Rungs Removed or Scaffolds Built? A Dynamic Theory of Autonomy, Review, and Learning in Early-Career Labor Markets

Georgios Petropoulos
University of Southern California

June 8, 2026

Abstract

How a firm configures generative AI determines whether the technology develops or hollows out its junior talent. We build a dynamic equilibrium theory of early-career skill formation in which tasks generate both current output and future expertise and AI either performs work autonomously or assists a worker who reviews its output. Because reviewing AI output can itself teach, the effect of delegation on skill depends on how it is configured. Two results are central. First, whether AI raises or erodes a workforce's long-run expertise turns not on how capable the AI is but on how intensively humans review its output: above a critical review intensity the same technology permanently builds expertise, and below it permanently erodes it. Second, a competitive market will not choose the configuration that develops the workforce. Because the expertise a worker gains is general and portable, and review is hard to contract on, a reviewing firm bears the cost but captures only part of the return. It therefore favors high autonomy and thin review, and competition sharpens rather than disciplines the resulting under-provision of review. This hold-up in the formation of general human capital falls hardest at career entry, where the foregone tasks are precisely those a novice learns from. The remedy is configuration, not slower adoption: the optimal policy withholds autonomy from juniors until they can verify and grants it to seniors, so that autonomy builds expertise rather than removing rungs. AI deployment is thus a talent-development decision, to be set by task and seniority rather than as a uniform automation target.

Keywords: Generative AI, human–AI collaboration, human capital, autonomy, early-career workers, apprenticeship, workforce development, technology management.

1 Introduction

The first large-scale evidence on generative AI and jobs points at the bottom of the career ladder. Using payroll microdata, Brynjolfsson et al. (2025a) document a sizable relative decline in employment for workers in their early twenties in the most AI-exposed occupations (Eloundou et al., 2024), concentrated where AI automates rather than augments labor. Using résumé and job-posting data on roughly 62 million workers across 285,000 firms, Hosseini and Lichtinger (2025) find the same seniority-biased pattern within firms that adopt generative AI. Junior employment falls sharply while senior employment continues to rise but most importantly for what follows, AI-exposed tasks are increasingly *removed from junior job postings*: the very rungs a novice would have climbed. Converging evidence points the same way: an age gradient in generative-AI-era hiring, with the youngest workers in exposed occupations falling behind (Lodefalk et al., 2026), a sharp slowdown in the growth of coder employment, one of the most exposed occupations (Crane and Soto, 2026), and broader labor-market evidence on AI exposure (Klein Teeselink, 2025).

The worry this raises should be analyzed in a dynamic setting. Many of the entry-level tasks AI is most apt to take over (drafting a first memo, cleaning a dataset, writing routine code) have modest production value but high learning value: they are how a novice acquires the primitives of a domain. If AI removes precisely these tasks, current output can rise even as the future stock of expertise falls. We ask whether this harm is an intrinsic property of capable AI or a consequence of *how firms choose to deploy it*. If the latter, a subsequent question is whether a competitive labor market can be expected to deploy it well.

Our answer turns on a distinction the debate has largely collapsed: between whether a firm adopts AI and how it configures it. The same technology can be run with the human kept on the work and assisted by the machine, or with the work handed to the machine and the human verifying the result. The prevailing intuition and the structure of existing models treat the second arrangement as pedagogically dead: if the worker does not perform the task, the worker does not learn from it, so learning can only be protected by keeping the human hands-on. We consider a channel this logic ignores. When AI completes a task, the worker can still learn by reviewing and verifying its output, and the output under review is often better than the worker could have produced alone. Reviewing an expert-grade solution can transfer more of a domain's knowledge than performing a routine version of the task oneself, provided the worker can understand what they are checking.

This single change reorganizes the problem. Whether reviewing teaches more than doing depends on the worker's ability to verify what the machine produced, which grows with seniority, and on the quality of that output, which grows with the AI's capability. The effect

of handing work to the machine therefore reverses sign over a career: for raw juniors, at the very start of their career, who cannot yet evaluate expert output, autonomy strips away the rungs they would have climbed. But for workers senior enough to review capable AI, the same autonomy *builds* skill, and restricting it would slow their development. The lever that matters is thus not how much to automate, but how to pair automation with review, task by task and seniority by seniority.

The paper's central result is what this implies for a market and firms' strategies. A configuration that develops the workforce rather than depleting it exists, but a competitive labor market will not choose it. Because the expertise a junior gains is general and portable, and because how intensively AI output is reviewed cannot be written into an enforceable contract, a firm that invests in review reaps only part of the return while bearing all of the cost. Each firm therefore favors the configuration that maximizes current output, high autonomy, thin review, which is exactly the configuration that starves learning. Competition, far from disciplining this, sharpens it: the market under-provides review relative to the social optimum, most starkly in the limiting case where review has no current-output value and the competitive choice collapses to full autonomy with no review. The missing rung, on this account, is not a property of the technology but a *deployment externality*: a hold-up problem in how AI is configured, which disappears only when the deployment itself can be contracted upon. The binding lever is the governance of workforce development, not limits on adoption.

