Belenzon, and Patacconi 2018). Against this backdrop, corporate AI-related publications present a notable counter-trend: they more than doubled between 2014 and 2019, rising from 11% to 29% of total corporate scientific output (Emerging Technology Observatory 2024). A long-standing question in the disclosure literature asks why firms voluntarily reveal valuable innovations when such disclosure may benefit competitors (Verrecchia 1983). Prior research offers several explanations, including signaling research quality to capital markets (Baruffaldi, Simeth, and Wehrheim 2023), attracting scientific talent (Rosenberg 2010; Zucker and Darby 1996), and accessing academic information networks that improve R&D productivity (Simeth and Cincera 2016; Polidoro and Theeke 2012). We propose an additional mechanism relevant to general-purpose technologies like AI: scientific publications can facilitate coordination with supply chain partners. Despite the growing importance of inter-firm production networks, little is known about how innovation disclosures shape supply chain relationships.

We conjecture that such disclosures facilitate technological upgrading among supply chain partners. As a general-purpose technology, AI transforms entire production networks rather than isolated firm operations—a firm cannot fully realize AI's benefits unless its suppliers also adopt AI capabilities. This creates a coordination problem: suppliers must make relationship-specific investments in AI—hiring specialized talent, developing systems tailored to the focal firm's requirements—but these investments lose value if the focal firm terminates the relationship or abandons its AI strategy. Because AI evolves rapidly and future needs are difficult to specify contractually, the incomplete contracting literature suggests that suppliers cannot fully protect themselves through formal agreements (Williamson 1979; Costello 2013). Anticipating these hold-up risks, suppliers underinvest in AI even when adoption would benefit the entire supply chain.

Patents and trade secrets do not resolve this problem—patents grant exclusive rights that create legal and financial barriers to supplier adoption, while trade secrets keep knowledge hidden entirely. Private bilateral disclosure is also inadequate: it is costly to scale across many suppliers, lacks the credibility of peer-reviewed research, and does not reach upstream tiers of the supply chain and offers no protection against misappropriation—suppliers could appropriate shared knowledge without attribution. Publications, by contrast, establish a public record of priority while placing knowledge in the public domain, allowing suppliers to freely build on disclosed AI capabilities. They also signal the focal firm's sustained commitment to its AI trajectory, reducing suppliers' uncertainty about whether their investments will retain value.

The incomplete contract framework generates two testable predictions. First, because publications provide both the knowledge for suppliers to develop AI capabilities and a credible signal of the focal firm's long-term commitment, they can help address coordination problems with suppliers. As such, focal firms' AI publications should predict subsequent AI adoption by upstream suppliers. Second, because supplier AI adoption is expected to benefits the production network, AI publications should also lead to improved future performance for the focal firm.

We obtain data on AI-related publications from the Emerging Technology Observatory (ETO) at Georgetown University, which identifies AI publications using machine learning models applied to a comprehensive corpus of scholarly articles. A publication is associated with a company if at least one author is affiliated with that company. The dataset provides comprehensive coverage of major U.S. publicly traded firms, including the S&P 500 from 2014 to 2019. We first show that among our sample firms, AI publications more than doubled from 2,592 in 2014 to 6,609 in 2019.

This growth is concentrated in Manufacturing and Services, which together account for over 85% of corporate AI research output. AI publications are also positively correlated with patent filings, suggesting that firms use scientific publication for a portion of their innovative portfolio while continuing to pursue traditional intellectual property protection for other inventions.

Having provided these descriptive patterns, we turn to our first prediction: whether focal firms' AI publications are associated with subsequent AI adoption by upstream suppliers. We measure supplier AI adoption following Babina et al. (2024), who develop a human-capital-based measure of firm-level AI investment that captures the number of AI-skilled employees at each firm. We find that a focal firm's AI publications predict subsequent AI adoption by its upstream suppliers. The effect is economically meaningful: a one-standard-deviation increase in AI publications is associated with a 7.1% increase in suppliers' AI employment in the following year. Consistent with our incomplete contracting framework, this effect is concentrated in supply chains characterized by shorter relationship duration, greater supplier network complexity, and longer supply chain length—conditions where coordination frictions and hold-up concerns are most severe.

We next examine how AI publications create value for disclosing firms. AI publications carry proprietary costs—competitors can freely access disclosed knowledge. However, if the benefits of facilitating supplier AI adoption outweigh these costs, we expect AI publications to lead to improved future performance. We proxy for private value creation using future sales revenue, with the expectation that the benefits of supply chain coordination ultimately manifest as increased revenue.[1] Specifically, we use a mediation framework to test whether supplier AI

[1] Although generative AI applications remain in early stages, many AI technologies were already delivering measurable value by 2019. In manufacturing—which accounts for over 30% of AI publications in our sample—applications such as predictive maintenance, quality control using computer vision, and demand forecasting were generating cost savings and efficiency gains (Bughin et al. 2018; McKinsey Global Institute 2018).

adoption serves as the channel through which AI publications affect future sales. We find that supplier AI adoption significantly predicts the focal firm's future sales, and that the direct effect of AI publications on future sales becomes insignificant when supplier AI adoption is included. This pattern is consistent with mediation—the economic value generated by AI publications materializes through the technological upgrading of the firm's suppliers. The estimated relative indirect effect is 7.1%, suggesting that a meaningful portion of the value from AI publications flows through the supply chain channel.

We further reinforce our main findings with two additional tests that examine alternative channels and mechanisms. Our theory predicts that AI publications facilitate coordination specifically with upstream suppliers, where hold-up concerns impede relationship-specific investments. However, publications are in the public domain—anyone can access them. If we observe spillovers only to suppliers and not to other parties, this suggests that firms strategically select which knowledge to disclose. Specifically, firms may publish AI research that is more relevant to upstream production processes than to downstream applications or competitive positioning. We test this by regressing the AI adoption of customers, business partners, and competitors on the focal firm's AI publications. We find no significant relationship between focal firms' AI publications and AI adoption by any of these groups. This pattern suggests that knowledge spillovers are not incidental—firms appear to strategically channel disclosed knowledge toward supply chain partners where coordination benefits are greatest, while withholding or protecting innovations that would primarily benefit competitors or are less relevant to downstream parties.

A key feature of our theory is that publications place knowledge in the public domain, allowing suppliers to freely adopt disclosed AI capabilities. Patents, by contrast, grant exclusive rights that create barriers to adoption. If this distinction matters, publications—not patents—should drive supplier AI adoption. We test this by including both AI publications and AI patents in the same specification. We find that only AI publications remain significant; AI patents have no incremental effect on supplier AI adoption. This result reinforces our proposed mechanism: it is the open nature of scientific disclosure, rather than general AI innovation activity, that facilitates technological upgrading among suppliers.

We face several endogeneity concerns. First, firms that publish AI research may systematically differ from non-publishers in ways that also affect supplier behavior. For example, better-managed firms or those with stronger supplier relationships might both publish more AI research and cultivate suppliers who invest in AI capabilities. We address this concern by including firm fixed effects, which absorb time-invariant differences across firms and identify the effect from within-firm variation in AI publications over time. Second, reverse causality could arise if suppliers' AI investments prompt focal firms to publish, rather than the reverse. Our research design mitigates this concern through the timing structure: we regress supplier AI adoption in year t+1 on focal firm publications in year t.

