Build, Apply, Operate: AI Role Complementarities and Firm Value

Completed Research Paper

Kim Melissa Bär
Vienna University of Economics and Business
Vienna, Austria
kim.melissa.baer@wu.ac.at

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

Abstract

Firms increasingly compete not on whether they employ artificial intelligence talent but on how they organize it. We distinguish three functionally distinct forms of AI work — Build (proprietary AI creation), Apply (domain-specific deployment), and Operate (governance and infrastructure) — and ask whether the value associated with one depends on the presence of another. Using Revelio Labs workforce data matched to Compustat for 1,927 publicly traded U.S. firms from 2019 to 2025, we estimate firm and year fixed-effects models of Tobin's Q. Build and Apply talent are associated with higher firm value jointly rather than separately, and this association is stronger where Operate intensity is high and outside the technology sector. Results hold under alternative role taxonomies, alternative valuation and productivity outcomes, and dynamic specifications. We contribute a functional-configuration perspective on AI value, showing that complementarities arise within AI human capital itself, so that configuration, not scale alone, shapes returns.

Keywords: AI workforce composition, complementarity, firm value, human capital

Introduction

As artificial intelligence becomes more broadly accessible through foundation models, external services, and off-the-shelf tools (Baird & Maruping, 2021; Berente et al., 2021; Brynjolfsson et al., 2025), a puzzle arises: if baseline AI capability is increasingly available to all firms, through what organizational mechanisms do firms capture differential value from it? Prior research has shown that the capabilities and organization of technology workers shape returns on technology investments (Cohen & Levinthal, 1990; Melville et al., 2004; Roberts et al., 2012; Tambe & Hitt, 2012), and this logic has been extended to artificial intelligence (Benbya et al., 2020; Benbya et al., 2021), where firms with stronger AI-capable workforces achieve higher market valuations (Babina et al., 2024). The commoditization of baseline AI capability, however, shifts the locus of advantage: competitive advantage depends less on whether firms invest in AI talent and more on how they organize, integrate, and govern AI work internally (Krakowski et al., 2023). Firms must decide which AI activities to build in-house, how to embed AI into products and business processes, and how to create the governance and infrastructure capacity needed to deploy AI reliably at scale. AI transformation is therefore not simply a technology adoption decision; it is a problem of workforce composition, cross-functional integration, and organizational governance.

Existing research already differentiates AI and data talent by expertise and by where it sits inside the organization. Abis and Veldkamp (2024) partition the data workforce by the vintage of analytical training and by data-management responsibility, and Tambe (2026) shows that algorithmic skills are rewarded

differently depending on whether they sit with domain experts or technical specialists, while other studies measure the quantity of AI workers (Babina et al., 2024; Tambe, 2014). What this work does not ask is how the functions those workers perform combine to create value. This is a consequential omission. Research on IT business value has long shown that separating technology investments by strategic purpose reveals performance effects that aggregate measures obscure (Aral & Weill, 2007; Kohli & Grover, 2008; Melville et al., 2004). Moreover, workforce capabilities and organizational practices generate returns through complementarities rather than simple addition (Milgrom & Roberts, 1990; Wu et al., 2020). AI capability itself has been recognized as multidimensional, spanning distinct skill domains, tangible resources, and intangible coordination capacities (Mikalef & Gupta, 2021), and as creating organizational value through different mechanisms that require distinct forms of organizational embedding (Benbya et al., 2021). Further, managing AI innovation demands different organizing logic than traditional IT, requiring distinct management practices and structural arrangements that make workforce configuration especially consequential (Berente et al., 2021; Lin & Maruping, 2025). Prior work has not directly tested whether creating, deploying, and governing AI are complementary within the AI workforce in their contribution to firm value.

We address this question by conceptualizing AI workforce composition as a portfolio of functionally distinct but potentially complementary organizational capabilities that correspond to three distinct organizational problems: creating proprietary AI assets, delegating tasks to AI within domain workflows, and governing AI artifacts that operate with degrees of autonomy once deployed. Specifically, we distinguish among Build roles oriented toward proprietary AI creation, Apply roles focused on domain-specific AI deployment, and Operate roles concerned with AI governance and infrastructure. Our central argument is that firms are more likely to create strategic value when they configure these roles as a coherent system rather than treat AI hiring as a simple aggregate investment. Build talent develops proprietary models, pipelines, and technical assets that competitors cannot easily replicate, a commitment that is qualitatively distinct from traditional IT investment because AI innovation demands different organizing logic and is less easily substituted through procurement alone (Lin & Maruping, 2025). Apply talent embeds those assets into products, decisions, and business processes where they can generate returns, a function that is more valuable when it draws on firm-specific rather than purely generic AI capabilities (Lou & Wu, 2021). Either capability in isolation is limited: Build without Apply leaves technical capability weakly embedded in organizational use, whereas Apply without Build constrains firms to deploying broadly available tools that are easier for rivals to imitate, amounting to a commoditized configuration with limited scope for firm-level differentiation (Hoffreumon et al., 2024). We therefore expect the value associated with one function to depend on the presence of the other, reflecting context-specific capabilities that are difficult for competitors to imitate (Bharadwaj, 2000).

We further argue that two contingencies shape the strength of this complementarity. First, Operate capabilities amplify the value of the Build–Apply combination by providing the governance, infrastructure, and integration capacity that connect AI creation to organizational deployment. Operate talent does not merely add another pool of AI labor; it strengthens the interface through which Build and Apply become a more integrated system of AI work, reducing coordination frictions and enabling firms to scale AI more reliably, accountably, and consistently across contexts (Aral et al., 2012; Benbya et al., 2024; Schneider et al., 2023; Sirmon et al., 2011). Second, the complementarity between Build and Apply should be stronger outside the technology sector. On average, proprietary AI creation is more common and therefore less distinctive in technology sectors, though substantial heterogeneity remains within both groups. In non-technology industries, by contrast, the choice to build proprietary AI rather than rely primarily on off-the-shelf solutions is more unusual and more likely to represent a distinctive, difficult-to-imitate organizational configuration when paired with domain deployment (Bharadwaj, 2000; Hoffreumon et al., 2024). In that sense, industry context shapes not only the level of AI investment, but also the strategic value of how firms configure and govern their AI workforce.

We test this theory by analyzing AI workforce composition across 1,927 publicly traded U.S. firms from 2019 to 2025 using Revelio Labs Workforce Intelligence data matched to Compustat financial data. The underlying workforce dataset covers more than 54 million position records and 34 million unique individuals over the same period, allowing us to construct a firm-level panel that tracks the stock of employees across 342 distinct AI-related occupational roles and captures not only how much AI talent firms employ, but also the functional role that talent performs. We develop a purpose-built taxonomy grounded in Revelio's standardized occupational classification that distinguishes Build roles involving proprietary

model development and AI system creation, Apply roles that embed AI into domain operations, and Operate roles that govern infrastructure, data quality, and deployment. We estimate two-way fixed effects panel regressions on Tobin's Q, exploiting an extensive-margin specification that interacts the discrete decision to invest in any Build capability with continuous Apply intensity to identify complementarity. The results support the hypothesized relationships and are robust across alternative taxonomy specifications, functional forms, and industry subsamples. Because the design exploits within-firm variation with lagged predictors, we interpret the estimates as within-firm associations rather than causal effects.