Two further results give the externality its force. First, whether AI raises or erodes a workforce's long-run expertise is not fixed by how capable the AI is but governed by a single margin the firm controls: above a critical review intensity the same technology permanently builds expertise, below it the identical autonomy permanently erodes it. Second, we microfound the review channel from verification, error detection, and explanation quality, making its conversion rate an output rather than a posit and identifying when reviewing teaches, fails to teach, or produces negative learning, including a region in which a junior held at high autonomy is trapped in a low-skill equilibrium. The microfoundation reinforces rather than threatens the headline: more capable AI can *reduce* expert learning by removing the errors experts learn most from, so the governable margin is review, not capability. The trap, moreover, is only a feature of bad deployment: under the optimal seniority-contingent rule (keep juniors on AI they review rather than AI that works for them, and hand over the controls only as they become able to check its work) it vanishes. Throughout, the deciding actor is a firm configuring AI across its workforce, so that AI configuration is a talent-development decision, not only a productivity one.

The framework is operational rather than only conceptual: every primitive maps to some-

thing a firm already logs or can cheaply elicit: the autonomy, review, and assist intensities sit in workflow telemetry, and a worker's ability to verify what the machine produced is recovered by seeded-error audits that plant known mistakes in AI output and record the detection rate by seniority. The resulting prescription is therefore defined over *(task, worker) cells* that a manager can actually populate, rather than over occupations or firms.

The remainder proceeds as follows. Section 2 places the paper in the literature and Section 3 sets up the environment and the learning technology. Section 4 establishes when autonomy teaches and derives the sign reversal. Section 5 solves the dynamic equilibrium, and Section 6 derives the deployment externality and the role of contractibility. Section 7 introduces the leading parametrization and calibrates. Section 8 microfounds the review channel and locates the conditions under which it teaches. Section 9 develops the managerial framework. Section 10 concludes. Appendices contain proofs and extensions.

2 Related Literature

Our paper contributes to 5 strands of literature:

Task-based technological change. In task-based models, technology reallocates tasks between labor and capital, and tasks generate current output (Autor et al., 2003; Acemoglu and Autor, 2011; Autor and Dorn, 2013; Goos, Manning, and Salomons, 2014; Acemoglu and Restrepo, 2018, 2019, 2020, 2022). Recent work brings AI into this frame, measuring task-level exposure and tracing it to within-firm labor demand (Hampole et al., 2025) and to firm growth and product innovation (Babina et al., 2024). We add a second value to each task, its learning content, and let the technology act on both. The novelty relative to this tradition is that a statically efficient deployment can be dynamically inefficient, and that the relevant policy unit is a task–seniority pair rather than an occupation. Our learning technology and its review channel are new to this literature. Our novel consideration implies that the augmentation–automation dichotomy is incomplete: a delegated task is not pedagogically dead, because reviewing the machine's output is itself a learning channel, which is why autonomy can *build* skill rather than only erode it.

Knowledge hierarchies and AI autonomy. A long tradition models how a given distribution of expertise is organized into hierarchies of producers and problem-solvers (Garicano, 2000; Garicano and Rossi-Hansberg, 2006). Within it, Antràs et al. (2006) show how heterogeneous agents sort into hierarchical teams, with less skilled agents specializing in production and more skilled in problem-solving as well as how communication technology reshapes that

assignment. Bloom et al. (2014) document that information and communication technologies have *distinct* effects on firm organization, the former pushing problem-solving authority down the hierarchy and the latter up and Fuchs et al. (2015) characterize the contracts that implement the efficient knowledge hierarchy, showing a firm-like arrangement is required where spot contracting fails. Ide and Talamàs (2025) bring AI into this frame, with agents that operate autonomously or only as copilots, and show autonomous AI benefits the most knowledgeable while nonautonomous AI benefits the least. Ide and Talamàs (2026) extend this to global knowledge work, mapping a basic-advanced capability distinction that parallels the capability margin in our model. Two concurrent papers are closest, and the contrast with each is sharp. Xu et al. (2025) study how generative AI reorganizes a knowledge-work hierarchy across the same design space we use, automation versus augmentation crossed with the worker versus the expert layer, with fallible AI output that must be verified. Their central finding is that as AI improves the firm optimally *lowers the knowledge requirement* of the worker role (deskilling) and may *narrow* its span of control. Friebel et al. (2026) model how productivity and learning shocks move a firm's span of control, generating transitions between pyramid and diamond structures and oscillating inequality across junior cohorts. We share with both the premise that AI is simultaneously autonomous and fallible and that verification is central. We differ on the object and, decisively, on the normative conclusion. Both take AI's effect on skill as a *primitive*, a deployment that fixes the role's knowledge requirement in Xu et al. (2025), an exogenous "faster-learning" shock in Friebel et al. (2026), and ask how the *organization* (composition, layers, span) adjusts around it. We endogenize the skill itself: reviewing AI output is the technology by which a worker's own expertise accumulates over a career, so AI configuration acts on the *production* of the future expertise distribution rather than the *allocation* of the current one, and the same act that delivers reliable output today produces tomorrow's skilled worker.

