

Value-Maxxing and the New Economics of AI Labor

The Economy Research Editorial^{1,2}

¹The Economy Research, 71 Lower Baggot Street, Dublin 2, Co. Dublin, D02 P593, Ireland

²Swiss Institute of Artificial Intelligence, Chaltenbodenstrasse 26, 8834 Schindellegi, Schwyz, Switzerland

Abstract

This article argues that the current AI debate is framing the question wrong. Far too much of it is focused on token volume, content generation and its adoption, when perhaps the more relevant topic is whether or not AI provides a tangible net value to the economy. Essentially, this article reinterprets AI as a new factor of production which is valuable in proportion to the extent to which firms optimize it by integrating tokenized computing power, human evaluation, workflow architecture and desired friction. Its primary policy recommendation is not that AI will render workers jobless causing mass unemployment but that it would it revalue and reconfigure work and the associated labor value unevenly. In this approach, firms, nations and individuals that lack equal access to AI skills, computing infrastructure and institutions would suffer from widening inequality.

1 Introduction - Introduction - From AI Usage to Value-Maxxing

The lexicon of "maxxing" has crossed over from online slang into the center of AI discussion because it captures an important organizational shift, where not just firms, workers and platform vendors are merely "using" generative AI, but actively trying to optimize around it. Merriam-Webster now defines "-maxxing" as "to make as effective and fruitful as possible,"^[1] while Brookings, in its recent contribution on "context-maxxing," contends that the dominant discourse of AI is caught in a bind between "using AI more" and "producing more AI,"^[2] missing the bigger questions around "is AI working in a way that sustains human cognitive agency?". At the same time, the business community is now bifurcating "token-maxxing" from "outcome-maxxing,"^[3] while a counter trend towards "friction maxxing" has started gaining traction^[4] as a defense of strategic opposition to smooth automation. None of this is purely semantic refinement: these phenomena are symptomatic of a deeper conflict about what exactly firms are optimizing when they incorporate AI into production, coordination and decision-making processes.

Too often, conventional framing misses the mark because it conflates proxy indicators with economic objectives. Output is not an objective; token-use is not an objective; even context control, as important as it is, is merely an indicator and not an end-in-itself. From an economic perspective, the goal is net value creation - a greater capacity to produce usable and profitable output at an acceptable trade-off of labor cost, token-use and, what is increasingly pertinent, compliance burden and error rate. It is for this reason that a more cohesive organising concept is not token-maxxing, content-maxxing or even context-maxxing but value-maxxing, an umbrella term coined in this report to frame the process of optimizing AI to maximize its economic value, rather than its "throughput". Token-maxxing and context-maxxing cease to stand opposed but become complementary means to higher value per token, worker, or unit of time.

This repatterning has become all the more urgent as generative AI is no longer a novel laboratory artifact but a structuring force on the labor market, which it has become between 2023 and 2026. The 2025 AI Index report from Stanford has stated that 78% of organizations were already using AI in 2024 (compared to 55% in 2023)^[5] while generative AI alone was projected to attract US\$33.9 billion of global private investment in 2024-an 18.7% increase from 2023.^[6] Official data from OECD now indicates that the rate of firm-level adoption of AI across member countries jumped from 8.7% in 2023 to 20.2% in 2025,^[7] much higher for larger companies than small ones, and according to World Bank data, job vacancies mentioning generative AI jumped by ninefold from 2021 to 2024,^[8] even if many of these jobs remained situated in high-income countries and amongst workers with better digital skills. AI's influence is beyond dispute; the question is rather what kinds of human skills, organizational models and positions of power it is rewarding.

The emerging narrative is not one of simple labor-saving automation but of a much more complex and politically significant reordering of work. Analysis from the OECD suggests that nearly a quarter of all workers across member countries already work in roles exposed to generative AI^[9] meaning that at least 20% of their job tasks could be significantly expedited or completed at least 50% faster. The ILO's updated global study confirms that 25% of the world's workforce now holds jobs with some exposure to generative AI,^[10] with administrative workers continuing to hold the largest share and female workers in high-income countries at

particular risk. IMF's broader measure of AI-exposed jobs finds that nearly 40% of global employment and almost 60% of employment in developed economies will be impacted.^[11] In light of these figures, the primary policy question transcends simply whether AI can boost productivity at the task-level; rather it is about whether the organizations integrating AI preserve judgment, sustain skill development and diffuse gains sufficiently broadly to prevent a deeper stratification of firms, jobs and nations.

The central thesis is simple but consequential: AI should be treated as a new factor of production whose value depends on how well it is integrated with labor, organizational context and friction. Low-context task optimization, for instance, relies primarily on strict adherence to tokens and standardized workflows. Higher-context tasks rely more heavily on human intelligence, namely, specification, verification and judgment. In all cases, the relevant economic logic remains: immediate gains in speed or output volume are not meaningful unless coupled with appropriate cost management across labor inputs, token costs, governance burden and error potential. This is the underlying economic rationale for why certain forms of friction might become a source of value, why AI-native labor can claim rising wage premiums, and why, in the context of usage-priced tokenized compute, there is greater incentive to invest in more selective and smaller labor forces. The rise of AI does not signal the disappearance of labor, but its repricing, re-ranking, and reorganization.