Third, one might worry that the correlation reflects an industry-wide AI adoption trend rather than a causal supply chain effect—if AI is simply "hot," both focal firms and their supply chain partners may adopt AI simultaneously. Several features of our results are inconsistent with this alternative explanation. We show that AI publications predict adoption by suppliers but not by customers, business partners, or competitors. If industry trends were driving our results, we would expect similar spillovers to all connected parties. Moreover, year fixed effects absorb aggregate time trends affecting AI adoption economy-wide. Our cross-sectional tests provide additional support: the effect of AI publications on supplier adoption is stronger in supply chains with shorter relationship duration, more suppliers, longer supply chain length, and greater supplier industry diversity—precisely the conditions where our theory predicts coordination frictions are most severe. An industry trend cannot explain this systematic variation. Finally, we conduct a horse-race analysis between AI publications and AI patents, finding that only publications predict supplier adoption. This distinction reinforces our proposed mechanism: it is the open, public nature of scientific disclosure—not general AI innovation activity—that facilitates supplier upgrading. While we cannot fully rule out all sources of endogeneity absent a clean natural experiment, the totality of evidence supports a supply chain coordination interpretation.

We contribute to the literature on coordination failures in technology adoption. Prior work in economics establishes that general purpose technologies face adoption challenges when complementary investments must be made across organizational boundaries (Bresnahan and Trajtenberg 1995; Helpman and Trajtenberg 1998). The accounting literature has examined how information sharing and disclosure affect supply chain relationships (Costello 2013; Crawford, Huang, Li, and Yang 2020), but has not explored how disclosure might coordinate technology adoption across firms. We extend this literature by identifying scientific publications as a mechanism to resolve coordination failures in supply chains. When AI adoption requires compatible investments by multiple suppliers, each party faces uncertainty about others' actions, potentially leading to underinvestment. We show that public scientific disclosure creates common knowledge about a focal firm's technological commitments, enabling suppliers to coordinate their AI investments with confidence. This mechanism offers a scalable alternative to private bilateral communication, which cannot create the common knowledge necessary for network-wide coordination.

Our study also advances the literature on corporate scientific publication. Prior research documents various motives for firms to publish scientific findings, including signaling research quality to capital markets (Baruffaldi, Simeth, and Wehrheim 2022), attracting scientific talent (Stern 2004; Roach and Sauermann 2010), and strategically deterring competition through defensive publishing (Johnson 2014). We identify a new motive: firms publish scientific research to coordinate technological upgrading among supply chain partners. Unlike prior explanations that focus on capital market or labor market benefits, our mechanism operates through real production relationships. We show that the value of scientific disclosure can extend beyond the disclosing firm to shape investment decisions throughout the supply chain.

Finally, our findings inform research on innovation disclosure and intellectual property strategy. Prior work examines how patent disclosures generate both knowledge spillovers and proprietary costs (Kim and Valentine 2021; Dyer, Glaeser, Lang, and Sprecher 2023). We add to a nascent literature showing that public disclosure of innovation—rather than protection—can generate private value for disclosing firms. Coleman, Fronk, and Valentine (2023) document that firms benefit from sharing software code on GitHub through crowdsourced improvements and widespread adoption. We extend this insight to scientific publications and identify a distinct mechanism: public disclosure can resolve coordination problems by creating common knowledge across the supply network. While patents protect innovations by granting exclusive rights—and in doing so create financial and legal barriers to complementary adoption—publications facilitate coordinated adoption by placing knowledge in the public domain without exclusionary intent. Our findings suggest that public disclosure and traditional IP protection serve complementary roles, with firms strategically choosing which innovations to publish based on where coordination benefits are greatest.

2. Literature Review

2.1 Corporate Scientific Publications

A substantial body of research examines why profit-maximizing firms engage in open science activities that seemingly give away valuable knowledge. The traditional view holds that firms protect intellectual property through patents, trade secrets, and other appropriability mechanisms to capture returns from innovation (Teece, 1986; Cohen, Nelson, and Walsh, 2000). Scientific publication, by placing knowledge in the public domain, appears puzzling as it undermines these strategies.

Recent scholarship identifies several rationales for corporate scientific disclosure. Firms may publish to establish priority and deter competitors from patenting in adjacent areas (Johnson, 2014; Valentine, Zhang, and Zheng, 2025). Publications can signal research capabilities to capital markets, potentially reducing information asymmetry and lowering the cost of capital (Baruffaldi, Simeth, and Wehrheim, 2022). Scientific output also aids in recruiting talented researchers who value academic recognition and the ability to publish (Stern, 2004; Roach and Sauermann, 2010).

These benefits must be weighed against signficant costs. Verrecchia (1983) establishes the foundational framework for analyzing proprietary costs of disclosure: when disclosed information can be used by competitors to the detriment of the disclosing firm, managers rationally withhold information even when disclosure would otherwise benefit investors. Scientific publications represent a particularly severe form of proprietary cost exposure because the disclosed knowledge is detailed, technical, and directly applicable to competitive activities. Arora, Belenzon, and Sheer (2021) document that firms have become increasingly cautious about scientific disclosure, with corporate publications declining even as patenting has increased. They argue that the growing strategic importance of technology has raised the proprietary costs of open science, leading firms to shift toward more protective appropriability strategies. Arora, Belenzon, and Patacconi (2018) document a broader secular decline in corporate scientific research, attributing this trend to reduced vertical integration and increasing specialization of knowledge production. Together, these findings suggest that proprietary cost considerations weigh heavily in firms' publication decisions and that the trend has been toward less, not more, scientific openness.

A related literature examines “open innovation”—the strategic use of knowledge inflows and outflows to accelerate internal innovation and expand markets for external use of innovation (Chesbrough, 2003, 2006). Coleman, Fronk, and Valentine (2023) study a prominent form of open innovation: firms’ voluntary disclosure of software source code on the GitHub platform. They find that open innovation generates private value for disclosing firms through three channels: crowdsourced improvement of disclosed projects, widespread adoption of the firm’s technologies by third parties, and enhanced investment efficiency from managerial learning. Importantly, they document that firms use open innovation as a complement to, rather than a substitute for, traditional intellectual property protection—firms active on GitHub also patent extensively. Their findings demonstrate that voluntary disclosure of proprietary knowledge can benefit firms even when it weakens traditional appropriability mechanisms.

This paper extends the literature on corporate scientific publications by examining a previously unexplored function: the role of publications in coordinating technology adoption across organizational boundaries. While prior work focuses on capital market benefits, competitive dynamics, or labor market signaling, we propose that scientific publications can serve as coordination mechanisms that enable ecosystem-wide adoption of complex technologies. Our framework explains why firms might publish even when the direct proprietary costs are substantial:

the coordination benefits that flow from creating common knowledge across the supply network can outweigh the costs of knowledge dissemination to competitors.

2.2. General Purpose Technologies and Coordination Challenges

Artificial intelligence exhibits the defining characteristics of a general purpose technology (GPT). Bresnahan and Trajtenberg (1995) characterize GPTs as technologies with pervasive applicability across sectors, inherent potential for technical improvement, and strong complementarities with application-sector innovations. Brynjolfsson, Rock, and Syverson (2021) provide extensive evidence that AI satisfies these criteria, documenting both AI's broad applicability and the emerging productivity effects of AI adoption across industries.

A central challenge with GPTs is that realizing their value requires coordinated co-invention across multiple parties. Helpman and Trajtenberg (1998) model this dynamic formally, showing that GPT adoption involves a coordination problem between the GPT sector and application sectors. Because investments are interdependent, each party's incentive to invest depends on expectations about others' behavior. When parties are uncertain about whether complementary investments will materialize, they may rationally underinvest, leading to slower adoption and diminished returns relative to what coordinated investment would achieve.

