AI Workforce Architecture and Digital Innovation: Complementarity, Governance, and the Value of Balance

Completed Research Paper

Christian Schumacher
Vienna University of Economics and Business
Vienna, Austria
christian.schumacher@wu.ac.at

Can Tihanyi
Vienna University of Economics and Business
Vienna, Austria
can.tihanyi@wu.ac.at

Abstract

Firms with comparable artificial intelligence headcounts can distribute that talent very differently across the functions AI work requires. We distinguish Invention, Application/Design, and Scaling/Governance roles and link Revelio Labs workforce data to Compustat and KPSS patent data for 3,236 publicly traded U.S. firms from 2012 to 2022. In firm and year fixed-effects Poisson pseudo-maximum likelihood (PPML) models, the Invention × Application/Design staffing interaction is associated with broader digital technology scope and entry into new digital subclasses, but not with greater patent volume. The relationship with forward citations is stronger where Scaling/Governance staffing is higher, and firms with more balanced role composition exhibit broader exploratory search. Alternative taxonomies and specifications yield the same pattern. For research on digital innovation and AI capability, this locates one source of innovation heterogeneity in the architecture of the AI workforce: which functions a firm covers, and how they combine.

Keywords: AI workforce composition, digital innovation, complementarity, human capital

Introduction

As artificial intelligence (AI) tools, cloud services, and off-the-shelf solutions lower experimentation costs and diffuse baseline AI capability, firms increasingly confront a similar opportunity set for AI-enabled innovation (Brynjolfsson et al., 2025; Nambisan et al., 2019). Yet firms do not innovate in the same way. Some use AI to open new digital search paths across products and processes, whereas others generate narrower or less exploratory outcomes, and many fail to convert AI adoption into meaningful innovation (e.g., Brynjolfsson et al., 2021). This heterogeneity suggests that one important differentiator is not AI adoption alone, but how firms organize AI work internally (Benbya et al., 2021; Lin & Maruping, 2025). Firms vary not only in the technical possibilities AI affords them but also in how they connect those possibilities to new value paths (Henfridsson et al., 2018; Nambisan et al., 2017; Yoo et al., 2010; Yoo et al., 2024). Those connections depend on a system of complementary organizational investments—among them process redesign, data resources, and managerial commitment—in which the AI workforce is a central element (Bresnahan et al., 2002; Mikalef & Gupta, 2021).

Prior research speaks to this heterogeneity only partially. Prior work shows that technology labor and complementary organizational capabilities shape returns to digital investment (Bresnahan et al., 2002; Tambe & Hitt, 2012), while recent studies examine the quantity, distribution, and labor-market value of AI

talent (Alekseeva et al., 2021; Babina et al., 2024; Tambe, 2026). Yet this literature largely treats AI talent as a stock rather than a functionally differentiated set of occupational roles. We therefore know little about whether the internal mix of AI roles is associated with the breadth, exploration, and impact of patent-based digital innovation. This gap matters because AI work inside firms is heterogeneous: some roles advance technical frontiers, others translate AI into domain-specific products and processes, and still others deploy, govern, and stabilize AI systems. IS research suggests that AI innovation differs systematically from other IT innovation and that organizational arrangements shape the kinds of innovation firms produce (Lin & Maruping, 2025). Related work likewise conceptualizes AI capability as multidimensional and as value creating through organizational, not merely technical, mechanisms (Benbya et al., 2021; Mikalef & Gupta, 2021). At the same time, digital innovation research emphasizes that exploratory recombination depends on combining diverse knowledge elements and linking them to unfamiliar application domains (Fleming, 2001; Henfridsson et al., 2018; Yoo et al., 2024). Together, these insights point to a question that current research has not addressed: how does the internal composition of AI talent relate to the kind of digital innovation firms produce?

We approach this question by treating the AI workforce as an internal architecture of functionally distinct but potentially complementary occupational roles. This workforce-level construct is related to, but distinct from, AI human capital and broader organizational AI capability, which also encompasses routines, structures, platforms, data resources, and coordination mechanisms (Mikalef & Gupta, 2021). Rather than treating AI talent as a homogeneous stock, we argue that firms differ in the kinds of AI work they are able to perform internally and that these differences matter for the direction of digital innovation. We distinguish three functional role categories. AI Invention roles are oriented toward frontier technical creation, including model development, experimentation, and algorithmic problem solving (Agrawal et al., 2019). AI Application/Design roles focus on embedding AI in domain contexts by translating technical possibilities into products, services, and processes and by framing use opportunities around user and organizational needs (Faraj et al., 2018; Lebovitz et al., 2021). AI Scaling/Governance roles focus on deployment, integration, reliability, documentation, quality assurance, and compliance, enabling firms to operationalize and stabilize AI-enabled innovations (Berente et al., 2021). These categories capture different but interdependent knowledge inputs: invention expands what is technically possible, application links technical possibilities to concrete use contexts, and scaling/governance helps embed innovations in organizational systems so that they can be deployed, maintained, and appropriated.

Drawing on complementarity theory (Athey & Stern, 1998; Ennen & Richter, 2010; Milgrom & Roberts, 1995) and the recombinant innovation lens (Fleming, 2001; Henfridsson et al., 2018), we argue that firms convert AI talent into digital innovation not through headcount alone, but through combinations of specialized roles that relax distinct innovation bottlenecks. Invention expands what is technically possible, but without strong links to domain problems it can remain disconnected from opportunities for digital value creation. Application/Design embeds AI in user settings, business processes, and product architectures, but without sufficiently strong invention capability it may be constrained to familiar technical repertoires and incremental use cases. Their joint staffing should therefore widen the firm's digital search space and support broader exploratory recombination. Scaling/Governance staffing plays a distinct conditional role: by supporting integration, documentation, monitoring, deployment, and oversight, it can shape whether novel combinations become reliable and organizationally consequential rather than remaining isolated experiments. This logic also suggests that workforce architectures without severe functional staffing gaps should be associated with broader exploratory search than equally large but highly concentrated workforces.

We test this theory in an observational matched firm-year panel of 3,236 publicly traded U.S. firms observed from 2012 to 2022, yielding a total of 18,010 firm-year observations. We combine Revelio Labs Workforce Intelligence data with Compustat and KPSS patent data to capture both the internal composition of firms' AI workforces and the character of their patented digital innovation. Revelio assigns each employment spell to a standardized occupational unit, so we see not only the scale of a firm's AI workforce but the kind of work it does. Building on these data, we develop a role-based AI workforce taxonomy and construct firm-level stocks of employees in AI Invention, AI Application/Design, and AI Scaling/Governance roles. We then link these workforce measures to patent-based indicators of digital innovation novelty and impact: digital Cooperative Patent Classification (CPC) breadth and new-to-firm digital subclass entry as indicators of recombinant novelty and exploratory search, and forward citations and Kogan–Papanikolaou–Seru–Stoffman (KPSS) economic value as indicators of innovation impact (Hall

et al., 2005; Kogan et al., 2017). Empirically, we find that the combination of AI Invention and AI Application/Design staffing is associated with broader digital technology scope and greater entry into new digital knowledge areas, but not with greater digital patent volume overall. We also find that this complementarity is linked to stronger innovation impact only among firms with greater AI Scaling/Governance staffing, and that firms with more balanced AI workforce compositions display broader and more novel digital innovation even after accounting for total AI intensity.

