# Generative AI and Organizational Structure in the Knowledge Economy

**(Authors' names blinded for peer review)**

Generative AI (GenAI) is rapidly transforming knowledge work, yet its implications for organizational hierarchies remain poorly understood. Unlike earlier automation technologies, GenAI can both perform tasks autonomously and assist human workers, while its intrinsic fallibility, the tendency to produce confident but incorrect outputs, demands continuous human oversight. We develop a theoretical model to study how GenAI reshapes workforce composition and organizational structure in knowledge-based hierarchies. Our analysis highlights two deployment dimensions, namely mode (automation vs. augmentation) and location (worker vs. expert layer), which generate a $2 \times 2$ design space whose organizational implications are not predicted by traditional technology adoption theories. We obtain three main findings. First, GenAI's effect on entry-level skill requirements is critically mode-dependent. Worker-level automation leads firms to hire fewer but more skilled workers who validate AI outputs and limit costly escalation to experts. Worker-level augmentation, by contrast, expands workers' effective capability, allowing firms to relax entry-level knowledge requirements while sustaining performance. The decline in junior employment documented in recent studies therefore reflects deployment choices favoring automation over augmentation, not an inevitable consequence of GenAI itself. Second, expert-level deployment uniformly lowers entry-level skill requirements, regardless of whether GenAI automates or augments. By expanding experts' capacity to support downstream workers, it enables organizations to employ a broader base of less specialized workers, thereby broadening entry-level access to knowledge work. Third, organizational structure evolves non-monotonically as GenAI improves: across all four deployment architectures, the span of control initially contracts before eventually expanding. Demand for senior expertise may therefore remain stable or even increase during early-to-intermediate stages before hierarchies ultimately flatten. Together, these results demonstrate that GenAI's organizational consequences depend on deployment design rather than adoption intensity alone, and that the same technology can upskill or deskill the workforce depending on how and where it is deployed.

*Key words*: Generative AI, knowledge economy, organizational structure, workforce, knowledge-based hierarchy, intrinsic fallibility, deployment flexibility

## 1. Introduction

The knowledge economy, in which expertise serves as a key economic resource, has long relied on organizational hierarchies to efficiently allocate specialized knowledge (Powell and Snellman 2004). A canonical form is the "knowledge-based hierarchy," in which lower-level workers handle routine problems while escalating more complex cases to higher layers of expertise, such as those in analyst–manager, attorney–partner, and software engineer–tech lead structures. When workers encounter problems beyond their capacity, they turn to experts at higher levels for guidance (Garicano 2000). The advent of Generative AI (GenAI), however, is fundamentally challenging this established order by enabling machines to perform complex knowledge-based tasks through natural language

interactions, potentially transforming how organizations structure their hierarchies and deploy their workforce.

Emerging deployment patterns across industries reveal systematic variation in how GenAI is integrated into daily workflows. In software development, AI coding agents such as Claude Code and Cursor autonomously generate code, shifting the developer's role from execution toward validation and integration. In creative fields, tools like Adobe Firefly and Canva's AI features expand users' capabilities, allowing non-designers to produce content that previously required specialized expertise. GenAI is also transforming expert-level roles. In professional services, AI platforms like McKinsey's Lilli function as virtual knowledge repositories, providing direct answers to junior employees' queries and reducing their routine reliance on senior staff. Meanwhile, organizations are deploying AI-powered synthesis tools to assist senior consultants and partners directly, enabling them to rapidly retrieve firm expertise and respond to complex escalations more efficiently. This heterogeneity reveals that GenAI is not a uniform technology; rather, it is being integrated in fundamentally different ways that vary both in how it interacts with human ability and where it is positioned within the organizational hierarchy.

Despite widespread adoption and intense public interest, we lack rigorous theoretical frameworks for understanding how GenAI reshapes organizational hierarchies and the distribution of skills within firms. Recent empirical evidence reveals striking workforce composition effects: early-career employment has declined in the most AI-exposed occupations, while senior employment has remained stable or even risen (Brynjolfsson et al. 2025a, Hosseini and Lichtinger 2025). Furthermore, these entry-level declines are concentrated in settings where GenAI primarily automates tasks and are much more muted where GenAI augments human work (Brynjolfsson et al. 2025a). Yet existing organizational theories cannot explain why the same technology produces simultaneous junior displacement and senior stability, nor why these effects vary systematically across deployment contexts. Without clear conceptual foundations, it remains impossible to predict whether a given GenAI deployment will raise or lower skill requirements, steepen or flatten hierarchies, or displace or preserve senior roles.

Two features distinguish GenAI from prior workplace technologies and drive the organizational dynamics we characterize. The first is *intrinsic fallibility*: regardless of how GenAI is deployed, it produces probabilistically incorrect outputs that appear authoritative—a phenomenon known as *hallucination*—thereby creating a novel organizational burden of continuous human validation (Bubeck et al. 2023, Zhou et al. 2024). Unlike rule-based automation, which either produces a correct output or fails transparently, or earlier AI systems that attach explicit uncertainty estimates to their recommendations, GenAI's errors are indistinguishable from correct outputs without domain expertise, making validation central to organizational design. Critically, validation demands

domain knowledge: workers can verify outputs only within their own area of expertise. GenAI therefore does not simply substitute for human knowledge; it transforms the nature of knowledge work from execution to oversight, requiring organizations to retain the very expertise that GenAI ostensibly replaces, not to perform tasks, but to verify them.

The second distinctive feature is *deployment flexibility*. GenAI can be integrated into knowledge work along two dimensions: deployment mode and deployment location. Deployment mode determines how GenAI effectiveness relates to human ability, specifically through automation or augmentation. Under automation, GenAI autonomously executes tasks within a defined scope, and its effectiveness is independent of users' baseline skill. Under augmentation, GenAI extends the user's capability boundary beyond their intrinsic expertise, with effectiveness inherently coupled to baseline ability. Deployment location specifies which role within the hierarchy GenAI reshapes. In knowledge-based hierarchies, workers focus on task execution, whereas experts provide guidance on exceptions rather than routine execution (Garicano 2000). The organizational consequences of GenAI adoption are therefore inherently layer-dependent. At the worker level, GenAI either executes tasks that workers would otherwise perform (automation) or operates alongside workers as a capability amplifier (augmentation). At the expert level, GenAI either substitutes for experts as a first-line solution provider (automation) or accelerates experts' consultation efficiency (augmentation). Together, this two-dimensional flexibility generates a $2 \times 2$ design space (Table 1), each cell carrying distinct implications for workforce composition.

**Table 1 GenAI Deployment Design Space**

| | **Automation** | **Augmentation** |
|---|---|---|
| **Worker Level** | Task Executor | Capability Amplifier |
| **Expert Level** | Virtual Supervisor | Communication Accelerator |

The interaction between intrinsic fallibility and deployment flexibility generates organizational dynamics that traditional technology-adoption theories do not predict. These two features play complementary but distinct roles. Intrinsic fallibility establishes *why* organizational design matters under GenAI: human oversight remains essential regardless of deployment configuration. Deployment flexibility determines *how* organizational consequences vary: the choice of mode and location shapes who validates, what is validated, and how validation reshapes skill requirements and workforce composition. Traditional automation executed rule-based tasks without probabilistic errors requiring cognitive oversight; earlier AI systems provided predictions or recommendations but left execution to human workers; and information and communication technologies reduced coordination frictions without performing cognitive work. GenAI uniquely combines autonomous cognitive

execution with intrinsic fallibility and unprecedented deployment flexibility, creating organizational design challenges and opportunities that existing frameworks cannot address.

This paper asks: *How does GenAI adoption reshape organizational structure in knowledge-based hierarchies?* To answer this question, we develop a formal analytical model building on Garicano (2000), where hierarchies emerge endogenously to economize on knowledge acquisition costs. We incorporate three GenAI-specific elements. First, a capability parameter determines the scope of tasks that GenAI can address. Second, intrinsic fallibility is captured by a hallucination rate that is not mechanically tied to capability, consistent with evidence that more advanced models may exhibit comparable or higher error rates across tasks (Zhou et al. 2024). Third, we model explicit validation and correction time requirements that draw on scarce worker or expert labor. We analyze four deployment architectures spanning the $2 \times 2$ design space and characterize their effects on optimal workforce composition, skill requirements, and managerial spans of control (i.e., the ratio of low-level workers to high-level experts).[1] Our analysis yields findings that challenge conventional wisdom about technology-driven organizational change.

First, the effect of GenAI on entry-level skill requirements is highly mode-dependent. Under worker-level automation, GenAI takes over routine execution, yet the firm still relies on workers to validate AI outputs and resolve residual tasks. The firm optimally raises entry-level skill requirements, hiring fewer but more knowledgeable juniors to curb costly escalation to senior experts and to ensure reliable oversight of AI outputs. Under worker-level augmentation, in contrast, GenAI enlarges workers' effective problem-solving range, allowing the firm to lower the knowledge requirement for junior roles while sustaining performance through AI-assisted workflows; deskilling is optimal. This distinction suggests that the decline in junior employment is not an inevitable consequence of GenAI per se, but is more consistent with deployment choices that tilt toward automation rather than augmentation, in line with emerging evidence on early-career impacts (Brynjolfsson et al. 2025a).

Second, deployment location can be as consequential as deployment mode. Under expert-level automation, AI-based substitutes for expert input lower the cost of escalation, so the firm optimally reduces entry-level skill requirements while preserving problem-solving capacity through AI plus expert oversight. A distinctive implication is what we term *post-adoption capability insensitivity*: once the firm adopts expert-level automation, further improvements in GenAI capability

[1] We emphasize that this $2 \times 2$ framework is an analytical decomposition of distinct organizational mechanisms rather than a rigid product taxonomy. In practice, a single technology deployment may blend elements across cells. For example, GitHub Copilot operates primarily as worker-level automation when generating code blocks autonomously, but shifts toward worker-level augmentation when providing in-line suggestions during active editing. Our analytical separation is precisely what enables the key insight: these blended deployments produce qualitatively different, and sometimes opposing, organizational consequences depending on which mechanism dominates.

shift queries between GenAI and human experts but leave the optimal entry-level skill threshold unchanged within the AI-covered region. Under expert-level augmentation, GenAI accelerates experts' handling of consultations and lowers the effective cost of accessing high-level knowledge, with deskilling strengthening as the technology advances. Taken together, expert-level adoption yields a clean first-order prediction: it uniformly reduces entry-level skill requirements, broadening access to knowledge work regardless of deployment mode.

Third, departing from the conventional view that technological progress monotonically flattens hierarchies, we uncover a robust non-monotonic evolution of the span of control across all deployment modes: the span contracts initially and expands later as GenAI advances. This implies that even if organizations ultimately flatten, the demand for senior expertise can remain stable or even rise during early-to-intermediate deployment stages, consistent with large-scale evidence (Brynjolfsson et al. 2025a, Hosseini and Lichtinger 2025). The mechanisms, however, differ by deployment mode and location. Under worker-level automation, early automation gains reduce frontline labor demand while leaving expert demand largely unchanged, narrowing the span of control. As GenAI matures, induced upskilling curbs upward referrals, compresses expert demand, and widens the span. In the other deployment modes, deskilling at the entry level takes precedence in the early stage, increasing the volume of tasks that require expert intervention and thereby tightening the span of control. At later stages, GenAI's direct efficiency gains substantially reduce reliance on expert input, and the span of control eventually widens.