2 Value-Maxxing as the Economic Objective of AI Adoption

The context-maxxing frame usefully corrects AI strategies that focus on usage quotas or output volume by shifting attention away from model consumption itself and toward the controlled information environment in which human and machine cognition interact. In the context of context-maxxing, the user relies upon user-controlled infrastructure and three intertwined skills: specification, orchestration, and exploration, to increase their cognitive efficiency and command over the task. Nevertheless, this advanced notion of context-maxxing remains incomplete as a political-economic framework because it still leaves control as the main yardstick of success. The degree of control is only a means, not an end, of the production process, in that it is meaningful only to the degree that it lowers costs of error, repetition and dependency. What truly matters in the economic sense is not the volume of control but whether such control has improved the ratio between output value and cost incurred.

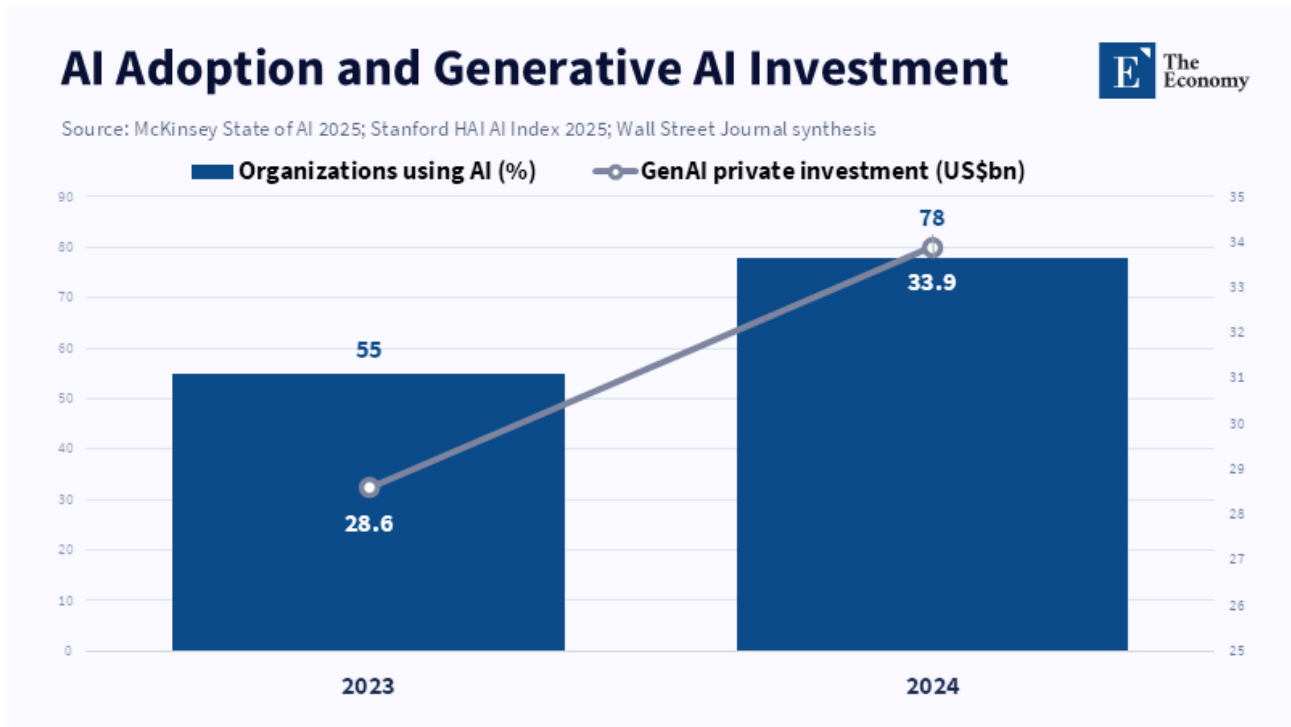


Figure 1: AI has moved from experimentation to economic infrastructure; the policy question is whether adoption converts into measurable value.

It follows that token-maxxing and content-maxxing are essentially two faces of the same coin, as content is nothing other than tokens expressed as words, images or code. A document rapidly produced by an AI could be worthless if it does not achieve a sales target, does not improve a decision-making process, does not generate an acceptable recommendation or does not contribute to enduring organizational capability. Conversely, a more costly or slower workflow producing outputs of higher value, due to reduced hallucinations, hidden limitations, or preserved professional judgment in high-stake decisions, offers a greater economic return. The appropriate objective function, in simple economic terms, therefore defines value-maxxing as the optimization over the net value added, considering labor inputs, token costs, compliance costs, organizational and coordination costs and estimated risk costs. Under this framework, debates around whether one should maximize tokens or output amounts are analytically insufficient and conceptually incomplete, as these are merely factors contributing to a larger optimization problem.