This study makes three contributions. First, we contribute to research on digital innovation heterogeneity. Whereas Lin and Maruping (2025) show that AI-enabled innovation differs from other IT innovation at the aggregate organizational level, we disaggregate the AI workforce into three functional roles and show that composition, not only scale, is associated with the character of patent-based digital innovation. Our discriminating tests distinguish breadth from count and novelty from derivative output, indicating that AI workforce architecture is one important factor conditioning which kinds of innovation emerge. More broadly, this contributes to digital innovation research that conceptualizes innovation as recombinant by showing how distinct AI roles are associated with firms' ability to generate new digital search paths (Hanelt et al., 2021; Henfridsson et al., 2018; Kohli & Melville, 2019; Yoo et al., 2010; Yoo et al., 2024). Second, we contribute to the emerging literature on AI capability by disaggregating AI human capital into distinct functional roles and showing that composition matters beyond scale (Aral & Weill, 2007). Whereas prior studies examine the amount of AI talent firms hire, where it is distributed, or how capital markets value it (Alekseeva et al., 2021; Babina et al., 2024; Tambe, 2026), we show that the internal mix of AI Invention, AI Application/Design, and AI Scaling/Governance roles helps explain why firms with comparable AI investments generate different innovation trajectories. This also responds to calls to unpack AI capability as a multidimensional organizational phenomenon rather than a monolithic technical asset (Benbya et al., 2021; Mikalef & Gupta, 2021). Third, we contribute to complementarity research by showing that complementary AI roles are associated with innovation quality, not just input scale. Specifically, the joint presence of invention-oriented and application-oriented roles co-varies with a broader digital search space, and scaling/governance staffing is linked to greater translation of novelty into impact, extending complementarity arguments in organizational and IT research to the internal composition of AI work (Bresnahan et al., 2002; Ennen & Richter, 2010; Milgrom & Roberts, 1995).

Theoretical Background and Hypotheses

AI Innovation as an Organizational Design Problem

AI innovation differs from conventional IT innovation in ways that make workforce configuration especially consequential. AI innovations are more process-oriented and less radical than comparable IT, with heterogeneous organizational approaches correlating systematically with innovation outcomes (Lin & Maruping, 2025). Managing autonomy, learning, and inscrutability requires dedicated capabilities for deployment, monitoring, and quality control (Berente et al., 2021), while hidden organizational work—ground truth creation, model evaluation, quality assurance—spans technical, domain, and governance expertise (Lebovitz et al., 2021). The two-tier structure from frontier research to domain adaptation further underscores that firms must decide not only whether to invest but at which tier and how to coordinate across tiers (Cockburn et al., 2019). Workforce composition is nevertheless only one part of the organizational setting in which innovation occurs: process redesign (Brynjolfsson et al., 2021), data resources and the capabilities built around them (Gregory et al., 2021; Mikalef & Gupta, 2021), and competitive conditions (Melville et al., 2004) all bear on innovation outcomes. We focus on the workforce element because it is theoretically central to the coordination of AI work and observable at scale.

The IS complementarity tradition establishes that returns on technology investment depend on which kinds of technical work a firm staffs, not on headcount alone: productivity gains emerge not from technology alone but from simultaneous investments in organizational change and skilled labor (Bresnahan et al., 2002; Melville et al., 2004). Dispersing algorithmic literacy across domain workers rather than concentrating it in isolated technical teams is likewise associated with higher returns on AI investments (Tambe, 2026), and integrated human capital systems that manage talent across recruitment, development, and assignment can deliver performance gains individual skills alone cannot (Ployhart et al., 2014; Tambe & Hitt, 2012). What remains less understood is how these principles translate into innovation outcomes when AI talent is functionally differentiated across roles.

AI Workforce Architecture: Three Functional Role Categories

Three functions recur across process accounts of innovation: frameworks that separate idea generation from conversion and diffusion (Hansen & Birkinshaw, 2007), acquisition from assimilation, transformation, and exploitation (Zahra & George, 2002), and the structuring, bundling, and leveraging of resources (Sirmon et al., 2011). Specializing this logic to AI, work distinguishing frontier development from application-oriented activity (Cockburn et al., 2019) and work on the deployment, monitoring, and coordination required to make AI systems organizationally usable (Berente et al., 2021; Lebovitz et al., 2021) motivate a corresponding separation of creation, domain embedding, and operationalization. We therefore distinguish Invention, Application/Design, and Scaling/Governance as the three functional categories of AI work.

AI Invention encompasses frontier technical creation: the development of novel algorithms, model architectures, training techniques, and proprietary AI assets. Workers in this domain include AI research scientists, machine learning engineers, deep learning engineers, and generative AI engineers whose primary function is to push the technical frontier. AI Application/Design encompasses domain embedding, product translation, and opportunity framing: the identification and integration of AI capabilities into specific business problems, products, and decision contexts. Workers in this domain include AI product owners, data analysts, and domain specialists who translate technical possibilities into organizational applications. AI Scaling/Governance encompasses staffing in AI-related roles whose primary functional contribution is to deploy, operationalize, monitor, scale, and govern AI systems in production environments, making AI innovations production-grade, reproducible, and protectable. Workers in this domain include machine-learning operations (MLOps) engineers, enterprise AI architects, data governance specialists, and compliance professionals whose collective function is to connect AI development to reliable organizational use. The Scaling/Governance label identifies a workforce-role category; it is not a direct measure of organizational governance capability.

This decomposition maps onto established theoretical distinctions: Invention corresponds to the frontier AI tier (Cockburn et al., 2019); Application/Design to the dispersed algorithmic literacy and organizational embedding emphasized in recent IS research (Lin & Maruping, 2025; Tambe, 2026); and Scaling/Governance to the management and coordination challenges identified as distinctive to AI (Berente et al., 2021; Lebovitz et al., 2021; Tiwana & Konsynski, 2010). The framework is firm-level: it captures not the presence of individual AI workers but the collective role configuration that emerges from how firms distribute AI human capital across functionally distinct roles. It parallels the layered modular architecture of digital innovation (Yoo et al., 2010), where innovation emerges from coupling modular layers: Invention produces the technical foundations, Application/Design connects them to domain contexts and products, and Scaling/Governance sustains deployment and integration (Agrawal et al., 2019; Benbya et al., 2021). The theoretical significance lies in predictions about patent-based digital innovation outcomes. Digital innovation is recombinant: it emerges through combining digital resources, knowledge, and capabilities into new value paths (Henfridsson et al., 2018; Yoo et al., 2024), with novelty rising as combinations draw from distant domains (Fleming, 2001). In our setting, one observable set of recombination inputs is the workforce composition across specialized roles carrying knowledge for distinct stages of AI-related work. Different workforce configurations therefore imply different patterns of digital recombination, as we develop below.

Recombinant Novelty through Invention and Application/Design Complementarity

The recombinant innovation perspective holds that digital innovation value comes from novel combinations of digital resources, knowledge components, and organizational capabilities (Henfridsson et al., 2018). Innovation novelty increases when firms search across distant knowledge domains rather than recombining familiar components within established fields (Cassiman & Veugelers, 2006; Fleming, 2001; March, 1991). We therefore expect the joint staffing of AI Invention and AI Application/Design roles to support cross-domain recombination in the AI innovation context.

Invention talent generates technical knowledge at the AI frontier: new algorithms, model architectures, and training techniques. On its own, however, frontier technical knowledge tends to produce innovation in familiar technical domains. Machine learning researchers solve problems they know, generating technically

sophisticated but narrowly clustered innovation in established AI patent classes. This is the natural search behavior of specialized technical experts whose local search routines exploit existing knowledge rather than exploring unfamiliar application domains (March, 1991; Nelson & Winter, 1982). A firm that invests exclusively in Invention may produce deep technical innovations (better classifiers, faster training methods, improved architectures) but these remain concentrated in the technical domains the firm already knows. The resulting innovation portfolio is technically deep but recombinantly narrow.

