From AI Access to Organizational Capability: Pricing the Corporate AI Transition

Keith Lee

Swiss Institute of Artificial Intelligence, Chaltenbodenstrasse 26, 8834 Schindellegi, Schwyz, Switzerland

Abstract

Corporate AI use is commonly measured through access, licenses or reported use, but these proxies tell us little about whether firms have embedded AI in the production function. This article argues that the economic value of AI today is increasingly derived from its implementation: its linkage to data, workflows, systems and organizational decision processes. Using recent asset-pricing evidence from the AI Premium literature, the analysis distinguishes nominal access from substantive consumption and explain why intensive, experienced and deeply integrated AI use provides a meaningful market signal. The current pattern is consistent with a temporary deployment bottleneck, but as enterprise tools, integrations, practices and standards emerge, competitive advantage will likely move toward workflow transformation, analytic substitution, validation, decision integration and measurable benefit. AI reduces the cost of machine-produced insights, whereas judgment, validation, structuring the organization and holding people accountable will remain relatively valuable. The future AI premium will thus depend less on owning AI assets and more on converting abundant machine-generated capacity into lower costs, better decisions, improved performance and sustained returns.

1 Introduction - The Gap Between AI Access and Economic Capability

Corporate AI is being measured with a category error. A firm can be described as an AI adopter because employees have access to a general-purpose chatbot, because one business function uses an AI application, because a production process has been redesigned around model outputs or because models have been connected programmatically to internal data and operational systems. These states differ economically even though conventional adoption statistics can place them under the same label. The measurement problem is already visible in official data. Eurostat reported that 19.95 percent of European Union enterprises with at least ten employees used at least one AI technology in 2025.[1] A narrower comparison of production use found AI in production at about 7 percent of United States firms and 4 percent of European Union firms.[2] United States survey estimates have ranged much higher when questions encompass any business function or employment-weighted adoption. The Federal Reserve warns that experimental, incidental and economically insignificant use is difficult to separate from meaningful adoption in current statistics.[3]

That distinction changes how the corporate AI transition should be interpreted. Purchasing access through Copilot, ChatGPT, Claude, Gemini or another service gives an organization an option to use increasingly capable models. The purchase itself doesn't reveal much information about where the models enter production, which data they can reach, whether employees use them repeatedly, whether workflows have changed, who verifies outputs or whether any resulting capacity becomes lower cost, higher output or better decisions. A cross-country survey of more than 5,000 senior executives illustrates the gap. Around 70 percent of surveyed firms reported active AI use, while more than 80 percent reported no effect on productivity or employment during the preceding three years.[4] The coexistence of widespread reported use and limited realized firm-level impact matches an implementation lag. It also shows why license counts and broad adoption declarations provide weak measures of economic capability.

This paper argues that the present AI premium is best understood as the pricing of exposure to an incomplete organizational transition. Access has become relatively easy. Implementation remains difficult. Value capture remains harder still. Firms have to connect models with data, software and work processes, specify access rights, establish controls, train employees, redesign interfaces, change task allocation, build evaluation routines and persuade workers and managers to rely on new systems where reliance is justified. These investments resemble the complementary intangible capital associated with earlier general-purpose technologies.[5] Their scarcity during the early diffusion period can make implementation capacity economically significant before the eventual productivity effects are fully visible.

The strongest recent market evidence for this interpretation comes from a 2026 research constructs an AI factor from realized consumption recorded by OpenRouter between January 2024 and April 2026. The dataset contains roughly 380 trillion tokens across more than 400 language models and represents approximately 2 percent of current global monthly AI token consumption according to their estimate. The factor combines weekly growth in tokens, expenditure and active users, after which firm-level AI betas are calculated from the comovement of stock returns with innovations in that factor. Firms with higher exposure subsequently earned substantially higher returns. The value-weighted high-minus-low portfolio earned 64.1 basis points per week in the baseline result.[6]

The result has two meanings that need to be kept separate. The first concerns depth of consumption. The return signal is much stronger when the AI factor is constructed from frontier closed-source models, established paying users, seasoned users and longer prompts. Casual and extensive-margin consumption produces weaker estimates. This supports a distinction between nominal access and economically substantive use. The data measure realized inference rather than corporate license ownership, so the study cannot establish that buying a Copilot or ChatGPT subscription produces exactly zero market effect. It does establish something more defensible and more useful: variation associated with intensive and sophisticated consumption contains substantially more asset-pricing information than variation associated with new or casual use.[7]

The second meaning concerns asset pricing. A positive AI premium is an expected-return result. It cannot automatically be converted into a statement that investors assign higher present valuations to AI-exposed firms. The premium is defined as compensation investors require for an additional unit of market-implied AI exposure. In standard finance, a higher required return can accompany stronger expected growth opportunities, greater systematic risk or both. Holding expected cash flows fixed, a higher discount rate reduces present value. Technological revolutions can alter systematic uncertainty precisely because investors are learning which businesses will gain and which will lose as a recent technology diffuses.[8] The correct language, accordingly, is that markets price exposure to the AI transition. Claims about a mechanical valuation bonus go beyond the evidence.

The early character of the transition is central to interpreting another result. In the occupation and skill mapping, positive market-implied AI exposure is concentrated in skills associated with installation, repair, programming, persuasion, instruction and systems-oriented work.[9] Analytical, scientific and operations-control skills are associated with more negative exposure.[10] This does not establish a permanent hierarchy in which installation remains more valuable than analysis. It may describe a particular phase of diffusion. Systems still need to be built, connected and accepted before massive quantities of analytical work can be shifted onto them. Implementation specialists, managers capable of changing workflows, employees able to coordinate adoption and workers able to resolve the institutional frictions of deployment are temporarily complementary to diffusion. Once deployment becomes routine, that complementarity can weaken.

The argument developed here extends the market result through time. During the present phase, implementation removes a deployment constraint. That creates productive capacity. Productive capacity eventually must become output. As APIs, model-routing systems, security practices, enterprise connectors and implementation expertise become more standardized, the scarcity value of basic deployment should fall. Competitive differentiation should migrate toward the quality of use: which tasks are delegated, how model output is evaluated, how decisions are changed, how workflows are redesigned and how much expensive professional work can be performed at lower marginal cost. The long-run corporate AI question is accordingly different from the current one. In the present phase, investors need to identify firms capable of crossing the implementation bottleneck. In the mature phase, they will need to identify firms capable of converting abundant machine-generated analysis into organizational action and cash flow.

2 What Markets Are Pricing in the AI Transition

2.1 The Baseline Result and the Intensive Margin

The analysis makes an unusually useful contribution because its central explanatory variable is realized model consumption. OpenRouter connects users and developers to hundreds of models through one API endpoint and the dataset records paid inference requests rather than survey answers about whether a company claims to use AI.[11] The records are aggregated into growth in token volume, dollar expenditure and distinct users. This provides a high-frequency measure of changes in AI consumption intensity. It also imposes an interpretive boundary. OpenRouter represents only a fraction of global inference and its API-oriented ecosystem can capture sophisticated users differently from consumer-facing applications. The study should be read as an unusually rich measure of one segment of realized AI consumption rather than a census of enterprise deployment.