The theoretical contributions of this study are threefold. First, we extend the complementarity tradition in technology value research (Bresnahan et al., 2002; Melville et al., 2004; Milgrom & Roberts, 1990) by showing that complementarities exist within AI human capital itself. The contribution is not that AI talent is heterogeneous, which prior work already shows (Abis & Veldkamp, 2024; Babina et al., 2024; Tambe, 2026), but that functionally distinct forms of AI work may be complementary, such that the value associated with one function depends on another. Second, we conceptualize Operate capabilities as an enabling mechanism rather than only a risk safeguard, extending resource orchestration theory (Sirmon et al., 2011) to an infrastructure capability whose primary role is relational rather than stand-alone. Third, we identify an industry boundary: the Build–Apply relationship is stronger in non-technology firms, where proprietary AI creation is less common and therefore more distinctive (Bharadwaj, 2000; Hoffreumon et al., 2024).

Theoretical Background and Hypotheses

AI Human Capital and the Build/Apply/Operate Architecture

As artificial intelligence becomes more widely accessible through foundation models, external AI services, and off-the-shelf tools, the central strategic question shifts from whether to adopt AI at all to whether firms have the internal capability to recognize, assimilate, and apply these externally available assets (Cohen & Levinthal, 1990; Roberts et al., 2012). In the AI setting this absorptive-capacity problem takes a distinctive form: AI artifacts are opaque, act with a degree of autonomy, and adapt after deployment (Berente et al., 2021), and they increasingly function as agentic entities to which work is delegated rather than passive tools that execute predefined routines (Baird & Maruping, 2021). Designing the AI workforce is therefore less a staffing question than a question of how firms govern the delegation of work between human workers and agentic systems. This emphasis is consistent with the broader information systems literature showing that the composition and organization of technology workers, not merely their quantity, shape returns on technology investments (Bresnahan et al., 2002; Kohli & Grover, 2008; Melville et al., 2004; Tambe & Hitt, 2012). Recent evidence extends this logic to artificial intelligence, showing that firms with stronger AI-capable workforces exhibit higher market valuations (Babina et al., 2024) and that AI-related market value is greater when algorithmic expertise is distributed across domain specialists rather than concentrated in isolated technical teams (Tambe, 2026). The value of AI human capital, in short, is not primarily additive; it is configurational and relational.

As noted above, existing firm-level research differentiates AI talent by quantity (Babina et al., 2024; Tambe, 2014), by the composition of data work (Abis & Veldkamp, 2024), by organizational location (Tambe, 2026), and by the wage premia attached to AI skills (Alekseeva et al., 2021), but not by the kind of AI work performed inside firms. Conceptual research in IS suggests this is a consequential omission. Mikalef and Gupta (2021) show that AI capability is multidimensional, spanning tangible, human, and intangible resources; we extend this logic from capability dimensions to workforce-function dimensions and ask which combinations of AI work complement one another. Berente et al. (2021) provide the mechanism: because AI artifacts are opaque and act with a degree of autonomy, the handoff between creating AI assets and embedding them in business processes imposes coordination costs that traditional IT does not, which is why the governance and infrastructure work that supports reliable deployment cannot be folded into a generic IT function. Because AI systems can act on organizational tasks rather than merely executing fixed routines (Baird & Maruping, 2021), embedding AI is itself a delegation-governance capability, distinct from building the artifact and from running generic IT. Tambe (2026) asks where AI skills sit across organizational units; we ask what functions AI workers perform. The data managers in Abis and Veldkamp (2024) come closest to our Operate category through their data-stewardship element, but Operate spans the wider deployment and governance layer — model operations and monitoring, platform and infrastructure engineering, security, and compliance — and, where prior work characterizes the

composition and rewards of such categories, we test whether this governance capacity conditions the value of other AI functions.

Within this functional lens, we conceptualize AI human capital as a portfolio of distinct but interdependent forms of work whose value may depend on how they are combined. Building on this delegation view (Baird & Maruping, 2021), we argue that firms must solve three distinct organizational problems: creating the artifact, governing its delegation into business processes, and providing the scaling and governance infrastructure through which it runs. Build talent includes workers whose primary function is to create proprietary AI assets such as models, training pipelines, and algorithmic components. Apply talent includes workers whose primary function is to embed AI into products, decisions, and business processes within specific domains. Operate talent includes workers whose primary function is to provide the governance, infrastructure, monitoring, and deployment environment that connects AI development to organizational use. This partition is distinct from, but informed by, Mikalef and Gupta's (2021) tangible/human/intangible cut, which decomposes capability rather than workforce, and it differs from explore/exploit (March, 1991; O'Reilly & Tushman, 2013), which partitions learning modes rather than work functions. The three categories are analytically distinct but interdependent subsystems through which AI, as a technical system, becomes coupled to organizational structure and work practices (Bostrom & Heinen, 1977; Sarker et al., 2019).

Build–Apply Complementarity (H1)

The information systems literature has long argued that technology creates value through complementarity. Investments in technology, organizational redesign, and human capabilities often generate returns that exceed the sum of their separate effects because each increases the marginal value of the others (Aral & Weill, 2007; Bresnahan et al., 2002; Kohli & Grover, 2008; Melville et al., 2004; Milgrom & Roberts, 1990; Wu et al., 2020). In the AI setting, a key channel through which this complementarity operates is absorptive capacity: a firm's ability to extract value from external knowledge depends on its stock of related internal knowledge (Cohen & Levinthal, 1990; Roberts et al., 2012). As foundation models and external AI services make baseline AI capability widely available, the firms best positioned to absorb value from these external resources are those that already possess the internal technical knowledge embodied in Build roles. We argue that this logic applies within the AI workforce itself. In particular, the strategic value of AI human capital should depend not only on how much AI talent a firm employs, but on whether proprietary AI creation and domain-specific deployment are configured as a coherent organizational system.

Build talent creates proprietary AI assets such as trained models, data pipelines, and predictive algorithms that embody firm-specific technical knowledge difficult for competitors to replicate. These assets can be understood as agentic IS artifacts whose organizational value is realized only through deliberate embedding in domain workflows (Baird & Maruping, 2021). On its own, however, Build capability can remain unrealized potential. Technical assets do not automatically become organizational outcomes; they must be connected to concrete products, workflows, and decision settings. A firm that accumulates AI research scientists and machine learning engineers without corresponding domain-embedded AI users possesses technical capability disconnected from revenue-generating business processes. A firm can therefore possess sophisticated AI development capability yet capture limited business value if those capabilities remain weakly embedded in domain operations. Research on technical skill concentration supports this concern: firms with abnormally high concentrations of technical talent exhibit inflated current valuations followed by negative future returns, suggesting that markets eventually discount technical capability that is not translated into sustainable business performance (Fedyk & Hodson, 2023).

Conversely, Apply talent lacking access to internally developed AI assets is constrained to deploying generic, commercially available tools equally accessible to competitors. This amounts to a commoditized configuration that confers no durable advantage (Hoffreumon et al., 2024). Because AI innovation requires different organizing logic than traditional IT (Lin & Maruping, 2025), proprietary AI creation represents a qualitatively distinct strategic commitment that is less easily substituted through procurement alone during a period in which AI capabilities are rapidly diffusing through commercial channels (Brynjolfsson et al., 2025). Yet Build without Apply leaves proprietary capabilities unembedded in the domains where they generate returns.