Application/Design talent introduces a qualitatively different form of knowledge: domain problems, user needs, market opportunities, and process inefficiencies. These workers do not push the AI frontier directly; rather, they identify where frontier technical capability connects to unsolved problems in unfamiliar application domains. On their own, Application/Design workers cannot build novel technical solutions. They can identify opportunities but lack the means to create frontier AI capabilities to address them. But when combined with Invention, they redirect technical search toward domains the firm has not previously patented in, expanding the recombinant search space. Critically, this interaction is not unidirectional. Invention provides the technical building blocks that Application/Design redirects toward novel contexts, but Application/Design also generates demand-side problem signals that pull Invention toward unfamiliar technical challenges. Domain problems identified by Application/Design workers create specification pressure that directs frontier technical work toward new areas, while the availability of proprietary technical capability shapes which domain opportunities Application/Design workers can credibly pursue. The complementarity therefore operates through mutual redirection rather than a simple technology-push mechanism.

Under complementarity logic, the marginal contribution of each functional role should be greater when the other is also present. Proprietary AI techniques may become more useful for exploratory recombination as domain-embedded workers identify applications and problem specifications, while domain embedding may be more consequential when it can draw on internally developed technical knowledge that can be customized to firm-specific data and processes (Choudhury et al., 2020; Tambe, 2026). The resulting interaction is consistent with design recombination (Henfridsson et al., 2018) and with the broader complementarity literature in which the value associated with one input depends on the configuration of other inputs (Bresnahan et al., 2002; Ennen & Richter, 2010; Krakowski et al., 2023; Milgrom & Roberts, 1995). We therefore hypothesize:

Hypothesis 1: The Invention × Application/Design staffing interaction is positively associated with more recombinant patent-based digital innovation, reflected in broader digital technology scope and entry into new digital technology domains.

From Novelty to Impact through Scaling/Governance

Recombinant novelty is necessary but not sufficient for impactful patent-based digital innovation. Novel AI applications that remain prototypes that are unstable, unvalidated, and undocumented do not generate the concrete technical implementations that become well-cited patents or economically valuable innovations. The gap between demonstrating that a novel AI application works in a development environment and building a production system with documented architecture, reproducible results, and sufficient implementation specificity is what Scaling/Governance staffing is oriented toward closing.

The theoretical argument draws on two streams. First, work on resource orchestration directs attention to the steps between holding a knowledge input and obtaining an outcome from it: the same portfolio of technical and domain knowledge yields different results depending on how it is assembled and put to work (Sirmon et al., 2011). In a patenting setting this shifts the question from whether a firm can generate a novel Invention × Application/Design combination to whether that combination survives the passage from working prototype to a reproducible, documented, production-tested system, the form in which it becomes disclosable as a patent and citable thereafter. Scaling/Governance roles are the workforce input that operates at this passage: they build the deployment pipelines, monitoring, and documentation through which a working prototype becomes a reproducible production system. Second, research on AI management has established that deploying AI reliably requires dedicated attention to autonomy, learning, and inscrutability, all of which demand organizational capability for monitoring, quality control, and accountability (Berente et al., 2021; Lebovitz et al., 2021; Raisch & Krakowski, 2021). The complementary

investments in process redesign, training, governance, and organizational restructuring also shape returns from technology (Brynjolfsson et al., 2021; Gregory et al., 2021).

Three channels motivate the conditional prediction. First, workers in Scaling/Governance roles can support reproducible systems through model registries, deployment pipelines, and documentation, increasing the specificity and implementability of the firm's technical innovations—and patents describing production-tested systems generate more forward citations than proof-of-concept disclosures (Hall et al., 2005). Second, codification enables iterative building on prior innovations, creating cumulative knowledge assets of greater technical depth and economic value. Third, governance capability—auditing, bias monitoring, compliance, access controls—enables deployment in higher-stakes domains, expanding the range of contexts in which the Invention×Application/Design complementarity can generate impact. Accordingly, we expect the positive Invention × Application/Design staffing association with patent-based digital innovation impact to be stronger at higher levels of Scaling/Governance staffing. We therefore hypothesize:

Hypothesis 2: The positive Invention × Application/Design staffing interaction is more strongly associated with patent-based digital innovation impact at higher levels of AI Scaling/Governance staffing.

Architecture Over Headcount: Functional Coverage and Exploratory Breadth

The preceding hypotheses concern specific staffing interactions. We now turn to a broader architectural implication: whether the distribution of AI workers across the three functional categories is associated with exploratory breadth after accounting for total AI workforce scale. Complementarity arguments imply that severe weakness in one necessary function can constrain what a system can accomplish even when other functions are well staffed (Ennen & Richter, 2010; Milgrom & Roberts, 1995). The prediction is therefore about functional coverage: configurations with major staffing gaps should show narrower search than configurations that maintain meaningful staffing across Invention, Application/Design, and Scaling/Governance.

We operationalize this idea as balance across the three role categories. Balance does not mean equal numbers of workers, nor does it identify an optimal configuration relative to firm strategy, industry, size, or AI maturity. It captures the absence of severe functional staffing gaps: whether AI work is highly concentrated in one category or distributed across the distinct functions involved in developing, applying, and operationalizing AI. The Herfindahl-Hirschman concentration measure therefore serves as a proxy for functional coverage, not for a theoretically optimal proportional allocation. The bottleneck rationale is consistent with the innovation-value-chain logic of Hansen and Birkinshaw (2007), which separates idea generation, conversion, and diffusion and emphasizes that a severe weakness in one stage can constrain the overall innovation process. In our setting, very limited or absent staffing in one functional category may restrict the range of technological opportunities that firms can pursue and combine. The prediction is thus about missing functional complements and the extensive margin of role presence more than about fine-grained proportional evenness once all three functions are already staffed.

The alternative hypothesis is specialization: firms that concentrate AI talent in one domain may develop deeper expertise and produce more impactful innovation through depth rather than breadth (Krakowski et al., 2023; Tambe et al., 2020). Our prediction is therefore discriminating rather than general: balance should be associated with exploratory breadth, whereas concentration may still support depth within established technical domains. Digital CPC Breadth captures the technological scope of digital patenting, while New Digital Subclasses captures entry into domains in which the firm has not previously patented. This framing also implies an extensive-margin interpretation: if the association mainly reflects the absence of missing role categories, fine-grained concentration need not remain consequential among firms that already staff all three categories. We therefore hypothesize:

Hypothesis 3: Balanced AI workforce composition is associated with broader exploratory search across digital technology domains.

Empirical Setting and Data

Our sample consists of 3,236 publicly traded U.S. firms observed from 2012 to 2022, yielding 18,010 firm-year observations after applying the one-year lag structure used in our design. The main analysis uses

outcomes through 2021 to mitigate patent truncation bias, as later patents may not yet have been granted or accumulated sufficient citations (Hall et al., 2005; Lerner & Seru, 2022; Lin & Maruping, 2025). We combine three data sources. To measure workforce composition, we use Revelio Labs Workforce Intelligence data. Revelio reconstructs employment spells for individual workers from publicly posted professional profiles and maps each spell onto a single standardized occupational unit. This data source has gained increasing traction in IS and management research (Marchetti & Puranam, 2026; Tambe, 2026), supported by validation evidence showing high correlations with firm-reported disclosures. For 2012 to 2021 the extract underlying this study contains more than 84 million employment spells covering more than 48 million distinct workers. Each record contains standardized occupational classifications for 17,000 fine-grained role categories, allowing us to move beyond broad firm-level AI workforce proxies and directly measure the internal architecture of AI labor within firms. Following prior research (e.g., Lin & Maruping, 2025), we measure digital innovation using patent-based indicators, which are widely used in IS research to capture the direction, novelty, and economic significance of firm-level innovation. Specifically, we draw on the KPSS patent database (Kogan et al., 2017; December 2024 release, updated through the end of 2024), which provides patent-level forward citations and market-value estimates for all U.S. utility patents. We link patents to Compustat firms using the Center for Research in Security Prices (CRSP) permno identifier and the CRSP-Compustat merged (CCM) linking table. We identify digital patents using CPC technology class codes, specifically prefixes G06, G11, G16, and Y04, and construct firm-year measures of digital innovation novelty and impact. Finally, we use Compustat North America to obtain firms' financial characteristics. The observation window predates the broad diffusion of foundation-model tools, a scope condition we return to in the Discussion.