These coordination challenges are particularly salient in supply chain contexts. When a focal firm adopts AI, realizing productivity gains often requires that suppliers also develop compatible AI capabilities. A substantial accounting literature examines how information asymmetries shape supply chain relationships. Costello (2013) demonstrates that hold-up concerns and information asymmetry are important determinants of supply contract design, with parties using financial covenants and contract duration to mitigate incentive conflicts when relationship-specific assets are exchanged. Raman and Shahrur (2008) show that relationship-specific investments influence earnings management behavior in supply chain contexts.

The disclosure literature documents important interactions between supply chain relationships and firms' disclosure policies. Crawford, Huang, Li, and Yang (2020) find that customer concentration reduces the likelihood of management earnings forecasts, consistent with large customers having private channels to access supplier information. Ellis, Fee, and Thomas (2012) show that proprietary cost considerations shape disclosure about customer relationships. Pandit, Wasley, and Zach (2011) document information externalities along supply chains, with suppliers' stock prices responding to customers' earnings announcements. These studies establish that information flows matter for supply chain coordination, though none examines how disclosure might coordinate technology adoption specifically.

The challenge of coordinating technology adoption across supply chains parallels that faced by platform leaders who must encourage complementary investments by ecosystem participants (Gawer and Cusumano, 2014; Jacobides, Cennamo, and Gawer, 2018). Unlike platforms, however, supply chain relationships typically lack formal governance structures and standardized interfaces. This raises the question of how coordination can be achieved in these more loosely coupled settings where relationships are bilateral and technologies are firm-specific.

2.3 Network Effects and Equilibrium Selection

The economics of network effects provides a theoretical foundation for understanding coordination failures in technology adoption. Katz and Shapiro (1985, 1986) develop models in which the value of adopting a technology increases with the number of other adopters. These network externalities create strategic complementarities: each party’s optimal action depends positively on the actions of others. When a supplier's AI investment is more valuable if the focal firm and other suppliers also invest in AI, the adoption decision becomes interdependent across the network.

Strategic complementarities give rise to multiple equilibria. In one equilibrium, all parties adopt the technology and realize network benefits. In another, no party adopts because each anticipates that others will not adopt. Both equilibria can be self-fulfilling. Farrell and Saloner (1985, 1986) show that this multiplicity creates "excess inertia"—beneficial technologies may fail to achieve adoption because no party is willing to move first.

The role of information in equilibrium selection has received considerable attention. Morris and Shin (2002) demonstrate that public signals have disproportionate effects in coordination games because they create common knowledge. When a signal is publicly observed, each party knows the content of the signal, knows that others know, knows that others know that others know, and so on. This higher-order knowledge enables coordination in ways that private information cannot.

A parallel accounting literature examines how innovation disclosures create knowledge spillovers. Kim and Valentine (2021) demonstrate that mandatory patent disclosures generate both spillover benefits for rivals and proprietary costs for disclosing firms. Dyer, Glaeser, Lang, and Sprecher (2023) show that patent disclosure quality affects follow-on innovation. Kim and Valentine (2023) find that public firm disclosures facilitate patent market transactions by reducing information frictions. This literature establishes that innovation disclosures can have real effects beyond the disclosing firm, though the mechanisms examined—knowledge spillovers and information provision—differ from the coordination mechanism we propose.

Our framework emphasizes that scientific publications can coordinate technology adoption through common knowledge creation. Publications are inherently public: once released, they are observable to all interested parties simultaneously. A focal firm's AI publications create common knowledge about its technological commitments across the entire supplier network. Each supplier observes the publications, knows that other suppliers observe them, and can therefore coordinate investment decisions with confidence. Unlike private bilateral communication, which Crawford, Huang, Li, and Yang (2020) show can substitute for public disclosure with concentrated customers, publications reach all supply chain participants simultaneously, enabling network-wide coordination.

3. Hypothesis Development

3.1 AI Publications and Supplier AI Adoption

We propose that focal firms' AI-related scientific publications facilitate AI adoption among their suppliers by serving as coordination mechanisms. Our argument proceeds in three steps.

First, AI adoption in supply chains exhibits strong complementarities. The value of a supplier's AI investment depends on whether the focal firm has compatible AI capabilities, and vice versa. Consider a supplier investing in AI-enabled quality control systems. The returns to this investment are substantially higher if the focal firm can integrate quality data into its own AI-driven processes. These complementarities create strategic interdependence: optimal investment decisions depend on expectations about partners' investments. This interdependence is well-documented in the supply chain literature, with Costello (2013) showing that relationship-specific investments shape contracting terms and Raman and Shahrur (2008) demonstrating that such investments influence financial reporting choices.

Second, these complementarities give rise to coordination challenges. When multiple suppliers must make AI investments that are valuable only if the focal firm and other suppliers also invest, each party faces uncertainty about others' actions. A supplier may be willing to invest if confident that the focal firm is committed to AI and that other suppliers will develop compatible capabilities. But absent this confidence, the supplier rationally withholds investment. When all suppliers reason similarly, the outcome is an equilibrium in which beneficial AI adoption fails to occur. This dynamic mirrors the excess inertia problem identified by Farrell and Saloner (1985, 1986).

Third, AI publications can resolve coordination failure by creating common knowledge about the focal firm's AI strategy. Unlike private communications with individual suppliers, publications are observable to all supply chain participants simultaneously. As Morris and Shin (2002) demonstrate, such public signals transform private beliefs into common knowledge, enabling coordination in ways that private information cannot. Each supplier knows the focal firm's AI approach, knows that other suppliers know, and can therefore invest with confidence.

Several features distinguish publication from alternative mechanisms. Patent filings, while publicly disclosing technical information, signal an intent to exclude others from using the technology. This exclusionary intent may discourage complementary investment by creating uncertainty about whether suppliers will be permitted to develop compatible capabilities (Kim and Valentine, 2021). Moreover, patents create direct financial and legal barriers—suppliers seeking to develop compatible AI capabilities may face licensing fees, royalty obligations, or infringement risk, reducing expected returns to relationship-specific investments.

Private bilateral communication also falls short as a coordination mechanism. Crawford, Huang, Li, and Yang (2020) document that large customers access information through private channels, but such communication reaches only direct parties and cannot create common knowledge across the supplier network. Scientific publications address both limitations by placing technical knowledge in the public domain without exclusionary intent, enabling suppliers to freely develop compatible capabilities while simultaneously creating common knowledge across the entire network.

Additionally, the scientific content of publications provides technical detail that enables suppliers to develop genuinely compatible capabilities rather than merely increasing investment levels. Dyer, Glaeser, Lang, and Sprecher (2023) demonstrate that disclosure quality matters for follow-on innovation; we expect similar dynamics in supply chain coordination, where detailed technical disclosure facilitates compatible investment. We therefore hypothesize:

H1: Focal firms' AI-related scientific publications are positively associated with subsequent AI adoption among their suppliers.

This prediction is consistent with the coordination mechanism we propose, though we acknowledge that observational data cannot definitively distinguish coordination effects from knowledge spillovers or other mechanisms. We develop cross-sectional predictions that help differentiate our proposed mechanism from alternatives.

3.2 Supply Chain Characteristics and Coordination Needs

If AI publications function as coordination mechanisms, their effects should vary with supply chain characteristics that affect coordination difficulty such as relationship duration, the number of suppliers, and supply chain complexity.

First, coordination challenges are more acute in shorter-duration relationships. Partners in long-standing relationships develop mutual understanding through repeated interaction, substituting for public signals. Costello (2013) documents that contract duration reflects parties' accumulated relational capital. Suppliers in newer relationships face greater uncertainty about the focal firm's strategic direction and cannot rely on relationship history to inform expectations. Public disclosure through AI publications therefore provides greater incremental coordination value when relationships are relatively short.