Our analysis yields actionable insights for stakeholders navigating the GenAI transformation. For firms, technology procurement cannot be decoupled from hiring strategy. Under worker-level automation, the role of entry-level employees shifts from primary task execution to oversight and validation of AI outputs. In this setting, organizations benefit from hiring workers with sufficient domain expertise and judgment to detect subtle errors and escalate problematic cases appropriately. Training should therefore strengthen workers' ability to critically evaluate AI-generated outputs, recognize potential failure modes, and apply domain knowledge when validating or correcting automated results. Conversely, under augmentation, firms can broaden recruitment toward candidates with strong foundational traits but without deep technical mastery, pivoting HR criteria from peak technical execution toward adaptability and human–AI collaborative aptitude.

Deployment decisions at the expert layer carry different implications. By expanding experts' capacity to supervise and support entry-level workers, expert-level GenAI enables organizations to sustain performance while lowering entry-level knowledge requirements. Critically, firms should anticipate that demand for senior expertise may remain stable or even increase during early-to-intermediate stages of GenAI deployment. Preemptive hierarchy flattening is therefore counterproductive, as it risks creating bottlenecks in exception handling and quality control precisely when expert oversight becomes more critical.

For policymakers and educators, our findings suggest that the labor-market consequences of GenAI depend less on the overall rate of adoption than on how the technology is deployed within organizational hierarchies. Policies that encourage augmentation-oriented applications can help broaden participation in knowledge work, while deployments that expand expert capacity may similarly enable firms to support a larger base of less specialized workers. Furthermore, workforce development programs should move beyond generic AI literacy toward mode-specific curricula that emphasize deep domain expertise for validation in automation settings versus adaptive problem-solving in augmentation contexts.

The rest of the paper is organized as follows. Section 2 compares and contrasts our paper with the related literature. Section 3 outlines the model setup for pre-GenAI organizations. Sections 4 and 5 analyze worker-level automation and augmentation, respectively. Section 6 examines GenAI adoption at the expert level. We conclude in Section 7. All proofs and supplemental materials are presented in the appendices.

## 2. Literature Review

Our work contributes to three research streams: the economics of organizational design and knowledge-based hierarchies, the literature on AI's labor market and economic impacts, and related work on human–AI collaboration.

Our work builds most directly on the foundational literature examining how organizations structure knowledge work and allocate problem-solving across hierarchical layers.[2] Garicano (2000) establishes the seminal framework in which hierarchy emerges endogenously to economize on knowledge acquisition costs: workers handle routine problems independently while escalating exceptional cases to more knowledgeable experts, and the optimal span of control balances the benefits of knowledge specialization against communication costs. This framework has been extended to incorporate product positioning (Wu 2015), heterogeneous communication technologies (Gumpert 2018), and variations in decision complexity (Bloom et al. 2014). Garicano and Rossi-Hansberg (2006) and Garicano and Rossi-Hansberg (2015) provide comprehensive surveys.

A key insight from this literature is that technologies altering communication costs can shift the balance between centralized and decentralized problem-solving, thereby reshaping spans of control (Bloom et al. 2014). However, these earlier information and communication technologies (ICTs) primarily reduced coordination and communication frictions without substituting for cognitive labor or introducing systematic unreliability. GenAI differs on both dimensions: it can autonomously

[2] The relationship between technology adoption, skill demand, and organizational structure is a foundational inquiry in economics. Traditional frameworks—the skill-biased technical change hypothesis (Katz and Murphy 1992, Autor et al. 1998), the routine-biased technical change refinement (Autor et al. 2003, Goos and Manning 2007), and the task-based framework (Acemoglu and Autor 2011)—generally predict uniform directional effects for a given technology, typically upskilling or the displacement of routine tasks.

generate substantive task outputs, yet does so with non-trivial error risk that requires human validation.

The paper most closely related to ours is Ide and Talamas (2025), which embeds AI in a competitive economy of knowledge hierarchies, modeling AI as agents with an exogenously fixed knowledge level and showing that equilibrium outcomes hinge on whether agents act autonomously as "co-workers" or only as advisory "co-pilots." Our paper shares this knowledge-hierarchy foundation but introduces a mechanism absent from perfect-substitute agent models: GenAI's intrinsic fallibility creates an endogenous validation-and-escalation workload that reshapes workforce predictions. This mechanism yields workforce predictions—specifically, how entry-level skill thresholds and spans of control respond to deployment mode and location—that do not emerge from frameworks treating AI as an error-free additional layer of knowledge.

Our theoretical predictions connect to a rapidly growing empirical literature documenting AI's heterogeneous effects on workforce composition and job quality. Early research on pre-GenAI automation technologies established that AI adoption can simultaneously displace workers in routine cognitive tasks while increasing demand for workers in complementary roles (Acemoglu and Restrepo 2019, Acemoglu and Johnson 2023, Autor 2024). Recent empirical studies of GenAI adoption document substantial productivity gains but reveal considerable heterogeneity across skill levels and contexts. Noy and Zhang (2023) and Brynjolfsson et al. (2025b) find that GenAI tools substantially accelerate task completion, with particularly pronounced benefits for initially lower-performing workers. Peng et al. (2023) and Cui et al. (2024) document that GitHub Copilot enables developers to code significantly faster, though with quality concerns that require heightened code review. Dell'Acqua et al. (2023) show that consultants using GenAI exhibit significant productivity improvements but face what they term a "jagged frontier": some tasks benefit dramatically while others see minimal gains or quality degradation.

Two recent large-scale studies provide convergent evidence of GenAI's seniority-biased labor market effects that are particularly relevant to our paper. Brynjolfsson et al. (2025a) report that early-career employment declined by approximately 13% in the most AI-exposed occupations after 2022, while senior employment in those occupations remained stable or rose. Complementing this reduced-form evidence with a firm-level adoption design, Hosseini and Lichtinger (2025) use direct measures of GenAI adoption and show that adopting firms reduce junior employment by roughly 9–10% within six quarters while senior employment remains largely unchanged, with effects concentrated in high-exposure occupations and driven primarily by reduced hiring rather than separations. Taken together, these papers establish a robust empirical regularity—entry-level opportunities can contract even as senior roles remain stable—but they are not designed to pin

down why the effects differ across contexts. Our paper provides the missing organizational mechanism by endogenizing deployment architecture within the firm: not only whether GenAI is used for automation versus augmentation, but also where it is deployed in the hierarchy and who bears the validation and rework burden created by intrinsic fallibility. This framework offers a lens for interpreting the seniority-biased pattern by tracing how deployment architecture shapes entry-level skill requirements and spans of control.

Lastly, our analysis relates to a growing literature examining how human workers and AI systems interact. Fügener et al. (2021) study "humans-in-the-loop" configurations where workers oversee algorithmic decision-making, documenting both benefits and potential cognitive pitfalls. Wang et al. (2024) find that less experienced workers benefit more from AI assistance, a pattern consistent with our augmentation mode. Bastani et al. (2021) and Ibrahim et al. (2021) examine optimal designs for human–algorithm collaboration in prediction and decision tasks. More broadly, recent work has begun to investigate how the modality and structure of human–AI collaboration shape creative and knowledge-intensive outcomes (e.g., Chen and Chan 2024, Hou et al. 2025, Lu and Zhang 2025). Bastani and Cachon (2025) show that as AI errors become rarer, motivating human vigilance can become paradoxically more expensive. Zhang et al. (2025) study a platform setting in which human creators can delegate tasks to AI and show that hidden effort distortions can induce excessive delegation. Dai and Taylor (2026) further show that enterprise GenAI design choices (e.g., temperature) interact with hidden human effort under moral hazard to endogenously shape the optimal hallucination–creativity trade-off. Siderius et al. (2026) study a dynamic principal-agent model in which reliance on organizational AI can erode worker expertise over time, showing that incentive frictions can cause firms to underinvest in oversight and tolerate skill atrophy. Together, these studies reinforce that organizations' skill and oversight demands depend as much on how GenAI is deployed and governed as on its raw capability. Our work complements this literature by embedding deployment flexibility within a knowledge-based hierarchy where intrinsic fallibility makes validation a binding organizational constraint, showing that the same GenAI technology can produce opposite workforce effects depending on where and how it is deployed.

## 3. Benchmark: Pre-GenAI Organizational Structure

We construct a parsimonious model of knowledge-based hierarchy and establish the pre-GenAI organizational structure as a benchmark. We consider a two-layer firm consisting of low-level junior workers and high-level senior experts. While our model formally depicts a two-layer hierarchy, the "senior expert" role structurally proxies for middle-management functions centered on exception handling, quality assurance, and oversight. The firm faces one problem per unit of time, a normalization that captures a steady-state workflow without loss of generality. The difficulty of each

problem $d$ is ex-ante unknown and uniformly distributed over the interval $[0,1]$. If the problem is solved, the firm can produce one unit of output; otherwise, no output is produced. For simplicity, we normalize the value of this output to 1. We assume that each worker can process one problem per unit of time. Therefore, the firm employs one worker and $n$ experts, which gives a *span of control* (the ratio of workers to experts) of $s = \frac{1}{n}$.

The labor market consists of individuals with varying levels of expertise, quantified by their *knowledge level* $x \in [0,1]$, which represents an individual's ability to solve problems up to difficulty $x$. The firm perfectly observes each candidate's knowledge level and, in an efficient labor market, selects the precise knowledge levels needed to maximize profitability.

The worker is employed to independently handle the easier problems and seeks assistance for more complex ones. Specifically, a worker with knowledge level $x$ independently solves problems with difficulty $d \in [0,x]$, where $0 \leq x \leq 1$. If the problem's difficulty exceeds the worker's knowledge level, i.e., $d \in (x,1]$, the worker requests help from experts, who then provide solutions, allowing the worker to handle the problem and produce one unit of output. Importantly, the worker remains the sole producer of output. For problems exceeding their knowledge boundary, the worker relies on expert consultation to obtain the solution rather than delegating the task itself to the expert for execution.

Following Bloom et al. (2014), we assume that experts possess complete knowledge over the entire difficulty spectrum $[0,1]$. In the context of a two-layer hierarchy, this assumption simply implies that all problems under consideration fall within the firm's collective problem-solving capacity, so that escalations represent transfers of work upward rather than unresolved uncertainty. Importantly, it ensures that changes in organizational structure are driven entirely by endogenous reorganization rather than by limits to top-layer competence, cleanly isolating our focal mechanism: how GenAI's fallibility reallocates verification and escalation burdens across layers. We relax this assumption in Appendix E and show that our insights remain robust.