From financing training to providing learning exposure. Training theory, since Becker (1964), asks how investment in transferable skill is financed. Because such skill can walk out the door, the question is who pays for it and how the return is appropriated. A large literature shows that labor-market frictions and private information can make firms willing to finance even general skills (Acemoglu and Pischke, 1998; Katz and Ziderman, 1990; Stevens, 1994; Loewenstein and Spletzer, 1998), and that apprentice contracts and firm commitment can sustain training where spot markets cannot (Malcomson et al., 2003; Dustmann and Schönberg, 2012). Garicano and Rayo (2026) bring AI into this tradition with a dynamic apprenticeship model in which careers remain viable when an *expertise leverage ratio*, the value of a trained graduate relative to AI's standalone output, exceeds a critical threshold,

and otherwise training compresses toward collapse. This entire tradition, however, takes the *learning technology as given* and asks whether its returns can be financed and appropriated. We ask a logically prior question that AI makes pressing: does the work environment still *deliver the learning* to be financed? Our object is therefore a *learning* viability condition: does the junior receive enough learning exposure to accumulate expertise? We do not ask: can the senior profit from training the junior? In Appendix B.1 we show that the learning and the financial viability conditions are logically independent and can move in opposite directions, so each captures a failure the other misses, and we view our contribution and the financing literature as complementary halves of a single requirement for robust ladders.

Expertise, human capital and early-careers. The paper connects to human-capital accumulation and early-career wage growth (Ben-Porath, 1967; Mincer, 1974; Topel and Ward, 1992; Rubinstein and Weiss, 2006), to the internal wage and promotion structure of firms (Baker et al., 1994), to early-career scarring (Kahn, 2010) and to the role of expertise in shielding workers from automation (Autor and Thompson, 2025). That the early career is where the stakes concentrate is underscored by Deming (2023), who shows that nearly all life-cycle wage growth occurs *within* jobs and that it is sorting into complex, learning-intensive work that drives the steeper trajectories of educated workers, so a technology that strips the learning content from entry-level work bears directly on lifetime outcomes. In this tradition the slope of the early-career wage profile reflects the accumulation of human capital and the gradual revelation of ability, a property of the worker and the job. Our contribution makes that slope a function of deployment: because reviewing AI output can build skill or hollow it out depending on configuration, the same technology can steepen or flatten the early-career profile even when entry wages or current output rise, acting through the production of skill rather than its pricing. That the accumulated expertise is largely general—portable across employers (Gathmann and Schönberg, 2010) is precisely why the firm cannot appropriate its return and privately underprovides it, tying the early-career wage profile directly to the deployment externality. Methodologically, the paper builds on the career-dynamics tradition linking human-capital accumulation and employer learning to wage and promotion paths (Farber and Gibbons, 1996; Gibbons and Waldman, 1999, 2006). We share its general dynamic structure but add the channel through which AI configuration governs the rate of accumulation itself.

Human–AI collaboration and delegation. A body of work studies how organizations combine human and machine effort. Shrestha et al. (2019) characterize how decisions can be fully delegated to algorithms, split sequentially between human and machine, or aggregated

across them. Fügener et al. (2022) show experimentally that effective delegation is bounded by humans' *metaknowledge*, their ability to judge what they themselves do not know. On the adoption side, Kawaguchi (2021) studies when workers actually follow algorithmic recommendations and Parasuraman and Manzey (2010) documents the monitoring failures that arise when a human is nominally kept "in the loop" over automated output. Field evidence on generative AI at work (Brynjolfsson et al., 2025b; Cui et al., 2026; Noy and Zhang, 2023; Dell'Acqua et al., 2023; Peng et al., 2023) shows large contemporaneous productivity gains from such collaboration, though Dell'Acqua et al. (2023) also document a "jagged frontier": AI raises quality and speed on tasks within its capability but degrades performance on tasks beyond it, a capability boundary that workers cannot easily locate. Dell'Acqua et al. (2025) push this further, showing in a field experiment that an AI "teammate" lets less-experienced workers reach performance levels that previously required collaboration with or supervision by more senior colleagues. Closest to our concern, Fügener et al. (2021) find that AI advice raises each worker's current accuracy while *eroding* the unique knowledge that distinguishes workers from one another, collaboration that improves output today can deplete human capability. This literature, however, measures that depletion contemporaneously and treats the human's role as an input to *current* performance. We add two things it does not contain. First, the same configuration choice that optimizes current human–AI performance (e.g, how much to automate, how intensively to review, how much to assist), also sets the rate at which the worker *accumulates* expertise over a career, so the configuration is simultaneously a performance lever and a workforce-development lever, and the two need not point the same way. Second, and most importantly, because the skill is portable and review is non-contractible, a competitive labor market selects too much autonomy and too little review, a market failure that single-firm design frameworks cannot see.