AI Workforce Taxonomy

The explanatory variables rest on a role taxonomy we built specifically for this study. Revelio independently standardizes tens of millions of firm-specific title variants into an activity-based hierarchy ranging from approximately 17,000 fine-grained occupations to ten broad groups. The clusters are learned from résumé and profile descriptions and from the responsibilities sections of job postings after removing seniority-related wording; at the individual level, occupation inference also draws on skills, job descriptions, prior titles, professional connections, and company and title context. The displayed occupation title is therefore a representative label for an activity-based cluster rather than the basis for membership. We work at the finest level so that roles such as machine learning engineer, data engineer, and business intelligence analyst remain distinct. From approximately 17,000 standardized occupations, combined keyword- and skill-based filtering identified 369 AI-related candidates. Three independent expert coders assigned these candidates to the three functions, which were specified before any occupation was coded, without inspecting outcome data, yielding 42 Invention, 231 Application/Design, and 96 Scaling/Governance occupations (Fleiss' κ = 0.84). Representative titles include AI research scientist, machine learning engineer, and deep learning engineer (Invention); AI product manager, AI data strategist, and domain analyst (Application/Design); and enterprise AI architect, MLOps engineer, and data governance specialist (Scaling/Governance). The coders applied one assignment rule to ambiguous cases: each standardized occupation was classified by its primary functional contribution—what the role principally produces—rather than by isolated words in its title. Remaining disagreements were resolved in discussion by returning to the same rule, so every final assignment reflects coder consensus. The same criterion governs all fourteen architect-titled occupations and analyst-titled occupations. In a narrower robustness taxonomy, 25 Invention, 43 Application/Design, and 29 Scaling/Governance occupations meet more restrictive inclusion criteria.

Measures

Dependent Variables

Our dependent variables capture patent-based digital innovation rather than AI innovation in general or the full range of firms' digital innovation activity. They therefore exclude AI-enabled process improvements, service innovations, product enhancements, and other digital innovations that do not result in patent filings. Following prior research (Benner & Tushman, 2003; Fleming, 2001; Katila & Ahuja, 2002; Lin & Maruping, 2025), we use two measures of exploratory breadth for H1 and H3. Digital CPC Breadth is the number of distinct four-character CPC subclasses (e.g., G06F, G06N, G11B, Y04S) represented in a firm's digital patent filings in a given year and thus reflects the technological scope of digital innovation. New

Digital Subclasses counts the CPC subclasses a firm enters for the first time, capturing exploratory search into unfamiliar digital technology domains. As a discriminating test, we also examine total digital patent count and non-digital patent count. Our theory implies that this complementarity concerns the character of digital innovation, its breadth and exploratory nature, rather than its volume, and that this pattern should be specific to digital rather than non-digital innovation. To test H2, which predicts that the Invention × Application/Design staffing interaction is more strongly associated with innovation impact at higher levels of AI Scaling/Governance staffing, we measure innovation impact using two indicators. Forward Citations captures the citations received by a firm's digital patents and is a widely used proxy for innovation influence and quality (Hall et al., 2005; Trajtenberg, 1990). KPSS Economic Value estimates patent value based on stock market reactions at grant (Kogan et al., 2017), providing a market-based complement to the citation-based measure. To test H3, which predicts that balanced AI workforce composition is associated with broader exploratory search across digital technology domains, we again use Digital CPC Breadth and New Digital Subclasses as our measures of innovation novelty.

Independent Variables

We measure role categories using Revelio's standardized occupational units rather than raw title strings, following prior AI-workforce research (Acemoglu et al., 2022; Alekseeva et al., 2021; Babina et al., 2024). We follow Babina et al. (2024) in constructing workforce intensity measures. For each AI workforce category Invention, Application/Design, Scaling/Governance, we compute AI workforce intensity as the share of firm employees in the respective category. Because role counts and headcount totals come from the same Revelio records, the ratio is internally consistent and requires no cross-source scaling. AI staffing could nonetheless proxy for a firm's general technical depth, so every specification carries a non-AI IT intensity term built from 790 IT occupations disjoint from the AI taxonomy (e.g., Software Engineer, Database Administrator, Network Engineer). The within-firm correlation between Application/Design intensity and non-AI IT intensity is modest ($r = 0.193$), consistent with related but non-identical forms of technical staffing. All continuous independent variables are winsorized at the 1st and 99th percentiles.

To test H1, we construct the interaction Invention Presence × Application/Design Intensity, where Invention Presence is a binary indicator equal to one if the firm employs at least one AI Invention worker. We use this extensive-margin specification because AI Invention roles are sparse, with roughly 70% of firm-year observations containing no Invention workers. The interaction estimates whether the association between Application/Design intensity and the outcome differs when Invention staffing is present, which is the observable implication of complementarity in a design of this kind. In robustness checks, we replace the binary indicator with continuous Invention Intensity. For H2, the main analysis splits Invention-presence observations at the median of Scaling/Governance Intensity and compares the Invention × Application/Design interaction across subsamples. Because the median threshold is not unique, we also report a pooled continuous three-way interaction. For H3, we compute the Herfindahl-Hirschman Index (HHI) of AI workforce concentration across the three categories, where higher values indicate a more concentrated and less balanced AI workforce architecture. We include total AI intensity as a control for overall AI workforce scale. As a robustness check, we also use an entropy index as an alternative measure of workforce balance.

Control Variables

All control variables are measured with a one-year lag. We control for firm size (log total assets), liquidity (cash holdings scaled by total assets), book leverage, growth opportunities (log Tobin's Q, winsorized), intangible asset intensity, capital intensity (net property, plant, and equipment scaled by total assets), industry concentration (three-digit North American Industry Classification System [NAICS] HHI), and non-AI IT workforce intensity. Together, these controls account for firm resources, financial structure, investment profile, industry conditions, and broader digital labor capacity that may independently shape patent-based digital innovation. All continuous control variables are winsorized at the 1st and 99th percentiles.