Each time a worker consults an expert, the communication process consumes $t_c$ units of the expert's time, where $t_c > 0$ is the *communication cost*. Since the worker escalates problems with probability $1-x$ and each consultation takes $t_c$ units of expert time, the firm requires $n_0 = (1-x)t_c$ experts. We assume that the communication cost $t_c$ lies within the range $(0, \bar{t}_c)$, where $\bar{t}_c = \frac{2k}{k+2w}$. If $t_c$ reaches or exceeds $\bar{t}_c$, the communication cost becomes high enough that the firm would find it more profitable to employ a worker with complete knowledge, eliminating the need for expert support.

Drawing on the literature on hierarchical organization in the knowledge economy (e.g., Garicano 2000, Antràs et al. 2006), we assume that the firm pays an individual with knowledge level $x$ a wage of $w + \frac{1}{2}kx^2$, where $w$ is the market price of a unit of work time and $\frac{1}{2}kx^2$ represents the

compensation for the individual's knowledge, with $k$ as the knowledge premium.[3] The firm's profit function is thus:

$$\Pi_0(x) = 1 - \left(w + \frac{1}{2}kx^2\right) - \left(w + \frac{1}{2}k\right)(1-x)t_c. \tag{1}$$

The presence of experts with comprehensive knowledge ensures that any problems the firm encounters can be resolved, allowing the firm to maintain one unit of output per unit of time. The term $\left(w + \frac{1}{2}kx^2\right)$ represents the wage the firm must pay to hire one worker with knowledge level $x$, while $\left(w + \frac{1}{2}k\right)(1-x)t_c$ reflects the total wage required to employ $(1-x)t_c$ experts. Here, since the expert's knowledge level is 1, the firm pays each expert a wage of $w + \frac{1}{2}k$. The firm's objective is to determine the optimal knowledge level $x^*$ for its workers to maximize profit per unit of time.

The following Lemma 1 reveals the firm's optimal organizational structure in the pre-GenAI environment. Throughout this paper, "organizational structure" refers specifically to two endogenous outcomes: the optimal knowledge level of entry-level workers ($x^*$), which determines workforce composition, and the span of control ($s^* = 1/n$), which captures hierarchical shape.

LEMMA 1. *In the pre-GenAI organization, the optimal worker knowledge level is* $x_0^* = \frac{(k+2w)t_c}{2k}$ *and the span of control is* $s_0^* = \frac{1}{(1-x_0^*)t_c}$.

Organizations face a fundamental trade-off when determining the optimal knowledge level of workers. While lower knowledge requirements reduce worker wage costs, they simultaneously increase the need for worker-expert communications, necessitating more senior positions. As shown in Lemma 1, the optimal worker knowledge level $x_0^*$ responds differently to changes in knowledge premium $k$ and communication costs $t_c$. A lower knowledge premium reduces the cost of hiring more knowledgeable workers, increasing worker autonomy. Conversely, lower communication costs make it more efficient to rely on managerial consultation. This finding aligns with previous research by Garicano (2000) and Bloom et al. (2014) on how information and communication technology (ICT) development influences organizational hierarchies.

## 4. Worker-Level Automation

Having established the pre-GenAI benchmark, we now examine how the four GenAI deployment modes reshape this organizational structure. In this section, we begin with a common deployment mode in which GenAI autonomously executes low-level knowledge tasks that would otherwise be performed by junior workers. We refer to this mode as *worker-level automation* (denoted by subscript $t$). Specifically, GenAI automatically completes tasks in the interval $[0, r_t]$, where $0 < r_t < 1$ represents GenAI's *automation capability*, while the remaining tasks in $(r_t, 1]$ are handled by workers (see Figure 1).

[3] This wage structure reflects the convex cost of acquiring knowledge at increasing levels. To satisfy participation constraints, the firm aims to set wages that align with the costs employees incur to attain their knowledge levels.

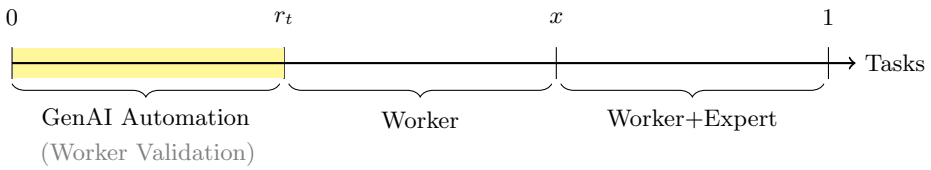



**Figure 1 Delegation of tasks under worker-level automation**

To formalize the presence of hallucinations, we assume that each GenAI-automated task is solved incorrectly with probability $h$, where $0 < h < 1$ denotes the *hallucination rate*. Because erroneous AI outputs may trigger losses far exceeding the value of the task itself, such as reputational damage or strained business relationships, we assume that the firm mandates human validation for all AI-processed tasks to identify and resolve potential hallucinations. We treat the hallucination rate $h$ as an exogenous primitive independent of GenAI's capability $r_t$. This reflects growing evidence that improvements in model scale or capability do not necessarily translate into lower error rates, and that more advanced language models may exhibit equal or even higher hallucination rates across tasks (Zhou et al. 2024). In Appendix C, we show that our main results remain robust when $h$ is inversely correlated with $r_t$.

Since GenAI automates tasks in the interval $[0, r_t]$, the firm requires only $(1 - r_t)$ workers per unit of time to process the remaining tasks in $(r_t, 1]$. To mitigate the potentially severe losses caused by hallucinations, workers are also responsible for validating all AI-completed outputs and reworking tasks whenever hallucinations are detected. Consequently, the firm must hire workers whose knowledge level satisfies $x \geq r_t$. We assume that each validation consumes $t_v$ units of worker time, while reworking a task following a detected hallucination consumes $t_r$ units of worker time. We impose $t_v < 1$ and $t_r < 1$, implying that either validation or rework alone is less time-consuming than manual task execution (normalized to one). However, we assume $t_v + t_r > 1$, so that when a hallucination occurs, the combined effort of validation and rework strictly exceeds the time required for manual processing. This condition captures the realistic scenario where catching and correcting an AI error is more costly than handling the task manually from the outset, which is precisely what makes the hallucination risk organizationally consequential. Absent this condition (i.e., if $t_v + t_r \leq 1$), GenAI adoption would always be optimal regardless of the hallucination magnitude. Therefore, after GenAI adoption, the firm's profit function is given by

$$\Pi_t(x) = 1 - \left(w + \frac{1}{2}kx^2\right)\left[(1 - r_t) + (t_v + ht_r)r_t\right] - \left(w + \frac{1}{2}k\right)(1 - x)t_c. \tag{2}$$

Since worker validation eliminates the impact of hallucinations on the final output, the firm's output per unit of time remains normalized to one. The total worker time requirement $[(1 - r_t) + (t_v + ht_r)r_t]$ corresponds to the number of workers the firm must employ per unit of time. The

term $[(1-r_t)+(t_v+ht_r)r_t]$ consists of three components: $(1-r_t)$ units of time devoted to non-automated tasks; $t_v r_t$ units allocated to validating AI-completed tasks; and an expected rework time of $ht_r r_t$ when hallucinations occur. By contrast, the demand for senior experts remains at $(1-x)t_c$, as the probability that a worker encounters a problem beyond their knowledge level is unchanged at $1-x$, with each consultation requiring $t_c$ units of expert time.[4]

### 4.1. Impact of GenAI Advancements on Worker Knowledge

We first examine when firms should implement worker-level automation and how GenAI technology advancements reshape organizational knowledge hierarchies.

LEMMA 2. *The firm should adopt worker-level automation if and only if the hallucination rate satisfies* $h < \overline{h}_t$, *where* $\overline{h}_t = \frac{1-t_v}{t_r}$ *for* $r_t \le x_0^*$ *and* $\overline{h}_t < \frac{1-t_v}{t_r}$ *for* $r_t > x_0^*$.

Lemma 2 establishes that the firm's adoption of worker-level automation is contingent upon a critical hallucination threshold, $\overline{h}_t$. When GenAI capability is relatively low ($r_t \le x_0^*$), the threshold is given by $\overline{h}_t = \frac{1-t_v}{t_r}$. This condition ensures that automating tasks in $[0, r_t]$ yields a net reduction in worker demand: the labor saved exceeds the validation and rework burden, i.e., $r_t > (t_v + ht_r)r_t$. However, as GenAI capability expands beyond the firm's pre-adoption worker knowledge level ($r_t > x_0^*$), the adoption decision becomes more restrictive. While higher $r_t$ enables the automation of increasingly complex tasks, it simultaneously requires that the firm employ workers with a knowledge level at least equal to $r_t$ to effectively validate AI-generated outputs. This requirement inflates worker wage costs, making adoption more sensitive to inefficiencies caused by hallucinations. Consequently, the hallucination threshold for adoption tightens as GenAI becomes more capable ($\overline{h}_t < \frac{1-t_v}{t_r}$). This result highlights that GenAI adoption is governed by the strategic interaction between the automation capability and the resulting organizational burden of human oversight.

PROPOSITION 1. *The optimal worker knowledge level under worker-level automation is* $x_t^* = \max\{r_t, \min\{1, x_t\}\}$, *where* $x_t = \frac{x_0^*}{(1-r_t)+(t_v+ht_r)r_t}$. *Moreover,* $x_t^*$ *increases in GenAI's automation capability* $r_t$ *and decreases in hallucination rate* $h$, *and* $x_t^* > x_0^*$.

Proposition 1 shows that worker-level automation drives an *upskilling* effect: the firm optimally raises entry-level skill requirements after GenAI adoption ($x_t^* > x_0^*$).[5] The mechanism operates

[4] Our baseline treats capability $r_t$ and hallucination rate $h$ as exogenous and abstracts from direct monetary adoption costs (e.g., licensing, integration) to isolate the organizational mechanism. The organizational costs of GenAI integration, including validation, rework, and consultation burdens, are modeled explicitly. Appendix D endogenizes capability and reliability via convex investment costs and confirms that our main insights are robust.