The combination of Build and Apply addresses this value-realization gap, and we expect it to do so more than additively. Build creates the proprietary assets, and Apply converts them into realized business value

through domain-specific embedding in products, processes, and decisions. Critically, this is not merely an argument that firms need both capabilities, but that each makes the other more productive at the margin. Proprietary models become more valuable as more domain-embedded users find applications for them, because each new application amortizes the fixed cost of model development while generating context-specific adaptations that further differentiate the firm. Simultaneously, domain users become more effective when they can draw on proprietary rather than generic tools, because internally developed models can be customized to firm-specific data, processes, and decision contexts in ways that off-the-shelf solutions cannot. Domain deployment also generates feedback, use cases, and data that improve the relevance of internally developed AI assets over time, creating a reciprocal reinforcement loop that is the essence of complementarity (Choudhury et al., 2020; Lou & Wu, 2021; Milgrom & Roberts, 1990). Our prediction is therefore that firms exhibit greater market valuation when Build and Apply are jointly present than when either is present alone. This distinguishes our argument from the aggregate thesis that simply accumulating AI talent increases firm value. Formally:

Hypothesis 1: Build capability positively moderates the relationship between Apply intensity and firm market valuation, such that the contribution of Apply talent to firm value is greater when firms possess Build capability.

Operate Talent as Governance and Integration Capacity (H2)

Hypothesis 1 proposes that Build and Apply talent are complementary and that their joint presence is associated with greater firm value than either in isolation. We now argue that the strength of this complementarity varies across firms because the translation of AI assets from creation to organizational use depends on governance, infrastructure, and integration capacity. In our framework, Operate talent amplifies the Build–Apply complementarity by providing the organizational mechanisms through which AI assets move from development into reliable deployment.

Resource orchestration theory suggests that possessing valuable resources is necessary but not sufficient to realize value from them (Sirmon et al., 2011). What matters is whether firms can structure, bundle, and leverage those resources so that potential value becomes realized value. In IT-enabled contexts, this orchestration function has been characterized as a set of sensing, learning, integrating, and coordinating capabilities through which firms convert digital resources into organizational outcomes (Pavlou & El Sawy, 2011). Operate roles embody this orchestration in the AI context: model monitoring and observability perform the sensing function, MLOps and deployment pipelines perform the integrating function, and model risk, compliance, and AI governance perform the coordinating function. When orchestration mechanisms are more developed, the same resource portfolio can generate stronger returns because the process linking resources to outcomes becomes more efficient and complete. We adopt orchestration because Operate in our setting sits between Build and Apply as a bridging capability rather than as a pole of a single tension (Sirmon et al., 2011). We conceptualize Operate as a workforce-level manifestation of this orchestration capacity: firms do not govern and scale AI in the abstract, but through workers such as MLOps engineers, infrastructure architects, and data governance specialists whose collective function is to connect AI creation to AI deployment.

This amplifying role operates through three interconnected channels. First, Operate talent reduces deployment friction by building repeatable pipelines that move AI assets from development environments into production settings, reducing the time, cost, and failure rate of the Build-to-Apply handoff. Second, Operate talent increases reliability through monitoring, versioning, rollback, and performance-management capabilities that allow Apply talent to embed AI outputs into consequential business decisions with greater confidence. Third, Operate talent provides governance capacity through auditing, bias monitoring, regulatory compliance, and access controls that enable firms to deploy AI in higher-stakes organizational contexts (Benbya et al., 2024; Schneider et al., 2023; Thiebes et al., 2021). These functions are particularly consequential in AI because such systems differ from traditional IT in their opacity, autonomous action, and ongoing adaptation, all of which intensify governance challenges as AI is deployed at scale and in more consequential domains (Berente et al., 2021; Sunyaev et al., 2025). Firms that invest in this capacity can therefore extend the range of contexts in which the Build–Apply combination can be deployed productively, reliably, and accountably (Sunyaev et al., 2025).

The distinctive role of Operate in our framework is relational rather than stand-alone. Operate may have some direct value on its own, but its primary contribution is to strengthen the interface through which Build

and Apply become an integrated system of AI work. By lowering coordination costs, reducing risk, and making deployment more reliable, Operate allows firms to realize more of the strategic value embedded in their existing portfolio of Build and Apply capabilities. Critically, we do not argue that Build–Apply complementarity requires Operate in order to exist. Rather, Operate should make this complementarity stronger because the movement from AI creation to organizational deployment becomes more efficient, scalable, and governable. Formally:

Hypothesis 2: The positive association between the Build–Apply interaction and firm market valuation is stronger in firms with higher Operate talent intensity.

Industry Context: Complementarity Beyond the Technology Sector (H3)

The preceding hypotheses propose that AI workforce composition and governance infrastructure relate to firm value through complementarity and orchestration mechanisms. We now argue that the strength of the Build–Apply complementarity depends on competitive context. Research on IT capability, drawing on the resource-based view, holds that the value of a capability depends on its distinctiveness within a given environment: capabilities that are common within an industry may be necessary for competition, but are less likely to generate above-normal returns on their own (Bharadwaj, 2000; Melville et al., 2004). The same AI workforce configuration should therefore not be equally valuable across sectors if the strategic distinctiveness of Build capability differs by industry. This expectation is consistent with integrative models of IT business value showing that competitive environment moderates the returns to technology investments (Melville et al., 2004).

In technology firms, proprietary AI creation is closer, on average, to an expected capability. These firms compete through the development of more capable models, novel algorithms, and proprietary AI architectures, and sustained investment in internal AI creation is often part of normal competitive participation. Because Build capability is closer to an industry norm, combining it with Apply talent may still be valuable, but it is less likely to represent a strongly differentiating capability configuration. The Build–Apply combination is therefore less distinctive in settings where strong internal AI creation is already widespread.

Non-technology firms face a different strategic logic. For these firms, AI is more often embedded into existing operations, customer interfaces, and decision processes rather than sold as the focal product. The dominant AI opportunities in these industries lie in automation and insight, where AI extracts efficiency and informational gains from existing business processes (Benbya et al., 2021). In this context, proprietary AI creation is less common and therefore more strategically revealing. When paired with Apply talent, Build reflects not only technical capability but a willingness to move beyond generic tools toward a firm-specific, distinctive, and difficult-to-imitate AI configuration (Bharadwaj, 2000), precisely because many firms in these industries can rely on commercially available alternatives.

This argument is an average one about relative distinctiveness, not a categorical claim; there is meaningful variation within both technology and non-technology sectors. Existing evidence is consistent with this asymmetry: complementarity between internal and external AI sourcing is weaker in the IT sector than in other industries, suggesting that internally developed AI is less differentiating where strong Build capabilities are already common (Hoffreumon et al., 2024). More broadly, analytics capabilities generate greater value in firms oriented toward process improvement and recombinatory innovation, the dominant mode in many non-technology sectors, than in firms pursuing novel technological breakthroughs (Wu et al., 2020). Figure 1 summarizes the research model. Formally:

Hypothesis 3: The positive association between the Build–Apply interaction and firm market valuation is stronger in non-technology firms than in technology firms.