3 Environment

We begin by setting up the environment and learning technology in a form chosen for closed-form tractability, deferring to Section 8 a microfoundation of the learning by reviewing AI output channel on which the central results turn.

Time is continuous, $t \geq 0$. We describe a single worker's career. The worker is summarized by human capital $h \in [0, 1]$, interpreted as seniority, judgment, and domain expertise. A worker of type h is assigned a flow of tasks indexed by complexity z, drawn from a distribution $F(\cdot \mid h)$ with density $f(\cdot \mid h)$. We assume task assignment shifts toward harder problems with seniority:

Assumption 1 (Seniority)**.** *$F(\cdot \mid h)$ is ordered by first-order stochastic dominance in h: for $h' > h$, $F(z \mid h') \leq F(z \mid h)$ for all z.*

Each task carries learning content $\ell(z) \geq 0$ which is interpreted as the *marginal, transferable* increment to general human capital a worker obtains by engaging with a task of complexity z.

The (generative AI) technology is described by the following objects:

- *Capability* $\theta \in [0, 1]$ is the complexity frontier: AI can complete tasks with $z \leq \theta$.
- *Output quality* is $A(\theta)$, with $A' > 0$ and $A(0) = 0$.
- *Autonomy* $a \in [0, 1]$ is the share of AI-capable tasks the system completes without the worker performing them.
- *Review intensity* $r \in [0, 1]$ is how intensively the worker verifies autonomously produced output.
- *Copilot intensity* $m \geq 0$ is the assistance the worker receives on tasks they perform themselves.

A capable task ($z \leq \theta$) is either delegated autonomously (fraction a) or kept in copilot mode (fraction $1 - a$). Tasks beyond the frontier ($z > \theta$) are always performed by the worker.

The worker absorbs knowledge from reviewing AI output only to the extent they can evaluate it. *Verification ability* is $p(h)$, with $p' > 0$ and $p(0) = 0$. A novice who cannot tell a correct solution from a plausible-looking one learns little from reviewing. An expert extracts the reasoning.

Performing a task teaches by doing. Reviewing a completed task teaches by example. We posit three per-task learning increments:

- **Hands-on, beyond the frontier** ($z > \theta$): the worker performs the task and learns $\alpha\, \ell(z)$.
- **Copilot** (capable task, performed with assistance): the worker performs the task and the assistant amplifies comprehension, yielding $\alpha\, \ell(z)\big(1 + m\, A(\theta)\, p(h)\big)$.
- **Autonomous + review** (capable task, delegated and verified): the worker reviews an expert-grade completion and absorbs $\beta\, r\, A(\theta)\, p(h)\, \ell(z)$.

Here $\alpha > 0$ is the hands-on conversion rate and $\beta > 0$ the review conversion rate. The economically substantive premise is $\beta > \alpha$: *per task, reviewing an expert-grade solution can transfer more knowledge than producing a routine version oneself.* This is what makes the review channel capable of dominating hands-on learning, and it is the formal content of the "virtual apprenticeship" intuition.

This agrees with the *worked-example effect*, a widely replicated finding in the science of learning: for a learner with the prerequisites to follow it, studying a fully worked expert solution yields better subsequent *independent* performance, at lower time and mental effort, than solving the equivalent problem unaided (Sweller and Cooper, 1985; Cooper and Sweller, 1987; Sweller et al., 1998), with a recent meta-analysis confirming sizable retention and transfer gains in mathematics (Barbieri et al., 2023). Three mechanisms supply the premium that β aggregates: the active processing of a model solution (the *self-explanation effect*, Chi et al., 1989); the high transfer value of detecting and correcting errors in worked output (error-management training raises complex-transfer performance with a large effect size, Keith and Frese, 2008); and the observation of expert performance a novice cannot yet produce, which is the mechanism of apprenticeship (Collins et al., 1989).

Aggregating over the task distribution, we define the capable and frontier learning masses as integrals over the capable and beyond-frontier regions of the support of $F(\cdot \mid h)$:

$$\mathcal{L}_\theta(h) \equiv \int_{\{z \le \theta\}} \ell(z)\, \mathrm{d}F(z \mid h), \qquad \bar{\mathcal{L}}_\theta(h) \equiv \int_{\{z > \theta\}} \ell(z)\, \mathrm{d}F(z \mid h). \tag{1}$$

The *measure* of tasks with complexity below the AI frontier (capable tasks without weighting them by learning content) is:

$$Q_\theta(h) \equiv F(\theta \mid h) = \int_{\{z \le \theta\}} \mathrm{d}F(z \mid h). \tag{2}$$

The gross learning rate is

$$\Phi(h; \theta, a, r, m) = \alpha\, \bar{\mathcal{L}}_\theta(h) + (1 - a)\, \alpha\big(1 + mA(\theta)p(h)\big)\mathcal{L}_\theta(h) + a\, \beta\, r\, A(\theta)p(h)\, \mathcal{L}_\theta(h). \tag{3}$$

With no AI ($\theta = 0$, hence $\mathcal{L}_\theta = 0$ and all tasks are performed hands-on), $\Phi_0(h) = \alpha \int \ell(z)\, \mathrm{d}F(z \mid h)$, the integral taken over the support of $F(\cdot \mid h)$. For $h \ge \theta$ the worker faces no AI-capable tasks ($\mathcal{L}_\theta = 0$) and $\Phi = \Phi_0$: AI cannot touch the learning of a worker whose entire task queue lies beyond its frontier.