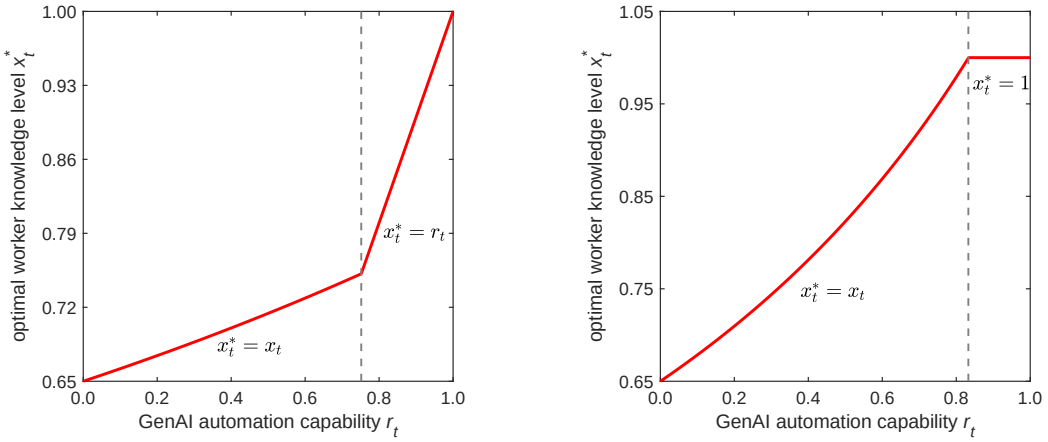
[5] Throughout the paper, "upskilling" and "deskilling" refer specifically to changes in the entry-level knowledge threshold $x^*$: upskilling means the firm raises the worker knowledge requirement, deskilling means it lowers it. These terms do not characterize overall skill demand in the firm, as changes in span of control can simultaneously move demand for senior expert roles in the opposite direction.

through a cost-ratio channel distinct from traditional skill-biased technical change. Automation reduces the firm's total worker-time requirement from 1 to $(1 - r_t) + (t_v + ht_r)r_t$, lowering the effective cost of worker knowledge per unit of output. The marginal cost of raising $x$, captured by the convex knowledge premium $\frac{1}{2}kx^2$, is now borne by fewer worker units, while the marginal benefit from reducing expert consultations remains unchanged at $(w + \frac{1}{2}k)t_c$ per unit of escalation probability. As a result, the firm's optimal worker knowledge level tilts upward. Intuitively, when automation handles routine execution, each unit of worker knowledge is leveraged over a smaller labor base, making it cost-effective for the firm to invest in higher-quality frontline labor that reduces costly upward referrals.

This upskilling intensifies as GenAI's automation capability $r_t$ increases, which further reduces the worker-time base over which knowledge costs are amortized. Furthermore, a decline in the hallucination rate $h$ strengthens the upskilling effect: a lower $h$ reduces expected rework time per automated task, reducing total worker demand and lowering the marginal cost of knowledge investment. This comparative static may seem paradoxical: lower error rates lead to higher, not lower, skill requirements. The driver is the cost channel: more reliable automation requires fewer worker-hours, making the convex knowledge premium less expensive to bear, so the firm responds by raising the skill threshold.

Beyond this cost-ratio channel, the validation mandate introduces a direct skill-floor effect: because workers must be capable of detecting errors in all AI-automated tasks, the firm is constrained to hire workers with $x \geq r_t$. Figure 2 illustrates how these two channels interact to produce distinct upskilling trajectories depending on GenAI's reliability. Under moderate hallucination rates (Figure 2a), as GenAI capability increases, the optimal worker knowledge level shifts from the unconstrained optimum to exactly matching GenAI's automation capability ($x_t^* = r_t$). The skill-floor constraint becomes binding: the labor effort required to validate and rework AI-generated outputs is substantial enough that hiring workers with knowledge beyond $r_t$ would incur prohibitively high wage costs relative to the marginal benefit. The firm therefore sets the worker knowledge level at the minimum threshold required for validation. In this sense, the validation requirement amplifies automation-induced upskilling: expanding the automation scope enlarges the set of tasks that workers must be able to verify and correct. Conversely, in a high-reliability environment (Figure 2b), reliable GenAI automation drives the optimal worker knowledge level to its upper bound of one. The traditional hierarchy collapses into a single-layer organization of fully knowledgeable workers.

Our theoretical finding that worker-level automation induces an upskilling effect aligns closely with the empirical evidence discussed earlier (Brynjolfsson et al. 2025a, Hosseini and Lichtinger 2025): the documented decline in junior employment among GenAI-adopting firms is consistent




(a) $h = 0.4$ (b) $h = 0.1$

**Figure 2 Impact of GenAI capability enhancement on worker knowledge level under worker-level automation ($k = 0.8$, $w = 0.25$, $t_c = 0.8$, $t_v = 0.5$, $t_r = 0.8$)**

with our mechanism: automating routine execution reduces demand for low-skilled executors, prompting firms to hire fewer but more knowledgeable workers who oversee AI outputs and reduce costly escalations.

This mechanism differs sharply from traditional automation-driven upskilling (e.g., Acemoglu and Restrepo 2019). In those frameworks, automation eliminates routine tasks from workers' portfolios, and upskilling arises because the remaining tasks are inherently more complex. Under GenAI, by contrast, automated tasks do not leave the worker's domain of responsibility: they return as a validation tax that the worker must pay on every AI-executed task. The upskilling imperative arises not from the removal of routine work but from the oversight of AI-executed work, driven by intrinsic fallibility. Furthermore, as we show in Section 5, this upskilling outcome is specific to the automation mode; the same technology deployed for augmentation produces the opposite effect. This mode-dependence has no counterpart in traditional automation theories, which predict a uniform directional skill effect for a given technology.

### 4.2. Impact of GenAI Advancements on Organizational Structure

We next examine how advancements in GenAI technology reshape organizational structure under worker-level automation. In particular, we study how improvements along two fundamental dimensions of GenAI (its capability to automate a broader range of tasks and its reliability in generating correct outputs) affect the firm's managerial span of control, defined as the ratio of junior workers to senior experts.

PROPOSITION 2. *After the adoption of worker-level automation,*

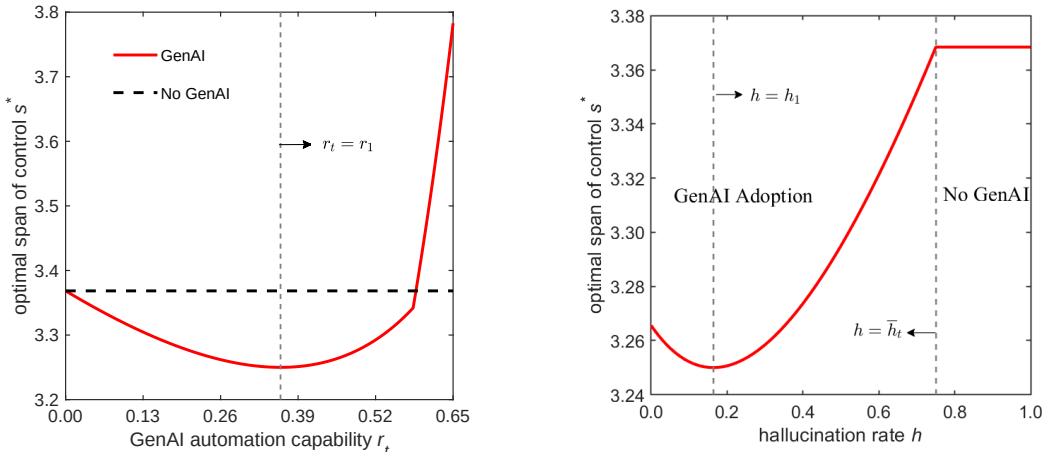
(a) *When $h \geq h_0$ or $r_t < r_0$, the optimal span of control is given by $s_t^* = \frac{1 - r_t + (t_v + h t_r) r_t}{(1 - x_t^*) t_c}$. Furthermore, the span of control $s_t^*$*

(1) *decreases in automation capability* $r_t$ *on* $(0, r_1)$ *and increases in* $r_t$ *thereafter;*

(2) *decreases in hallucination rate* $h$ *on* $(0, h_1)$ *and increases in* $h$ *thereafter.*

(*b*) *Otherwise, when* $h < h_0$ *and* $r_t \geq r_0$*, the traditional two-layer organizational hierarchy collapses into a single-layer structure composed solely of workers with complete knowledge, i.e.,* $x_t^* = 1$.

Proposition 2 demonstrates a non-monotonic evolution of organizational structure in response to GenAI advancements. Specifically, the optimal span of control exhibits a distinct "U-shaped" trajectory as the automation capability $r_t$ increases (see Figure 3a). This dynamic stems from the shifting dominance of two competing forces: a direct *task-automating effect* and an indirect *organization-upskilling effect*. When automation capability is limited (i.e., at low $r_t$), the task-automating effect prevails: since the modest reduction in worker demand provides little incentive for upskilling, increases in $r_t$ primarily contract the worker headcount while expert demand remains stable, narrowing the span of control. Conversely, at high $r_t$, the *organization-upskilling effect* takes precedence. As the marginal labor savings from further automation diminish within a significantly contracted workforce, the firm shifts its adjustment margin toward worker knowledge. To minimize costly referrals, the firm substantially elevates the worker knowledge level. As a result, the reduction in expert demand outpaces the decline in worker headcount, thereby widening the span of control. Notably, as illustrated in Figure 3a, the span of control surges once $r_t$ reaches a sufficiently high level. This is because the firm raises the worker knowledge threshold to $x_t^* = r_t$ to ensure validation and correction of increasingly complex AI outputs. Such intense upskilling effectively internalizes the vast majority of problem-solving within the worker layer, causing the demand for senior experts to collapse and leading to a rapid flattening of the organization.




(a) $h = 0.1$ (b) $r_t = 0.4$

**Figure 3 Impact of GenAI advancements on span of control under worker-level automation** **($k = 0.8$, $w = 0.25$, $t_c = 0.5$, $t_v = 0.4$, $t_r = 0.8$)**

A similar non-monotonic pattern can arise as GenAI becomes more reliable (Figure 3b), although the underlying mechanism is subtly different from that for capability improvements. When hallucination rates are relatively high ($h > h_1$), a marginal reduction in $h$ mainly alleviates workers' correction burden on AI-automated tasks. We refer to this channel as the direct *correction-reduction effect*. This effect reduces worker labor demand more than expert demand, thereby narrowing the span of control. Once hallucinations become sufficiently rare ($h < h_1$), further reliability gains operate primarily through the organization-upskilling channel: workers can internalize more problem solving, reduce referrals, and drive a faster contraction of expert demand, which widens the span of control. Importantly, the initial contraction at high hallucination rates hinges on the presence of worker oversight. Without these frictions, improved reliability would chiefly amplify the upskilling channel that substitutes away from experts, so decreases in $h$ would increase the span of control monotonically rather than producing an initial decline.

Finally, Proposition 2(b) identifies an extreme organizational regime. When GenAI is both highly capable and sufficiently reliable, the upskilling effect becomes so strong that the firm optimally hires workers with complete knowledge ($x_t^* = 1$). In this case, workers can independently handle all non-automated tasks while validating GenAI outputs, eliminating the need for expert assistance. The traditional two-layer hierarchy therefore collapses into a single-layer organizational structure composed solely of fully knowledgeable workers.

Overall, worker-level automation induces a staged reconfiguration of organizational structure as automation expands. At low levels of automation, the hierarchy becomes steeper as routine tasks are automated while expert roles remain intact. As automation deepens, firms respond by upgrading worker knowledge, compressing the expert layer and flattening the organization. In the limit, sufficiently capable and reliable GenAI renders expert assistance unnecessary, collapsing the hierarchy into a single-layer organization of fully knowledgeable task executors.

## 5. Worker-Level Augmentation

While worker-level automation centers on substituting human effort in routine execution, GenAI can also reshape organizations through a fundamentally different mechanism: it is embedded directly into workers' workflows to expand their effective problem-solving capabilities beyond their intrinsic skill endowments. In this deployment mode, GenAI operates alongside workers during task execution, providing real-time suggestions and refinements that workers assess and incorporate, so that outcomes are jointly produced rather than autonomously generated. Conceptually, this pattern resembles a copilot-style use of GenAI, in which the technology augments workers' capabilities without displacing their role in the production process. We refer to this capability-amplifying deployment as our second deployment mode, *worker-level augmentation* (denoted by subscript $g$).