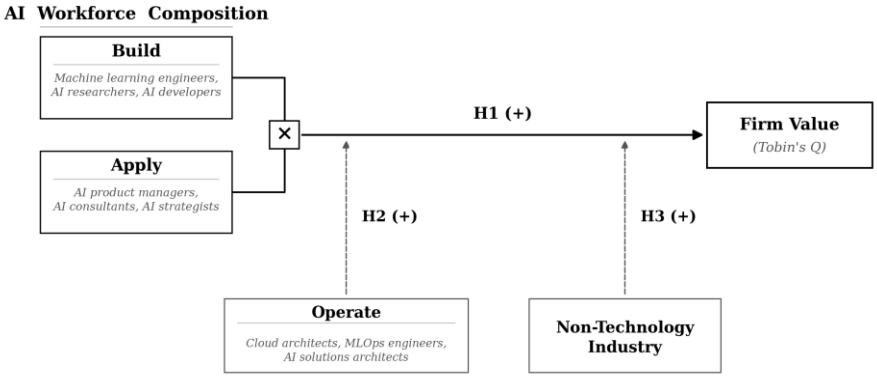


Figure 1. Research Model

Empirical Setting and Data

We examine how firms' internal AI workforce composition relates to firm value using a matched panel that combines Revelio Labs workforce data with Compustat financial data. Revelio provides individual-level employment histories derived from public professional profiles and standardized into approximately 17,000 fine-grained occupational categories, yielding over 54 million position records and 34 million unique individuals from 2019 to 2025. This granularity allows us to move beyond external AI-exposure proxies based on occupation-task mappings (Felten et al., 2021), and to directly measure the composition of AI labor inside firms. Because occupational structure aggregates across thousands of profiles per firm-year, Revelio-based measures are comparatively stable and consistently observed across firms and over time, and they capture realized workforce composition rather than stated intentions or task-level exposure scores. This data source has gained traction in strategy, finance, and IS research (Babina et al., 2024; Marchetti & Puranam, 2026; Tambe, 2026) and has been validated against firm-reported disclosures and official labor statistics (Cai et al., 2024; Liang et al., 2025).

We aggregate Revelio entities to the ultimate parent firm and match them to Compustat North America using a cascading identifier procedure that prioritizes parent-level gvkeys and uses CUSIP and ticker-based fallbacks when needed. From the firms with overlapping fiscal years (3,347), we exclude financial firms (SIC 6000–6999) and utilities (SIC 4900–4999) because their valuations are shaped primarily by regulatory capital requirements and rate-setting rather than operating capability (Bharadwaj, 2000; Tambe, 2014), and we require a minimum of 200 Revelio-tracked employees so that workforce composition measures rest on a meaningful profile base rather than a handful of sampled individuals. After further requiring non-missing lagged values for all independent variables and controls, the final estimation sample comprises 1,927 firms and 8,645 firm-year observations spanning fiscal years 2019 to 2025. The panel is broad in sectoral coverage, spanning manufacturing (47% of firm-years), services (26%), retail, wholesale, transportation, and mining/construction, and in firm size, with a median Revelio-tracked workforce of approximately 2,350 employees.

AI Workforce Taxonomy

Our key explanatory variables are derived from a purpose-built AI role taxonomy grounded in Revelio's standardized occupational classification system. We begin with Revelio's v3 role taxonomy of approximately 17,000 distinct occupations, which provides a consistent, high-resolution mapping of job titles across firms and over time and avoids inconsistencies inherent in firm-reported labels. Using a combination of skill- and keyword-based filtering and an ensemble coding approach, we first identify the universe of potentially AI-related roles. Whereas prior AI occupational taxonomies measure tasks' exposure to AI-driven automation (Felten et al., 2021), our aim is to distinguish functionally distinct roles within firms' internal AI workforce. We iteratively refine the candidate set through manual validation by two independent academic experts in artificial intelligence and data science, external to the author team.

Interrater agreement between the two coders was high (Cohen's κ = 0.82), indicating very good agreement (Landis & Koch, 1977); remaining disagreements were resolved through discussion to produce the final classification. This process yields a final set of 342 AI-related occupations. For the main analyses, we employ a functional taxonomy that distinguishes among Build, Apply, and Operate roles. Build roles (25 occupations) capture proprietary AI creation and model-development activities and include titles such as AI Research Scientist, Machine Learning Engineer, Deep Learning Engineer, and Generative AI Engineer. Apply roles (257 occupations) capture the embedding of AI into products, decisions, and business processes and include titles such as AI Product Owner, AI Data Analyst, and Business Intelligence Analyst. Operate roles (60 occupations) capture the infrastructure, deployment, and governance capabilities that connect AI development to organizational use and include titles such as Enterprise AI Architect, AWS DevOps Architect, and Data Governance Specialist. We conduct robustness checks using alternative taxonomies that reclassify the same role universe into Specialists, Infrastructure, and Domain Experts, and, separately, into AI Core, AI Infrastructure and Governance, and AI-Enabled Analytical categories. The results remain substantively unchanged, reducing concerns that our findings are driven by any single classification scheme. The complete taxonomy with role-level classifications is available from the authors upon request.

To validate that these functional categories correspond to distinct bundles of capability rather than to labeling conventions, we conducted a skills-based construct validation using Revelio's individual-level skill records. For the 904,582 workers who held one of our 342 AI-classified roles at a sample firm during the panel window, we retrieved 39.8 million person-skill records, an average of 44 skills per worker, and worked at Revelio's 1,500-cluster skill level. Skill signatures were estimated on the 872,969 workers (96.5%) whose roles fall unambiguously within a single functional category, discarding skills held by fewer than 1% of AI workers as too rare to estimate reliably and those held by more than 50% as too generic to discriminate, and ranking the remainder by the log-odds of appearing in one category relative to the unweighted average of the other two. The fifty highest-scoring skills form each category's signature, and the three signatures share no skills: Build emphasizes model-development and computational capabilities, Operate emphasizes infrastructure, cloud, network, and security capabilities, and Apply emphasizes domain expertise. A worker is treated as holding a category's skill profile when at least three of that category's signature skills appear on their profile, and firm-year skill intensities are the count of such workers per 100 employees, scaled identically to our occupational-title measures. Because reporting dates identify when a skill was first observed rather than when it was acquired, we treat skill profiles as time-invariant worker attributes, so within-firm variation in the skill-based measures arises from employees entering and leaving the firm rather than from changes in individual skill listings.

Measures

Dependent Variable

Our primary dependent variable is Tobin's Q, measured as the natural logarithm of the ratio of a firm's market value to its book value of total assets:

$$ln(Q)i,t = ln((AT - CEQ + PRCC_F \times CSHO) / AT)$$

where AT is total assets, CEQ is common equity, and PRCC_F × CSHO is market capitalization. This is the simplified approximation commonly used in IT value research and related work on digital and human-capital complements (Babina et al., 2024; Bharadwaj, 2000; Tambe, 2014). We log-transform Q because the distribution is right-skewed and winsorize at the 1st and 99th percentiles to mitigate outlier influence. Tobin's Q is well suited to our setting because it captures the forward-looking value of firm-specific intangible assets relative to tangible assets. This aligns with our theory, as the value of AI workforce composition should emerge first through organizational capabilities before it is fully reflected in accounting performance. Our results are robust to an alternative debt-adjusted Tobin's Q that incorporates both short-term and long-term debt (Build–Apply interaction $\beta = 0.071$, $p = 0.029$) and to revenue per employee as an accounting-based performance indicator ($\beta = 17{,}785$, $p = 0.003$). This suggests that the findings capture not only investor valuation but also underlying operating performance.

Independent Variables

Our key independent variables measure the share of each AI workforce category in a firm's total employee base per year. For each category c ∈ {Build, Apply, Operate}, we compute:

$$AI_intensity_{c,i,t} = (AI_headcount_{c,i,t} / Total_employees_{i,t}) \times 100$$

where both numerator and denominator are drawn from Revelio to ensure measurement consistency. To separate AI-specific workforce composition from general digitalization, we construct a non-AI IT workforce control as the share of employees in 790 non-AI IT occupations identified within the same Revelio occupational classification used for our AI taxonomy (e.g., Java Developer, Network Administrator, Database Architect), scaled as headcount per 100 employees. This control enters all specifications and ensures that our AI workforce measures do not merely proxy for broader technology investment.