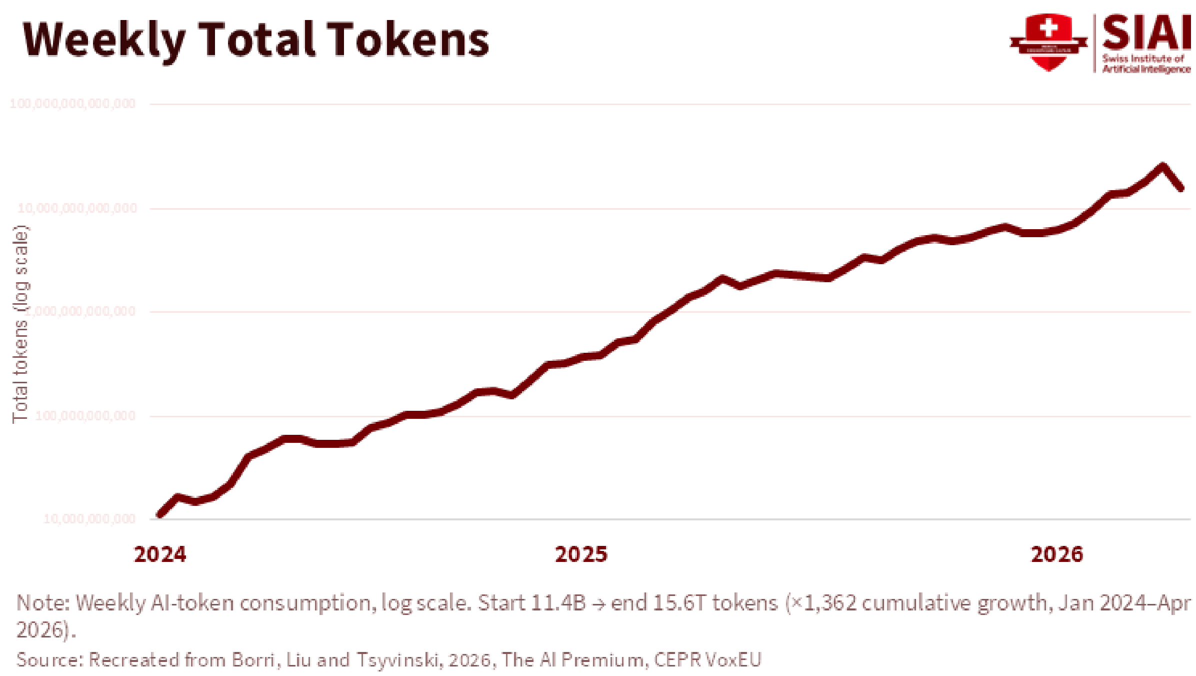

Figure 1: Realized AI consumption rose from billions to trillions of tokens per week, showing the scale of use behind the market-based AI factor.

The composition of realized consumption offers a complementary view of what the AI factor is measuring. Categorized OpenRouter tokens in the final sample week are concentrated in programming and technology-related prompts, which together account for more than three-fifths of categorized volume, while roleplay and science follow at a considerable distance and consumer-facing categories such as marketing, translation and trivia remain marginal. This concentration is consistent with the intensive-margin results reported below: the factor draws disproportionately on technical, production-oriented use rather than diffuse consumer experimentation.

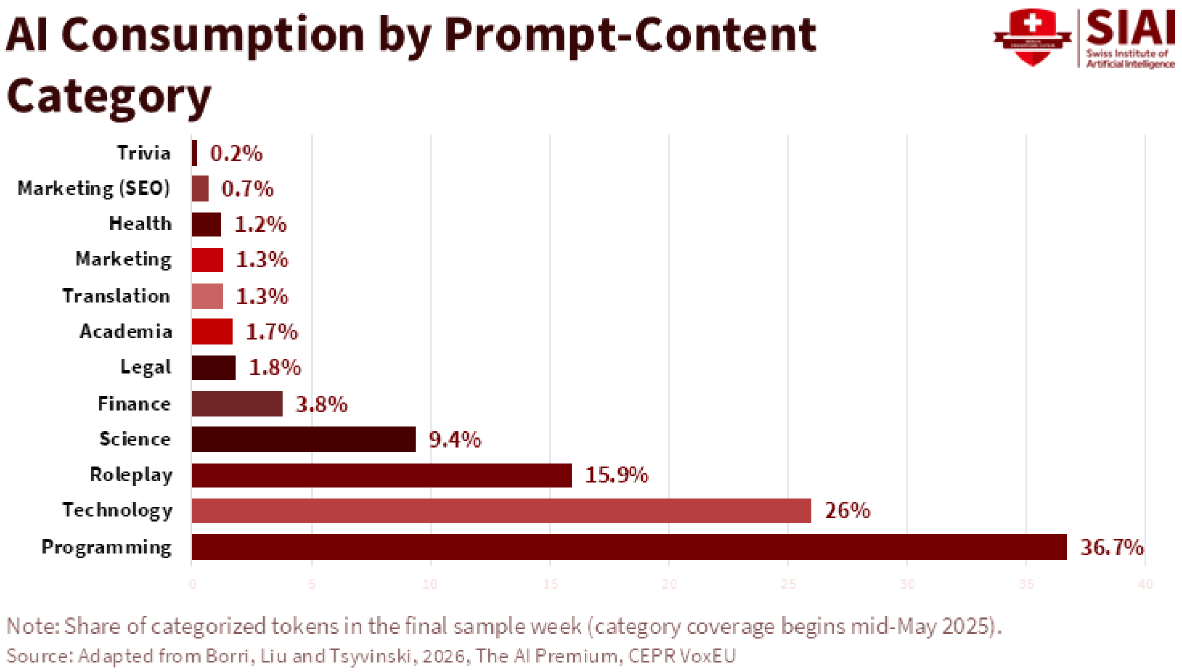

Figure 2: Programming and technology account for almost two-thirds of categorized token use, reflecting the technical concentration of current AI consumption.

2.2 The Baseline Result and Its Robustness

The baseline market result is economically substantial. Firms are sorted according to rolling AI betas estimated from thirteen weeks of return data, with market returns included in the first-stage regression. The value-weighted spread between the high and low AI-beta portfolios is 64.1 basis points per week. The result survives controls designed to separate the AI factor from technology-sector exposure, semiconductor returns, AI-themed exchange-traded funds, industry effects and public attention measured through search interest. Industry-demeaned exposure generates a high-minus-low spread of 60.3 basis points per week in one principal specification, close to the baseline estimate. The finding appears, as a result, to capture something beyond a simple technology-stock rally.[12]