Under worker-level augmentation, workers integrate GenAI into their task execution process and operate as a tightly coupled human–AI collaborative unit. This in-workflow integration expands a worker's effective capability boundary beyond their intrinsic knowledge level $x$, but the magnitude of this expansion depends jointly on the worker's own skill and the augmentation capability of GenAI.[6] Specifically, we assume that an augmented worker can resolve problems within the difficulty range $[0, x + r_g(1-x)]$, where $r_g \in (0,1)$ denotes the GenAI's augmentation capability. Problems with difficulty exceeding this augmented boundary, i.e., $d \in (x + r_g(1-x), 1]$, still require upward consultation with senior experts (see Figure 4).

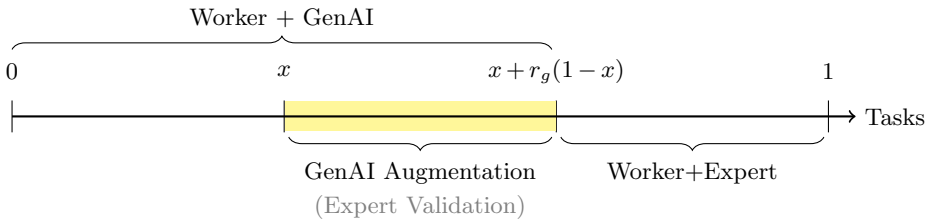



**Figure 4 Delegation of tasks under worker-level augmentation**

This formulation captures a key feature of the copilot-style augmentation: the benefits of GenAI are inherently heterogeneous across workers. As a worker's intrinsic knowledge level $x$ increases, the incremental value of GenAI assistance diminishes, because highly skilled workers already possess much of the relevant problem-solving capability. Conversely, less-skilled workers experience a disproportionately larger expansion of their effective capability boundary from the same augmentation technology. In this sense, worker-level augmentation exhibits an equalizing effect, consistent with empirical evidence that GenAI disproportionately enhances the performance of lower-skilled workers (Noy and Zhang 2023, Brynjolfsson et al. 2025b, Caplin et al. 2025).

However, this expansion of capability boundaries introduces a critical validation challenge. For tasks in the augmented interval $(x, x + r_g(1-x)]$, problem difficulty exceeds the worker's intrinsic knowledge level, preventing the worker from independently verifying the accuracy of the human–AI collaborative output. As in the automation setting, we assume that solutions in this interval are subject to hallucinations, with probability $h$ of being incorrect. To eliminate such hallucinations, senior experts must validate AI-assisted outputs. When errors are detected, experts intervene to assist workers in correcting them. Accordingly, after adopting worker-level augmentation, the firm's profit function is given by

$$\Pi_g(x) = 1 - \left(w + \frac{1}{2}kx^2\right) - \left(w + \frac{1}{2}k\right)\left[r_g(1-x)(t_v + ht_c) + (1-r_g)(1-x)t_c\right]. \tag{3}$$

[6] Beyond expanding the capability boundary, worker–AI collaboration may also improve execution efficiency for tasks that already fall within the worker's intrinsic knowledge range $[0, x]$. We formally incorporate such a productivity improvement in Appendix B.

In contrast to worker-level automation, workers remain responsible for executing all tasks within each unit of time. Therefore, the firm continues to employ one worker and incurs a total worker-hiring cost of $w+\frac{1}{2}kx^2$. The total expert time requirement, $[r_g(1-x)(t_v+ht_c)+(1-r_g)(1-x)t_c]$, determines the number of experts needed and reflects multiple managerial responsibilities. The first component, $r_g(1-x)(t_v+ht_c)$, captures the oversight burden associated with human–AI collaboration within the augmented task range: experts devote $r_g(1-x)t_v$ units of time to validating AI-assisted outputs and incur an additional expected time cost of $r_g(1-x)ht_c$ to assist workers in correcting hallucination-related errors. The second component, $(1-r_g)(1-x)t_c$, represents conventional expert intervention for tasks that exceed the combined capabilities of the worker and GenAI, i.e., those in the interval $(x+r_g(1-x),1]$. We further assume that the worker's hallucination-induced rework cost is negligible under augmentation, since the worker is already engaged in execution and the primary rework friction is the expert-worker communication time.

### 5.1. Impact of GenAI Advancements

LEMMA 3. *The firm should adopt worker-level augmentation if and only if the hallucination rate satisfies* $h<\overline{h}_g=1-\frac{t_v}{t_c}$.

Lemma 3 shows that firms should adopt worker-level augmentation only when the hallucination rate $h$ lies below a critical threshold $\overline{h}_g$. Notably, this threshold is determined solely by validation and communication costs. The condition $h<\overline{h}_g$ can be rewritten as $r_g(1-x)(t_v+ht_c)<r_g(1-x)t_c$. The right-hand side of this inequality can be interpreted as the expert time when the worker directly consults them about problems within the difficulty range $(x,x+r_g(1-x)]$. The left-hand side represents experts' working time for the same problems under GenAI augmentation, specifically the time spent validating AI-assisted outputs and correcting hallucination-induced errors. Accordingly, worker-level augmentation is profitable to firms if and only if the expert time devoted to mitigating hallucinations is less than the time required for direct assistance in problem-solving. Unlike worker-level automation, where adoption hinges mainly on the net change in workers' workload, augmentation is pinned down by experts' validation and communication efficiency. The reason is that augmentation shifts the oversight burden from workers to experts.

PROPOSITION 3. *The optimal worker knowledge level under worker-level augmentation is* $x_g^*=\frac{(k+2w)[r_gt_v+(1+r_gh-r_g)t_c]}{2k}$. *Moreover,* $x_g^*$ *decreases in GenAI's augmentation capability* $r_g$ *and increases in hallucination rate* $h$, *and* $x_g^*<x_0^*$.

Proposition 3 characterizes how worker-level augmentation reshapes the firm's optimal skill demand. The central result is a *deskilling effect*: as firms deploy GenAI to augment frontline workers, they optimally reduce the intrinsic knowledge required of entry-level employees ($x_g^*<x_0^*$). Under

augmentation, GenAI operates as a capability amplifier that extends a worker's effective problem-solving boundary beyond their own knowledge level, thereby substituting for human expertise. By enabling less specialized workers to handle problems that previously required higher skill levels, the firm can sustain operational performance while economizing on labor costs. As a result, the firm's optimal hiring strategy shifts toward workers with lower intrinsic knowledge, with GenAI filling the resulting capability gap.

The proposition further shows that this deskilling effect strengthens as GenAI technology advances. As augmentation capability increases, workers can rely more on AI support during execution, enabling them to handle a broader range of problems relative to their intrinsic knowledge and reducing the marginal value of higher human expertise. Meanwhile, a lower hallucination rate reduces the expected cost of AI-assisted work by decreasing the frequency of errors that require expert intervention. Together, stronger augmentation capability and higher reliability reinforce the deskilling effect, pushing the optimal worker knowledge level downward.

Under the skill-biased complementarity view, new technologies raise the productivity of highly skilled workers and increase firms' demand for expertise (Bartel et al. 2007). In contrast, our research suggests that GenAI augmentation enables a different dynamic. Rather than driving upskilling, GenAI allows for what recent research terms "*deskilling*," where specialized human expertise is increasingly substituted by AI capabilities (Kinder et al. 2024, Liu et al. 2024). Agrawal et al. (2023) formalize this logic as the "Turing Transformation": AI equips workers to perform tasks that previously required extensive education and experience, shifting the competitive margin from human expertise to human–AI collaboration.

PROPOSITION 4. *After the adoption of worker-level augmentation, the optimal span of control is* $s_g^* = \frac{1}{(1-x_g^*)[r_g(t_v+ht_c)+(1-r_g)t_c]}$.

(*a*) $s_g^*$ *decreases in GenAI's augmentation capability* $r_g$ *on* $(0, r_2)$ *and increases in* $r_g$ *on* $(r_2, 1)$.

(*b*) $s_g^*$ *decreases in hallucination rate* $h$ *on* $(0, h_2)$ *and increases in* $h$ *on* $(h_2, \overline{h}_g)$.

Proposition 4(a) characterizes how worker-level augmentation reshapes the firm's managerial span of control as augmentation capability increases. Classic theories and recent empirical evidence suggest that stronger augmentation should monotonically widen the span of control, as AI-assisted workers need less expert guidance (Garicano and Rossi-Hansberg 2006, Bloom et al. 2014, Babina et al. 2023, Hoffmann et al. 2024). Proposition 4 overturns this prediction: the relationship is non-monotonic. When augmentation capability is limited (i.e., $r_g < r_2$), improvements in GenAI primarily affect organizations through indirect worker skill adjustment rather than through immediate gains in worker autonomy. In this regime, even modest augmentation allows firms to lower

worker knowledge requirements to reduce labor costs. As workers become less skilled, their dependence on expert guidance increases, which raises the volume of upward referrals and narrows the span of control. We refer to this channel as the indirect *organization-deskilling effect*.

As augmentation capability becomes sufficiently strong (i.e., $r_g > r_2$), the dominant adjustment margin shifts. At high levels of augmentation, AI-assisted workflows substantially enhance the effective problem-solving capacity of workers, especially when intrinsic skill requirements are low. In this regime, further reductions in worker skill no longer generate large increases in expert dependence. Instead, the ability of augmented workers to resolve problems independently improves, reducing the need for expert involvement. This direct *knowledge-augmenting effect* dominates, leading to a widening of the span of control.

A non-monotonic pattern also arises with respect to improvements in GenAI reliability, as described in Proposition 4(b). Reductions in the hallucination rate affect organizational structure through two countervailing forces. First, higher reliability lowers the expected cost of AI-assisted workflows, encouraging firms to further relax worker skill requirements to economize on wages; this amplifies the indirect *organization-deskilling effect* and increases workers' reliance on expert support. Second, because experts must intervene to correct AI-induced errors, a lower hallucination rate directly reduces the time experts spend on corrections; this direct *correction-reduction effect* frees managerial capacity and enables experts to supervise more workers. When hallucination rates are relatively high ($h > h_2$), reductions in $h$ mainly operate through the deskilling margin. Firms aggressively lower worker skill requirements, which increases the frequency of expert consultations and narrows the span of control. By contrast, when GenAI becomes sufficiently reliable ($h < h_2$), the correction-reduction effect dominates: expert time savings from fewer errors outweigh the incremental consultation burden from lower-skilled workers, resulting in a broader span of control. Notably, under worker-level augmentation, validation and correction are borne by experts, so the correction-reduction effect reduces expert demand, thereby widening the span of control. This contrasts sharply with worker-level automation, where this burden is borne by workers, so greater correction efficiency reduces worker demand and narrows the span of control.