To test the complementarity hypothesis (H1), we construct the product Build_intensity × Apply_intensity. Because a number of firm-year observations have zero Build workers, we additionally specify an extensive-margin interaction using Build_any (a binary indicator equal to one if the firm employs any Build workers) interacted with Apply_intensity. This specification isolates the value premium that Apply intensity commands conditional on the firm possessing any frontier capability-building capacity. We show robust results for both specifications.

To examine boundary conditions, we leverage two additional design features embedded in the data. First, for the governance-amplification test, we split the sample at the median of Operate intensity, thereby distinguishing firms with stronger versus weaker AI-related infrastructure, deployment, and governance capacity. Second, for the industry-contingency test, we partition firms into technology and non-technology sectors. A firm is classified as a technology firm if its primary SIC code falls within industrial machinery and computers (SIC 35), electronic equipment (SIC 36), instruments (SIC 38), communications (SIC 48), or business services including software (SIC 73), or if its NAICS code indicates information services (NAICS 51) or computer systems design and related services (NAICS 5415).

Control Variables

All control variables enter with a one-year lag to mitigate simultaneity concerns and allow time for workforce composition to affect firm outcomes. Following standard practice in IT value research (Bharadwaj, 2000; Tambe, 2014), we include firm size (log total assets) and firm age since IPO to capture scale and maturity; return on assets, cash holdings, and operating cash flow over total assets to capture profitability and liquidity; R&D intensity, capital intensity (net PP&E over total assets), and reported intangible asset intensity to capture investment in growth and intangibles; market leverage and SG&A intensity to capture financial and operational structure; and lagged sales growth to capture recent momentum.

Empirical Model

We estimate firm fixed-effects panel regressions of the following general form:

$$ln(Q)_{i,t} = \beta' X_{i,t-1} + \gamma' Z_{i,t-1} + \mu_i + \tau_t + \varepsilon_{i,t}$$

where X is the vector of AI workforce intensity variables and their interactions, Z is the vector of controls, μi are firm fixed effects, and τt are year fixed effects. Standard errors are clustered at the firm level to account for serial correlation within firms; clustered standard errors appear in parentheses in all tables, and significance is denoted as * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Firm fixed effects absorb all time-invariant unobserved heterogeneity, including industry, founding conditions, and stable organizational characteristics, ensuring that our estimates are identified from within-firm changes in AI workforce composition across the 2019 to 2025 period. Year fixed effects control for common temporal shocks such as aggregate AI labor market trends and macroeconomic conditions, so the identifying variation is the deviation of each firm's composition from its own average after partialling out year shocks. The one-year lag structure on all independent and control variables provides temporal ordering and reduces concerns about reverse causality, while allowing time for workforce investments to be reflected in market valuation. We do not claim causal identification: unobserved time-varying shocks that jointly shape workforce composition and firm value cannot be fully ruled out. The design establishes within-firm associations that

are temporally ordered and robust to observed controls, and we interpret all results as associational. Workforce composition is itself an endogenous, realized configuration that can reflect product strategy, technical opportunities, organizational structure, and team-level staffing decisions; we do not interpret it as an exogenously chosen or optimal firm-wide ratio.

Results

Table 1 presents descriptive statistics for the dependent variable, the focal AI workforce measures, and the main firm-level controls. The descriptives show that Build roles are relatively sparse compared with Apply roles, which motivates our extensive-margin specification using Build presence interacted with Apply intensity. We also examined pairwise correlations and variance inflation factors; the results do not indicate material multicollinearity, so we do not report them in the main text for space reasons.

Variable	**Mean**	**SD**	**Min**	**Max**
Tobin's Q (ln)	0.70	0.64	−0.42	2.65
Build Intensity (per 100 emp.)	0.14	0.45	0	12.58
Apply Intensity (per 100 emp.)	1.39	1.35	0	23.73
Operate Intensity (per 100 emp.)	0.33	0.53	0	7.96
Build Presence (binary)	0.51	0.50	0	1
Non-AI IT Intensity (per 100 emp.)	10.85	11.03	0	76.19
Firm Size (ln total assets)	7.90	1.72	1.92	12.76
Return on Assets	−0.02	0.22	−1.88	0.28
R&D Intensity	0.33	2.25	0	27.53
Market Leverage	0.26	0.22	0	0.93
Sales Growth	0.08	0.29	−1.14	1.61
N = 8,645 firm-years				

Table 1. Descriptive Statistics

Table 2 reports tests of Hypothesis 1 using firm and year fixed-effects models with one-year-lagged predictors and standard errors clustered at the firm level. Model (1) presents the additive specification and shows no robust positive standalone association between AI workforce categories and firm market valuation, consistent with the argument that AI talent is associated with higher firm value through complementarity rather than simple accumulation. Models (2) and (3) test this complementarity directly. In Model (2), the continuous interaction between Build Intensity and Apply Intensity is positive and significant ($\beta = 0.030$, $p < 0.001$). Because many firm-year observations have no Build workers, Model (3) provides our preferred specification by interacting a binary indicator for Build Presence with continuous Apply Intensity. Consistent with H1, this interaction is positive and significant ($\beta = 0.089$, $p = 0.005$), indicating that the association between Apply talent and firm value is stronger when firms possess internal Build capability. The negative coefficient on Build Presence ($\beta = -0.142$, $p < 0.001$) is conditional on Apply Intensity equaling zero and should not be interpreted as a general negative effect of Build; the focal quantity for H1 is the conditional Build × Apply relationship. Substantively, a one-unit increase in Apply Intensity is associated with a 6.7% increase in Tobin's Q among firms with Build Presence (conditional slope $\beta = 0.065$, $p = 0.005$), whereas the corresponding effect is negative and statistically insignificant among firms without Build workers ($\beta = -0.025$, $p = 0.246$). The net effect of Build Presence becomes non-negative once Apply Intensity reaches 1.59, above the sample median (1.04) but below the 75th percentile (1.80). Relative to the within-firm standard deviation of ln(Tobin's Q) (0.245), these conditional effects are economically meaningful. The complementarity is also robust to an alternative taxonomy that reclassifies the same role universe into Specialists (proprietary AI creators), Domain Experts (domain-embedded deployers), and Infrastructure (governance and deployment): both the continuous ($\beta = 0.076$, $p < 0.001$) and extensive-margin ($\beta = 0.111$, $p < 0.001$) Specialist × Domain interactions are positive and significant. Taken together,

the results support H1: the combination of proprietary AI creation and domain-specific deployment is associated with higher market valuation than either capability in isolation.