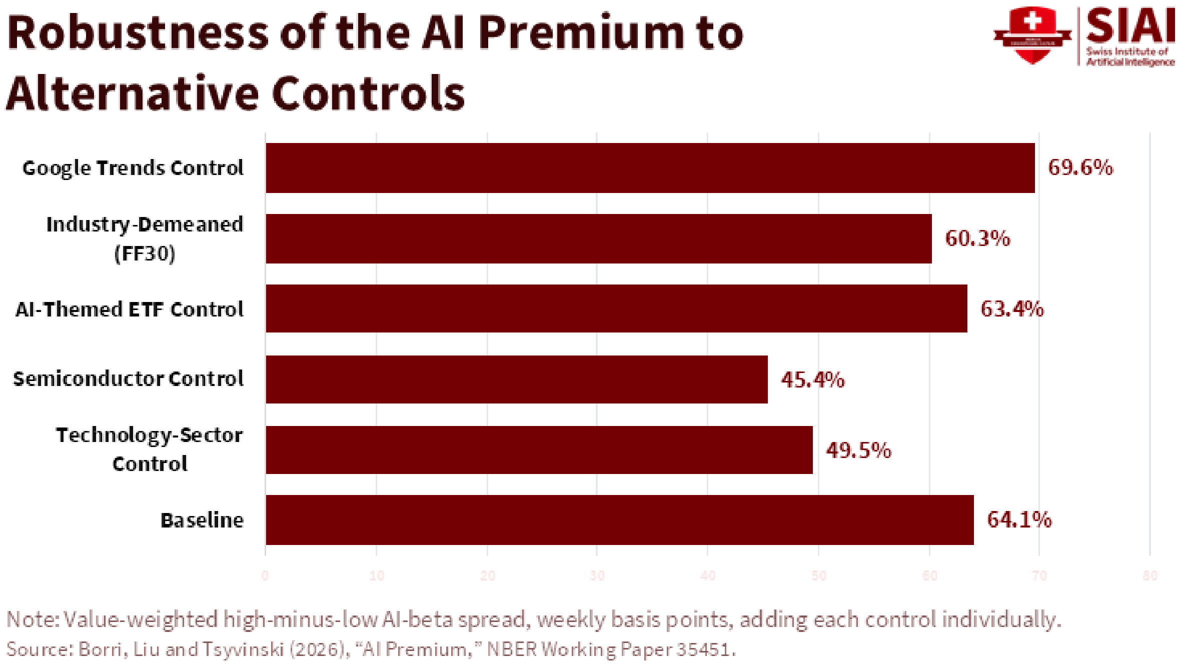

Figure 3: The AI premium remains sizeable after controls for technology exposure, semiconductors, AI-themed funds, industry composition and public attention.

Its most relevant finding for corporate strategy comes from the decomposition of the consumption factor. Exposure based on closed-source model consumption produces a high-minus-low return spread of 53.4 basis points per week, compared with 32.3 basis points for open-weight consumption. The associated statistics are 2.47 and 1.76. Model category alone should not be treated as a permanent indicator of quality. The economic content is narrower. During this sample, consumption closer to the frontier carried a clearer market signal.[13] One plausible interpretation is that frontier consumption contained more information about commercially relevant capability, demanding applications and proximity to the technological frontier. User characteristics strengthen the same conclusion. The weekly high-minus-low spread is 66.8 basis points when the factor is based on paid or core accounts and 21.9 basis points for new accounts. It is 59.3 basis points for seasoned users and 29.3 for non-seasoned users. Long-prompt consumption generates a 54.1 basis-point spread compared with 33.0 for short prompts. These estimates do not reveal exactly what any corporate user was doing. They establish a consistent pattern across several proxies: the market signal strengthens as observed consumption moves from entry toward repeated, experienced and complex use.[14]

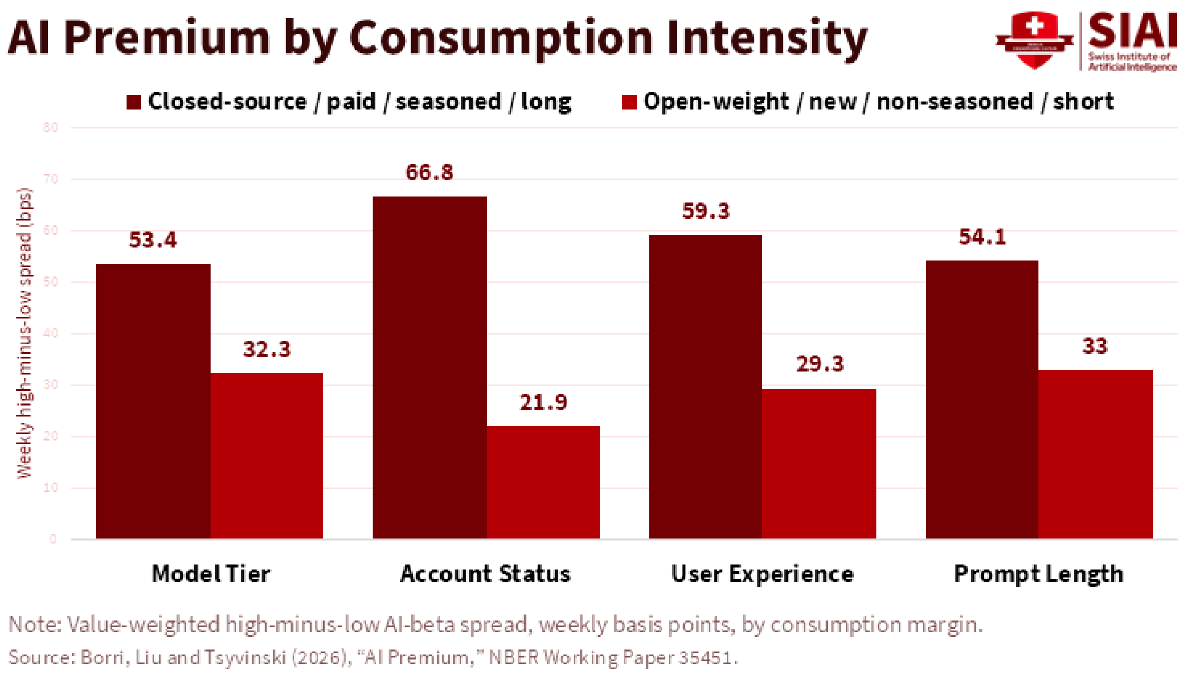

Figure 4: Repeated, experienced and more complex AI consumption carries a stronger market signal than entry-stage or casual use.

That pattern provides a better empirical foundation for criticizing superficial corporate AI claims than a categorical assertion about subscriptions. A company may purchase thousands of seats and later build a valuable deployment around them. Another may purchase the same number and see usage decay after an initial trial. A third may achieve significant effects through a small number of heavily used API integrations. Seat counts cannot distinguish among these cases. Registered user counts have the same weakness because registration records the extensive margin. The Borri evidence gives considerably more weight to the intensive margin. Experienced users, continued consumption and complex requests appear to contain information that initial access does not.

The distinction also explains why a generalized label such as AI company is analytically weak. The AI factor identifies positive and negative exposures across industries rather than reproducing the technology-sector classification. The market-return relationship remains when broad industry effects are removed. Around nineteen frontier-model release events, the high AI-beta portfolio outperformed the low AI-beta portfolio by 1.9 percent over a five-trading-day window in the study's event analysis. Yet the positive spread also persists after release weeks are excluded, which means the effect is not confined to short bursts of model-release enthusiasm.[15] Geography supplies another clue. The high-minus-low spread in developed markets is 17.9 basis points per week, with a statistic of 2.77, while the corresponding emerging-market estimate is 5.0 basis points and statistically indistinguishable from zero in the reported specification. This difference fits the idea that market pricing is strongest where firms, investors, infrastructure and production systems are closer to frontier development and commercial deployment. It does not prove that developed-market firms are inherently better adopters. It identifies where the transition factor had stronger pricing relevance during the sample.[16]