### 5.2. Worker-Level Comparison: Automation vs. Augmentation

Having characterized worker-level automation and augmentation, we now compare the two deployment modes and summarize their implications for entry-level skill requirements and hierarchy. The central difference lies in how the two deployment modes change the knowledge required at the entry level. Under automation, GenAI reduces the amount of entry-level labor needed for task execution, which raises the marginal value of worker expertise: firms respond by hiring fewer but more knowledgeable junior workers who can validate AI outputs and limit costly upward referrals. Under augmentation, by contrast, GenAI partially substitutes for workers' problem-solving

capability during execution, allowing firms to relax entry-level knowledge requirements and keep less-skilled junior workers productive with AI augmentation. This mode-dependent logic offers a parsimonious interpretation of emerging evidence: for example, Brynjolfsson et al. (2025a) document that early-career employment declines are concentrated in occupations where GenAI is used in more automation-oriented ways, whereas entry-level employment changes are muted (and can even turn positive) in more augmentative settings. Together, the theory and evidence suggest that the observed decline in junior jobs is not an inevitable consequence of GenAI per se, but is more consistent with deployment choices that push GenAI toward automation rather than augmentation.

Despite their opposite implications for entry-level skill requirements, the two deployment modes share a common and counterintuitive prediction for hierarchy: as GenAI advances, the optimal span of control can exhibit a U-shaped trajectory along both the capability and reliability dimensions. Under automation, the span of control initially narrows because expert time remains necessary for complex exceptions, while entry-level labor demand falls as the automated scope expands or as higher reliability reduces rework frictions. As GenAI matures, stronger indirect upskilling curbs upward referrals and compresses expert demand, widening the span of control. Under augmentation, early-stage progress encourages deskilling that increases reliance on experts, contracting the span of control; once GenAI becomes sufficiently capable and reliable, firms rely less on experts for consultation and oversight, and the span of control expands. Importantly, this early-stage contraction under either deployment mode is consistent with empirical evidence that junior employment can contract while demand for senior expertise remains stable or even rises (Brynjolfsson et al. 2025a, Hosseini and Lichtinger 2025).

Finally, the two deployment modes imply different long-run endpoints. Under sufficiently capable and reliable automation, extreme upskilling can render expert oversight redundant, potentially collapsing the organization toward a single-layer structure. Under augmentation, by contrast, hierarchy persists: AI-augmented workers still require expert validation and correction when errors arise, so augmentation yields flatter but enduring hierarchies.

## 6. Expert-Level GenAI Adoption

Having examined how GenAI reshapes the organizational foundation through worker-level adoption, we now turn to deployment at the expert layer. In traditional knowledge-based hierarchies, senior experts primarily function as exception handlers rather than frontline executors, serving as repositories of advanced knowledge that support junior workers in resolving problems beyond their expertise. Because experts' core role is to provide support rather than to perform routine tasks directly, our analysis focuses on how GenAI alters this support mechanism.

We consider two distinct adoption pathways. Under *expert-level automation*, GenAI operates as an autonomous knowledge source that substitutes for human experts in responding to workers'

queries. Under *expert-level augmentation*, GenAI acts as a capability multiplier that enhances the efficiency of human experts in assisting junior workers.

### 6.1. Expert-Level Automation

We model expert-level automation (denoted by subscript $e$) by treating GenAI as an autonomous knowledge repository with knowledge range $[0, r_e]$, where $0 < r_e < 1$ captures GenAI's automation capability. When a worker with intrinsic knowledge level $x$ faces a problem of difficulty $d > x$, the worker turns to the expert layer. For problems with $d \in (x, r_e]$, GenAI provides guidance; for problems with $d \in (r_e, 1]$, GenAI fails and human experts supply the solutions through the traditional communication channel. A key distinction from worker-level augmentation is that the difficulty range that can be resolved without human-expert involvement is determined entirely by GenAI capability $r_e$, independent of the worker's intrinsic knowledge level $x$ (see Figure EC.2).

Because enterprise GenAI systems remain probabilistic, we incorporate hallucination risk through a hallucination rate $h$. To mitigate potential losses from authoritative but incorrect outputs, we assume the firm mandates human-in-the-loop validation for all AI-guided work. Specifically, whenever a worker completes a task with GenAI guidance, a human expert reviews the output and intervenes if an error is detected. Accordingly, experts spend $t_v$ per AI-guided task for validation and, with probability $h$, incur an additional $t_c$ to communicate with the worker to correct hallucination-induced errors (detailed profit functions are provided in Appendix A.3).

Lemma 4. *The firm should adopt expert-level automation if and only if the hallucination rate $h < \overline{h}_e$ and GenAI's automation capability $r_e > \overline{r}_e$, where $\overline{h}_e = 1 - \frac{t_v}{t_c}$ and $\overline{r}_e = \frac{(k+2w)(ht_c+t_c+t_v)}{4k}$.*

Lemma 4 characterizes the adoption conditions for expert-level automation. Adoption is contingent upon two thresholds: a reliability threshold $h < \overline{h}_e$ and a capability threshold $r_e > \overline{r}_e$. The reliability condition reflects an efficiency requirement: AI-mediated support necessarily triggers expert validation and potential hallucination correction; hence, the expected per-task overhead $t_v + ht_c$ must be strictly lower than the cost of direct expert consultation $t_c$. Otherwise, routing tasks through GenAI fails to generate time savings and the firm optimally refrains from adoption. The capability condition reflects a scale requirement: if GenAI's knowledge coverage is limited, a substantial mass of tasks still lies beyond its capability range and continues to require direct expert intervention. In that case, GenAI adds an intermediate step without materially reducing expert workload. Only when GenAI covers a sufficiently broad portion of the task space does the reduction in direct referrals render adoption profitable.

Proposition 5. *The optimal worker knowledge level under expert-level automation is $x_e^* = \frac{(k+2w)(t_v+ht_c)}{2k}$. Furthermore, $x_e^*$ is independent of GenAI's automation capability $r_e$ but increases in hallucination rate $h$, and $x_e^* < x_0^*$.*

Proposition 5 characterizes the firm's optimal choice of worker knowledge under expert-level automation. The firm reduces the knowledge requirement for frontline workers, implying a deskilling effect ($x_e^* < x_0^*$). This contrasts sharply with worker-level automation, which drives upskilling (Proposition 1). Intuitively, once expert knowledge can be accessed through GenAI, the firm can economize on worker knowledge (and thus wages) while maintaining problem-solving capacity via AI-mediated support.

A striking feature is that the optimal worker knowledge level is independent of GenAI's automation capability $r_e$. This capability insensitivity is unique to expert-level automation and has no parallel in worker-level adoption, where capability enhancements continuously reshape skill requirements (Propositions 1 and 3). To see why, note that lowering $x$ by one unit shifts one unit of tasks from worker-independent execution to AI-guided execution requiring expert validation. The wage savings from this deskilling depend on the current skill level but not on $r_e$; the added validation cost per shifted task, $t_v + ht_c$, is likewise independent of $r_e$ because each AI-guided task incurs the same oversight cost regardless of how broad the AI-covered region is. Since neither side of the marginal trade-off involves $r_e$, the optimal $x_e^*$ is invariant to capability. By contrast, reliability directly enters the trade-off: a lower hallucination rate $h$ reduces the per-task validation cost $t_v + ht_c$, lowers the marginal cost of deskilling, and therefore decreases $x_e^*$.

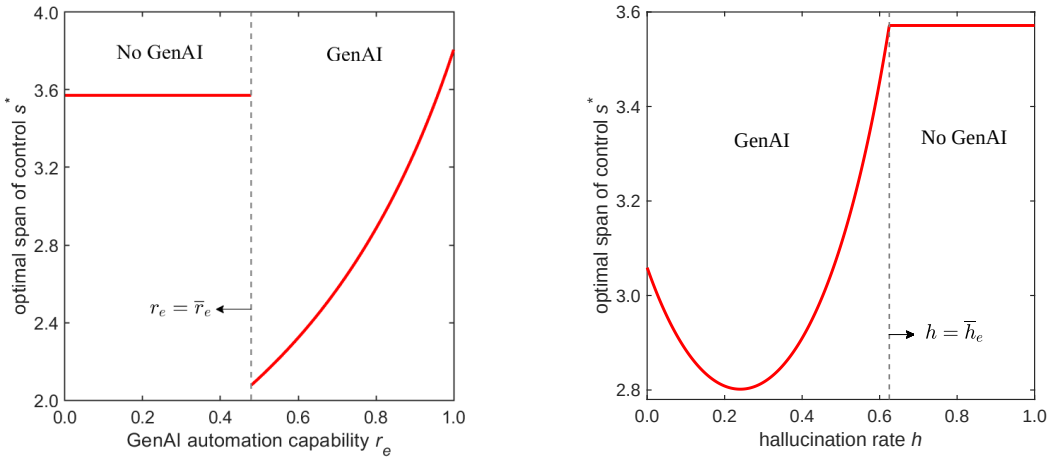
PROPOSITION 6. *After the adoption of expert-level automation, the optimal span of control is* $s_e^* = \frac{1}{(1-r_e)t_c + (r_e - x_e^*)(t_v + ht_c)}$.

(*a*) $s_e^*$ *exhibits a discontinuous decline at* $r_e = \overline{r}_e$ *and strictly increases in GenAI's automation capability* $r_e$ *on* $(\overline{r}_e, 1)$.

(*b*) $s_e^*$ *decreases in hallucination rate* $h$ *on* $(0, h_3)$ *and increases in* $h$ *on* $(h_3, \overline{h}_e)$.

Proposition 6 characterizes the organizational implications of expert-level automation. Along the capability dimension, the key feature is a discontinuity at adoption: adopting GenAI induces a discrete drop in worker knowledge level, increasing the mass of AI-guided tasks that require expert oversight and then causing a sudden decline in the span of control (Figure 5a). Conditional on adoption, since $x_e^*$ is insensitive to $r_e$, higher GenAI capability mainly shifts tasks from direct worker-expert consultation to AI guidance. Under the low hallucination rate ($h < \overline{h}_e$), the resulting reduction in direct consultation time dominates the added oversight burden, lowering expert demand and making $s_e^*$ strictly increase in $r_e$.

Along the reliability dimension, Proposition 6(b) implies a non-monotonic response driven by two forces. The direct correction-reduction effect lowers experts' expected correction time as $h$ falls, which tends to expand the span of control. Offsetting this, the indirect organization-deskilling effect strengthens as reliability improves, inducing lower $x_e^*$ and increasing the mass of AI-guided tasks




(a) $h = 0.1$ (b) $r_e = 0.8$

**Figure 5 Impact of GenAI advancements on span of control under expert-level automation ($k = 0.8$, $w = 0.25$, $t_c = 0.8$, $t_v = 0.3$)**

that require expert validation. When hallucinations are high, the deskilling channel dominates and $s_e^*$ decreases as $h$ falls; when hallucinations are low, the correction-reduction channel dominates and $s_e^*$ increases with further reductions in $h$. This mechanism parallels the reliability comparative statics under worker-level augmentation.