This notion becomes especially clear when the evidence on productivity is read through the article’s central argument. The finest existing studies indeed report significant improvements in task-level outputs, where for instance Noy and Zhang reported that ChatGPT assistance decreased time spent on intermediary professional writing by 40% and simultaneously enhanced output quality by 18%;^[12] and Brynjolfsson, Li and Raymond found a similar 14% overall improvement in customer support productivity^[13] (with novice and lower skilled employees seeing 34% gains)^[14] with simultaneous improvements in customer satisfaction and loyalty. Such effects are considerable and cannot be ignored in policy discussion as they signal clear gains at the level of the production frontier for certain work tasks.

Nevertheless, such results do not offer grounds for simplistic productivity narratives. Indeed, the same studies also highlight how the value produced by AI is heavily influenced by task characteristics, worker inputs and level

of oversight. As evidenced in a randomized field experiment with 758 consultants on "the jagged technological frontier," assistance from GPT-4 increased the speed, volume, and quality of work by 12.2%, 25.1% and more than 40%, respectively on tasks falling within the frontier. However, if asked to tackle an issue outside the frontier, consultants provided correct answers 19 percentage points less frequently.^[15] Thus, AI enhances or hinders work in the abstract but creates value only when correctly paired with tasks and adequately regulated by human oversight; it is precisely these conditionalities that value-maxxing seeks to render central.

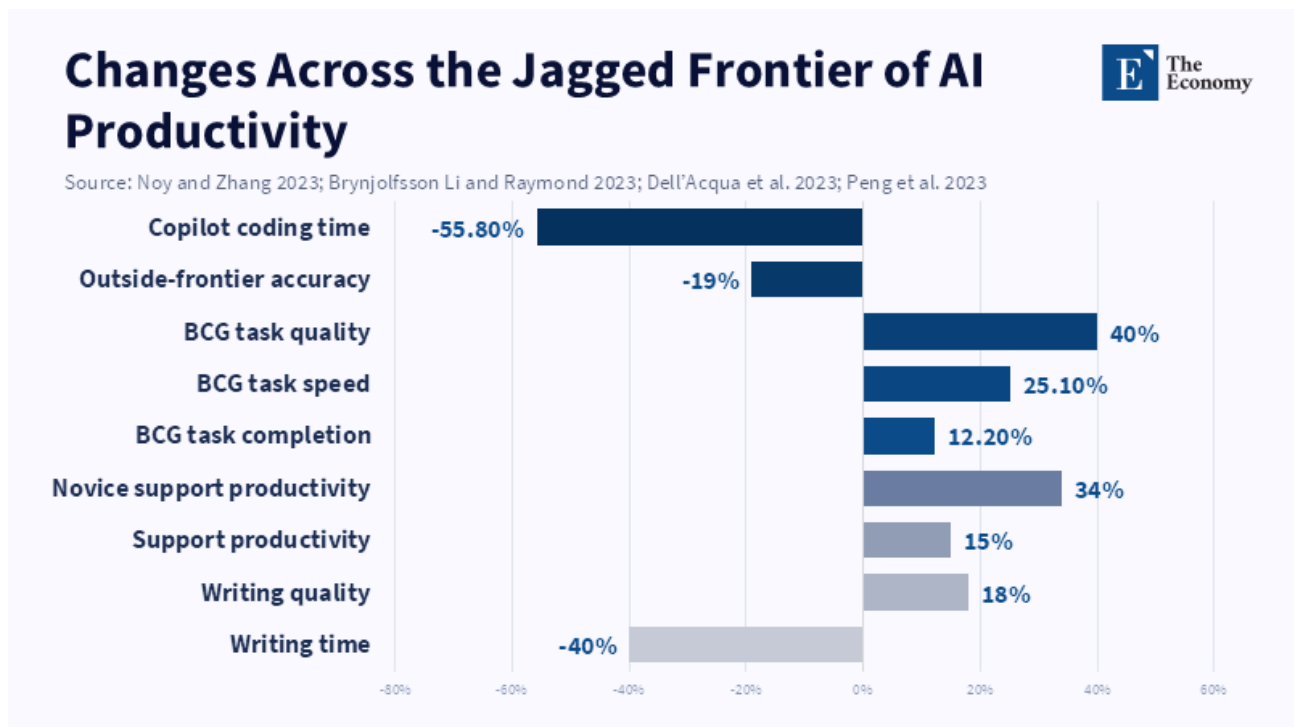


Figure 2: AI produces large gains inside its performance frontier, but the same system can reduce accuracy when applied beyond it.

This helps clarify why "friction maxxng" should not be viewed as irrational or anti-technical. Friction represents the deliberate insertion of pauses, verifications, explanatory text and individual judgment into an AI-assisted process. These frictions prevent an artificial intelligence system from compromising the value of its output. Studies examining the link between AI-use confidence, worker self-confidence and critical thinking levels report that higher user trust in GenAI leads to lower levels of critical evaluation^[16] while higher user self-confidence in the system boosts it; other works examine how automation bias and the human tendency toward over-reliance on automated systems can lead workers to accept inaccurate AI-driven recommendations.^[17] In these scenarios, frictions are not inefficiency but a control mechanism to preserve the value of final outputs by maintaining human evaluative control.