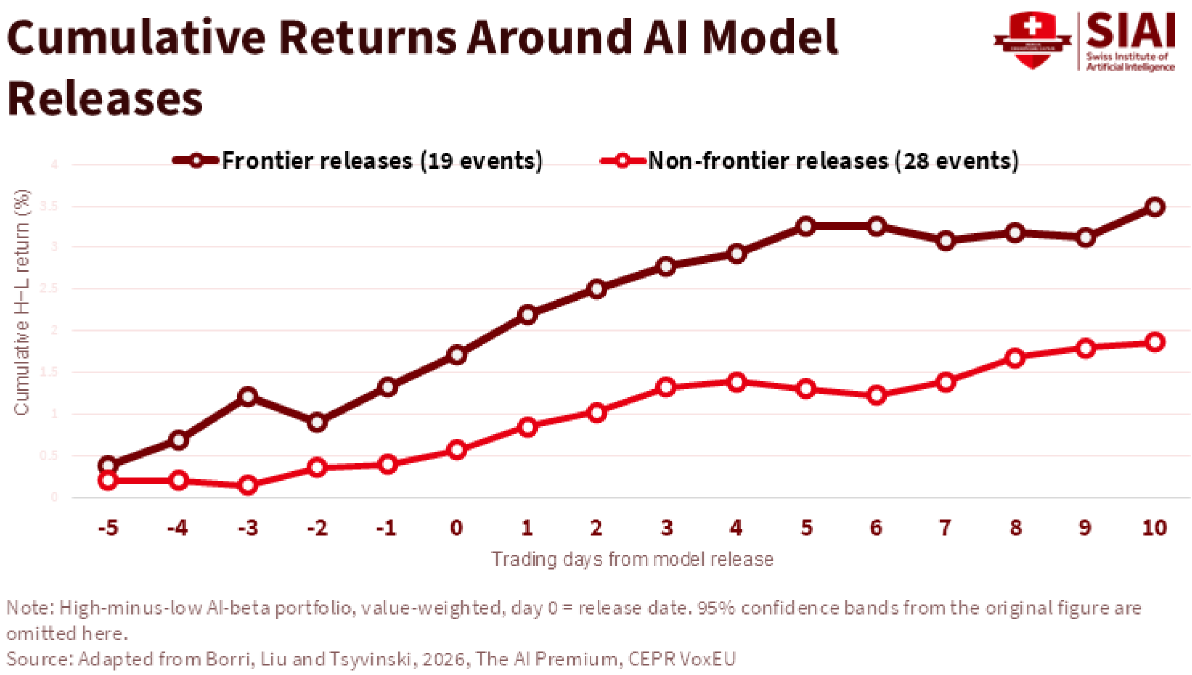

Figure 5: High-AI-beta firms respond more strongly around frontier-model releases, although the premium extends beyond release events.

Geography supplies another clue. The high-minus-low spread in developed markets is 17.9 basis points per week, with a statistic of 2.77, while the corresponding emerging-market estimate is 5.0 basis points and statistically indistinguishable from zero in the reported specification. This difference fits the idea that market pricing is strongest where firms, investors, infrastructure and production systems are closer to frontier development and commercial deployment. It does not prove that developed-market firms are inherently better adopters. It identifies where the transition factor had stronger pricing relevance during the sample.[16]

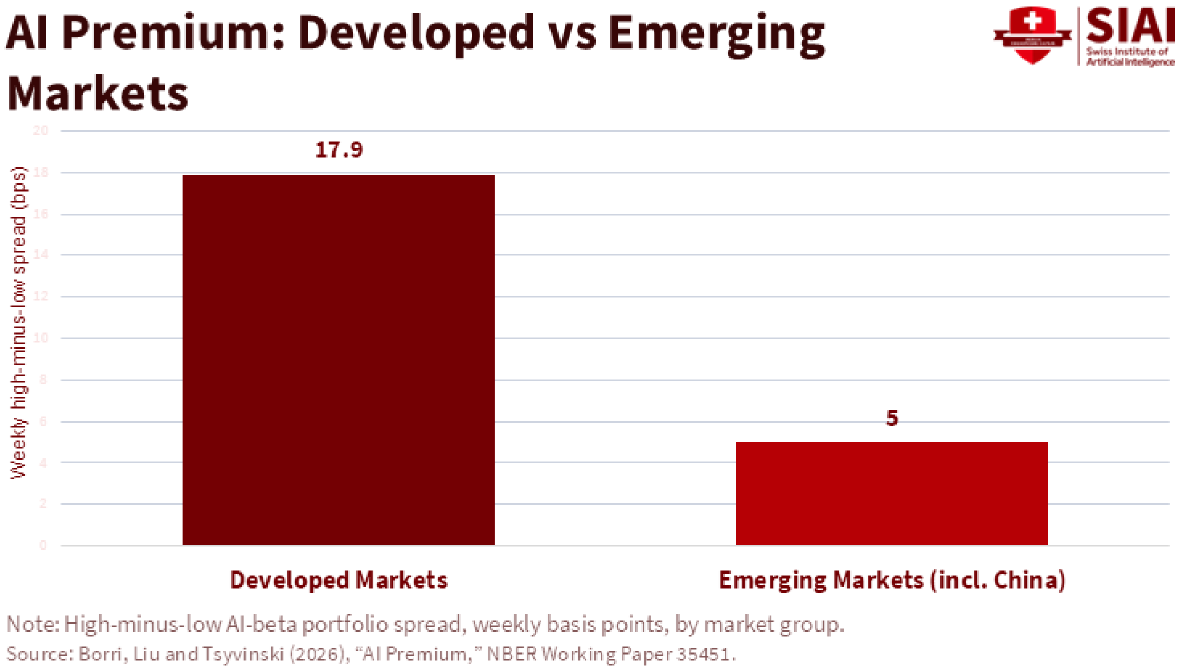

Figure 6: The AI premium is substantially stronger in developed markets, where firms and investors are closer to frontier infrastructure and commercial deployment.

2.3 Nominal Adoption, Disclosure and Occupational Exposure

Nominal adoption measures appear especially weak when compared across surveys. In the United States, the Census Business Trends and Outlook Survey measured roughly 18 percent of firms using AI at the end of 2025 after its question broadened to any business function. The employment-weighted Survey of Business Uncertainty produced a much higher figure of around 78 percent, while an individual worker survey found work-related generative AI use at about 41 percent. The Federal Reserve attributes much of this dispersion to differences in units of analysis, weighting and question wording and explicitly identifies the difficulty of separating experimental use from economically meaningful adoption.[17] The implication for asset pricing is immediate. Any factor built from a binary adoption variable risks combining firms that occupy quite various positions in the deployment process.

This is also why announcements have limited informational content. A statement that a firm has launched an AI strategy may describe budget allocation, an experimental program, a procurement decision or a completed production redesign. Regulatory filings can contain genuine information but mention counts alone cannot distinguish between those states. The disclosure test is valuable because it finds that the consumption-based signal survives the inclusion of portfolios formed from AI mentions in regulatory filings. Corporate rhetoric and realized consumption exposure should not be used interchangeably.