### 6.2. Expert-Level Augmentation

Finally, we examine expert-level augmentation (denoted by subscript $u$). In a knowledge-based hierarchy, senior experts primarily create value by supplying solutions when frontline workers face problems beyond their expertise, making worker-expert consultation time the key bottleneck. Expert-level augmentation captures the use of GenAI to streamline this consultation process rather than replacing experts as the solution provider. For example, GenAI can retrieve relevant information, draft proposed solutions, or structure recommendations, enabling experts to respond more quickly.

To formalize this setting, we model expert-level augmentation as a reduction in consultation time from $t_c$ to $(1 - r_u)t_c$, where $r_u \in (0, 1)$ measures GenAI's augmentation capability. This proportional reduction implies an equalizing effect, where GenAI delivers the greatest time savings for consultations that are initially the most costly. However, this efficiency gain is subject to hallucination risk $h$. When hallucination occurs, the expert must revert to the manual process, incurring the original time cost $t_c$. Therefore, the expected expert time per consultation becomes $(1 - r_u)t_c + ht_c$ (the detailed profit function is provided in Appendix A.4).

LEMMA 5. *The firm should adopt expert-level augmentation if and only if the hallucination rate* $h < r_u$.

The firm adopts expert-level augmentation if and only if the expected consultation time under GenAI is lower than the pre-GenAI benchmark, yielding the threshold $h < r_u$. The threshold highlights a capability-accuracy trade-off: greater augmentation capability expands the set of hallucination rates for which adoption is profitable; equivalently, a lower hallucination rate renders adoption viable even when augmentation capability is modest. The adoption criteria differ qualitatively across the four deployment architectures: capability can relax the reliability requirement (expert-level augmentation), tighten it (worker-level automation), or leave it unchanged (worker-level augmentation; expert-level automation), and may itself need to clear a minimum threshold (expert-level automation). This heterogeneity implies that there is no one-size-fits-all adoption rule; firms must assess adoption in a deployment-specific manner.

PROPOSITION 7. *After the adoption of expert-level augmentation:*

(*a*) *The optimal worker knowledge level is* $x_u^* = \frac{(k+2w)(1-r_u+h)t_c}{2k}$*, which decreases in GenAI's augmentation capability* $r_u$*, increases in hallucination rate* $h$*, and satisfies* $x_u^* < x_0^*$.

(*b*) *The optimal span of control is* $s_u^* = \frac{1}{(1-x_u^*)[(1-r_u)t_c+ht_c]}$*:*

(1) $s_u^*$ *decreases in GenAI's augmentation capability* $r_u$ *on* $(0, r_4)$ *and increases in* $r_u$ *on* $(r_4, 1)$.

(2) $s_u^*$ *decreases in hallucination rate* $h$ *on* $(0, h_4)$ *and increases in* $h$ *on* $(h_4, \overline{h}_u)$.

Conditional on adoption, expert-level augmentation induces a deskilling effect: the firm optimally lowers the knowledge requirement for frontline workers, because augmentation reduces the marginal cost of relying on experts for exception handling. Moreover, as higher augmentation capability and lower hallucination rates further reduce the expected consultation time, the incentive to deskill becomes stronger (Proposition 7(a)). Proposition 7(b) implies a non-monotonic effect of GenAI advancements on organizational structure. Although augmentation ultimately allows each expert to serve more workers by lowering the effective consultation time, early-stage improvements in GenAI can initially narrow the span of control. The reason is that deskilling increases the volume of referrals to experts, which may dominate the per-referral efficiency gain when GenAI's augmentation capability is still moderate or the hallucination rate remains high. As GenAI becomes sufficiently capable and reliable, the direct reduction in consultation time dominates, reducing total expert workload and expanding the span of control, producing a U-shaped trajectory.

### 6.3. Comparison and Cross-Layer Implications

The organizational consequences of expert-level GenAI adoption highlight that deployment location can be as consequential as deployment mode. Although worker-level and expert-level implementations rely on the same underlying technology, their interactions with knowledge-based hierarchies generate distinct organizational signatures.

**Automation Asymmetry:** Expert-level automation exhibits a clear cross-layer reversal relative to worker-level automation. At the worker level, automation drives upskilling by shifting junior roles toward validation and escalation control. At the expert level, however, automation induces deskilling. By providing an autonomous substitute for human-expert consultation, expert-level automation lowers the marginal value of worker knowledge and allows firms to economize on entry-level wages while preserving problem-solving capacity through AI-mediated support.

A further distinctive feature is post-adoption capability insensitivity. Under worker-level automation, capability improvements expand the scope that workers must validate and therefore continuously push skill requirements upward. Under expert-level automation, by contrast, higher capability mainly reallocates queries between GenAI and human experts and does not change the firm's optimal worker-knowledge choice within the AI-covered region. As a result, once adoption occurs, the optimal worker knowledge level is invariant to further capability gains.

**Augmentation Symmetry:** In contrast to the asymmetries of automation, augmentation produces qualitatively symmetric effects across both layers of the hierarchy. Whether it expands workers' effective capability boundary or accelerates experts' handling of consultations, it consistently lowers the effective cost of accessing high-level expertise. This leads to a uniform relaxation of frontline knowledge requirements, generating a deskilling effect that intensifies as the technology advances.

Taken together, expert-level adoption delivers a clean first-order skill prediction: both automation and augmentation deskill the frontline, lowering the entry-level skill threshold. The adjustment paths differ: automation produces a discrete drop in skill requirements followed by capability insensitivity, while augmentation generates a more gradual, capability-sensitive response. Both modes also imply similar hierarchical dynamics: direct efficiency gains and endogenous deskilling jointly drive a non-monotonic span of control, contracting initially before expanding as GenAI advances.

## 7. Conclusion and Managerial Implications

This paper develops an analytical framework showing that GenAI's organizational consequences depend critically on deployment design: not just whether the technology is adopted, but how and where. Building on knowledge-based hierarchy theory, the model incorporates two defining features of GenAI: intrinsic fallibility, which necessitates continuous human validation, and deployment flexibility, which allows GenAI to be integrated across both deployment modes and hierarchical layers. The resulting predictions challenge conventional wisdom about technology-induced organizational change while offering practical guidance for stakeholders navigating the GenAI transformation.

### 7.1. Summary of Key Findings

Our analysis generates three main findings, summarized in Table 2. First, GenAI's impact on entry-level skill requirements is mode-dependent. Under worker-level automation, GenAI takes over routine execution but introduces a validation burden due to intrinsic fallibility; the firm responds by hiring fewer but more knowledgeable workers who can verify AI outputs and reduce costly upward referrals to senior experts. Under worker-level augmentation, GenAI instead operates as a capability amplifier that extends workers' effective problem-solving boundaries; the firm responds by relaxing frontline knowledge requirements and sustaining performance through AI-assisted workflows. The observed decline in entry-level employment is therefore not an inevitable consequence of GenAI per se, but reflects deployment choices that tilt GenAI toward automation rather than augmentation.

Second, deployment location can be as consequential as deployment mode. Expert-level GenAI adoption uniformly lowers entry-level skill requirements regardless of whether GenAI serves as an autonomous knowledge source (automation) or a consultation accelerator (augmentation). In both cases, expert-level deployment reduces the marginal value of frontline knowledge by providing an alternative channel for resolving problems beyond workers' expertise, allowing firms to economize on entry-level wages while preserving problem-solving capacity through AI-mediated or AI-enhanced expert support. This uniform deskilling effect stands in sharp contrast to the mode-dependent outcomes at the worker level.

Third, organizational structure evolves non-monotonically as GenAI advances. Across all four deployment architectures, spans of control first decrease and then increase, driven by the shifting dominance between direct efficiency gains and indirect workforce restructuring. This pattern implies that demand for senior expertise may remain stable or even increase during early-to-intermediate stages of deployment, even before hierarchies ultimately flatten, a prediction consistent with recent large-scale evidence on GenAI's seniority-biased labor market effects.

**Table 2 Predicted organizational impacts across GenAI deployment architectures**

| | **Automation Mode** | **Augmentation Mode** |
|---|---|---|
| **Worker Level** | Task executor<br>**Upskilling** ($x_t^* > x_0^*$)<br>Capability-sensitive<br>U-shaped span of control | Capability amplifier<br>**Deskilling** ($x_g^* < x_0^*$)<br>Capability-sensitive<br>U-shaped span of control |
| **Expert Level** | Virtual supervisor<br>**Deskilling** ($x_e^* < x_0^*$)<br>Capability-insensitive<br>Initial drop, then increase | Communication accelerator<br>**Deskilling** ($x_u^* < x_0^*$)<br>Capability-sensitive<br>U-shaped span of control |

Notably, we do not assume that firms can frictionlessly choose a single optimal GenAI architecture in the abstract. In practice, deployment location and mode are shaped by constraints such

as compliance regimes, liability and risk tolerance, legacy systems, and the existing distribution of expertise. Our goal is therefore to provide a conditional design guide: given a firm's choice of where GenAI is embedded in the hierarchy and whether it is used primarily for automation or augmentation, the model characterizes how skill thresholds and spans of control adjust as AI capability and fallibility change. This structure yields transparent "if–then" prescriptions that firms can use to anticipate second-order organizational consequences of a chosen deployment strategy.

### 7.2. Managerial and Policy Implications

For firms, our analysis offers a strategic framework for aligning technology procurement with human capital planning. The decision to adopt GenAI cannot be decoupled from hiring strategy. First, firms deploying GenAI for *automation* (e.g., autonomous coding agents or customer service bots) face a counterintuitive skill imperative: the optimal response is to raise, not lower, entry-level knowledge requirements. Under automation, the worker's primary role shifts from execution to validation of AI outputs. Firms should therefore raise technical qualification requirements for entry-level roles, ensuring that workers possess the domain expertise to detect subtle AI hallucinations. Failing to upskill the intake pipeline risks a proliferation of undetected errors or an overwhelming escalation load on senior experts.

Conversely, firms prioritizing *augmentation* (e.g., co-pilots for writing or research) can broaden their recruitment funnels. By deploying GenAI as a capability amplifier, firms can hire candidates who possess strong foundational traits but lack specialized technical mastery, as the AI tool bridges the execution gap. This suggests a strategic pivot for HR: under augmentation strategies, hiring criteria should shift from testing for peak technical execution to assessing adaptability and human-AI collaborative aptitude.

Finally, firms should anticipate that middle-management demand will not shrink in the near term and may even increase during early-to-intermediate stages of GenAI deployment. Organizations that preemptively flatten their hierarchies in anticipation of future efficiencies risk creating significant bottlenecks in exception handling and quality control during the transition period.

For junior workers, our model speaks directly to career strategy. Under worker-level *automation*, the returns to validation skills (the ability to detect and correct AI errors) rise sharply, because the hallucination risk makes human verification the binding constraint on AI-automated production. Entry-level workers in automation-heavy environments should invest in domain expertise sufficient to evaluate AI outputs, not merely to execute tasks that AI can perform. This contrasts with the conventional advice to develop AI-complementary soft skills: our model suggests that deep technical knowledge becomes more valuable, not less, precisely because someone must verify what the machine produces. Under worker-level *augmentation*, the strategic calculus

reverses. Because augmentation extends workers' effective capability boundary, the marginal return to intrinsic knowledge declines. Workers in augmentation-heavy environments face a different risk: skill atrophy, as reduced demand for deep expertise may erode the human capital that currently differentiates them from less experienced peers (Siderius et al. 2026).