This also explains why some organizations might appear maximally efficient on a given monthly dashboard but may actually be destroying value long term. When routine drafts, triage or analysis are automated away and replaced without retaining moments for human understanding or validation, some short-term gains come at the cost of longer-term stocks of judgment. Brookings' critique of "usage maxxng" becomes critical in this context as an oversimplified view of AI-assisted productivity risks undermining human cognition through its reliance on seemingly productive but shallow interaction with systems. The concept of "mechanized convergence" discussed

by Microsoft researchers points to how unmediated use of the same AI system may decrease the diversity of output^[18] in contexts where differentiated judgment is precisely the value driver. Homogeneity of output in such instances cannot be an acceptable side-effect; it is an economic cost.

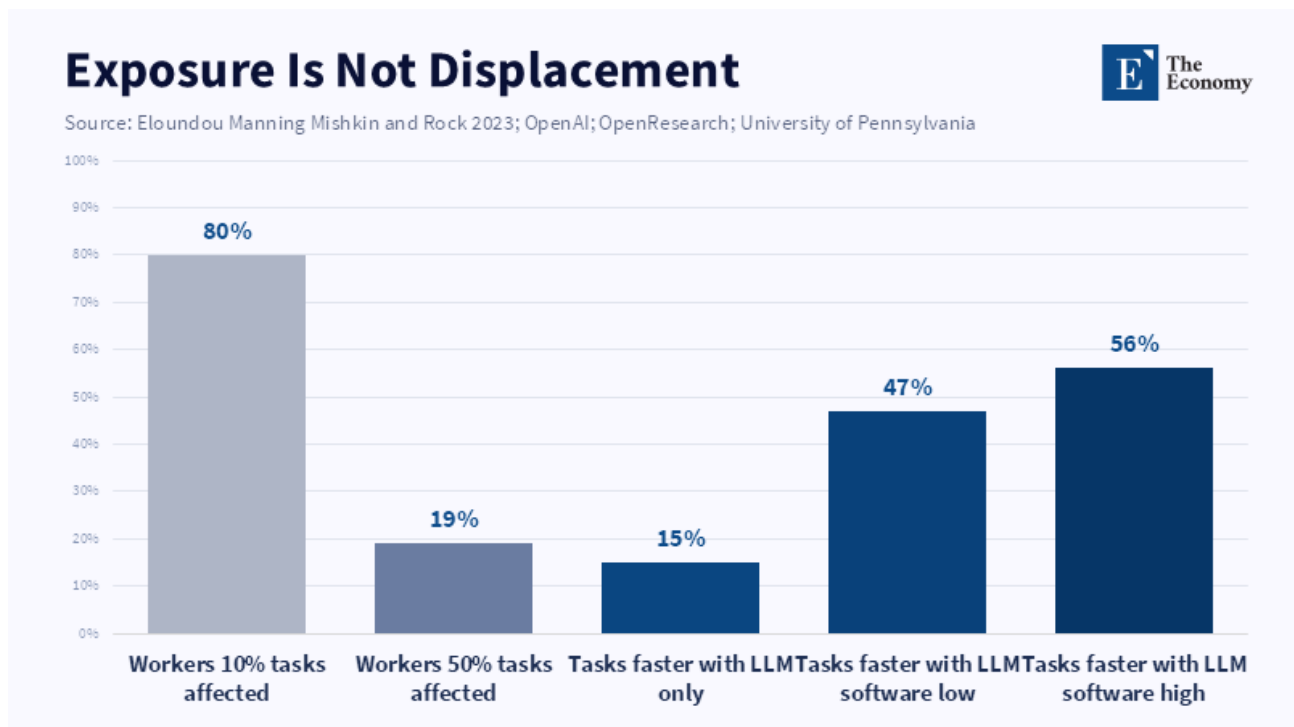


Figure 3: Task exposure shows where AI can reshape work, but exposure should not be mistaken for automatic job loss.

The consequence for practice is that organizations can no longer afford to monitor their AI usage via superficial metrics such as installed seats, number of prompts generated or pages-written, or solely the speed of task completion. Instead, companies must monitor the downstream impact, rework rates, factual accuracy of output, times-to-validated completion, revenue or decision quality, litigation events and changes in the capacity of workers to perform tasks without AI assistance. Value-maxxing is therefore not merely a buzzword; it is a managerial imperative and a recovery of a fundamental economic principle: efficiency gains must be judged by the value created.

3 Human Customization and the Rise of AI-Native Labor

Accepting the AI model as a generalized production factor rather than an all-knowing oracle brings human customization to the foreground. AI should not be compared with human intelligence but it would be more reasonable to consider it a device for the expansion of human intellect like data bases, spreadsheet or search engine. The Brookings notion of specification, orchestration and exploration is useful here because it emphasizes how humans can usefully enhance AI’s output value by mastering the definition of the problem, the sequence and execution of tasks, and the redeployment of AI-facilitated efficiencies towards tasks requiring higher levels of human expertise.