Empirical Model

We estimate Poisson pseudo-maximum likelihood (PPML) models with high-dimensional fixed effects (Correia et al., 2020):

$$E[DV_{i,t} \mid X] = \exp(\beta' X_{i,t-1} + \gamma' Z_{i,t-1} + \mu_i + \tau_t)$$

where X denotes the vector of AI workforce variables and interactions, Z denotes the vector of controls, μ_i are firm fixed effects, and τ_t are year fixed effects. We cluster standard errors by firm. PPML is well suited to our setting because the dependent variables are non-negative outcomes with a substantial share of zeros, and the estimator remains consistent under heteroskedasticity without requiring strong distributional assumptions (Santos Silva & Tenreyro, 2006). This property is also useful for the KPSS economic value measure, which is continuous and highly right-skewed rather than count-valued. PPML remains appropriate in this case because it requires only correct specification of the conditional mean and accommodates both the mass of zeros and the extreme right tail that characterize the variable (Santos Silva & Tenreyro, 2011). Firm fixed effects absorb time-invariant unobserved heterogeneity across firms, while year fixed effects account for common temporal shocks, such as shifts in AI labor markets and patent office processing cycles. The PPML firm-fixed-effects sample retains 5,225 (Digital CPC Breadth), 4,465 (New Digital Subclasses), and 7,119 (non-digital) firm-years of 18,010; firms with zero within-firm variation in an outcome are dropped by construction, so identification concentrates on firms that actively patent digitally and exhibit AI-capability variation over 2012–2022. Pairwise correlations among the three focal AI workforce measures do not exceed 0.15, and all variance inflation factors in the H1/H2 and H3 specifications are below 2.8 (mean VIF = 1.89 and 1.84), indicating no collinearity concern. Relative to firms dropped because an outcome has no within-firm variation, retained firms are, on average, larger, more technology-oriented, and more patent-active. Several features of the research design help mitigate endogeneity concerns. First, the one-year lag between workforce composition and innovation outcomes reduces simultaneity bias. Second, firm fixed effects eliminate time-invariant confounders, such as persistent organizational culture, industry positioning, and managerial orientation toward innovation. Third, year fixed effects absorb economy-wide and technology-wide shocks that may affect all firms simultaneously. In robustness analyses, we additionally control for patent stock intensity (total patent stock scaled by total assets), following Lin and Maruping (2025), to account for pre-existing innovation capability. Although these design choices cannot fully rule out reverse causality, the combination of lagged predictors, within-firm estimation, and extensive controls substantially reduces the scope for alternative explanations.

Results

Table 1 reports descriptive statistics on our main variables.

Variable	Mean	SD	Min	Max
Digital CPC Breadth	2.07	2.48	0	18
New Digital Subclasses	0.27	0.62	0	6
Fwd Citations (Digital)	2,023.66	43,079.16	0	1,819,542
KPSS Econ. Value (Digital)	8,923.45	144,388.60	0	7,643,977
AI Invention Intensity (per 100)	0.12	0.44	0	9.54
AI App./Design Intensity (per 100)	1.24	1.23	0	25.00
AI Scaling/Gov. Intensity (per 100)	0.78	0.92	0	13.14
AI Workforce HHI	0.58	0.16	0.33	1.00
Non-AI IT Intensity (per 100)	17.91	15.11	0	106.46

Table 1. Descriptive Statistics

Table 2 reports PPML estimates testing H1. The focal term is the Invention Presence × Application/Design Intensity interaction. A positive coefficient indicates that the association between Application/Design

staffing and patent-based digital innovation is stronger when Invention staffing is present, the pattern H1 predicts.

In Model (1), where Digital CPC Breadth is the dependent variable, the interaction is positive and statistically significant (β = 0.118, SE = 0.031, p < 0.001). For New Digital Subclasses, it is likewise positive and significant (β = 0.357, SE = 0.126, p < 0.01). The main effects are not statistically significant in these models, so the empirical pattern is concentrated in the conditional staffing association. A 25th-to-75th-percentile increase in Application/Design Intensity among firms with Invention Presence is associated with 19.1% greater Digital CPC Breadth (p = 0.001) and 70.0% more New Digital Subclasses (p = 0.028). Models (3) and (4) provide discriminating evidence about the scope of this association. The interaction is small and statistically insignificant for total digital patent count (β = −0.030, SE = 0.048) and effectively zero for non-digital patent count (β = −0.004, SE = 0.040). Thus, the H1 pattern concerns the breadth and exploratory character of patent-based digital innovation rather than patent volume generally. The number of observations differs across models (5,225 for digital outcomes, 4,465 for new digital subclasses, and 7,119 for non-digital patents) because PPML with firm fixed effects drops firms without within-firm variation in the relevant outcome. Taken together, the results support H1: the pattern concerns the breadth and direction of digital patenting rather than its volume, and it does not extend to non-digital patenting.

	(1)	**(2)**	**(3)**	**(4)**
	CPC Breadth	**New Subclasses**	**Digital Patents**	**Non-Digital Pat.**
	(Novelty)	*(Exploration)*	*(Volume)*	*(Placebo)*
Invention Presence × App./Design Int.	0.118***	0.357**	−0.030	−0.004
	(0.031)	(0.126)	(0.048)	(0.040)
Invention Presence	0.015	−0.176	0.195	0.137
	(0.052)	(0.170)	(0.144)	(0.087)
App./Design Intensity	−0.011	0.072	0.008	−0.031***
	(0.028)	(0.077)	(0.059)	(0.008)
Scaling/Governance Intensity	0.014	−0.098	0.182**	−0.119
	(0.024)	(0.087)	(0.062)	(0.084)
Controls	Yes	Yes	Yes	Yes
Firm Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
Observations	5,225	4,465	5,225	7,119

Table 2. Invention × Application/Design Complementarity (H1)
*† p < 0.10; * p < 0.05; ** p < 0.01; *** p < 0.001*

Table 3, Panel A examines H2 using the split-sample specification. We divide Invention-presence observations at the median of Scaling/Governance Intensity and estimate the full coefficient structure separately in the high- and low-Scaling/Governance subsamples. Because the median threshold is not unique and the subsamples may differ in the firms contributing identifying variation, we report a pooled continuous three-way interaction as a robustness check below.

For Forward Citations, the Invention Presence × Application/Design interaction is positive and significant in the high-Scaling/Governance subsample (β = 0.578, p < 0.001) and not significant in the low subsample

(β = −0.476, p = 0.281); the cross-subsample difference is significant (z = 2.26, p = 0.024). For KPSS Economic Value, the corresponding estimates are β = 0.354 (p < 0.001) and β = 0.008 (p = 0.934), again with a significant difference (z = 2.55, p = 0.011).

These split-sample patterns provide conditional evidence consistent with H2: the positive Invention × Application/Design staffing association with patent impact is concentrated in the high-Scaling/Governance subsample. As in any interaction model, the Invention Presence coefficient is evaluated at Application/Design staffing of zero and therefore does not describe firms that employ both role types; the quantity of interest for H2 is the interaction itself.

Panel A	**Fwd Citations**	**Fwd Citations**	**KPSS Value**	**KPSS Value**
	High Gov.	*Low Gov.*	*High Gov.*	*Low Gov.*
Invention Presence × App./Design Int.	0.578***	−0.476	0.354***	0.008
	(0.149)	(0.442)	(0.100)	(0.092)
Invention Presence	−1.622**	0.376	−1.228†	−0.033
	(0.623)	(0.307)	(0.681)	(0.172)
App./Design Intensity	4.111***	1.914**	1.629***	−0.015
	(0.269)	(0.573)	(0.196)	(0.036)
Controls & Firm & Year Fixed Effects	Yes	Yes	Yes	Yes
Observations	1,329	2,631	1,699	3,217
Panel B	***CPC Breadth***		***New Subclasses***	
AI Workforce HHI (concentration)	−0.637***		−1.639***	
	(0.161)		(0.467)	
Log Total AI Intensity	0.231***		1.259***	
	(0.054)		(0.225)	
Controls & Firm & Year Fixed Effects	Yes		Yes	
Observations	5,225		4,465	

Table 3. Governance Amplification (H2) and Workforce Balance (H3)
*† p < 0.10; * p < 0.05; ** p < 0.01; *** p < 0.001*

Table 3, Panel B tests H3 by relating AI workforce concentration to exploratory breadth while controlling for total AI intensity. The HHI of workforce concentration is negative and significant for Digital CPC Breadth (β = −0.637, SE = 0.161) and New Digital Subclasses (β = −1.639, SE = 0.467), while total AI intensity is positive for both outcomes. A 75th-to-25th-percentile decrease in concentration is associated with 36.6% greater Digital CPC Breadth (p = 0.001) and 123.2% more New Digital Subclasses (p = 0.016). These estimates support H3's prediction about broader exploratory search. The pattern also has an important extensive-margin component: among firms already staffing all three role categories, the concentration coefficient is statistically indistinguishable from zero (reported with the robustness tests below), so avoiding major functional staffing gaps is more consequential in these data than fine-grained proportional evenness once all three functions are present.