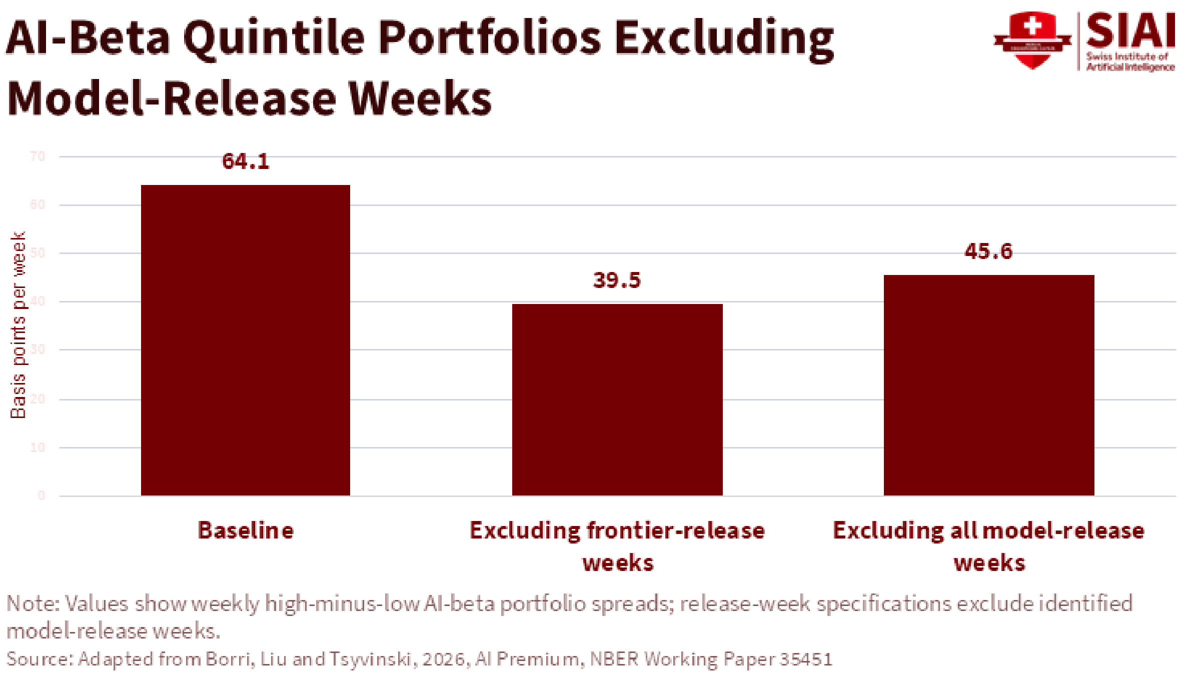

Figure 7: The premium remains positive after model-release weeks are removed, showing that priced AI exposure is not confined to announcement effects.

The deeper lesson is that adoption should be decomposed. Access measures whether a firm can use AI. Consumption intensity measures whether people or systems use it. Implementation measures whether models are connected to organizational processes. Process redesign measures whether the allocation of work changes. Value capture measures whether those changes appear in costs, output, quality, innovation, working capital, decision speed or some other economically relevant outcome. The present evidence is strongest on the gap between the first two categories. It is increasingly suggestive on implementation. Firm-level causal evidence on the complete chain from model integration to realized profits remains much thinner.[18]

The occupational skill mapping is especially relevant to implementation. Installation and repair, programming, persuasion, instruction and systems integration appear toward the positive side of the skill exposure ranking, while science and operations-control skills appear toward the negative side. The variable records skill content. It does not observe an enterprise connecting a model API to a customer database or redesigning an approval system. The pattern provides suggestive evidence for an implementation-complementarity interpretation rather than a causal estimate of the return on installing AI.[19] That distinction matters. Treating the figure as proof that technical deployment automatically raises firm value would reproduce the same measurement error that the intensive-margin results help correct.

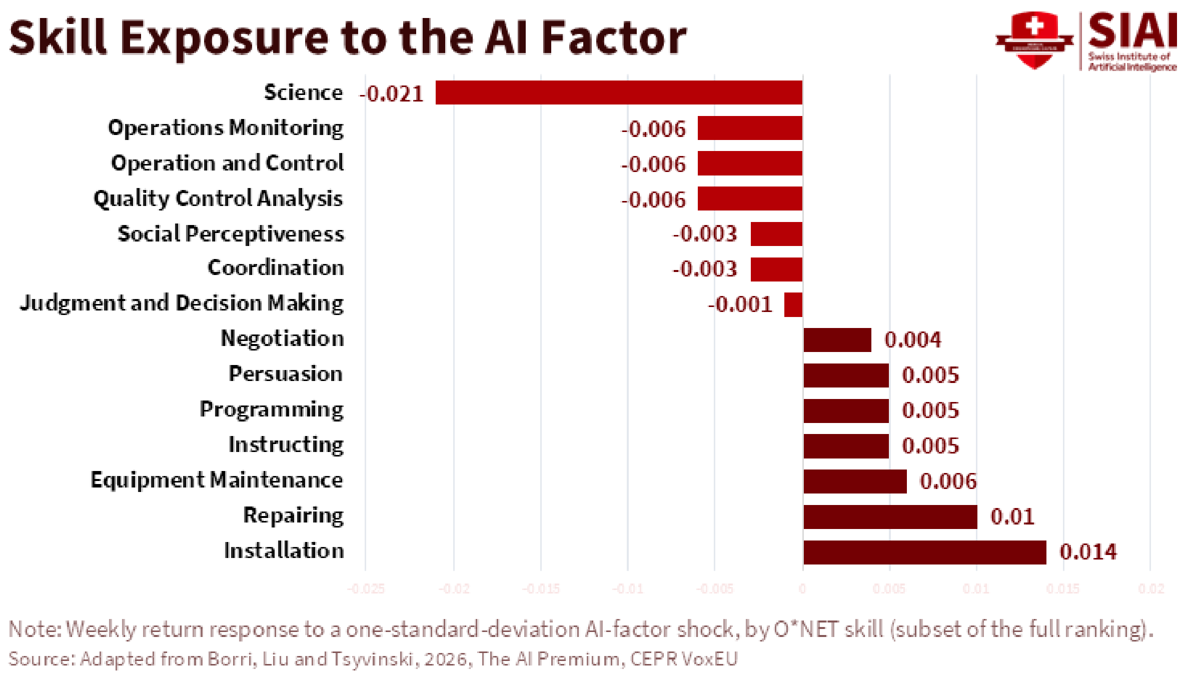

Figure 8: Implementation-oriented skills currently carry more positive market-implied AI exposure, while scientific and operations-control skills sit on the negative side.

3 Why Implementation Carries the Strongest Signal Today

3.1 Implementation as Organizational Capital

Implementation is the organizational layer between access and production. Its economic significance comes from the fact that a general-purpose model arrives without the firm-specific information, permissions, process rules, interfaces and decision rights required to operate inside a particular business. A subscription can deliver a powerful model to an employee. A production system requires the model to receive the right information at the right time, produce outputs in a usable form, route those outputs to the appropriate worker or software process, expose failures to review and preserve responsibility for consequential decisions. These are complementary investments in technology rather than properties of the model itself.

Productivity literature provides a useful historical analog. The productivity J-curve describes how general-purpose technologies require substantial complementary intangible investment before the resulting productivity gains are fully measured. Organizational redesign, new processes and accumulated knowledge can initially appear as costs while the productive asset is being constructed. Benefits arrive later when the complementary capital becomes usable at scale. Their framework predates the present generative AI diffusion but its logic fits the current implementation problem closely.[20] The model endpoint can be purchased immediately. The surrounding organization cannot.