H2a: The positive association between focal firms' AI publications and supplier AI adoption is stronger for supply chains characterized by shorter average relationship duration.

Coordination becomes more difficult as the number of parties increases. With few suppliers, bilateral communication may suffice to align expectations. Crawford, Huang, Li, and Yang (2020) find that customer concentration affects disclosure policy, consistent with private channels substituting for public disclosure when few parties must be informed. As the supplier base expands, bilateral coordination becomes impractical. Public disclosure scales efficiently: a single publication reaches all suppliers simultaneously and creates the common knowledge necessary for coordination across a large network.

H2b: The positive association between focal firms' AI publications and supplier AI adoption is stronger for firms with a larger number of suppliers.

Supply Chain Complexity. More complex supply chain structures create additional coordination challenges. Industry diversity increases coordination difficulty because suppliers in different industries may have different technological requirements. Supply chain depth creates challenges because information must propagate across multiple tiers. Pandit, Wasley, and Zach (2011) document that information flows along supply chains; public disclosure facilitates this flow by being equally observable to all tiers.

H2c: The positive association between focal firms' AI publications and supplier AI adoption is stronger for firms with more complex supply chain structures.

3.3 Economic Consequences

If AI publications facilitate supplier AI adoption, and if supplier AI adoption enhances supply chain performance, then supplier AI adoption should mediate the relationship between focal firm publications and economic outcomes.

The logic connecting supplier AI adoption to focal firm performance operates through supply chain efficiency. When suppliers develop AI capabilities that complement the focal firm's AI investments, the supply chain operates more effectively—through improved demand forecasting, better quality control, or faster response to changing conditions. These efficiency gains ultimately support the focal firm's ability to generate revenue.

Importantly, our framework predicts that the effect of AI publications on focal firm performance operates through supplier AI adoption rather than through direct effects on the focal firm's own capabilities. The coordination mechanism benefits the focal firm by enabling complementary investments throughout the supply chain.

H3: Supplier AI adoption mediates the relationship between focal firms' AI publications and subsequent firm performance.

4. Data and Summary Statistics

4.1. Data

The data on AI-related publications are from the Emerging Technology Observatory (ETO), a research initiative housed at Georgetown University's Center for Security and Emerging Technology. ETO identifies AI-related publications using machine learning models applied to its Merged Academic Corpus, which contains over 280 million scholarly articles worldwide.[2] A publication is associated with a company if at least one of its authors lists an affiliation with that company. The dataset provides comprehensive coverage of major U.S. publicly traded firms, including full coverage of the S&P 500.

We adopt the measure of firms' investments in AI proposed in Babina et al. (2024), which is a human-capital-based measure of firm-level AI investments. It captures each firm's employees who are AI-skilled. The information on employees' AI-related skills is from Cognism, an aggregator of employment profiles for lead generation and client relationship management services. Stock returns are from the Center for Research on Security Prices (CRSP). Accounting data and firm characteristics are from the Compustat database. Our sample consists of U.S. publicly traded firms from 2014 to 2019. We begin in 2014, the first year for which AI-related publication data are available, and end in 2019 because the AI investment data from Babina et al. (2024) conclude in 2020, allowing one additional year to capture firms' future AI investments.

We obtain detailed firm-level supply-chain data from FactSet Revere. The FactSet Supply Chain Relationships dataset was built to identify business relationship interconnections among companies globally. FactSet Revere creates a relationship classification structure consisting of four main categories: Supplier, Customer, Competitor, and Strategic Business Partner.[3] Both Direct and Reverse Relationships are identified for each company, as available. Direct relationship is

[2] For details, please see https://eto.tech/dataset-docs/private-sector-ai-indicators/.
[3] Strategic business partners, for example, include joint venture investors, firms with research collaboration, or firms with integrated product offerings.

disclosed by a company of interest, and a reverse relationship is disclosed by another company with the company of interest. As a result, FactSet Supply Chain Relationships offers a comprehensive network of relationship interconnections, many of which would not have been captured by simply analyzing a company's public filings.

4.2. Descriptive Statistics

We begin by summarizing the distribution of AI knowledge production and key firm characteristics. Table 1 presents the descriptive statistics of our sample. Detailed variable definitions and Compustat item codes are provided in Appendix A. Panel A reports the temporal distribution of AI-related publications. We observe a monotonic increase in AI research output over our sample period. The number of AI publications more than doubled from 2,592 in 2014 to 6,609 in 2019, representing a shift from 11.25% to 28.70% of the total sample volume. This trend highlights the increasing strategic importance that firms are placing on fundamental scientific research in artificial intelligence.

Panel B breaks down AI publications by industry. The distribution is highly concentrated, with the Manufacturing (20-39) and Service (70-89) sectors accounting for the vast majority of AI research output (30.61% and 56.04%, respectively). This heterogeneity suggests that AI applications are particularly salient in industries characterized by complex production processes and high intangible asset intensity.

Panel C provides summary statistics for the variables used in our analysis. We start with the average number of AI-related employees of suppliers, customers, strategic partners, or competitors for a focal firm. Each supplier firm for a focal firm in our sample, on average, has 50 AI-related employees (*Supplier_AIEmp_t+1*), which is about 1.6% of their workforce. Firms classified as customers, on average, hire 54 employees with AI skills. The values are 45 employees and 65 employees for firms classified as strategic partners and competitors, respectively. The number of AI-skilled employees for all four categories is highly skewed, with the 75th percentile all zero, suggesting the majority of companies in the global supply chains have not materially embraced AI-related technologies during our sample period. And the average number of AI-related employees for a focal firm is 13. Turning to the AI-related publication, we find that, on average, a focal firm contributes 5 AI-related research publications and 6 AI-related patents per year.

5. The Dissemination of AI Knowledge in Supply Chains

5.1 Empirical Design and Baseline results

Our primary hypothesis posits that a focal firm's engagement in AI research facilitates the adoption of AI technology among its upstream suppliers. To test this, we regress the supplier's future AI employment (*Supplier_AIEmp_t+1*) on the focal firm's current AI publications (*AI Pub_t*), controlling for a vector of firm characteristics and fixed effects:

Eq. (*1*):

$$Supplier\_AIEmp_{t+1} = \beta_1 AI\ Pub_t + Controls + \ Firm\ FE + Year\ FE\ + \epsilon_{i,t}$$

Table 2 presents the baseline results. Column (1) estimates the relationship without controls and fixed effects, yielding a positive and statistically significant coefficient. In Column (2), we include a battery of important firm characteristics as control variables. In Column (3), we include control variables and firm and year fixed effects. Our control variables include firm size, number of suppliers, leverage, book-to-market ratio, capital expenditure ratio, R&D scaled by assets, fixed investments scaled by assets, cash ratio, ROA, the percentage change in sales, and an indicator variable that equals one if a firm reports negative net income. Firm fixed effects control for unobserved firm characteristics that do not vary across time. Time fixed effects control for unobserved factors that vary over time, such as macroeconomic conditions.

We find that the coefficients on the focal firm's AI publications are positive and statistically significant after adding control variables and fixed effects. Column (3) shows that firms with larger assets, more suppliers, higher book-to-market, higher capital expenditure, or lower ROA are associated with higher AI investments among their suppliers. In Column (4), we further control for the lagged number of AI-skilled employees by both focal firms and their suppliers in year t. The number of AI-skilled employees at focal firms is not related to suppliers' future AI employment. As expected, lagged suppliers' AI employment is strongly and positively related to their future AI employment. Importantly, after including these two additional control variables, the relationship between focal firms' AI publications and their suppliers' AI adoptions remains positive and significant. The coefficient on *AI Pub_t* in Column (4) is 0.181 (p-value = 0.01), suggesting that a one-standard-deviation increase in *AI Pub_t* is associated with a 7.1% increase in suppliers' AI employment in year t+1, indicating that the effect is economically meaningful.[4]

Overall, these findings suggest a spillover effect: knowledge creation at the customer level appears to induce technological upgrading upstream in the supply chain.