The empirical evidence also supports the argument that quality of work varies widely by user. In customer support, novice employees saw much greater increases in productivity gains than expert workers, as they were

more able to leverage the tacit knowledge captured in an AI's underlying model. In the BCG field experiment, although lower performing employees also saw benefits in the range of 12.2–25.1% and 40% improvement in speed and quality, they remained subject to dramatic declines when the task was deemed outside the system's performance frontier. Researchers at MIT, investigating random assignments of Copilot to 1,974 developers, reported similar gains in pull-request completions ^[19] while also noting significant heterogeneities tied to work and organization design and system use. Such evidence does not undermine the role of individual expertise or the existence of a skills gap. Quite the opposite, it redefines the locus of expertise from content creation toward its specification, management and evaluation.

It is this concept of worker input that underpins what might be termed "superhuman labor." Superhuman workers are not those who excel beyond human capacity at any task, but those who are best able to scale machine generation of inputs and outputs without ceding the definition of the problem, the validation of results or the responsibility for the outcome. OECD data on demand for skills shows this clearly: skills demanded for tasks with the highest levels of AI exposure, across ten European nations, include not just AI-specific know-how but also business-process management, personal and social-interpersonal and generalist cognitive skills. 72% of jobs with high levels of AI exposure demanded at least one management skill,^[20] whereas 67% included at least one business process skill.^[21] The scarcity in labor markets is not that of the best prompt engineers but of those who can embed the insights of the AI system into organizational decision-making.

This explains the rising demand and wage premiums for AI-native workers. The 2025 PwC Global AI Jobs Barometer found the average wage premium for AI-related skills to be 56%^[22] and notes that job postings demanding these skills now change at rates 66% higher in high-exposure professions compared to low-exposure ones.^[23] This aligns with LinkedIn's Work Change report, which noted the number of job postings including AI literacy as a desired skill has sextupled from the prior year,^[24] and AI-based hiring continues to grow at rates substantially higher than general employment. This global trend of expanding demand for specialized AI labor is also corroborated by World Bank data, which shows that generative AI skill-demanding job postings quadrupled in only three years and remain concentrated in high-income economies.^[25] This signals not merely 'conversational literacy' with new AI-driven tools, but 'native fluency' across the AI stack.

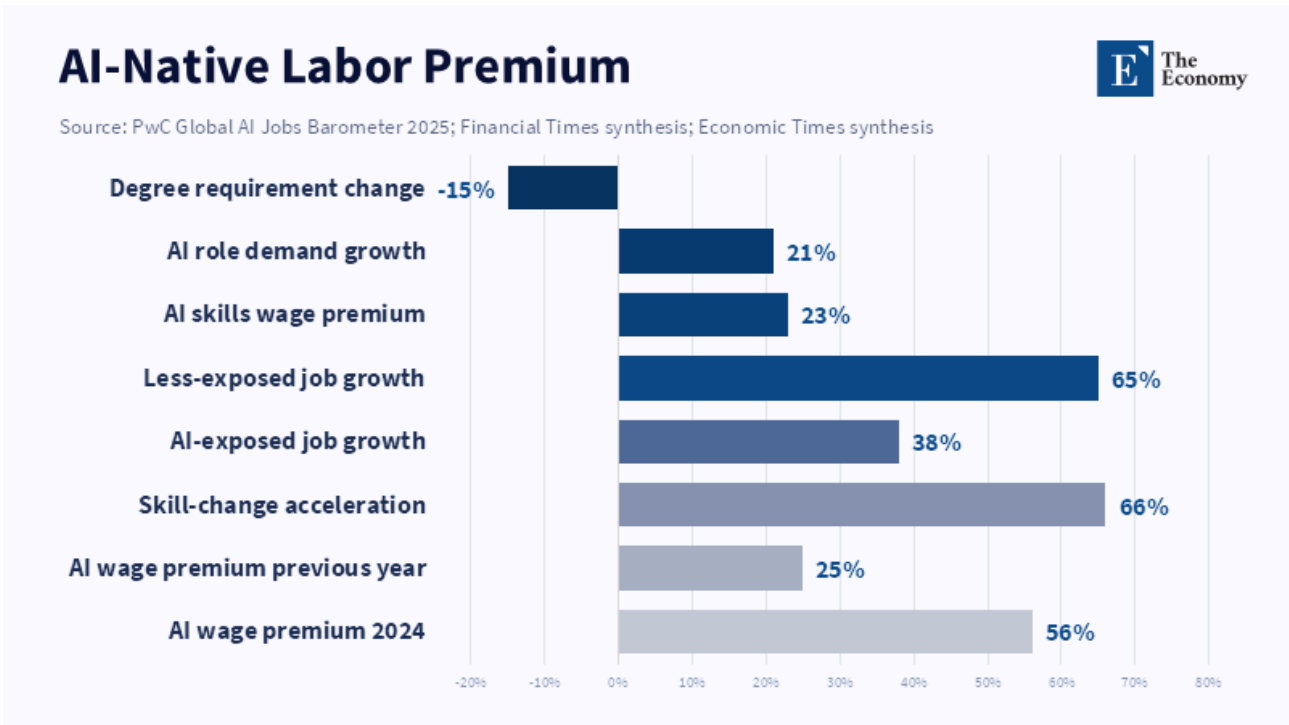


Figure 4: The labor-market premium is shifting toward workers who combine AI fluency with judgment, verification and organizational execution.