Flow output combines the worker and the AI:

$$Y(h; \theta, a, r, m) = y_0(h) + \pi(a; h, \theta) - C_r(r)\, a\, Q_\theta(h) - C_m(m)\, (1 - a)\, Q_\theta(h), \tag{4}$$

where $y_0(h)$ is baseline output increasing in h, $\pi(a; h, \theta)$ is the throughput gain from delegating capable tasks (increasing in a, so $\pi_a > 0$: autonomy raises current output), and $C_r, C_m \ge 0$ are the per-task costs of review and assistance, normalized so that $C_r(0) = C_m(0) = 0$ (cost is incurred only for positive intensity, there is no fixed setup cost). These costs scale

with the *measure* of capable tasks $Q_\theta(h)$ rather than with the learning-content mass $\mathcal{L}_\theta(h)$: a task is not cheaper to review or assist merely because the worker happens to learn little from it, so the operating cost depends on how many capable tasks there are, not on how much skill they transfer. The throughput and cost schedules are left general here (only $\pi_a > 0$ and $C_r(0) = C_m(0) = 0$ are imposed).

Human capital evolves as

$$\dot{h} = \Phi(h; \theta, a, r, m) - \delta h, \tag{5}$$

with depreciation/obsolescence rate $\delta > 0$. The state is confined to $[0, 1]$, so (5) is read as constrained dynamics on the interval.[1]

4 When Does Autonomy Teach? The Scaffolding Ratio

Whether handing a task to AI helps or harms a worker's learning depends on a single comparison: does reviewing the machine's output teach more than performing the task would have? This section makes that comparison precise.

We first isolate the per-task comparison at the heart of the paper, holding h fixed; Section 5 lets h evolve.

Definition 1 (Scaffolding ratio)**.** *For a capable task, the* scaffolding ratio *is the ratio of autonomous-with-review learning to copilot learning,*

$$\psi(h, \theta) = \frac{\beta\, r\, A(\theta)\, p(h)}{\alpha\big(1 + mA(\theta)p(h)\big)}. \tag{6}$$

In other words, the scaffolding ratio $\psi(h, \theta)$ compares two ways a worker of human capital h can learn from a single task, by reviewing what the AI (of capability θ) produced versus by doing the task themselves with the AI assisting. Put simply, it asks, per task, how much does reviewing the machine's work teach the worker, relative to doing the work hands-on? When the ratio is above 1, reviewing teaches more than doing, so handing the task to the AI and having the worker check it builds skill. When it's below 1, doing teaches more, so delegating the task strips away learning the worker would otherwise have gained.

Substituting (3), the effect of AI on learning relative to the no-AI benchmark decomposes

[1] Formally $\dot{h} = \Pi_{[0,1]}\big[\Phi(h; u) - \delta h\big]$, the projection onto the tangent cone of $[0, 1]$: the drift evolves as $\Phi(h; u) - \delta h$ on the interior and is truncated so that the state never leaves $[0, 1]$.

cleanly (the algebra is in Appendix A). The net learning effect is

$$\Phi - \Phi_0 = \mathcal{L}_\theta(h)\Big\{\underbrace{\alpha\, mA(\theta)p(h)}_{\text{copilot gain } \geq 0} + \underbrace{a\,\alpha\big(1 + mA(\theta)p(h)\big)\big(\psi(h,\theta) - 1\big)}_{\text{autonomy contribution}}\Big\}. \tag{7}$$

Differentiating in a, the copilot term drops out and the autonomy margin is isolated:

$$\begin{aligned} \frac{\partial \Phi}{\partial a} &= \mathcal{L}_\theta(h)\,\alpha\big(1 + mA(\theta)p(h)\big)\big(\psi(h,\theta) - 1\big), \qquad \text{so,} \\ \operatorname{sign}\left(\frac{\partial \Phi}{\partial a}\right) &= \operatorname{sign}\big(\psi(h,\theta) - 1\big) \quad (\text{for } h < \theta, \text{ where } \mathcal{L}_\theta(h) > 0). \end{aligned} \tag{8}$$

Equation (8) is the formal statement that autonomy is good or bad for learning according to whether reviewing beats doing the task. The copilot channel (m) always adds learning (≥ 0), but the autonomy channel switches sign at $\psi = 1$ and operates at all only where capable tasks remain, since the common factor $\mathcal{L}_\theta(h)$ vanishes once $h \geq \theta$.