Robustness Tests

We conduct several additional analyses to assess robustness. First, replacing the binary Invention Presence indicator with continuous Invention Intensity leaves the H1 pattern qualitatively unchanged. Second, extending outcomes through 2022 produces substantively similar results. Third, adding patent-stock intensity (total patent stock scaled by total assets) following Lin and Maruping (2025) leaves the focal H1 interaction similar. Fourth, replacing year fixed effects with NAICS × year fixed effects and, separately, adding lagged sales growth does not alter the main conclusions.

Taxonomy boundary checks produce the same substantive H1 pattern. Under the narrower taxonomy, the Digital CPC Breadth interaction moves from $\beta = 0.118$ to $\beta = 0.124$ (SE = 0.040, $p = 0.002$), and the New Digital Subclasses interaction moves from $\beta = 0.357$ to $\beta = 0.341$ (SE = 0.145, $p = 0.019$). A demanding generic-analyst sensitivity test removes 146 of the 231 Application/Design occupations—retaining 85 explicitly AI/ML-marked analyst roles and removing roughly three-quarters of Application/Design headcount—yet H1 remains positive and significant for Digital CPC Breadth ($\beta = 0.137$, SE = 0.050, $p = 0.006$; N = 4,538) and New Digital Subclasses ($\beta = 0.472$, SE = 0.195, $p = 0.015$; N = 3,840). Because the exclusion materially changes the scale of the regressor, we interpret these coefficients as robustness rather than as stronger effects. The within-firm Application/Design–non-AI IT correlation is modest ($r = 0.193$). Finally, using individual-level skill records that played no role in occupation assignment, the fifty most characteristic skills of each category share no overlap: Invention is characterized by computational, simulation, and algorithm-development skills; Application/Design by domain expertise concentrated in finance and accounting; and Scaling/Governance by data-center, cloud, virtualization, and security skills.

Timing checks bound the horizon of the association. In a lead-lag specification, the lagged Invention × Application/Design interaction remains positive and significant for Digital CPC Breadth ($\beta = 0.108$, $p < 0.001$) and New Digital Subclasses ($\beta = 0.314$, $p = 0.023$), while the corresponding one-period lead interaction is statistically indistinguishable from zero ($\beta = 0.049$, $p = 0.248$; $\beta = 0.080$, $p = 0.541$). This pattern is evidence against simple short-run reverse timing in which patenting moves before the observed workforce configuration, but it cannot rule out anticipation horizons longer than one year or time-varying strategic changes that jointly affect hiring and patenting. At two- and three-year lags, the H1 association attenuates: for Digital CPC Breadth, from 0.117 ($p < 0.001$) at L1 to 0.062 ($p = 0.016$) at L2 and 0.059 ($p = 0.039$) at L3; for New Digital Subclasses, from 0.306 ($p = 0.010$) at L1 to statistical zero at L2 and L3.

For H2, a pooled continuous three-way interaction provides a threshold-free robustness check. In the full Forward Citations sample, Invention Presence × Application/Design Intensity × Scaling/Governance Intensity is positive and significant ($\beta = 0.284$, SE = 0.116, $p = 0.014$), indicating that the positive Invention × Application/Design association strengthens as Scaling/Governance staffing increases. The analogous KPSS Economic Value PPML specification did not converge to a stable solution, so we do not interpret an unstable coefficient. The pooled Forward Citations result reduces concern that the split-sample pattern is solely an artifact of the median threshold, but it does not provide the same pooled evidence for the second H2 outcome.

For H3, results are robust to restricting the analysis to firms with nonzero AI employment and to an entropy-based balance measure. They also remain robust among firms with more than ten AI employees, addressing the mechanical link between HHI and very small AI workforces: the concentration coefficient is −0.806 (SE = 0.270) for Digital CPC Breadth and −2.244 (SE = 0.783) for New Digital Subclasses, compared with baseline coefficients of −0.637 and −1.639. By contrast, among firms already staffing all three role categories, the concentration coefficient is statistically indistinguishable from zero (Digital CPC Breadth: $\beta = -0.120$, $p = 0.74$; New Digital Subclasses: $\beta = 0.304$, $p = 0.79$). This pattern reinforces the interpretation of H3 as functional coverage and avoidance of severe staffing gaps rather than proportional evenness.

Discussion and Conclusion

This study examines whether the internal composition of AI talent, not merely its scale, is associated with patent-based digital innovation. The results support a compositional view of the AI workforce. The Invention × Application/Design staffing interaction is associated with broader digital technology scope and entry into new digital subclasses, but not with higher digital patent volume overall. For patent impact, the

same staffing interaction is positive in the high-Scaling/Governance subsample and the pooled Forward Citations three-way interaction indicates that it becomes stronger as Scaling/Governance staffing increases. Firms with less concentrated AI workforces also exhibit broader exploratory search, although the attenuation of the concentration coefficient among firms already staffing all three categories indicates that the H3 pattern operates substantially through the extensive margin of functional coverage. These patterns describe a pre-generative-AI, early-diffusion period.

Theoretical Contributions

Three aspects of the findings deserve particular theoretical attention.

First, the H1 results indicate that the Invention × Application/Design staffing interaction is associated with the breadth and exploratory quality of patent-based digital innovation rather than its volume. The interaction is associated with broader digital technology scope and entry into new digital knowledge areas, but not with total digital patent count, and it does not spill over to non-digital patents. This specificity matters: AI workforce composition does not simply coincide with higher patent volume but co-varies with a broader recombinant search space through which firms appear to discover and pursue new digital value paths (Fleming, 2001; Henfridsson et al., 2018; Yoo et al., 2010). Whereas Lin and Maruping (2025) establish that AI innovation differs from other IT innovation at the aggregate organizational level, our results locate part of that difference in how AI work is internally configured, with the internal coupling of frontier technical roles and domain-embedding roles providing one observable organizational correlate of more exploratory digital recombination.

Second, the H2 findings highlight a distinctive role for AI Scaling/Governance staffing. The Invention × Application/Design staffing interaction is associated with measurable innovation impact (forward citations and economic value) among firms with high AI Scaling/Governance staffing; in the low-Scaling/Governance subsample the same interaction is insignificant, and the pooled three-way interaction for forward citations points the same way. This asymmetry suggests that Scaling/Governance is not merely a downstream support function or a stand-alone innovation input but co-occurs with the conditions under which novel recombinations become deployable, citable, and economically consequential innovations rather than isolated prototypes. Research on managing AI identifies deployment, monitoring, and quality control as distinctively demanding (Berente et al., 2021; Lebovitz et al., 2021); our results place that work on the innovation path rather than only downstream of it, consistent with orchestration accounts in which the same knowledge inputs yield different outcomes depending on how they are assembled and put to work (Sirmon et al., 2011) and with evidence that organizational practice conditions the returns to technology and skill (Bresnahan et al., 2002). The returns to AI staffing therefore appear to depend not only on pairwise complementarities but on a broader configuration in which deployment, documentation, monitoring, and oversight accompany the translation of novelty into impact.