	Model (1)	**Model (2)**	**Model (3)**
	ln(Tobin's Q)	ln(Tobin's Q)	ln(Tobin's Q)
Build Intensity × Apply Intensity		0.030***	
		(0.008)	
Build Presence × Apply Intensity			0.089**
			(0.032)
Build Intensity	0.022	−0.204**	
	(0.046)	(0.063)	
Build Presence			−0.142***
			(0.037)
Apply Intensity	0.020	0.010	−0.025
	(0.015)	(0.016)	(0.021)
Operate Intensity	−0.042	−0.044	−0.054
	(0.032)	(0.033)	(0.033)
Non-AI IT Intensity	−0.006	−0.005	−0.007
	(0.004)	(0.004)	(0.004)
Controls	Yes	Yes	Yes
Firm Fixed Effects	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes
Observations	8,645	8,645	8,645

Table 2. AI Workforce Composition and Firm Value

Table 3 examines two boundary conditions of the Build–Apply complementarity. Columns (4) and (5) split the sample at the median of Operate Intensity to test H2. The Build Presence × Apply interaction is positive and significant in the high-Operate subsample ($\beta = 0.114$, $p = 0.040$) but insignificant in the low-Operate subsample ($\beta = -0.045$, $p = 0.350$). A pooled triple-interaction specification confirms that the difference is statistically significant: the three-way interaction between Build Presence, Apply Intensity, and Operate Intensity is positive and significant ($\beta = 0.108$, $p = 0.028$), indicating that the Build–Apply complementarity is significantly stronger in firms with high governance and infrastructure capacity. This pattern is consistent with Operate talent providing the orchestration capacity through which Build and Apply become an integrated system. Columns (6) and (7) split the sample by technology sector to test H3. The Build Presence × Apply interaction is strongly significant in non-technology firms ($\beta = 0.124$, $p < 0.001$) but only marginally significant in technology firms ($\beta = 0.068$, $p = 0.086$). A pooled triple-interaction specification confirms that the difference is statistically significant: the three-way interaction between Build Presence, Apply Intensity, and the technology-sector indicator is negative and significant ($\beta = -0.066$, $p = 0.041$), indicating that the Build–Apply complementarity is significantly weaker in technology firms. The pattern is consistent with H3: outside the technology sector, where proprietary AI creation is uncommon, combining Build and Apply reflects a rarer, more firm-specific workforce configuration that is associated with higher market valuation. In technology firms, where Build capability is closer to a baseline expectation, the same combination yields weaker incremental returns.

	Model (4)	**Model (5)**	**Model (6)**	**Model (7)**
	ln(Tobin's Q)	ln(Tobin's Q)	ln(Tobin's Q)	ln(Tobin's Q)
	High Operate	Low Operate	Tech Firms	Non-Tech
Build Presence × Apply	0.114*	−0.045	0.068	0.124***
	(0.056)	(0.049)	(0.040)	(0.020)
Build Presence	−0.205*	−0.020	−0.116*	−0.172***
	(0.081)	(0.039)	(0.051)	(0.027)
Apply Intensity	−0.055	0.012	−0.025	−0.051**
	(0.041)	(0.026)	(0.032)	(0.019)
Operate Intensity	−0.042	−0.437*	−0.067	−0.005
	(0.034)	(0.171)	(0.039)	(0.034)
Non-AI IT Intensity	−0.003	−0.008	−0.005	0.006
	(0.005)	(0.010)	(0.005)	(0.008)
Controls	Yes	Yes	Yes	Yes
Firm & Year Fixed Effects	Yes	Yes	Yes	Yes
Observations	4,345	4,300	3,398	5,247

Table 3. Heterogeneity Analyses

Additional Analysis

We first validate the focal relationship using the skill-based measures described above. Reconstructing firm-year Build and Apply intensities from skill profiles rather than occupational-title counts, the Build × Apply interaction remains positive and significant in the full sample ($\beta = 0.0514$, SE = 0.0181, $p = 0.005$). To separate construction from validation, we derived and froze the signatures on one random half of firms and tested them on the held-out half, where the interaction also remains positive and significant ($\beta = 0.0617$, SE = 0.0261, $p = 0.018$). The focal relationship therefore reproduces without relying on occupational-title counts. Results are substantively unchanged when skills are alternatively dated using Revelio's first_reported field and restricted to those observed by each firm-year.

Three further diagnostics address measurement and timing concerns. First, within the estimation sample, 81.6% of entries into Build, Apply, or Operate involve employees newly joining the firm rather than incumbents appearing in a different role. Because genuine internal transfers cannot be separated from pure title relabeling among the remaining 18.4%, treating every such transition as relabeling provides a conservative upper bound on how much of the measured variation could be an artifact of titles. Second, the relationship is specific to AI labor: replacing Apply with general non-AI IT staffing measured on the same per-100-employee scale yields no complementarity (Build × non-AI IT: $\beta = 0.0006$, $p = 0.65$), and when both interactions are estimated jointly, Build × Apply remains positive and significant ($\beta = 0.0928$, SE = 0.0334, $p = 0.006$) while Build × non-AI IT remains insignificant ($\beta = -0.0016$, $p = 0.24$). Third, in a lead specification, future Build × Apply composition does not predict current valuation ($\beta = 0.0108$, $p = 0.49$). This reduces concerns about simple reverse timing, but does not establish causal ordering, since an unobserved contemporaneous shift in strategy could still move both.

We also examined the Build × Apply relationship across broader industry groups. The interaction is positive and significant in manufacturing ($\beta = 0.1224$, $p < 0.001$) and retail ($\beta = 0.1945$, $p < 0.05$), and smaller and statistically insignificant in services ($\beta = 0.0558$, $p = 0.16$). We report these as descriptive heterogeneity rather than as evidence of statistically significant differences across sectors, which we do not formally test. The pattern is broadly consistent with H3 while indicating that the technology and non-technology contrast is a coarse average distinction; the services group is itself internally heterogeneous, containing both highly regulated, AI-intensive financial and healthcare firms and firms with minimal AI staffing.

As an additional contextual robustness test, we examine whether the Build–Apply complementarity strengthened following the late-2022 acceleration in AI commoditization associated with ChatGPT and open-source foundation models (Storey et al., 2025). This period plausibly lowered barriers to generic AI deployment and made baseline Apply-style capabilities more broadly accessible, thereby increasing the strategic importance of combining domain deployment with proprietary internal Build capability. In triple-interaction specifications, the post-2022 increment is positive and significant both for the continuous Build Intensity × Apply Intensity interaction ($\beta = 0.018$, SE = 0.0024, $p < 0.001$) and for the extensive-margin Build Presence × Apply Intensity interaction ($\beta = 0.043$, $p = 0.021$). The relationship is therefore stronger after 2022 rather than absent before it. We treat this as contextual heterogeneity rather than a difference-in-differences design, because the timing is common to all firms and there is no treated and untreated contrast. This pattern is consistent with our theory that, as generic AI tools proliferated, markets placed an increasing valuation premium on firms that combined proprietary creation with domain deployment rather than relying solely on commoditized alternatives.

As additional robustness tests, we re-estimate the models using alternative dependent variables, alternative lag structures, and alternative measurement choices. Replacing our simplified Tobin's Q with a debt-adjusted measure that incorporates both short-term and long-term debt yields a positive and significant Build–Apply interaction ($\beta = 0.071$, $p = 0.029$). We also estimate models using revenue per employee as an accounting-based performance indicator; the Build–Apply interaction remains positive and significant ($\beta = 17{,}785$, $p = 0.003$), indicating that the complementarity pattern extends to realized firm productivity and is not solely an artifact of investor expectations. Extending the specification to two-year lags, the continuous Build × Apply interaction remains positive and significant ($\beta = 0.0249$, SE = 0.0117, $p = 0.033$), while the extensive-margin interaction attenuates and is no longer significant ($\beta = 0.0202$, SE = 0.0198, $p = 0.31$). The association is therefore not confined to a one-year horizon, although its extensive-margin manifestation appears less persistent. In a more demanding dynamic specification that includes lagged Tobin's Q as an additional control, the Build–Apply interaction remains positive and significant ($\beta = 0.079$, $p < 0.001$), suggesting that the result is not an artifact of persistence in firm valuation. We further obtain similar inferences under the alternative role-taxonomy mappings described in our measures section.