Proposition 1 (Autonomy reverses sign in capability and seniority)**.** *Fix (r, m) with $\beta r > 0$ and let Assumption 1 hold. The scaffolding ratio $\psi(h,\theta)$ is strictly increasing in worker seniority h and in AI capability θ. Because autonomy acts on learning only through the capable-task mass $\mathcal{L}_\theta(h)$, which is positive for $h < \theta$ and zero once the queue lies beyond the frontier, two regions must be distinguished:*

(i) Sign of the autonomy margin (capable region). *For a worker who still faces capable tasks with $h < \theta$,* $\operatorname{sign}(\partial\Phi/\partial a) = \operatorname{sign}(\psi(h,\theta) - 1)$*, and there is a frontier locus $x^* = A(\theta)p(h) = \alpha/(\beta r - \alpha m)$ (well defined when $\beta r > \alpha m$) such that $\psi \gtrless 1$ as $A(\theta)p(h) \gtrless x^*$: below it autonomy removes rungs, above it autonomy builds scaffolds. For $h \geq \theta$ autonomy is learning-*neutral *because there is nothing to delegate.*

(ii) Scaffolding region and its feasibility. *For fixed θ, let $h_\psi(\theta)$ solve $p(h_\psi) = x^*/A(\theta)$. Autonomy* strictly builds *learning exactly on the* scaffolding region $\{h : h_\psi(\theta) < h < \theta\}$*, where both the per-task review beats doing with $\psi > 1$ and capable tasks remain. This region is non-empty if and only if*

$$A(\theta)\, p(\theta)\, (\beta r - \alpha m) > \alpha.$$

When this fails, $\psi < 1$ for every $h < \theta$, so autonomy only removes rungs for $h < \theta$ or is neutral for $h \geq \theta$ and the scaffolding phenomenon does not arise.

(iii) Comparative statics. *The threshold h_ψ is decreasing in review intensity r and in AI quality $A(\theta)$ and increasing in copilot intensity m. The feasibility condition in (ii) is slackened by higher β/α, r, and θ and by lower m. Better review technology thus*

enlarges the scaffolding region and makes its existence more likely.

The result contradicts the engineering heuristic "keep the human hands-on to preserve learning": once AI is capable enough and the worker can verify, hands-on practice on tasks is the *inferior* teacher. The learning value of reviewing rises with the quality of what is reviewed ($A(\theta)$) and with the reviewer's capacity to absorb it ($p(h)$), whereas the learning value of doing a routine task does not.

The review channel is most valuable where verification ability is still low and fades as the worker's own competence grows. The prescription it motivates is the pedagogical counterpart of the seniority-contingent deployment derived below.

Definition 2 (Raw junior)**.** *A worker is a* raw junior *if their seniority lies below the scaffolding threshold, $h < h_\psi$, equivalently if their verification ability $p(h)$ is too low for reviewing to out-teach doing ($\psi(h, \theta) < 1$). At such h autonomy removes rungs rather than building skill. The polar case is the entrant at h_0. A worker with $h \geq h_\psi$, who can verify capable output well enough that review teaches more than practice, is* seniority-ready*.*

Proposition 1 previews the seniority profile developed formally in Appendix B.3: raw juniors have low ψ (they cannot yet verify) *and* by Assumption 1, a large capable-task mass $\mathcal{L}_\theta$ (much of their queue is automatable), so they bear the rung removal. Seniors have high ψ but small $\mathcal{L}_\theta$ (their queue lies beyond the frontier), so AI barely moves their learning; the scaffolding benefit is concentrated on intermediates.

5 Dynamic Equilibrium

We now let human capital evolve under (5) and close the model with a competitive labor market. We fix a deployment $u = (a, r, m)$ constant within a regime. Section 6 endogenizes it.

5.1 Career paths and steady state

We first characterize a worker's career path under a fixed deployment: when an interior steady-state level of expertise exists, whether it is unique, and how the worker approaches it. Existence requires only that learning not vanish at entry, which we now state before turning to the dynamics.[2]

[2]The corner is ruled out by Assumption 2: $\theta = 1$, $a = 1$, $p(0) = 0$. There, a raw entrant (a worker with $h \to 0$) faces no task beyond the AI frontier (none to do hands-on), no task kept in copilot mode, and no review learning (they cannot yet verify), so $\Phi(0) = 0$ and $h = 0$ is absorbing: the worker never accumulates and the career ladder has no bottom rung at all.

Assumption 2 (Positive entry-level learning)**.** *At least one entry channel is open:* $\theta < 1$*, or* $a < 1$*, or* $p(0) > 0$*. Equivalently, the deployment is not the universal-AI, full-autonomy, no-novice-verification corner.*

Under this assumption the dynamics are well behaved at the bottom of the ladder, and the steady state is characterized as follows.