Third, the H3 finding that workforce balance is associated with digital innovation novelty beyond total AI headcount speaks to the complementarity prediction that severe imbalances reduce the returns to any one component (Ennen & Richter, 2010; Milgrom & Roberts, 1995). This is an important counterpoint to pure specialization: while concentrated AI workforces may produce depth in established technical domains, balanced compositions are associated with the cross-domain recombination that characterizes exploratory digital innovation (Krakowski et al., 2023; Tambe et al., 2020). Our argument is not that equal allocation is universally optimal but that severe functional staffing gaps limit the benefits of complementarity, reframing AI workforce design as a portfolio problem in which avoiding missing complements may matter as much as deepening any single role category. This is what the innovation-value-chain logic of Hansen and Birkinshaw (2007) implies: a severe weakness in one stage constrains the whole process, so it is the presence of a stage rather than its exact size that binds. Additional analysis supports this reading: when we restrict the sample to firms that already staff all three AI role categories, the workforce-concentration coefficient attenuates. At this stage of AI diffusion the principal architectural distinction is whether firms establish presence across Invention, Application/Design, and Scaling/Governance, with fine-grained role balance playing a secondary role once a complete architecture is in place. Whether balance becomes more consequential as firms mature beyond initial staffing is an open question for future work. Our H3 measure captures balance in role composition, one observable facet of AI workforce architecture; it does not directly test whether firms with balanced compositions also coordinate across these functions, which we leave to future work.

Practical Implications

The findings offer managers a diagnostic view of AI workforce composition rather than a causal prescription. Moving from an AI workforce concentrated in a single role type to one distributed across the three functions is associated with roughly 12 percent broader digital technology scope and 29 percent higher entry into new digital knowledge areas. Likewise, the Invention × Application/Design association with patent impact is concentrated where Scaling/Governance staffing is higher. These patterns give managers benchmarks for identifying staffing gaps or overconcentration.

Practically, the results suggest treating MLOps, deployment infrastructure, model monitoring, compliance, and data governance as part of the innovation architecture rather than as downstream support, and they position workforce audits that flag overconcentration in one role category as a diagnostic lens for firms whose AI initiatives stall between experimentation and deployment. Such audits should be interpreted alongside strategy, process design, data infrastructure, managerial commitment, and other organizational investments that our workforce measures do not capture.

Limitations and Future Research

Our findings describe within-firm associations under an observational design and should be interpreted within several boundaries. Lagged predictors and firm fixed effects absorb time-invariant confounding and reduce contemporaneous simultaneity, but they cannot rule out unobserved time-varying changes in strategic priorities, innovation processes, managerial commitment, data infrastructure, or patenting strategy that jointly shape AI staffing and patent outcomes. Anticipatory hiring is the most plausible such alternative. Firm and year fixed effects, the non-AI IT control, the NAICS × year and patent-stock robustness checks, and the insignificant one-period leads constrain simple versions of it, but they cannot eliminate a firm-specific strategic shift set in motion more than one year in advance. Under either a complementarity or an anticipation reading, the joint presence of the three role types is informative about which firms move into broader digital domains, but stronger causal claims require exogenous variation in AI hiring. Patent-based measures capture only part of firms' digital innovation and exclude AI-enabled process improvements, service innovations, product enhancements, and other innovations that do not result in filings. Our workforce measures capture standardized occupational roles, not internal team boundaries, reporting structures, routines, coordination mechanisms, platforms, or governance practices. The three-category taxonomy necessarily simplifies within-role variation despite high coder agreement and the boundary validations reported above. The 2012–2022 observation window is a pre-generative-AI, early-diffusion setting and should not be extrapolated directly to the foundation-model era, when external models and services may change both workforce composition and the relative importance of the three role categories. In addition, the PPML fixed-effects estimates generalize most directly to larger, more technology-oriented, patent-active firms that exhibit within-firm variation in digital patenting and AI workforce composition; they should not be generalized to firms that do not patent digitally. Future work should exploit exogenous variation in AI hiring, event studies around discrete workforce changes, and data on internal collaboration to distinguish complementary workforce mechanisms from anticipatory strategic reconfiguration and to observe how formal role composition interacts with actual coordination.

Acknowledgements

Declaration of AI Use: In preparing and revising this manuscript, the authors used generative AI tools to improve language and readability and, to a limited extent, to help refine and debug analysis code. All AI-assisted text and code were reviewed, tested, and verified by the authors, who take full responsibility for the accuracy, integrity, and content of the final version.

References

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics, 40*(S1), S293–S340. https://doi.org/10.1086/718327

Agrawal, A., Gans, J. S., & Goldfarb, A. (2019). Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives, 33*(2), 31–50. https://doi.org/10.1257/jep.33.2.31

Alekseeva, L., Azar, J., Giné, M., Samila, S., & Taska, B. (2021). The demand for AI skills in the labor market. *Labour Economics, 71*, Article 102002. https://doi.org/10.1016/j.labeco.2021.102002

Aral, S., & Weill, P. (2007). IT assets, organizational capabilities, and firm performance: How resource allocations and organizational differences explain performance variation. *Organization Science, 18*(5), 763–780. https://doi.org/10.1287/orsc.1070.0306

Athey, S., & Stern, S. (1998). *An empirical framework for testing theories about complementarity in organizational design* (Working Paper No. 6600). National Bureau of Economic Research. https://doi.org/10.3386/w6600

Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics, 151*, Article 103745. https://doi.org/10.1016/j.jfineco.2023.103745

Benbya, H., Pachidi, S., & Jarvenpaa, S. L. (2021). Special issue editorial: Artificial intelligence in organizations: Implications for information systems research. *Journal of the Association for Information Systems, 22*(2), 281–303. https://doi.org/10.17705/1jais.00662

Benner, M. J., & Tushman, M. L. (2003). Exploitation, exploration, and process management: The productivity dilemma revisited. *Academy of Management Review, 28*(2), 238–256. https://doi.org/10.5465/amr.2003.9416096

Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly, 45*(3), 1433–1450. https://doi.org/10.25300/MISQ/2021/16274

Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics, 117*(1), 339–376. https://doi.org/10.1162/003355302753399526

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics, 140*(2), 889–942. https://doi.org/10.1093/qje/qjae044

Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics, 13*(1), 333–372. https://doi.org/10.1257/mac.20180386

Cassiman, B., & Veugelers, R. (2006). In search of complementarity in innovation strategy: Internal R&D and external knowledge acquisition. *Management Science, 52*(1), 68–82. https://doi.org/10.1287/mnsc.1050.0470

Choudhury, P., Starr, E., & Agarwal, R. (2020). Machine learning and human capital complementarities: Experimental evidence on bias mitigation. *Strategic Management Journal, 41*(8), 1381–1411. https://doi.org/10.1002/smj.3152

Cockburn, I. M., Henderson, R. M., & Stern, S. (2019). The impact of artificial intelligence on innovation: An exploratory analysis. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 115–148). University of Chicago Press. https://doi.org/10.7208/chicago/9780226613475.003.0004

Correia, S., Guimarães, P., & Zylkin, T. (2020). Fast Poisson estimation with high-dimensional fixed effects. *The Stata Journal, 20*(1), 95–115. https://doi.org/10.1177/1536867X20909691

Ennen, E., & Richter, A. (2010). The whole is more than the sum of its parts, or is it? A review of the empirical literature on complementarities in organizations. *Journal of Management, 36*(1), 207–233. https://doi.org/10.1177/0149206309350083