Analyzing tasks where human customization contributes to value, versus where it plays less critical role, is thus key to differentiating the use cases. In highly routinized tasks such as summarization of routine documents, internal search categorization, and templated customer service, value-maxxing primarily translates to optimal machine-use, strict scope limitation, and task-specific quality control, whereas human customization will focus more on the workflow’s architectural design rather than the creative production process itself. Conversely, in high-variance tasks characterized by regulatory considerations, tacit domain-specific knowledge, strategic uncertainty, or broad societal impact (such as consulting, law, finance, executive management or academic research), AI performance depends largely on human intervention to select relevant information, identify potential errors and frame outputs in relation to wider objectives. This is where AI-native labor plays a far more decisive role in producing surplus.

Critically, these distinct patterns both fall under the same general optimization logic: whether one is using AI on highly customized or highly standardized tasks, value-maxxing must also serve the goal of cost-minimization relative to output. The distinction is simply in the trade-off between labor input quality and token consumption. For low-variance tasks, the ideal is a minimization of human input coupled with sufficient human oversight. For high-variance tasks, investing in human input that significantly increases the output value of token inputs is the economically efficient approach. It is for this reason that the proliferation of AI need not flatten labor markets; indeed, it can exacerbate already rising inequality, offering higher returns for a narrow class of AI-native workers whose judgment adds value to computation.

This exacerbation is already compounding existing societal inequalities: OECD data indicates that while AI-related displacement risks affect low-skilled and older workers directly, the risks for others, including women, are more likely to take the form of lower opportunities and less productive work through limited access to AI-

enhancing tools and capabilities. Indeed, ILO’s latest estimates indicate that 3.3% of global jobs are at highest risk of augmentation, and women comprise a larger share of workers in this category in high-income countries (9.6%) versus men (3.5%),^[26] and that only 3.5% of men’s jobs are in high-risk sectors. Similarly, World Bank research shows high-income countries already capture over 70% of AI-skilled job vacancies while basic digital skills remain below 5% of the population in low-income nations and 66% in high-income countries.^[27] The emergent landscape is thus one where AI-native labor is developing in a deeply unequal global infrastructure.




Figure 5: AI exposure is widespread, but its risks and rewards are unevenly distributed across economies, occupations and gender groups.

Small firms represent a case in point. Generative AI is already used by 31% of OECD SMEs,^[28] and two-thirds of those SMEs report that they have experienced an increase in their employees’ productivity; yet 83% saw no change in total staffing^[29] while more than twice as many report that AI increases employee skill requirements as say that it does the opposite. It seems clear, therefore, that SMEs are not moving towards jobless production through automation, but reorganizing work in favor of more specialized human capital and a higher skill-effort ratio. The lesson for the workforce educators and the universities and training bodies is crystal clear. A curriculum in AI literacy is not simply a set of tools and features; rather, it needs to link subject matter knowledge with task specification, evidence verification and process orchestration. Administrators, on the other hand, should create protocols differentiating between exploratory and operational uses of AI. And policymakers, above all, must work to ensure that AI-native work does not devolve into an elite, access-limited stratum of the labor market.




Figure 6: For lower- and middle-income regions, the key constraint is whether AI exposure becomes augmentation or automation.

4 Tokenized Capital and the Cost Logic of AI Production

The next step concerns what the firm is actually purchasing when it acquires AI capacity. As a practical matter in the short run, firms don't actually buy intelligence as fixed capital. Instead, firms are renting tokenized access to models, retrieval, search, and tool invocations hosted in distant data centers. As a consequence, the manager's relevant cost input is the marginal cost of token consumption and other usage, rather than an underlying data center amortization schedule. But the market price of these tokens is itself the market expression of hardware and infrastructure investment on a grand scale. Estimates suggest it costs approximately \$78 million to train GPT-4, and \$191 million for Gemini Ultra;^[30] meanwhile the 2025 AI Index demonstrates a more than 280-fold decrease in training costs and, by Stanford's accounting, 30% annual reduction in hardware costs and 40% annual increase in energy efficiency.^[31] The interplay of collapsing costs and the immense capital expenditure required therefore do not conflict, but represent two dimensions of the same system: massive frontier fixed investments, which filter down as non-zero, dynamically changing marginal prices to users.

Official pricing guides illustrate the point. OpenAI charges \$5 per 1 million input tokens, and \$30 per million output tokens.^[32] Prompts longer than 272,000 tokens cost twice the input rate, and 1.5 times the output rate for the entire exchange; they also charge for separate tools such as web search. Anthropic prices its Claude Opus 4.7 at \$5/million input tokens and \$25/million output tokens,^[33] with extra charges for cache writes and reads. Gemini Developer API prices its usage at token-level,^[34] distinguishes between standard and batch processing, and charges for context caching and storage as well as for search grounding beyond a free allowance. The pricing model has shifted from subscription to usage. For high-volume enterprises, token-maxxing is a budget constraint because, even though a single query might not be prohibitively expensive, the scale of usage (millions

of queries, large contexts, frequent searches and calls) converts AI into a variable cost like electricity, one that requires active management.