Contemporary firm evidence shows that this surrounding capacity is scarce. The OECD's examination of AI adoption in firms identifies specialized skills and data maturity as barriers to uptake and implementation. Businesses surveyed for the report valued industry-specific training and practical work with relevant systems and datasets. Public institutions that assist digital diffusion also identified uncertainty about return on investment and inadequate data maturity as recurrent obstacles. Implementation capacity therefore includes technical competence, data organization and managerial ability to identify a production case with sufficient economic value.[21] European data reveal a large firm-size gradient consistent with fixed implementation costs. In 2025, 55.03 percent of large European Union enterprises reported using at least one AI technology, compared with 30.36 percent of medium enterprises and 17 percent of small enterprises. Eurostat explicitly identifies implementation complexity, scale economies and affordability as possible contributors to this pattern. Among enterprises that had considered using AI but had not done so, lack of relevant expertise has been reported as a major constraint in the underlying European statistics. The economic implication is that access to a model does little to remove complementarities involving data, systems and organizational expertise.[22]

This creates an implementation bottleneck with several layers. Data must be available in formats that can be retrieved or passed to models. Identity and access rules must limit exposure of sensitive information. Applications require connectors, APIs or other interfaces. Outputs need evaluation criteria. Employees need to know when to delegate a task, when to verify a response and when to escalate it. Managers must identify an owner for the process and define accountability when a model contributes to an outcome. These activities consume time and managerial attention. Their costs are especially visible before large-scale value capture begins.

Implementation therefore resembles installation capital. The concept should be used cautiously because much of the investment is intangible. The asset consists partly of software and technical configuration, partly of redesigned processes, partly of accumulated employee experience and partly of institutional knowledge about where the model succeeds and fails. A functioning production deployment can contain all four. The value of that capital lies in removing a constraint on later production. It creates the capacity to use models consistently on a scale.[23] It cannot guarantee profitable output.

This distinction resolves an apparent tension in the market evidence. Why should installation, systems integration, communication and persuasion appear positively exposed to the AI factor when analytical skills appear negatively exposed? At an early diffusion stage, the scarce input is frequently the ability to place AI inside an organization. Once models are available as external services, model access can expand far faster than firms can redesign data systems and work processes. The marginal economic contribution of implementation skill can rise during that interval because implementation determines whether cheap model capability can enter production. The occupational evidence is consistent with this mechanism, although it does not establish a causal relationship.[24]

3.2 Evidence From Field Experiments on Implementation

The distinction between productive capacity and productive output is also necessary because AI experiments show highly heterogeneous effects across tasks. Brynjolfsson, Li and Raymond studied 5,172 customer-support agents and found that AI assistance increased issues resolved per hour by 15 percent on average.[25] The gains were concentrated among less experienced and lower-skilled workers, while experienced workers experienced smaller gains and some quality deterioration. This is a genuine productivity effect in a well-defined production environment. The system supplied recommendations closely matched to a specific workflow and the organization could measure output. The result shows what implementation can enable once model assistance is tied to a task with observable performance.

Evidence from experienced software developers makes the limitation concrete. A 2025 randomized trial conducted by METR with experienced open-source developers working on repositories they knew well found that participants using then-current AI tools took 19 percent longer on the studied tasks.[26] The narrowness of the sample and identify several reasons their finding can coexist with strong benchmark performance and positive evidence elsewhere, including implicit project requirements, high-quality standards and differences between benchmark tasks and ordinary development work. The lesson for firms is neither that coding assistants fail nor that benchmarks are useless. The relevant lesson is that a capability demonstrated in isolation can impose coordination and verification costs when inserted into a demanding production setting.

Implementation is the process through which those hidden costs become visible. A firm can discover that drafting time falls while review time rises. A model can accelerate an analyst while increasing the burden on a senior employee responsible for verification. Automated customer communication can reduce labor input while creating new exception-handling work. Code generation can become cheaper while testing requirements increase. These outcomes can all occur alongside technically successful adoption. Productivity appears only after the entire process is measured rather than the model interaction alone.

The Microsoft Copilot field experiment illustrates the same principle from the other direction. Integration inside existing email, meeting and document applications reduced an important source of friction. Users did not have to move information manually into a separate chatbot for every task. The experiment found meaningful time-use changes among actual users, especially lower email time. It did not observe direct performance or firm-level productivity and broader task composition changed little. The result occupies a middle stage between access and full production redesign. Integration changes the cost of using AI, while economic value still depends on what employees do with saved time.[27]

The organizational effects can become larger when AI changes coordination itself. A field experiment at Procter and Gamble examined individuals and teams solving product-development problems. AI-enabled individuals and AI-enabled teams improved solution quality and individual users could achieve performance comparable with traditional two-person teams in the setting studied. The experiment also found more balanced solutions across commercial and technical domains, suggesting that AI can transmit some knowledge across professional boundaries. This provides a direct counterargument to a purely substitution-based account. AI can complement human work by reducing the cost of accessing expertise and coordinating knowledge. The same mechanism can alter organizational design if firms redesign teams around the new capability.[28]

A better internal measurement architecture would separate access from active use, active use from integration and integration from realized output. Firms need baseline measurements before deployment so that changes in labor time, cycle time, errors, rework, customer outcomes and unit costs can be compared with a pre-AI process. Usage measures need enough granularity to identify task classes rather than counting registered accounts. Integrated workflows need defined owners and review rules. Management then gains a way to stop low-value deployments rather than preserving them because they satisfy an AI adoption target.

Implementation should be treated as a transitional organizational complement. Its current scarcity can produce strong market relevance because much corporate AI activity remains between experimentation and production. The scarcity will erode. APIs, enterprise model gateways, retrieval systems, evaluation software and implementation practices are becoming easier to buy or reproduce. The strategic question then moves downstream. Once many firms can install comparable systems, installation itself ceases to separate leaders from followers. What the organization makes those systems do becomes decisive.

4 From Deployment to Analytical Substitution and Value Capture

The transition can be understood as a sequence of six organizational states: access, installation, substantive adoption, process redesign, analytical substitution and value capture. These states are an analytical synthesis rather than an empirically validated maturity model. Their purpose is to separate mechanisms that are routinely compressed into the single word adoption. A firm can advance unevenly and different business units can occupy different states at the same time. The sequence still clarifies why current market signals can differ from those likely to matter after deployment becomes common.

4.1 From Access to Process Redesign

Access is the cheapest stage. The organization purchases licenses, approves external tools or makes a model available through an existing software suite. Access matters because employees cannot experiment with prohibited or unavailable technology. Yet it creates little information about expected cash flows without evidence of use. The major differences among official adoption measures already show how easily access and occasional use can inflate headline diffusion. Seat ownership is therefore comparable to installed optionality. It gives the firm a chance to learn. Its financial significance depends on whether the option is exercised productively.