For senior experts, our model delivers a nuanced career implication. The near-term outlook is more stable than commonly feared: across all deployment modes, exception handling, validation, and oversight demands sustain or even expand expert workloads during early-to-intermediate stages of GenAI deployment, even as frontline execution becomes more efficient. As GenAI matures and direct efficiency gains accumulate, the demand for senior positions will eventually moderate, but the remaining roles do not vanish. They recenter on validation, error-correction, and judgment-intensive oversight—tasks that remain resistant to full automation. In sectors where automation reaches sufficient maturity to flatten the hierarchy substantially, senior experts can sustain their relevance by evolving beyond pure exception handling toward a hybrid role that combines high-skill execution with AI supervision, effectively serving as elite task performers who also bear responsibility for monitoring and correcting AI-generated outputs.

For policymakers, our findings refine the policy debate beyond the binary framing of "AI displaces workers" versus "AI augments workers." Two policy-relevant insights emerge. First, the distributional consequences of GenAI depend critically on deployment architecture, not just adoption intensity. The observed concentration of entry-level employment declines in automation-intensive settings (Brynjolfsson et al. 2025a) is not an inevitable consequence of GenAI but reflects specific, and potentially redirectable, deployment choices. Policies that incentivize augmentation-oriented development (Acemoglu et al. 2023) are well-targeted from our model's perspective, but our results add the nuance that expert-level deployment of any mode similarly lowers the entry-level skill threshold. Second, training programs should be mode-specific. Our model implies that generic "AI literacy" curricula may be insufficient: workers entering automation-heavy roles need deep domain expertise for validation, while those entering augmentation-heavy roles need adaptive problem-solving skills to leverage AI-expanded capability boundaries.

### 7.3. Limitations and Future Research

While our model provides a general framework for GenAI adoption in knowledge-based hierarchies, its predictions are sharpest in industries characterized by verifiable correctness and high error costs. Our mechanism hinges on the concept of hallucination, an objectively incorrect output that requires remediation. This dynamic maps closely to domains such as software engineering, legal services, and financial analysis, where outputs are functionally binary (correct vs. buggy; legally sound vs. invalid) and the cost of deploying an error is non-trivial. In these settings, the trade-off between execution efficiency and validation rigor is the binding constraint on organizational

structure. Conversely, in creative industries (e.g., graphic design or marketing copy generation), where hallucination may be interpreted as creativity and validation is subjective, the strict up-skilling/deskilling dynamics we predict may be less pronounced.

Several directions for future research emerge naturally from our framework. First, extending the analysis to multi-tier hierarchies could reveal how GenAI affects intermediate layers, particularly middle management roles that combine execution and supervision. Second, incorporating dynamics, such as learning curves, path dependence, and transition costs, would characterize whether early adoption patterns (e.g., entry-level hollowing under automation) are self-correcting or self-reinforcing as GenAI matures. Third, embedding our single-firm analysis in a competitive labor-market equilibrium would characterize industry-wide wage and employment effects, addressing the general-equilibrium feedback that our partial-equilibrium framework deliberately abstracts from. Finally, integrating worker agency, including strategic skill investment, technology resistance, and behavioral responses such as automation complacency, would enrich the model's predictions about the co-evolution of human capital and AI capability.

## References

Acemoglu, D., D. Autor. 2011. Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics*, vol. 4. Elsevier, 1043–1171.

Acemoglu, D., D. Autor, S. Johnson. 2023. Can we have pro-worker AI? Choosing a path of machines in service of minds. CEPR Policy Insight 123, CEPR Press, Paris and London.

Acemoglu, D., S. Johnson. 2023. Rebalancing AI. *International Monetary Fund* .

Acemoglu, D., P. Restrepo. 2019. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives* **33**(2) 3–30.

Agrawal, A., J.S. Gans, A. Goldfarb. 2023. The turing transformation: Artificial intelligence, intelligence augmentation, and skill premiums. *NBER Working Paper 31767* .

Antràs, P., L. Garicano, E. Rossi-Hansberg. 2006. Offshoring in a knowledge economy. *The Quarterly Journal of Economics* **121**(1) 31–77.

Autor, D. 2024. Applying AI to rebuild middle class jobs. *NBER Working Paper 32140* .

Autor, D.H., L.F. Katz, A.B. Krueger. 1998. Computing inequality: have computers changed the labor market? *The Quarterly Journal of Economics* **113**(4) 1169–1213.

Autor, D.H., F. Levy, R.J. Murnane. 2003. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* **118**(4) 1279–1333.

Babina, T., A. Fedyk, A.X. He, J. Hodson. 2023. Firm investments in artificial intelligence technologies and changes in workforce composition. *NBER Working Paper 31325* .

Bartel, A., C. Ichniowski, K. Shaw. 2007. How does information technology affect productivity? Plant-level comparisons of product innovation, process improvement, and worker skills. *The Quarterly Journal of Economics* **122**(4) 1721–1758.

Bastani, H., O. Bastani, W.P. Sinchaisri. 2021. Improving human decision-making with machine learning. *arXiv preprint arXiv:2108.08454* .

Bastani, H., G. Cachon. 2025. The human-AI contracting paradox. *Available at SSRN 5962739* .

Bloom, N., L. Garicano, R. Sadun, J. Van Reenen. 2014. The distinct effects of information technology and communication technology on firm organization. *Management Science* **60**(12) 2859–2885.

Brynjolfsson, E., B. Chandar, R. Chen. 2025a. Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Working paper, Stanford Digital Economy Lab.

Brynjolfsson, E., D. Li, L. Raymond. 2025b. Generative AI at work. *The Quarterly Journal of Economics* **140**(2) 889–942.

Bubeck, S., V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, et al. 2023. Sparks of artificial general intelligence: Early experiments with GPT-4. *arXiv preprint arXiv:2303.12712* .

Caplin, A., D. Deming, S. Li, D. Martin, P. Marx, B. Weidmann, K.J. Ye. 2025. The ABCs of who benefits from working with AI: Ability, beliefs, and calibration. *Management Science (Articles in Advance)* .

Chen, Z., J. Chan. 2024. Large language model in creative work: The role of collaboration modality and user expertise. *Management Science* **70**(12) 9101–9117.

Cui, Z.K., M. Demirer, S. Jaffe, L. Musolff, S. Peng, T. Salz. 2024. The effects of generative AI on high skilled work: Evidence from three field experiments with software developers. *Available at SSRN 4945566* .

Dai, T., T. Taylor. 2026. Designing enterprise AI systems: Hallucination, creativity, and moral hazard. *Available at SSRN 5996714* .

Dell'Acqua, F., E. McFowland III, E.R. Mollick, H. Lifshitz-Assaf, K. Kellogg, S. Rajendran, L. Krayer, 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 Technology & Operations Mgt. Unit Working Paper* (24-013).

Fügener, A., J. Grahl, A. Gupta, W. Ketter. 2021. Will humans-in-the-loop become borgs? Merits and pitfalls of working with AI. *MIS Quarterly* **45**(3) 1527–1556.

Garicano, L. 2000. Hierarchies and the organization of knowledge in production. *Journal of Political Economy* **108**(5) 874–904.

Garicano, L., E. Rossi-Hansberg. 2006. Organization and inequality in a knowledge economy. *The Quarterly Journal of Economics* **121**(4) 1383–1435.

Garicano, L., E. Rossi-Hansberg. 2015. Knowledge-based hierarchies: Using organizations to understand the economy. *Annual Review of Economics* **7**(1) 1–30.

Goos, M., A. Manning. 2007. Lousy and lovely jobs: The rising polarization of work in britain. *The Review of Economics and Statistics* **89**(1) 118–133.

Gumpert, A. 2018. The organization of knowledge in multinational firms. *Journal of the European Economic Association* **16**(6) 1929–1976.

Hoffmann, M., S. Boysel, F. Nagle, S. Peng, K. Xu. 2024. Generative AI and the nature of work. *Harvard Business School Strategy Unit Working Paper* (25-021).

Hosseini, S.M., G. Lichtinger. 2025. Generative AI as seniority-biased technological change: Evidence from U.S. résumé and job posting data. *Available at SSRN 5425555* .

Hou, J., L. Wang, G. Wang, H.J. Wang, S. Yang. 2025. The double-edged roles of generative ai in the creative process: experiments on design work. *Information Systems Research (Articles in Advance)* .

Ibrahim, R., S.-H. Kim, J. Tong. 2021. Eliciting human judgment for prediction algorithms. *Management Science* **67**(4) 2314–2325.

Ide, E., E. Talamas. 2025. Artificial intelligence in the knowledge economy. *Journal of Political Economy* **133**(12) 3762–3800.

Kalai, A.T., S.S. Vempala. 2024. Calibrated language models must hallucinate. *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 160–171.

Katz, Lawrence F, Kevin M Murphy. 1992. Changes in relative wages, 1963–1987: supply and demand factors. *The Quarterly Journal of Economics* **107**(1) 35–78.

Kinder, M., X. de Souza Briggs, M. Muro, S. Liu. 2024. Generative AI, the American worker, and the future of work. *Brookings Institution* .

Liu, Y., C. Liu, T. Huang. 2024. Is turing transformation happening? A study with a new measure of occupational AI exposure. *Available at SSRN 4515629* .

Lu, T., Y. Zhang. 2025. 1+ 1>2? Information, humans, and machines. *Information Systems Research* **36**(1) 394–418.

Magesh, V., F. Surani, M. Dahl, M. Suzgun, C.D. Manning, D.E. Ho. 2024. Hallucination-free? Assessing the reliability of leading AI legal research tools. *arXiv preprint arXiv:2405.20362* .

Noy, S., W. Zhang. 2023. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* **381**(6654) 187–192.

Peng, S., E. Kalliamvakou, P. Cihon, M. Demirer. 2023. The impact of AI on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590* .

Powell, W.W., K. Snellman. 2004. The knowledge economy. *Annual Review of Sociology* **30**(1) 199–220.

Siderius, J., R.A. Shumsky, A. Taylor. 2026. Use it or slowly lose it: Expertise atrophy with organizational ai usage. Working paper, Tuck School of Business, Dartmouth College.

Wang, W., G. Gao, R. Agarwal. 2024. Friend or foe? Teaming between artificial intelligence and workers with variation in experience. *Management Science* **70**(9) 5753–5775.

Wu, Y. 2015. Organizational structure and product choice in knowledge-intensive firms. *Management Science* **61**(8) 1830–1848.

Zhang, X., Y. Dou, D.J. Wu. 2025. Incentivizing the content creator when tasks can be delegated to artificial intelligence. *Available at SSRN 5118060* .

Zhou, L., W. Schellaert, F. Martínez-Plumed, Y. Moros-Daval, C. Ferri, J. Hernández-Orallo. 2024. Larger and more instructable language models become less reliable. *Nature* **634**(8032) 61–68.