Figure 7: Token pricing turns AI from a fixed software subscription into a variable production cost that firms must actively govern.

This hardware and infrastructure investment becomes clear at the level of individual firms. Alphabet’s 2025 cash flow statement reveals over \$91 billion in purchases of property and equipment.^[35] Meta reported \$72.2 billion in 2025 capital expenditures,^[36] and projected it could go higher in 2026. Microsoft’s 2025 annual report indicates an increase of \$20.1 billion in property and equipment,^[37] with continued massive capital expenditures projected due to cloud services and AI infrastructure. The capital required is too substantial to be treated as incidental; it places AI as a sector more akin to energy, telecommunications or heavy manufacturing than ordinary SaaS, and token usage is the retail form through which firms rent access to capital-intensive infrastructure.




Figure 8: Behind every low-friction AI query sits a capital-intensive infrastructure system of data centers, chips, energy and cloud capacity.

This distinction should be interpreted economically rather than literally. While not every token-priced generated output is more costly than every human-produced unit of output (for mundane tasks, this remains false), the shadow price of compute for frontier-grade and high-context applications is still high enough to govern business practices. Long contexts, data residency requirements, tools, search augmentation, and repeated agentic feedback loops drive usage costs higher than the displayed unit price for a short query; premier models cost more than the low-end ones. Therefore, the manager’s problem has shifted away from just adopting AI to designing production processes: deciding which tasks warrant frontier tokens, how best to handle others with less expensive models, and where human review is less costly than additional machine exploration.

This problem is fundamentally one of optimization as typically taught in microeconomics. A firm minimizing costs for a fixed level of output will equalize the marginal product per dollar spent across inputs: $\frac{MP_L}{P_L} = \frac{MP_T}{P_T}$. This means that firms do not over- or under-buy AI merely because it is novel, nor additional labor merely because it is familiar, but adjust the quantities of both until marginal return per dollar is equivalent. The concepts of value-maxxing, context-maxxing and token-maxxing are, in this usage-based AI environment, informal expressions of this microeconomic optimization problem. Value-maxxing represents the ultimate objective function, while the others are often intermediate goals.

This also highlights the labor market implications. In a usage-priced system with output highly dependent on the quality of user input, firms will favor hiring fewer, higher-quality workers whose domain expertise, context comprehension, and AI fluency command a high wage and justify expensive AI computation attached to their work. The PwC labor market data support this conclusion. While demand for AI-exposed occupations grew slower (38%) than for others (65%) between 2019 and 2024,^[38] wage growth was twice as rapid for highly AI-exposed occupations as for those least-exposed, and workers with AI skills received 56% higher compensation.^[39]

Jobs may not disappear outright but are instead revalued and sorted according to the ability to generate significant value under AI-intensive conditions.

The same economic reasoning explains why short-term returns to labor quality may increase even as AI capabilities advance. If frontier compute is expensive, human-input quality is amplified, as workers who demonstrate better context, better judgment on stopping points, or recognition of when AI has exceeded its capabilities waste less computation, guide the AI more efficiently, and manage the overall productivity. Labor shifts from mere task execution to "token stewardship" that justifies high compensation, especially in complex fields like law, deals, consulting, regulation, or engineering, where a single bad output can cascade into significant waste or liability. In these situations, "super human labor" is not a romantic ideal but rather labor with judgment so sharp that it increases overall value-per-token.

These trends are likely to play out unevenly across the economy. While OECD nations have more than doubled AI use since 2023, adoption remains high among large firms (52.0%) and lower among SMEs (17.4%).^[40] The latter cite concerns about legal uncertainty, data usage, lack of expertise, and incompatibility as obstacles to adoption and often have not yet invested sufficiently in staff training or company-specific AI policies. Value-maxxing thus also becomes an issue of competitive dynamics. Large companies have the ability to experiment with diverse models, withstand initial cost-inefficiency, negotiate better terms with vendors, and hire specialized AI talent; SMEs typically do not. Without policy intervention to mitigate this asymmetry, AI will continue to operate as another vector of inequality and market-power concentration.

5 Conclusion - $MC_L = MC_T$

Token-maxxing, content-maxxing, context-maxxing, and friction-maxxing, in essence, should be seen as alternative descriptions of a single economic issue in the new technical conditions. What has changed fundamentally is the form of capital. Traditional economics considers the allocation of inputs to be efficient when it's based on the relative marginal returns of labor and capital. In the AI economy, a growing share of capital comes to the firm in the form of tokenized access to models, contexts, tools, cached memory, search, and agentic workflows. Once this is understood, value-maxxing is no longer a fad but the return of economic fundamentals: using AI matters if it creates more value than cost, risk and institutional consequences.