Installation follows when the organization begins connecting models with systems, data and workflows. This is the current bottleneck emphasized by the market interpretation developed here. The OpenRouter study does not directly observe corporate installation projects, yet its intensive-margin findings, its early evidence on tool-using model activity and its positive exposure among systems-oriented skills are consistent with a market in which deeper operational use carries more information than casual consumption. Agentic requests invoking external tools rose from almost zero in 2024 to roughly half of token consumption by the end of their sample, while estimates of separate agentic premiums remained imprecise because this activity was very new.[29]

The transition toward APIs and tool use matters because programmatic integration changes the division of labor between person and model. Anthropic's September 2025 Economic Index provides provider-specific evidence from Claude usage. In its sample, 77 percent of first-party API transcripts were classified as automation patterns, compared with 12 percent classified as augmentation. Across tasks, 97 percent of tasks represented in API traffic showed automation-dominant patterns, compared with 47 percent in the sampled Claude.ai interactions.[30] These figures cannot be generalized to the entire AI economy because they concern one provider and one classification method. They still demonstrate a mechanism: API deployment can move AI from an interactive assistant toward a component that performs work inside a larger software process.

Substantive adoption comes after installation when use becomes consistent enough to change the ordinary production process. This stage is frequently slower than procurement because employees have to decide when AI is worth using and managers have to tolerate temporary experimentation. The Copilot experiment is instructive. More than 90 percent of treated workers tried the tool at least once,[31] while stabilized weekly use after the initial period remained just under 40 percent of treated workers. A corporate dashboard built around the ever-used metric would describe near-universal adoption in that group. A dashboard built around repeated weekly usage would report a vastly different state.

Process redesign begins when firms change task allocation and decision rights around AI capability. This is the point at which implementation can start to alter organizational productivity structurally. A model that saves an analyst forty minutes while every downstream approval remains unchanged creates a local efficiency gain. A redesigned process can remove a handoff, change staffing ratios, shorten a decision cycle or allow one team to handle a larger workload. The distinction between individual and firm productivity is becoming important here. Time saved by employees is economically valuable only if the organization reallocates it, converts it to additional output or changes the resource requirement for a process.

The implication for market pricing is that implementation can command a temporary premium while process redesign remains scarce. Investors observing a firm that repeatedly converts technical capability into changed production may rationally treat its return exposure differently from that of a firm buying the same underlying model through a subscription. The source of differentiation lies in complementary organizational capital. Once common software products automate much of integration, the price of that complement should decline. Firms will then be compared to what their redesigned processes accomplish.

Analytical substitution becomes increasingly important at that stage. Generative models have unusually direct applicability to knowledge tasks because they can produce text, code, classification, synthesis and structured recommendations at low marginal cost. The Borri skill mapping places nonroutine analytical and mathematical content on the negative side of market-implied AI exposure while interactive and social content is positive.[32] The result is consistent with a shift in relative scarcity. When the cost of producing a preliminary analysis falls, ownership of generic analytical production capacity becomes less distinctive. The valuable activities move toward determining which analysis deserves trust, choosing an action, integrating the decision with organizational constraints and accepting responsibility for the consequences.

This does not imply the disappearance of analysts. Task-based theories of technological change distinguish occupations from tasks because automation can substitute for parts of a job while leaving other parts complementary to new technology.[33] A financial analyst may delegate data extraction, scenario generation and first-draft commentary while devoting more effort to model selection, assumptions, client communication and accountability. A lawyer may reduce the time spent producing standard research while increasing review of sources and strategic interpretation. An engineer may obtain generated calculations while retaining responsibility for specification, safety and sign-off. The labor effect depends on whether the time saved creates additional demand for higher-level work or reduces the number of workers required to produce a given quantity of service.

The customer support evidence demonstrates one path toward democratized expertise. Less experienced workers obtained larger benefits from AI assistance, which is consistent with models transmitting patterns previously embodied in experienced workers.[34] The cybernetic-teammate experiment points in the same direction by showing that AI could help individuals integrate knowledge across professional domains.[35] These effects can raise productivity and reduce inequality in access to expertise inside firms. They can also reduce the scarcity rents associated with routine professional knowledge. The same mechanism that makes a junior employee more capable can weaken the economic value of selling standardized analysis at a high price.

The resulting competitive shift is subtle. Analytical capability becomes more widely available as model quality rises and inference costs fall. Advantage then depends increasingly on evaluation and execution. Producing an analysis becomes cheap relative to knowing which inputs matter, recognizing an invalid inference, reconciling the output with proprietary context and changing an organization in response. A company with superior internal data can obtain more useful analysis from the same general-purpose model. A company with clearer decision rights can act on reliable output faster. A company with strong evaluation systems can delegate a larger share of work without accepting unacceptable error rates. These complements determine whether abundant analytical production becomes economic value.

4.2 Governance, Value Capture and Strategic Implications

This creates a governance problem that differs from ordinary software automation. Deterministic software generally executes specified rules repeatedly. Generative systems can produce plausible outputs whose reliability varies with task, context and model. When their outputs feed decisions, firms need a rational allocation of human review. Universal human checking can eliminate much of the cost advantage. Minimal checking can accumulate operational risk. Efficient implementation therefore requires calibrated delegation, where the intensity of review reflects error cost, model reliability and reversibility of the decision.

Responsibility cannot be outsourced along with analysis. A model can recommend a credit decision, identify a contractual risk, produce a financial forecast or suggest a maintenance action. The organization still determines whether the recommendation is accepted and who owns the consequence. This is why human complementary skills remain valuable even as analytical generation becomes substitutable. Persuasion, coordination and judgment concern institutional action. Models can contribute to all three but organizations allocate authority through people, processes and legal structures.

As these controls mature, the transition reaches value capture. At this stage, the relevant question is whether the organization converts AI capacity into cash-flow-relevant outcomes. The measures vary by business. Labor-intensive service firms can examine cost per completed case, staffing requirements and throughput. Manufacturers can examine downtime, defects and inventory decisions. Financial firms can measure review time, false positives, customer acquisition costs and losses arising from poor decisions. Software firms can compare development cycle time with defects and maintenance burden. Management should be able to explain the channel through which AI changes an economic quantity.

This is the point at which the present installation emphasis should recede. Installation expertise currently has scarcity value because many firms are still building the basic organizational complement. Standardization will reduce that scarcity. Common enterprise platforms will make secure model route easier. Employees will arrive with more experience. Vendors will package connectors, evaluations and workflow tooling. Organizations will accumulate reusable deployment patterns. The frontier of competition will then move from whether models are connected to how effectively the resulting analytical abundance is exploited.

The future analytical premium is therefore unlikely to belong simply to organizations employing the greatest number of analysts or buying the largest number of model licenses. Firms that can substitute machine analysis for expensive professional time while preserving decision quality can realize cost advantages. Firms that use the same analytical abundance to make faster or better decisions can realize revenue, risk or capital-allocation advantages. Firms that generate enormous volumes of analysis without reliable evaluation may accumulate noise rather than capability.

For workers, the transition changes the composition of valuable expertise. Implementation specialists can command a premium while deployment knowledge is scarce. Their relative scarcity should decline as practices standardize. Domain experts who can verify model output remain important because responsibility and error detection cannot be inferred from generation speed. Workers whose role consists largely of producing standardized analytical artifacts face greater substitution pressure as models improve. Employees who combine domain understanding with the ability to redesign processes, coordinate across organizational boundaries and decide when machine output is dependable occupy a more complementary position.

The temporal sequence also provides a way to reconcile the positive market exposure of interactive skills with the negative exposure of analytical content. The current complement is associated with getting AI into the organization. The prospective substitution occurs after integration makes analytical production scalable. These forces can operate at the same time in different tasks. A consulting firm may require more systems architects and change managers during deployment while gradually reducing hours spent producing standardized research. A bank may expand AI governance and integration functions while automating portions of document review and reporting. There is no reason to expect every occupational effect to move in the same direction.