5.2. Characteristics of Supply Chain

To better understand the mechanisms driving this propagation, we examine how supply chain characteristics influence the relationship between the focal firm's AI-related research and

[4] The economic magnitude is obtained by solving for R in the equation IHS(13.541 * (1+R)) * 100 - IHS(13.541)*100 = 0.181* 38.417, where 13.541 is the mean value of untransformed suppliers' AI employment in year t+1.

supplier AI adoption. We measure supply chain characteristics at the SIC 3-digit level, ensuring that the measures are relatively exogenous to each individual firm.

First, we expect the effect to be stronger when supply chain relationships are relatively short in duration. Shorter business relationships tend to exacerbate incomplete contracting and heighten the risk of hold-up, as neither party has accumulated enough relational capital to fully mitigate opportunistic behavior. Under these conditions, firms may rely more heavily on public disclosures to signal credibility, commitment, and long-term intent to their partners. Columns (1) and (2) of Table 3 partition the sample based on the average duration of the supply chain relationship. Consistent with our expectation, the positive relationship is statistically significant only for industries with relatively short supply chain durations but not significant for those with long and stable relationships, with the difference being statistically significant.

Second, we expect the effect to be stronger when firms have a larger number of suppliers. Managing many supplier relationships increases coordination complexity and reduces the ability to maintain close, bilateral communication with each partner. As the supplier base expands, information asymmetry and uncertainty about the firm's long-term commitment also grow, making public disclosures a more efficient and scalable mechanism for communicating strategic priorities and assuring suppliers of the firm's continued investment and reliability. Columns (3) and (4) of Table 3 partition the sample based on the average number of suppliers. We find that the effect of AI publications on suppliers' subsequent AI adoption is stronger when firms have a greater number of suppliers.

Lastly, we expect the effect to be stronger when the focal firm's supply-chain relationships are more complex. Greater complexity raises coordination costs, increases information frictions, and makes it harder for firms to monitor and manage their supplier networks. We capture this complexity using two measures. The first is industry dispersion, the number of distinct industries of suppliers, which reflects the technological and operational heterogeneity the firm must manage. The second is supply-chain length, which captures how many tiers upstream a firm's supplier network extends. Longer chains imply more intermediaries and greater uncertainty (Li, 2024). When supplier networks are more diverse and complex, firms face a higher risk of misalignment and misunderstanding, making the value of clear and credible public disclosures even more significant.

Columns (5) and (6) report the subsample results based on supply-chain length, and Columns (7) and (8) report the results based on the number of distinct supplier industries. We find that the effect is stronger for firms with more diversified suppliers and longer supply-chain networks. This evidence suggests that knowledge dissemination through publications is especially valuable for coordinating AI adoption across complex and diverse supply chains.

5.3. Economic Consequences: Mediation Analysis

In this subsection, we examine whether suppliers' AI adoption translates into economic benefits for the supply chain by focusing on the focal firm's future sales. We use a mediation framework to test whether suppliers' AI adoption serves as a channel through which a firm's AI research influences its subsequent sales. Specifically, we estimate a structural equation model that links a firm's future sales to its current AI publications and to suppliers' AI employment in year t+1. We present the estimation results of the structural equation model in Table 4.

Column (1) of Table 4 shows that supplier AI employment (*Supplier_AIEmp_t+1*) significantly predicts the focal firm's future sales (*Sales_t+1*). Column (2) reiterates our baseline finding that focal firm AI publications (*AI Pub_t*) predict subsequent supplier AI employment.

Column (3) includes both variables in the regression for focal firms' future sales. We observe that the direct effect of the focal firm's AI publications on its future sales becomes statistically insignificant when supplier AI employment is included, consistent with full mediation of the effect of the focal firm's AI publications through suppliers' AI adoption on its subsequent sales. The estimated Relative Indirect Effect (RIT) is 7.1%, suggesting that a meaningful portion of the economic value generated by corporate AI scientific research materializes through the technological upgrading of the focal firm's suppliers.

5.4. Alternative Channels and Robustness Tests

In this subsection, we first examine whether the knowledge spillovers from a focal firm's AI research extend beyond the supply chain to other stakeholders, such as customers, business partners, or competitors. Table 5 presents regressions of the AI employment of these groups on the focal firm's AI publications. The results show that the coefficients for customers (Column 1), business partners (Column 2), and competitors (Column 3) are all statistically indistinguishable from zero. This indicates that the positive knowledge spillovers documented in Table 2 are specific to upstream suppliers and do not generalize to other parties.

One interpretation is that firms strategically choose the content and level of their disclosures through AI publications to influence AI adoption among their supply chain partners, aligning incentives and facilitating technological upgrading where it matters most for operational efficiency. In contrast, customers, business partners, and competitors face strategic or contractual barriers that limit their access to firm-specific knowledge, which limit them taking advantage of the selected public disclosures (e.g., customers may not be directly involved in the production process, other business partners may operate in complementary but loosely coupled roles, and competitors may be deliberately shielded due to competitive concerns).

Together, these patterns suggest that knowledge spillovers are not merely incidental but are intentionally channeled through the supply chain, highlighting the focal firm's role in coordinating AI adoption among strategically important upstream partners.

Second, Table 6 presents a "horse race" analysis between AI publications and AI patents. While patents represent proprietary innovation, publications represent open science. When both *AI Pub_t* and *AI Patent_t* are included in the specification, only *AI Pub_t* remains positive and significant. This suggests that the open disclosure inherent in scientific publications is the primary driver of AI-related knowledge spillovers to the upstream supply chain, rather than proprietary knowledge protected by patents or general AI-related technology advancement of the focal firm.

Finally, Table 7 verifies the robustness of our main findings using alternative variable definitions. We replace the main independent variable with a binary indicator (*Dummy AI Pub*) in Column (1) and a ratio scaled by total publications (*Scaled AI Pub*) in Column (2). In Column (3), we use an alternative dependent variable scaled by total supplier employees (*Scaled AI Emp*). Across all specifications, the relationship between focal firm AI research publications and supplier AI adoption remains positive and statistically significant, confirming that our results are not driven by specific functional forms or measurement choices.

5.5. AI Research Publication and Stock Returns

In this subsection, we examine whether the AI research publications of focal firms are associated with subsequent stock market returns. In Table 8, we link focal firms' AI publication intensity in a year (the number of AI publications scaled by total assets) to their monthly stock market returns in the following year. If investors understand and appreciate the implications of focal firms' AI publications on firm valuations, we should observe that AI publication intensity predicts subsequent monthly returns.

Table 8 explores the capital market implications using Fama-Macbeth regressions (Fama and MacBeth (1973)). We find that AI publication intensity is a strong positive predictor of future monthly stock returns. This premium persists even after controlling for standard common firm characteristics that are shown to predict stock returns, including market capitalization, book-to-market ratio, momentum, profitability, and investment (Fama and French (1993, 2015) and Carhart (1997)). In Column (3), we further control for firms' R&D-to-total-asset ratio. In Column (4), focusing specifically on firm-years with at least one AI publication, the effect remains robust. These findings suggest that equity market investors recognize and reward corporate AI research capabilities—and the resulting efficiency gains in the production network—but that this recognition occurs with a short time lag over the subsequent month.