Lemma 1 (Existence, uniqueness, and stability)**.** *Let* u *be fixed and* $\Phi(\cdot; u)$ *be given by* (3) *under Assumption 1. Suppose entry-level learning is positive,* $\Phi(0; u) > 0$ *under Assumption 2 and the learnable mass is exhausted at the top,* $\mathcal{L}_\theta(1) = \bar{\mathcal{L}}_\theta(1) = 0$*, so that* $\Phi(1; u) = 0$*. Then* $\dot{h} = \Phi(h; u) - \delta h$ *has an interior steady state* $h^* \in (0, 1)$*, since* $\Phi(0; u) > 0 = \delta \cdot 0$ *and* $\Phi(1; u) = 0 < \delta$*. The entry rate decomposes as*

$$\Phi(0;u) = \underbrace{\alpha\bar{\mathcal{L}}_\theta(0)}_{\text{frontier hands-on}} + \mathcal{L}_\theta(0)\Big[\underbrace{(1-a)\alpha\big(1+mA(\theta)p(0)\big)}_{\text{copilot}} + \underbrace{a\beta r A(\theta)p(0)}_{\text{review}}\Big], \tag{9}$$

The steady state is unique and globally asymptotically stable whenever average learning per unit of human capital is strictly decreasing,

$$\frac{\mathrm{d}}{\mathrm{d}h}\frac{\Phi(h;u)}{h} < 0 \qquad \Longleftrightarrow \qquad \varepsilon_\Phi(h) \equiv \frac{h\,\Phi'(h;u)}{\Phi(h;u)} < 1, \tag{10}$$

i.e. the elasticity of gross learning in human capital is below one.

Condition (10) accommodates a feature special to the review channel: because ψ rises with h, the learning rate can be *locally increasing* in human capital over an intermediate range, a "virtuous circle" in which becoming more senior makes the review channel more productive. The steady state remains unique as long as this local force does not overturn the eventual decline driven by the shrinking stock of learnable content.

5.2 Wage determination

Human capital is general and portable: a worker who separates carries h to the next employer. Consider a competitive market for general human capital with free entry of identical firms. A matched firm chooses deployment u, produces $Y(h; u)$, pays $w(h)$, and loses the worker at hazard s, after which the worker re-enters the market and is valued at the same type h. The firm's match value $J(h)$ and the worker's value $W(h)$ solve

$$(\rho+s)J(h) = Y(h;u) - w(h) + J'(h)\big[\Phi(h;u) - \delta h\big], \qquad \rho W(h) = w(h) + W'(h)\big[\Phi(h;u) - \delta h\big], \tag{11}$$

the worker being indifferent to separation because re-match is at the same type, so $W(h) = \int_0^\infty e^{-\rho t} w(h_t)\,\mathrm{d}t$ along the induced path.

A single firm is atomistic in the market for type h: its own deployment does not move the market price $w(h)$ of that type. Hence $\partial w/\partial u = 0$. Free entry drives firm value to zero, $J(h) \equiv 0$, so the firm's shadow value of human capital is $\mu_F \equiv J'(h) = 0$ and the equilibrium wage is, in closed form,

$$w(h) = Y\big(h; \bar{u}(h)\big), \tag{12}$$

the full current output net of AI operating costs at the deployment $\bar{u}(\cdot)$ that firms play. It is strictly increasing in h. Wage growth is then

$$\frac{\dot{w}}{w} = \frac{Y_h(h;\bar{u})}{Y(h;\bar{u})}\,\dot{h} = \frac{Y_h(h;\bar{u})}{Y(h;\bar{u})}\big(\Phi(h;\bar{u}) - \delta h\big), \tag{13}$$

still proportional to the net learning rate. The level of pay can be high. It includes the autonomy throughput term $a\,\pi(h,\theta)$, while the *slope* is governed by Φ: entry productivity rises even as the trajectory stalls.

Equation (13) treats $\bar{u}$ as fixed at the worker's seniority. When the market deployment is itself seniority-contingent, $\bar{u} = \bar{u}(h)$, the total derivative of $w(h) = Y(h;\bar{u}(h))$ carries an indirect term,

$$\frac{\dot{w}}{w} = \frac{Y_h(h;\bar{u}) + Y_u(h;\bar{u})\,\bar{u}'(h)}{Y(h;\bar{u})}\,\dot{h}, \tag{14}$$

and (13) is the special case $Y_u\bar{u}' = 0$. Two arguments deliver it. *(i) Envelope.* In the competitive equilibrium priced here the wage is the upper envelope $w(h) = \max_u Y(h;u)$, since free entry makes firms play $\bar{u} = \arg\max_u Y$ (here $\mu_F = 0$). At an interior maximizer $Y_u = 0$, so the envelope theorem kills the indirect term and (13) holds exactly even as $\bar{u}(h)$ moves. *(ii) Fixed regime.* On any interval where $\bar{u}$ is constant, in particular each piece of a bang-bang or piecewise seniority-contingent program, $\bar{u}'(h) = 0$ and (13) again holds, so the path is smooth within each piece. Where the program switches deployment (at h_ψ, or at $h = \theta$ where capable tasks vanish), the path is continuous with a slope *kink* if pay equals the competitive envelope of (i), and *jumps in level* if pay tracks the development-oriented deployment actually played. Proposition 10, in Appendix B.4, sizes and signs this step for the seniority-contingent program.