Discussion and Theoretical Contributions

Prior work already differentiates AI talent by quantity and by where it sits in the organization. This study adds a functional-configuration perspective on how firms create value from AI workforces. Drawing on complementarity theory and resource orchestration logic, we decomposed AI human capital into Build, Apply, and Operate functions and examined how these capabilities relate to firm market valuation. The results are consistent with complementarity, although they remain associational: firm value is not systematically higher where more AI workers are employed overall. It is higher where a particular internal architecture of AI capability is present, most notably, the combination of proprietary AI creation and domain-specific deployment, especially when that combination is accompanied by governance and infrastructure capacity. In this sense, the value of AI human capital is not primarily additive. It is configurational and relational.

Our first contribution is to the complementarity tradition in information systems and technology value research. Prior work has shown that returns to digital technology investments depend on complementarities across technology, organizational practices, and human capabilities (Aral & Weill, 2007; Bresnahan et al., 2002; Kohli & Grover, 2008; Melville et al., 2004; Milgrom & Roberts, 1990; Wu et al., 2020). We extend that logic by showing that strategically consequential complementarities can also exist within AI human capital itself. In our setting, the standalone associations of Build, Apply, and Operate are weak or inconsistent, whereas the interaction between Build and Apply is positive and robust. This pattern is consistent with an absorptive-capacity interpretation of the result: Build capability creates new technical knowledge inside the firm, but its contribution to firm value depends on whether the firm can assimilate and apply that knowledge to commercial ends through domain-embedded Apply talent (Cohen & Levinthal, 1990; Roberts et al., 2012). The value of AI labor therefore does not appear to rise mechanically with quantity. Rather, the strategic value of AI talent depends on whether distinct AI functions are assembled into a coherent organizational system. In that sense, our findings complement prior evidence that AI workforce investment predicts firm value in aggregate (Babina et al., 2024), that the composition of data work has shifted with the vintage of analytical training and data-management responsibility (Abis &

Veldkamp, 2024), and that the organizational distribution of algorithmic skills matters (Tambe, 2026). Neither of the latter two asks whether the value associated with one AI function depends on the presence of another, which is the question our interaction tests address, and we sharpen the argument by showing that the internal functional composition of AI labor is itself a source of heterogeneity in market valuation.

A second contribution is to research on AI governance and the organizational conditions under which AI investments generate value. Much of the emerging governance literature frames governance primarily as a safeguard against risk, emphasizing bias mitigation, compliance, accountability, and ethical oversight (Benbya et al., 2024; Schneider et al., 2023; Thiebes et al., 2021). Our findings support a broader interpretation. Operate capability does not appear to function merely as an administrative support layer added after innovation has already occurred. Instead, it appears to be part of the productive architecture through which AI value is realized. Higher Operate staffing is associated with a stronger Build–Apply relationship, consistent with a resource-orchestration interpretation in which governance and infrastructure capacity help connect technical creation to reliable organizational deployment (Aral et al., 2012; Pavlou & El Sawy, 2011; Sirmon et al., 2011). In this sense, Operate is not simply a third category of AI labor. It appears to function as the integrating capability that manages the boundary at which AI assets, once built, are delegated into domain processes and governed while they run (Baird & Maruping, 2021; Berente et al., 2021).

Our third contribution is to identify an important boundary condition on AI-based competitive advantage. The Build–Apply complementarity is materially stronger in non-technology firms than in technology firms. This pattern is consistent with the view that the strategic value of a capability depends in part on how distinctive it is within the relevant competitive field (Bharadwaj, 2000). In technology sectors, proprietary AI creation is closer to a baseline expectation. It may still be necessary for competition, but it is less likely to differentiate firms strongly at the margin. In non-technology sectors, by contrast, combining internal AI creation with domain-specific deployment appears to represent a more distinctive and difficult-to-imitate organizational configuration. Consistent with this, the interaction is positive and significant in manufacturing and retail but not in services, reinforcing that the technology and non-technology contrast is a coarse average distinction rather than a sharp boundary. The same workforce architecture therefore has different valuation implications depending on context. This result usefully qualifies any implicit assumption that the market valuation implications of AI investment are uniform across sectors and suggests instead that the strategic implications of AI workforce composition are contingent on the competitive environment (Hoffreumon et al., 2024; Wu et al., 2020).

The post-2022 strengthening of the Build–Apply complementarity reinforces this interpretation and ties it to the broader commoditization argument developed in the theory section. As ChatGPT, open-source foundation models, and API-based services lowered the cost of generic AI deployment, baseline Apply-style capability likely became easier to access and easier to imitate (Brynjolfsson et al., 2025; Storey et al., 2025). The premium shifts away from generic access to AI and toward firm-specific capability configurations that are harder to replicate, especially outside the technology sector (Hoffreumon et al., 2024).

Managerial Implications

These findings also generate a more nuanced managerial implication than the usual recommendation to invest heavily in AI talent. For managers, the central issue is not simply how much AI talent to hire, but what portfolio of AI work the firm is trying to build. Firms that invest in Build without sufficient Apply capacity may create technically sophisticated assets that are less strongly embedded in commercial use. Firms that invest only in Apply may become proficient users of generic tools without creating much defensible differentiation. The more useful strategic question is whether the firm has assembled a coherent system in which proprietary creation and domain deployment raise each other's marginal value. Build, Apply, and Operate can therefore be used diagnostically, to audit internal AI staffing and to identify major functional gaps; the design does not imply that mechanically copying the observed configuration will raise firm value. Importantly, our preferred extensive-margin specification suggests that firms may not need massive frontier AI organizations for internal Build capability to matter. For many non-technology firms, the implication is not that they should become frontier model labs, but that selective internal Build capability may become strategically valuable when it is tied closely to firm-specific data, business processes, and use cases.

A further managerial implication concerns governance and infrastructure. MLOps, model monitoring, data governance, compliance, and platform engineering roles are often treated as overhead that slows innovation. Our results point in the opposite direction. These capabilities are associated with stronger returns to AI investment, plausibly by making the handoff between creation and deployment more scalable, reliable, and auditable. Executives should therefore evaluate governance not only as insurance against downside risk, but also as part of the productive system through which AI investments generate upside (Benbya et al., 2024; Schneider et al., 2023).

Limitations and Future Research

Several limitations suggest directions for future research. First, the analysis remains associational despite the lagged fixed-effects design. Time-varying strategic shifts, changes in leadership or managerial quality, or unobserved organizational redesign could influence both workforce composition and firm value. Second, our measures are inferred from occupational data rather than direct observation of tasks, workflows, or project-level collaboration, so we observe the structure of AI labor more clearly than the quality of coordination across roles. Third, while the Build/Apply/Operate taxonomy is theory-informed and robust across alternative classifications, role boundaries are inevitably imperfect and may evolve as AI work itself changes. Fourth, the sample is limited to publicly traded U.S. firms, which may constrain generalizability to smaller firms, private firms, and non-U.S. settings. Fifth, we observe internal AI workforce composition only. External vendors, contractors, and foundation-model providers may supply equivalent Build, Apply, or Operate functions, so the absence of an internal category does not imply the absence of the underlying function. Sixth, the skills-based validation shows that the categories correspond to distinct bundles of reported capability, but the skills data do not observe collaboration between roles, and because reporting dates identify when a skill was first observed rather than when it was acquired, they cannot directly track drift in the task content of a title over time. Finally, Tobin's Q captures forward-looking investor expectations rather than realized operational outcomes, so future work should test whether the same complementarities predict downstream outcomes such as innovation, deployment speed, productivity, or revenue growth.