6. Conclusion

This paper examines whether corporate AI publications help address coordination problems in supply chains. We propose that AI publications mitigate hold-up concerns that would otherwise discourage suppliers from making relationship-specific investments in AI capabilities. Publications accomplish this through two channels: they place knowledge in the public domain, allowing suppliers to freely build on disclosed AI capabilities, and they signal the focal firm's sustained commitment to its AI trajectory, reducing suppliers' uncertainty about whether their investments will retain value.

Using data on AI publications matched with supply chain relationships from 2014 to 2019, we find support for this mechanism. Focal firms' AI publications predict subsequent AI adoption by upstream suppliers, and this effect is concentrated in supply chains characterized by shorter relationship duration, greater supplier network complexity, and longer supply chain length—precisely the conditions where incomplete contracting theory predicts coordination frictions are most severe. We further show that supplier AI adoption mediates the relationship between AI publications and focal firm future performance, suggesting that the economic value generated by AI publications materializes through the technological upgrading of the firm's suppliers. Consistent with our proposed mechanism, we find that publications—not patents—drive supplier AI adoption, and that knowledge spillovers flow specifically to suppliers rather than to customers, business partners, or competitors.

Our findings contribute to three literatures. First, we extend the accounting literature on incomplete contracting by identifying scientific publications as a novel mechanism to address hold-up problems in supply chains. Traditional solutions—vertical integration, long-term contracts, or private bilateral disclosure—are costly or impractical when supply chains are complex and technology evolves rapidly. We show that public scientific disclosure can serve as a credible commitment device that encourages relationship-specific investments by suppliers. Second, we identify a new motive for corporate scientific publication: facilitating technological upgrading among supply chain partners. Unlike prior explanations that focus on capital market or labor market benefits, our mechanism operates through real production relationships. Third, we contribute to research on innovation disclosure by showing that public disclosure and traditional IP protection serve complementary roles—firms strategically choose which innovations to publish and which to protect based on where value creation is greatest.

Our findings offer practical implications. For managers, the results suggest that strategic disclosure of AI research can serve as a coordination tool that encourages technological upgrading throughout the supply chain. For policymakers interested in promoting AI adoption, our evidence indicates that incentives encouraging corporate scientific publication could generate positive spillovers beyond the publishing firm. More broadly, our findings highlight that the value of innovation disclosure extends beyond traditional capital market benefits to include real effects on production relationships and supply chain coordination.

References

Agrawal, A., J. Gans, and A. Goldfarb. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence.* Boston, MA: Harvard Business Review Press.

Arora, A., S. Belenzon, and A. Patacconi. 2018. The decline of science in corporate R&D. *Strategic Management Journal* 39 (1): 3–32.

Arora, A., S. Belenzon, and L. Sheer. 2021. Knowledge spillovers and corporate investment in scientific research. *American Economic Review* 111 (3): 871–898.

Arora, A., and A. Gambardella. 1994. The changing technology of technological change: General and abstract knowledge and the division of innovative labour. *Research Policy* 23 (5): 523–532.

Babina, T., A. Fedyk, A. He, and J. Hodson. 2024. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics* 151: 103745.

Baker, G., R. Gibbons, and K. J. Murphy. 2002. Relational contracts and the theory of the firm. *Quarterly Journal of Economics* 117 (1): 39–84.

Baruffaldi, S. H., M. Simeth, and D. Wehrheim. 2023. Signaling research quality through scientific publications. Working paper, Copenhagen Business School.

Brynjolfsson, E., and A. McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies.* New York, NY: W.W. Norton & Company.

Bughin, J., J. Seong, J. Manyika, M. Chui, and R. Joshi. 2018. Notes from the AI frontier: Modeling the impact of AI on the world economy. McKinsey Global Institute Discussion Paper.

Carhart, M. M. 1997. On persistence in mutual fund performance. *Journal of Finance* 52 (1): 57–82.

Christensen, H. B., V. V. Nikolaev, and R. Wittenberg-Moerman. 2016. Accounting information in financial contracting: The incomplete contract theory perspective. *Journal of Accounting Research* 54 (2): 397–435.

Coase, R. H. 1937. The nature of the firm. *Economica* 4 (16): 386–405.

Cockburn, I. M., and R. M. Henderson. 1998. Absorptive capacity, coauthoring behavior, and the organization of research in drug discovery. *Journal of Industrial Economics* 46 (2): 157–182.

Coleman, B., K. Fronk, and K. Valentine. 2023. How does open innovation create private value for firms? Evidence from GitHub. Working paper, University of Georgia.

Costello, A. M. 2013. Mitigating incentive conflicts in inter-firm relationships: Evidence from long-term supply contracts. *Journal of Accounting and Economics* 56 (1): 19–39.

Fama, E. F., and K. R. French. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33 (1): 3–56.

Fama, E. F., and K. R. French. 2015. A five-factor asset pricing model. *Journal of Financial Economics* 116 (1): 1–22.

Fama, E. F., and J. D. MacBeth. 1973. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy* 81 (3): 607–636.

Gilje, E. P., T. A. Gormley, and D. Levit. 2020. Who's paying attention? Measuring common ownership and its impact on managerial incentives. *Journal of Financial Economics* 137 (1): 152–178.

Glaeser, S. 2018. The effects of proprietary information on corporate disclosure and transparency: Evidence from trade secrets. *Journal of Accounting and Economics* 66 (1): 163–193.

Glaeser, S. A., and W. R. Landsman. 2021. Deterrent disclosure. *The Accounting Review* 96 (5): 291–315.

Grossman, S. J., and O. D. Hart. 1986. The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy* 94 (4): 691–719.

Hart, O., and J. Moore. 1990. Property rights and the nature of the firm. *Journal of Political Economy* 98 (6): 1119–1158.

Hui, K. W., S. Klasa, and P. E. Yeung. 2012. Corporate suppliers and customers and accounting conservatism. *Journal of Accounting and Economics* 53 (1–2): 115–135.

Joskow, P. L. 1987. Contract duration and relationship-specific investments: Empirical evidence from coal markets. *American Economic Review* 77 (1): 168–185.

Klein, B., and K. B. Leffler. 1981. The role of market forces in assuring contractual performance. *Journal of Political Economy* 89 (4): 615–641.

Li, Y. 2024. Supply chain structure and corporate disclosure. Working paper.

McKinsey Global Institute. 2018. Notes from the AI frontier: Applications and value of deep learning. McKinsey & Company Discussion Paper.

Polidoro, F., Jr., and M. Theeke. 2012. Getting competition down to a science: The effects of technological competition on firms' scientific publications. *Organization Science* 23 (4): 1135–1153.

Rosenberg, N. 2010. Why do firms do basic research (with their own money)?. In Studies on Science and the Innovation Process: Selected Works of Nathan Rosenberg, edited by N. Rosenberg, 225–234. Singapore: World Scientific.

Simeth, M., and M. Cincera. 2016. Corporate science, innovation, and firm value. *Management Science* 62 (7): 1970–1981.

Stern, S. 2004. Do scientists pay to be scientists?. *Management Science* 50 (6): 835–853.

Valentine, K., J. L. Zhang, and Y. Zheng. 2025. Strategic publication by firms. Working paper.

Verrecchia, R. E. 1983. Discretionary disclosure. *Journal of Accounting and Economics* 5: 179–194.

Williamson, O. E. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications.* New York, NY: Free Press.

Williamson, O. E. 1979. Transaction-cost economics: The governance of contractual relations. *Journal of Law and Economics* 22 (2): 233–261.

Williamson, O. E. 1985. *The Economic Institutions of Capitalism.* New York, NY: Free Press.