Faraj, S., Pachidi, S., & Sayegh, K. (2018). Working and organizing in the age of the learning algorithm. *Information and Organization, 28*(1), 62–70. https://doi.org/10.1016/j.infoandorg.2018.02.005

Fleming, L. (2001). Recombinant uncertainty in technological search. *Management Science, 47*(1), 117–132. https://doi.org/10.1287/mnsc.47.1.117.10671

Gregory, R. W., Henfridsson, O., Kaganer, E., & Kyriakou, H. (2021). The role of artificial intelligence and data network effects for creating user value. *Academy of Management Review, 46*(3), 534–551. https://doi.org/10.5465/amr.2019.0178

Hall, B. H., Jaffe, A. B., & Trajtenberg, M. (2005). Market value and patent citations. *RAND Journal of Economics, 36*(1), 16–38. https://doi.org/10.2307/1593752

Hanelt, A., Bohnsack, R., Marz, D., & Antunes Marante, C. (2021). A systematic review of the literature on digital transformation: Insights and implications for strategy and organizational change. *Journal of Management Studies, 58*(5), 1159–1197. https://doi.org/10.1111/joms.12639

Hansen, M. T., & Birkinshaw, J. (2007). The innovation value chain. *Harvard Business Review, 85*(6), 121–130, 142. https://hbr.org/2007/06/the-innovation-value-chain

Henfridsson, O., Nandhakumar, J., Scarbrough, H., & Panourgias, N. S. (2018). Recombination in the open-ended value landscape of digital innovation. *Information and Organization, 28*(2), 89–100. https://doi.org/10.1016/j.infoandorg.2018.03.001

Katila, R., & Ahuja, G. (2002). Something old, something new: A longitudinal study of search behavior and new product introduction. *Academy of Management Journal, 45*(6), 1183–1194. https://doi.org/10.5465/3069433

Kogan, L., Papanikolaou, D., Seru, A., & Stoffman, N. (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics, 132*(2), 665–712. https://doi.org/10.1093/qje/qjw040

Kohli, R., & Melville, N. P. (2019). Digital innovation: A review and synthesis. *Information Systems Journal, 29*(1), 200–223. https://doi.org/10.1111/isj.12193

Krakowski, S., Luger, J., & Raisch, S. (2023). Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal, 44*(6), 1425–1452. https://doi.org/10.1002/smj.3387

Lebovitz, S., Levina, N., & Lifshitz-Assaf, H. (2021). Is AI ground truth really "true"? The dangers of training and evaluating AI tools based on experts' know-what. *MIS Quarterly, 45*(3b), 1501–1525. https://doi.org/10.25300/MISQ/2021/16564

Lerner, J., & Seru, A. (2022). The use and misuse of patent data: Issues for finance and beyond. *The Review of Financial Studies, 35*(6), 2667–2704. https://doi.org/10.1093/rfs/hhab084

Lin, Y.-K., & Maruping, L. M. (2025). Organizing for AI innovation: Insights from an empirical exploration of U.S. patents. *MIS Quarterly, 49*(3), 1095–1122. https://doi.org/10.25300/MISQ/2025/18765

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science, 2*(1), 71–87. https://doi.org/10.1287/orsc.2.1.71

Marchetti, A., & Puranam, P. (2026). Are less hierarchical firms organized around stronger cultures? Evidence from big data. *Strategic Management Journal, 47*(2), 463–493. https://doi.org/10.1002/smj.70020

Melville, N. P., Kraemer, K., & Gurbaxani, V. (2004). Review: Information technology and organizational performance: An integrative model of IT business value. *MIS Quarterly, 28*(2), 283–322. https://doi.org/10.2307/25148636

Mikalef, P., & Gupta, M. (2021). Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. *Information & Management, 58*(3), Article 103434. https://doi.org/10.1016/j.im.2021.103434

Milgrom, P., & Roberts, J. (1995). Complementarities and fit: Strategy, structure, and organizational change in manufacturing. *Journal of Accounting and Economics, 19*(2–3), 179–208. https://doi.org/10.1016/0165-4101(94)00382-F

Nambisan, S., Lyytinen, K., Majchrzak, A., & Song, M. (2017). Digital innovation management: Reinventing innovation management research in a digital world. *MIS Quarterly, 41*(1), 223–238. https://doi.org/10.25300/MISQ/2017/41:1.03

Nambisan, S., Wright, M., & Feldman, M. (2019). The digital transformation of innovation and entrepreneurship: Progress, challenges and key themes. *Research Policy, 48*(8), Article 103773. https://doi.org/10.1016/j.respol.2019.03.018

Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of economic change*. Belknap Press of Harvard University Press.

Ployhart, R. E., Nyberg, A. J., Reilly, G., & Maltarich, M. A. (2014). Human capital is dead; long live human capital resources! *Journal of Management, 40*(2), 371–398. https://doi.org/10.1177/0149206313512152

Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation-augmentation paradox. *Academy of Management Review, 46*(1), 192–210. https://doi.org/10.5465/amr.2018.0072

Santos Silva, J. M. C., & Tenreyro, S. (2006). The log of gravity. *The Review of Economics and Statistics, 88*(4), 641–658. https://doi.org/10.1162/rest.88.4.641

Santos Silva, J. M. C., & Tenreyro, S. (2011). Further simulation evidence on the performance of the Poisson pseudo-maximum likelihood estimator. *Economics Letters, 112*(2), 220–222. https://doi.org/10.1016/j.econlet.2011.05.008

Sirmon, D. G., Hitt, M. A., Ireland, R. D., & Gilbert, B. A. (2011). Resource orchestration to create competitive advantage: Breadth, depth, and life cycle effects. *Journal of Management, 37*(5), 1390–1412. https://doi.org/10.1177/0149206310385695

Tambe, P. B. (2026). Reskilling the workforce for AI: Domain expertise and algorithmic literacy. *Management Science, 72*(1), 515–537. https://doi.org/10.1287/mnsc.2022.03968

Tambe, P. B., & Hitt, L. M. (2012). The productivity of information technology investments: New evidence from IT labor data. *Information Systems Research, 23*(3-part-1), 599–617. https://doi.org/10.1287/isre.1110.0398

Tambe, P. B., Hitt, L. M., Rock, D., & Brynjolfsson, E. (2020). *Digital capital and superstar firms* (Working Paper No. 28285). National Bureau of Economic Research. https://doi.org/10.3386/w28285

Tiwana, A., & Konsynski, B. (2010). Complementarities between organizational IT architecture and governance structure. *Information Systems Research, 21*(2), 288–304. https://doi.org/10.1287/isre.1080.0206

Trajtenberg, M. (1990). A penny for your quotes: Patent citations and the value of innovations. *RAND Journal of Economics, 21*(1), 172–187. https://doi.org/10.2307/2555502

Yoo, Y., Henfridsson, O., Kallinikos, J., Gregory, R. W., Burtch, G., Chatterjee, S., & Sarker, S. (2024). The next frontiers of digital innovation research. *Information Systems Research, 35*(4), 1507–1523. https://doi.org/10.1287/isre.2024.editorial.v35.n4

Yoo, Y., Henfridsson, O., & Lyytinen, K. (2010). The new organizing logic of digital innovation: An agenda for information systems research. *Information Systems Research, 21*(4), 724–735. https://doi.org/10.1287/isre.1100.0322

Zahra, S. A., & George, G. (2002). Absorptive capacity: A review, reconceptualization, and extension. *Academy of Management Review, 27*(2), 185–203. https://doi.org/10.5465/amr.2002.6587995