More broadly, the paper opens a research agenda on the evolution of AI workforce portfolios. Future studies could examine how firms sequence Build, Apply, and Operate investments over time, whether different organizational structures strengthen these complementarities, and how internal workforce architecture interacts with external sourcing choices such as vendors, partnerships, and foundation-model providers. Future work could also examine the antecedents of composition itself, including product strategy, organizational structure, labor-market constraints, and team-level staffing needs.

Overall, the value of AI human capital appears to be associated with configuration as well as with scale. Firm value is higher not simply where firms participate in the AI labor market, but where proprietary creation, domain deployment, and governance capacity are present together. The organizational processes that produce these configurations, and their causal role, remain open questions. As AI capabilities become more broadly accessible, the strategic challenge shifts from acquiring AI talent in the abstract to organizing, integrating, and governing AI work in ways that are commercially consequential and difficult to imitate.

Acknowledgements

Declaration of AI Use: During the preparation and revision of this manuscript, the authors used GenAI to refine language and improve readability and for limited assistance with code refinement and debugging. All AI-assisted text and code were reviewed, tested, and verified by the authors. The authors retain full responsibility for the accuracy, integrity, and content of the work.

References

Abis, S., & Veldkamp, L. (2024). The changing economics of knowledge production. *The Review of Financial Studies*, *37*(1), 89–118. https://doi.org/10.1093/rfs/hhad059

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., Brynjolfsson, E., & Wu, L. (2012). Three-way complementarities: Performance pay, human resource analytics, and information technology. *Management Science*, *58*(5), 913–931. https://doi.org/10.1287/mnsc.1110.1460

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

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

Baird, A., & Maruping, L. M. (2021). The next generation of research on IS use: A theoretical framework of delegation to and from agentic IS artifacts. *MIS Quarterly*, *45*(1), 315–341. https://doi.org/10.25300/MISQ/2021/15882

Benbya, H., Davenport, T. H., & Pachidi, S. (2020). Special issue editorial: Artificial intelligence in organizations: Current state and future opportunities. *MIS Quarterly Executive*, *19*(4), ix–xxi. https://aisel.aisnet.org/misqe/vol19/iss4/4

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

Benbya, H., Strich, F., & Tamm, T. (2024). Navigating generative artificial intelligence promises and perils for knowledge and creative work. *Journal of the Association for Information Systems*, *25*(1), 23–36. https://doi.org/10.17705/1jais.00861

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

Bharadwaj, A. S. (2000). A resource-based perspective on information technology capability and firm performance: An empirical investigation. *MIS Quarterly*, *24*(1), 169–196. https://doi.org/10.2307/3250983

Bostrom, R. P., & Heinen, J. S. (1977). MIS problems and failures: A socio-technical perspective. Part I: The causes. *MIS Quarterly*, *1*(3), 17–32. https://doi.org/10.2307/248710

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

Cai, W., Chen, Y., Rajgopal, S., & Azinovic-Yang, L. (2024). Diversity targets. *Review of Accounting Studies*, *29*(3), 2157–2208. https://doi.org/10.1007/s11142-024-09831-x

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

Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, *35*(1), 128–152. https://doi.org/10.2307/2393553

Fedyk, A., & Hodson, J. (2023). Trading on talent: Human capital and firm performance. *Review of Finance*, *27*(5), 1659–1698. https://doi.org/10.1093/rof/rfac068

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, *42*(12), 2195–2217. https://doi.org/10.1002/smj.3286

Hoffreumon, C., Forman, C., & van Zeebroeck, N. (2024). Make or buy your artificial intelligence? Complementarities in technology sourcing. *Journal of Economics & Management Strategy*, *33*(2), 452–479. https://doi.org/10.1111/jems.12586

Kohli, R., & Grover, V. (2008). Business value of IT: An essay on expanding research directions to keep up with the times. *Journal of the Association for Information Systems*, *9*(1), 23–39. https://doi.org/10.17705/1jais.00147

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

Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, *33*(1), 159–174. https://doi.org/10.2307/2529310

Liang, C., Lourie, B., Nekrasov, A., & Shevlin, T. (2025). The gender position gap and firm performance. *Journal of Business Finance & Accounting*, *52*(5), 2464–2491. https://doi.org/10.1111/jbfa.70006

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

Lou, B., & Wu, L. (2021). AI on drugs: Can artificial intelligence accelerate drug development? Evidence from a large-scale examination of bio-pharma firms. *MIS Quarterly*, *45*(3), 1451–1482. https://doi.org/10.25300/MISQ/2021/16565

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., 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. (1990). The economics of modern manufacturing: Technology, strategy, and organization. *American Economic Review*, *80*(3), 511–528. https://www.jstor.org/stable/2006681

O'Reilly, C. A., & Tushman, M. L. (2013). Organizational ambidexterity: Past, present, and future. *Academy of Management Perspectives*, *27*(4), 324–338. https://doi.org/10.5465/amp.2013.0025

Pavlou, P. A., & El Sawy, O. A. (2011). Understanding the elusive black box of dynamic capabilities. *Decision Sciences*, *42*(1), 239–273. https://doi.org/10.1111/j.1540-5915.2010.00287.x

Roberts, N., Galluch, P. S., Dinger, M., & Grover, V. (2012). Absorptive capacity and information systems research: Review, synthesis, and directions for future research. *MIS Quarterly*, *36*(2), 625–648. https://doi.org/10.2307/41703470

Sarker, S., Chatterjee, S., Xiao, X., & Elbanna, A. (2019). The sociotechnical axis of cohesion for the IS discipline: Its historical legacy and its continued relevance. *MIS Quarterly*, *43*(3), 695–719. https://doi.org/10.25300/MISQ/2019/13747

Schneider, J., Abraham, R., Meske, C., & vom Brocke, J. (2023). Artificial intelligence governance for businesses. *Information Systems Management*, *40*(3), 229–249. https://doi.org/10.1080/10580530.2022.2085825

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

Storey, V. C., Yue, W. T., Zhao, J. L., & Lukyanenko, R. (2025). Generative artificial intelligence: Evolving technology, growing societal impact, and opportunities for information systems research. *Information Systems Frontiers*, *27*(5), 2081–2102. https://doi.org/10.1007/s10796-025-10581-7

Sunyaev, A., Benlian, A., Pfeiffer, J., Jussupow, E., Thiebes, S., Maedche, A., & Gawlitza, J. (2025). High-risk artificial intelligence. *Business & Information Systems Engineering*, *67*(6), 981–994. https://doi.org/10.1007/s12599-025-00942-6

Tambe, P. (2014). Big data investment, skills, and firm value. *Management Science*, *60*(6), 1452–1469. https://doi.org/10.1287/mnsc.2014.1899

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

Thiebes, S., Lins, S., & Sunyaev, A. (2021). Trustworthy artificial intelligence. *Electronic Markets*, *31*(2), 447–464. https://doi.org/10.1007/s12525-020-00441-4

Wu, L., Hitt, L. M., & Lou, B. (2020). Data analytics, innovation, and firm productivity. *Management Science*, *66*(5), 2017–2039. https://doi.org/10.1287/mnsc.2018.3281