Zucker, L. G., and M. R. Darby. 1996. Star scientists and institutional transformation: Patterns of invention and innovation in the formation of the biotechnology industry. *Proceedings of the National Academy of Sciences* 93 (23): 12709–12716.

Table 1. Descriptive Statistics
Panel A: AI Publications over Time

Year	Num. AI Pub.	%AI Pub
2014	2,592	11.25%
2015	2,635	11.44%
2016	2,929	12.72%
2017	3,450	14.98%
2018	4,815	20.91%
2019	6,609	28.70%
Total	23,030	100.00%

Panel B: AI Publications across Industry

Industry	Num. AI Pub.	%AI Pub
Agriculture, Forestry, and Fishing (01-09)	0	0.00%
Mining (10-14)	107	0.46%
Construction (15-17)	60	0.26%
Manufacturing (20-39)	7,050	30.61%
Transportation & Communications (40-49)	957	4.16%
Wholesale & Retail (50-59)	1,323	5.74%
Finance (60-67)	174	0.76%
Service(70-89)	12,905	56.04%
Others	454	1.97%
Total	23,030	100.00%

Panel C. Summary Statistics of Firm-Year Panel

	(1) N	(2) Mean	(3) STD	(4) P25	(5) Median	(6) P75
Supplier_AIEmp_t+1	43,137	49.856	144.442	0.000	0.000	0.000
Supplier_AIEmp_Scaled_t+1	43,137	0.016	0.304	0.000	0.000	0.000
Customer_AIEmp_t+1	43,137	54.106	133.570	0.000	0.000	0.000
Partner_AIEmp_t+1	43,137	46.935	136.613	0.000	0.000	0.000
Competitor_AIEmp_t+1	43,137	64.820	143.207	0.000	0.000	0.000
AI Pub_t	43,137	4.848	38.417	0.000	0.000	0.000
AI Patent_t	43,137	5.760	44.475	0.000	0.000	0.000
AI Pub_Dummy_t	43,137	0.021	0.144	0.000	0.000	0.000
AI Pub_Scaled_t	43,137	1.073	9.690	0.000	0.000	0.000
Size_t	43,137	5.637	3.219	3.625	6.006	7.859
PP&E_t	43,137	0.251	0.297	0.018	0.107	0.418
Leverage_t	43,137	0.477	1.534	0.016	0.179	0.403
BM_t	43,137	1.348	1.622	0.420	0.861	1.435
CAPX_t	43,137	0.041	0.067	0.002	0.017	0.050
XRD_t	43,137	0.069	0.197	0.000	0.000	0.029
Cash_t	43,137	0.223	0.273	0.032	0.100	0.307
ROA_t	43,137	-0.708	3.301	-0.188	0.005	0.044
Loss_t	43,137	0.468	0.499	0.000	0.000	1.000
SaleGrowth_t	43,137	0.125	0.602	-0.012	0.000	0.119
Focal AIEmp_t	43,137	13.296	54.285	0.000	0.000	0.000
Supplier_AIEmp_t	43,137	43.987	133.770	0.000	0.000	0.000
Num. Suppliers_t	43,137	27.753	65.771	0.000	0.000	0.000

Notes: This table presents the descriptive statistics of the key variables. Panel A provides the distribution of AI publications across the years. Panel B provides the distribution of AI publications across industries. Panel C presents the summary statistics of key variables of the firm-year panel.

Table 2. AI Publications and Suppliers' AI Adoption

	(1)	(2)	(3)	(4)
Dep. Var	Supplier_AIEmp_t+1			
AI Pub_t	**0.472*** **	**0.179*** **	**0.228*** **	**0.181** **
	(0.000)	**(0.000)**	**(0.003)**	**(0.010)**
Num. Suppliers_t		1.321***	0.725***	0.033
		(0.000)	(0.000)	(0.229)
Size_t		2.225***	2.247***	2.052***
		(0.000)	(0.008)	(0.003)
PP&E_t		-43.404***	3.028	2.623
		(0.000)	(0.380)	(0.390)
Leverage_t		-0.547***	0.055	0.116
		(0.004)	(0.774)	(0.466)
BM_t		-3.889***	1.675***	1.479***
		(0.000)	(0.002)	(0.002)
CAPX_t		51.870***	9.747*	9.843**
		(0.000)	(0.088)	(0.050)
XRD_t		1.511	0.072	0.247
		(0.598)	(0.981)	(0.925)
Cash_t		7.635**	-1.760	-0.002
		(0.026)	(0.642)	(0.999)
ROA_t		0.686***	-0.252**	-0.209***
		(0.000)	(0.011)	(0.009)
Loss_t		12.797***	1.122	1.654
		(0.000)	(0.465)	(0.228)
SaleGrowth_t		0.199	-0.867	-0.672
		(0.776)	(0.151)	(0.222)
Focal AIEmp_t				0.061
				(0.102)
Supplier_AIEmp_t				0.427***
				(0.000)
Constant	47.568***	6.683**	12.279**	13.009***
	(0.000)	(0.025)	(0.015)	(0.002)
Observations	43,137	43,137	42,032	42,032
R-squared	0.016	0.395	0.785	0.810
Firm FE	No	No	Yes	Yes
Year FE	No	No	Yes	Yes

Notes: This table presents regression results from the regression of the average number of AI employees among suppliers in year t+1 on the focal firm's AI publications in year t. P-values, based on standard errors clustered at the firm level, are reported below the coefficient estimates in parentheses. ***, ***, * indicates statistical significance at 1%, 5%, and 10%, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

Table 3. Subsample Analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. Var	Supplier_AIEmp_t+1							
	Relation Duration		*Num. Supplier*		*Num. Supplier Layers*		*Num. Supplier Industry*	
	High	Low	High	Low	High	Low	High	Low
AI Pub_t	**0.058**	**0.305*****	**0.236****	**0.079**	**0.284****	**0.050**	**0.240****	**0.038**
	(0.550)	**(0.002)**	**(0.024)**	**(0.291)**	**(0.013)**	**(0.434)**	**(0.015)**	**(0.600)**
Diff (High – Low)	-0.247*		0.157		0.233*		0.202*	
P-Value	(0.075)		(0.221)		(0.077)		(0.100)	
Observations	19,981	21,935	19,669	22,254	20,699	21,226	20,876	21,058
R-squared	0.816	0.808	0.824	0.731	0.821	0.753	0.818	0.740
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table presents results from the regression of the average number of AI employees among suppliers in year t+1 on the focal firm's AI publications in year t within the corresponding subsample. Columns (1) and (2) present the results for subsamples of firm-years based on the average duration of supply chain relationships for each SIC 3-digit industry. Supplier relationship duration is the cumulative sum of the number of years that the supplier relationship has existed. Columns (3) and (4) present the results for subsamples of firm-years based on the average number of suppliers per firm for each SIC 3-digit industry. Columns (5) and (6) present the results for subsamples of firm-years based on the average number of layers of supply chain relationship for each SIC 3-digit industry. Columns (7) and (8) present the results for subsamples of firm-years based on the average number of unique supplier industries per firm within each SIC 3-digit industry. P-values, based on standard errors clustered at the firm level, are reported below the coefficient estimates in parenthesis. ***, ***, * indicates statistical significance at 1%, 5%, and 10%, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