Investors face the same temporal problem. Current AI exposure metrics need to identify firms crossing the deployment bottleneck. Future metrics will need to identify value capture. License counts should give way to measures such as the complexity of tasks performed through models, depth of workflow integration, share of targeted analytical work conducted with AI, professional-labor cost reductions, changes in cycle time, error-adjusted output and changes in decision quality. Each metric should be tied to an economic mechanism. A percentage of employees using AI is meaningful only when the relevant production consequence is specified.

Governments also need to recognize a distributional consequence of the implementation bottleneck. Large enterprises can spread fixed costs of data engineering, security, legal review and change management across larger operations. Small firms may receive similar model capability from the same external provider while lacking the complementary capital needed to integrate it. Eurostat's sharp size gradient and OECD evidence on skills and data maturity support that concern.[36, 37] Public intervention is most defensible where it lowers shared implementation barriers through training tied to real production problems, interoperable standards, access to high-quality public data and institutions that help smaller firms assess viable use cases. Subsidizing undifferentiated subscriptions would target the least scarce part of the production stack.

The strongest argument for optimism remains intact. AI can raise measurable task productivity, broaden access to expertise, reduce time spent on routine communication and help individuals perform work that previously required more experienced colleagues or larger teams. The evidence from customer support, consulting, Copilot and collaborative product development establishes that these mechanisms are real in specific settings.[38, 39, 40, 41] The qualification is organizational. Task-level gains become durable corporate returns only when firms place them inside production systems that capture saved resources or improve outcomes.

The long-run AI premium should migrate. During installation, markets can price exposure to firms and skills complementary to deployment.[42] As deployment becomes common, differentiation shifts toward sophisticated consumption and process redesign. As analytical tasks become cheaper, rents can move away from standardized analytical production. Value ultimately accrues where organizations translate model capability into reduced resource requirements, faster high-quality decisions, new products or superior coordination. The precise path will differ by industry and asset prices will continue to incorporate both opportunity and risk. The direction of measurement is clearer: access metrics should become less important as the technology diffuses, while operational and economic outcome metrics should become more important.

5 Conclusion - The Long-Term Meaning of the AI Premium

The central asset-pricing conclusion is narrower and stronger than the claim that markets give an automatic valuation bonus to firms using AI. The observed AI premium is a positive expected-return relationship associated with market-implied exposure to innovations in realized AI consumption. A higher expected return can compensate investors for systematic transition risk. Present valuation depends simultaneously on expected future cash flows and the rate at which those cash flows are discounted. Within that discipline, the evidence distinguishes meaningful exposure from corporate AI rhetoric. Sophisticated, repeated and frontier-oriented consumption carries the clearer market signal. Corporate mention counts do not explain that signal well. The current positive exposure of installation, system-oriented and interactive skills is in line with an economy still building the complementary organizational capital required to place AI inside production.

That condition is temporary. Installation removes a deployment constraint. Durable return requires value capture. As basic implementation becomes standardized, investors will need to place greater weight on task complexity, workflow penetration, analytical substitution, professional-labor savings, decision speed, decision quality and measurable changes in operating performance. License counts and registered-user totals will progressively lose information because access itself is becoming ordinary. The strategic implication follows directly. Firms outside the frontier-model race can still acquire economically meaningful AI exposure by integrating models deeply enough to change production. Their longer-run differentiation will depend on the quality of that use after integration becomes common. Markets are beginning to price the transition. The eventual winners will be identified by what the installed capacity produces.

References

[1, 22, 36] Eurostat (2026) *The Use of Artificial Intelligence Technologies in the European Union: Key Results, 2026 Edition.* Luxembourg: Publications Office of the European Union.

[2] Bick, A., Blandin, A., Deming, D.J., Fuchs-Schündeln, N. and Jessen, J. (2026) *Mind the Gap: AI Adoption in Europe and the U.S.* NBER Working Paper 34995. Cambridge, MA: National Bureau of Economic

Research.

[3, 17] Allen, J.S. (2026) 'Monitoring AI Adoption in the U.S. Economy', *FEDS Notes*, 3 April. Washington, DC: Board of Governors of the Federal Reserve System.

[4] Yotzov, I., Barrero, J.M., Bloom, N., Bunn, P., Davis, S.J., Foster, K.M., Jalca, A., Meyer, B.H., Mizen, P., Navarrete, M.A., Smietanka, P., Thwaites, G. and Wang, B.Z. (2026) *Firm Data on AI*. NBER Working Paper 34836. Cambridge, MA: National Bureau of Economic Research.

[5, 20, 23] Brynjolfsson, E., Rock, D. and Syverson, C. (2021) 'The Productivity J-Curve: How Intangibles Complement General Purpose Technologies', *American Economic Journal: Macroeconomics*, 13(1), pp. 333–372.

[6, 7, 9, 12, 13, 14, 15, 16, 29, 32, 42] Borri, N., Liu, Y. and Tsyvinski, A. (2026) *AI Premium*. NBER Working Paper 35451. Cambridge, MA: National Bureau of Economic Research.

[8] Pástor, Ľ. and Veronesi, P. (2009) 'Technological Revolutions and Stock Prices', *American Economic Review*, 99(4), pp. 1451–1483.

[10, 19] Borri, N., Liu, Y. and Tsyvinski, A. (2026) 'The AI Premium', *VoxEU*, Centre for Economic Policy Research, 29 July.

[11] OpenRouter (2026) *Models: One API for Hundreds of Models*. OpenRouter Developer Documentation.

[18, 27, 31, 40] Dillon, E.W., Jaffe, S., Immorlica, N. and Stanton, C.T. (2025) *Shifting Work Patterns with Generative AI*. NBER Working Paper 33795. Cambridge, MA: National Bureau of Economic Research.

[21, 37] OECD, Boston Consulting Group and INSEAD (2025) *The Adoption of Artificial Intelligence in Firms: New Evidence for Policymaking*. Paris: OECD Publishing.

[24] Brynjolfsson, E. and Hitt, L.M. (2000) 'Beyond Computation: Information Technology, Organizational Transformation and Business Performance', *Journal of Economic Perspectives*, 14(4), pp. 23–48.

[25, 34, 38] Brynjolfsson, E., Li, D. and Raymond, L.R. (2025) 'Generative AI at Work', *Quarterly Journal of Economics*, 140(2), pp. 889–942.

[26] Becker, J., Rush, N., Barnes, B. and Rein, D. (2025) *Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity*. METR.

[28, 35, 41] Dell'Acqua, F., Ayoubi, C., Lifshitz, H., Sadun, R., Mollick, E., Mollick, L., Han, Y., Goldman, J., Nair, H., Taub, S. and Lakhani, K.R. (2025) *The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise*. NBER Working Paper 33641. Cambridge, MA: National Bureau of Economic Research.

[30] Anthropic (2025) *Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption*. Anthropic, September.

[33] Acemoglu, D. and Restrepo, P. (2019) 'Automation and New Tasks: How Technology Displaces and Reinstates Labor', *Journal of Economic Perspectives*, 33(2), pp. 3–30.

[39] Dell'Acqua, F., McFowland III, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K.C., Rajendran, S., Krayer, L., Candelon, F. and Lakhani, K.R. (2023) *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality.* Harvard Business School Working Paper 24-013.