The strongest counterargument is that generative AI could boost productivity, alleviate skill bottlenecks and democratize access to formerly scarce capabilities. To a degree, this is true. Evidence from writing, customer service, software development, and SMEs suggests that AI is increasing productivity in these environments. However, such increases come with preconditions and trade-offs. Unlike a calculator or search engine, generative AI does more than compute or recall. It creates, synthesizes, recommends, and, therefore, actively constructs the cognitive environment. Its benefits, in short, are balanced by a number of new risks.

Ultimately, those firms that will gain the most value from generative AI are those that manage the interplay of labor, compute, context, and friction in such a way as to create demonstrable value, without undermining the human capabilities that stable productivity relies on. In the immediate term, this suggests favoring labor that has a good sense of judgment, verification, and orchestration skills while tokenized compute continues to

represent an important and costly capital input. In the longer run, it suggests an understanding of AI as a structure, a force that can either expand or narrow economic opportunities, based on the governance choices made by firms, institutions and policymakers.

References

- [1] Merriam-Webster. “-maxxing.” *Merriam-Webster Slang & Trending*.
- [2] Taylor, J. and Krishna, K. (2026) “Context-maxxing: A path to cognitive agency with generative AI.” *Brookings Institution*.
- [3] Sanchez, C. (2026) “Token Maxxing vs Outcome Maxxing Is the Wrong Question.” *ChristopherSanchez.ai*.
- [4] Mind the Product. “Why Friction Maxxing changes product strategy in 2026.”
- [5] McKinsey & Company. *The State of AI: Global Survey 2025*.
- [6] Stanford Institute for Human-Centered Artificial Intelligence. *Artificial Intelligence Index Report 2025*.
- [7] OECD. *AI Use by Individuals and Firms: ICT Access and Usage by Businesses Dataset*. 2025.
- [8] World Bank. *Generative AI, Digital Skills and Labour-Market Demand*. 2024.
- [9] OECD. *Job Creation and Local Economic Development 2024*.
- [10] International Labour Organization. *Generative AI and Jobs: A Refined Global Index of Occupational Exposure*. 2025.
- [11] International Monetary Fund. *Gen-AI: Artificial Intelligence and the Future of Work*. 2024.
- [12] Noy, S. and Zhang, W. (2023) “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.” *Science*.
- [13] Brynjolfsson, E., Li, D. and Raymond, L. (2023) “Generative AI at Work.” *NBER Working Paper*.
- [14] Brynjolfsson, E., Li, D. and Raymond, L. (2023) “Generative AI at Work.” *NBER Working Paper*.
- [15] Dell’Acqua, F. et al. (2023) “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality.” *Harvard Business School / Boston Consulting Group*.
- [16] Drosos, I. et al. (2025) “It Makes You Think: Provocations Help Restore Critical Thinking in AI-Assisted Work.” *Microsoft Research / CHI*.
- [17] Parasuraman, R. and Riley, V. (1997) “Humans and Automation: Use, Misuse, Disuse, Abuse.” *Human Factors*.
- [18] Drosos, I. et al. (2025) “It Makes You Think: Provocations Help Restore Critical Thinking in AI-Assisted Work.” *Microsoft Research / CHI*.

- [19] Peng, S., Kalliamvakou, E., Cihon, P. and Demirer, M. (2023) “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot.”
- [20] OECD. *Artificial Intelligence and the Changing Demand for Skills in the Labour Market*. 2024.
- [21] OECD. *Artificial Intelligence and the Changing Demand for Skills in the Labour Market*. 2024.
- [22, 23, 38, 39] PwC. *Global AI Jobs Barometer 2025*.
- [24] LinkedIn Economic Graph. *Work Change Report 2025*.
- [25] World Bank. *Generative AI, Skills Demand and Global Job-Posting Trends*. 2024.
- [26] International Labour Organization. *Generative AI and Jobs: A Refined Global Index of Occupational Exposure*. 2025.
- [27] World Bank. *Digital Progress and Trends Report*. 2024.
- [28, 29] OECD. *Generative AI and the SME Workforce*. 2025.
- [30, 31] Stanford Institute for Human-Centered Artificial Intelligence. *Artificial Intelligence Index Report 2024*.
- [32] OpenAI. *API Pricing Documentation*, accessed 14 May 2026.
- [33] Anthropic. *Claude API Pricing Documentation*, accessed 14 May 2026.
- [34] Google AI for Developers. *Gemini API Pricing Documentation*, accessed 14 May 2026.
- [35] Alphabet. *Annual Financial Statements and AI Infrastructure Capital Expenditure Guidance, 2025–2026*.
- [36] Meta. *Capital Expenditure Results and AI Infrastructure Guidance, 2025–2026*.
- [37] Microsoft. *AI-Enabled Data-Centre Investment Plan, FY2025*.
- [38] PwC. *Global AI Jobs Barometer 2025*.
- [39] PwC. *Global AI Jobs Barometer 2025*.
- [40] OECD. *AI Use by Individuals and Firms: ICT Access and Usage by Businesses Dataset*. 2025.