Table 4. Mediation Analysis for Future Sales

	(1)	(2)	(3)
Dep. Var	Sales_t+1	Supplier_AIEmp_t+1	Sales_t+1
Supplier_AIEmp_t+1	0.006**	.	0.006**
	(0.037)	(.)	(0.040)
AI Pub_t		0.183**	0.015
		(0.018)	(0.133)
RIT (Indirect Effect / Total Effect)		**7.1%**	
Observations	33,753	33,753	33,753
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Notes: This table reports the mediating effect of suppliers' AI employees in the relationship between a focal firm's AI publications and its future sales, estimated using structural equation models. Column (1) shows the regression of sales in year t+1 on suppliers' AI employees in year t+1. Column (2) reports the regression of suppliers' AI publications in year t+1 on the focal firm's AI publications in year t. Column (3) presents the regression of sales in year t+1 on both suppliers' AI employees in year t+1 and the focal firm's AI publications in year t, including the latter as the mediator. P-values, based on standard errors clustered at the firm level, are reported below the coefficient estimates in parenthesis. ***, ***, * indicates statistical significance at 1%, 5%, and 10%, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

Table 5. AI Publication and Other Parties' AI Adoption

	(1)	(2)	(3)
	Customer	Business Partner	Competitor
Dep. Var		Supplier_AIEmp_t+1	
AI Pub_t	**-0.068**	**0.090**	**-0.004**
	(0.224)	**(0.180)**	**(0.916)**
Observations	42,032	42,032	42,032
R-squared	0.840	0.778	0.906
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Notes: This table presents regression results from the regression of the average number of AI employees among other parties in year t+1 on the focal firm's AI publications in year t. Column (1) presents the results on customers' AI employees. Column (2) presents the results on business partners' AI employees. Column (3) presents the results on competitors' AI employees. P-values, based on standard errors clustered at the firm level, are reported below the coefficient estimates in parentheses. ***, ***, * indicates statistical significance at 1%, 5%, and 10%, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

Table 6. AI Publications vs. AI Patents

	(1)
Dep. Var	Supplier_AIEmp_t+1
AI Pub_t	0.146**
	(0.038)
AI Patent_t	0.091
	(0.105)
Observations	42,032
R-squared	0.810
Controls	Yes
Firm FE	Yes
Year FE	Yes

Notes: This table presents the results of horse-race analysis of AI publications and AI patents on suppliers' AI employees. P-values, based on standard errors clustered at the firm level, are reported below the coefficient estimates in parentheses. ***, ***, * indicates statistical significance at 1%, 5%, and 10%, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

Table 7. Robustness Tests

	(1)	(2)	(3)
	Dummy AI Pub	Scaled AI Pub	Scaled AI Emp
Dep. Var		Supplier_AIEmp_t+1	
AI Pub_t	**26.755****	**0.270****	**0.000*****
	(0.015)	**(0.034)**	**(0.009)**
Observations	42,032	42,032	42,032
R-squared	0.810	0.810	0.651
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Notes: This table presents regression results from the regression of the average number of AI employees among suppliers in year t+1 on the focal firm's AI publications in year t using alternative measures. , using alternative measures. Column (1) uses a dummy variable equal to one if the firm publishes at least one AI article in a given year. Column (2) uses AI publications scaled by total publications. Column (3) uses suppliers' AI employees scaled by their total employees. P-values, based on standard errors clustered at the firm level, are reported below the coefficient estimates in parentheses. ***, ***, * indicates statistical significance at 1%, 5%, and 10%, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

Table 8. AI Publications and Stock Returns: Fama-Macbeth Regression

	(1)	(2)	(3)	(4)
		Full sample		Firm-year with AI pub
AI Pub intensity	0.003***	0.003***	0.003***	0.002***
	(0.000)	(0.000)	(0.000)	(0.003)
Market capitalization		0.000	0.000	-0.003***
		(0.600)	(0.475)	(0.000)
Book-to-market		-0.001	-0.001	-0.008**
		(0.435)	(0.597)	(0.036)
Momentum		0.000	0.000	0.006
		(0.859)	(0.887)	(0.369)
Profitability		0.001	0.001**	0.000
		(0.338)	(0.017)	(0.670)
Investment		0.001	0.001	-0.000
		(0.851)	(0.860)	(0.973)
R&D/Assets			0.013	0.053***
			(0.455)	(0.000)
Observations	325,230	280,198	280,198	16,137
No. of months	108	108	108	108
Adjusted R^2	0.0003	0.0289	0.0414	0.0893

Notes: This table presents the relationship between firms' AI publication intensity and their subsequent monthly stock returns. AI publication intensity is the number of AI publications contributed by a firm in a year, scaled by its total assets ($billion). Market capitalization is the firm's stock market capitalization. Book-to-market is the ratio of book value per share to the market price per share. Momentum is the cumulative past 12-month return, skipping the most recent month. Profitability is annual revenues minus cost of goods sold, interest expense, and selling, general, and administrative expenses divided by book equity. Investment is the annual percentage change in total assets. Market capitalization is taken the natural logarithm. Columns (1), (2), and (3) are based on the full sample. Column (4) focuses on firm-year observations with at least one AI publication. P-values based on Newey-West standard errors with 6 lags are reported below the coefficient estimates in parentheses. *, **, ***, represent significance at the 10%, 5%, and 1% levels, respectively.

Appendix A: Variable Definition

Variable Name	Definition
Supplier_AIEmp_t+1	The inverse hyperbolic sine (IHS) transformation of the average number of AI-related employees of suppliers in year t+1.
Supplier_AIEmp_Scaled_t+1	The average number of AI-related employees of suppliers scaled by average number of suppliers' total employees in year t+1.
Customer_AIEmp_t+1	The inverse hyperbolic sine (IHS) transformation of the average number of AI-related employees of customers in year t+1.
Partner_AIEmp_t+1	The inverse hyperbolic sine (IHS) transformation of the average number of AI-related employees of business partners in year t+1.
Competitor_AIEmp_t+1	The inverse hyperbolic sine (IHS) transformation of the average number of AI-related employees of competitors in year t+1.
AI Pub_t	The inverse hyperbolic sine (IHS) transformation of the number of AI-related publications by a firm in year t.
AI Patent_t	The inverse hyperbolic sine (IHS) transformation of the number of AI-related patents by a firm in year t.
AI Pub_Dummy_t	A dummy variable that equals 1 if a firm publish at least one AI-related article in year t.
AI Pub_Scaled_t	The number of a firm's AI-related publications scaled by its total number of publications in year t.
Size_t	The natural logarithm of total assets at the end of year t (Compustat item AT).
XRD_t	R&D expenses scaled by total assets at the end of year t (Compustat items (XRD / AT).
BM_t	Book-to-market ratio at the end of year t (Compustat items (AT - DLTT – DLC) / (PRCC_F*CSHO)).
CAPX_t	Capital expenditures scaled by total assets at the end of year t (Compustat items CAPX / AT).
PP&E_t	Net value of property, plant, and equipment scaled total assets at the end of year t (Compustat items PPENT /AT).
Leverage_t	Total liabilities scaled by total assets at the end of year t (Compustat items (DLTT + DLC) / AT).
Cash_t	Total cash and cash equivalents scaled by total assets at the end of year t (Compustat items CHE /AT).
ROA_t	Net income scaled by total assets at the end of year t (Compustat items NI / AT).

Loss_t	An indicator variable that equals one if a firm reports negative net income in year t (Compustat items NI).
Sales Growth_t	The percentage change in sales from year t-1 to t (Compustat item SALE).
Focal AIEmp_t	The inverse hyperbolic sine (IHS) transformation of the average number of AI-related employees of the focal firm in year t+1.
Supplier_AIEmp_t	The inverse hyperbolic sine (IHS) transformation of the average number of AI-related employees of suppliers in year t.
Num. Suppliers_t	The inverse hyperbolic sine (IHS) transformation of the number of suppliers in year t.