Note that the law of motion (5) contains no wage term. Given a deployment, the human-capital path and steady state are invariant to how wages are determined. The autonomy reversal (Proposition 1) and the existence, uniqueness, and stability of the steady state (Lemma 1) concern only Φ, ψ, and h^*, and are therefore *unchanged* by endogenizing w. What the wage block determines is *which* deployment $\bar{u}$ the market selects and the level and

slope of the pay path through (12)–(13).

5.3 A watershed in the long run

Proposition 1 concerns a fixed h. We now consider workers that travel through the seniority distribution.

Proposition 2 (Rungs removed early, ceiling raised late)**.** *Consider a high-autonomy deployment with sufficient review so that $\psi(h,\theta) < 1$ for h below a threshold $h_\psi \in (h_0, \theta)$ and $\psi(h,\theta) > 1$ for $h \in (h_\psi, \theta)$. Let also Assumptions 1 and 2 hold. Then, relative to the no-AI baseline:*

(i) (Early-career loss.) *Define the* loss threshold h_c *as the unique seniority at which the two just balance,*

$$\alpha m A(\theta) p(h_c) \;=\; a\,\alpha\big(1 + mA(\theta)p(h_c)\big)\big(1 - \psi(h_c,\theta)\big), \tag{15}$$

equivalently $\psi(h_c,\theta) = 1 - \frac{mAp}{a(1+mAp)} \le 1$, so that $h_c = h_\psi$ when $m = 0$ and $h_c < h_\psi$ when $m > 0$. Then $\Phi < \Phi_0$ exactly on $\{h < h_c\}$: for entrants with $h_0 < h_c$ the net learning rate is below baseline over $[h_0, h_c)$ and the early-career wage slope (13) is flatter than under no AI. For $h \in (h_c, h_\psi)$ the copilot gain outweighs the autonomy loss and AI raises *learning despite $\psi < 1$.*

(ii) (Higher ceiling.) *If the steady state satisfies $h^* \in (h_\psi, \theta)$, then h^** exceeds *the no-AI steady state h_0^*: in the scaffolding region the review channel delivers more learning per unit time than hands-on practice, raising the long-run stock of expertise.*

The same deployment therefore lowers early-career growth and raises the expert ceiling.

Proposition 2 is the dynamic counterpart of the static reversal of Proposition 1. The same autonomous deployment that strips rungs from the youngest workers raises the expert ceiling for those who survive into the review-productive region: the cost lands on early-career growth, the benefit on the long-run stock. The youngest cohort is losing rungs while the settled cohort is gaining ceiling. The two effects are not in tension but sequential: a worker first loses, on $\{h < h_c\}$, then gains, once past h_ψ into the scaffolding region. The condition in part (ii), whether the steady state lands in the scaffolding region, is the substance of the result. It says the *sign* of AI's long-run effect on the workforce's expertise is itself a deployment choice.

Proposition 3 (The review watershed)**.** *Fix capability $\theta > 0$ and autonomy $a > 0$, let Assumptions 1 and 2 hold, and let $h^*(r)$ denote the unique stable steady state of Lemma 1.*

The map $h^*(r)$ is continuous and nondecreasing, with $\mathrm{d}h^*/\mathrm{d}r = \Phi_r/(\delta - \Phi_h) \geq 0$ at the stable steady state. It is strictly *increasing on the range where the steady state lies below the frontier, $h^*(r) < \theta$, and $p(h^*) > 0$. Where instead $h^*(r) \geq \theta$, the worker faces no capable tasks and review is long-run neutral. Maintaining the advanced-AI case $h_0^* < \theta$ places the crossing in the capable region. Suppose the two endpoints straddle the no-AI benchmark h_0^*:*

$$h^*(0) < h_0^* < h^*(1). \tag{16}$$

Then there is a unique critical intensity $r^* \in (0,1)$ at which $h^*(r^*) = h_0^*$, and

$$h^*(r) \gtrless h_0^* \quad \Longleftrightarrow \quad r \gtrless r^*. \tag{17}$$

Review intensity then sets the sign *of AI's long-run effect on expertise: $r > r^*$* permanently raises *the workforce's expertise above its no-AI level, $r < r^*$* permanently lowers *it.*

The lower inequality in (16) holds whenever autonomy is high enough that delegating without review loses learning, $a > \frac{mA(\theta)p}{1+mA(\theta)p}$ at $h^*(0)$. The upper inequality holds only if the review channel is strong enough with higher β/α, AI quality $A(\theta)$, and verification slope, and lower depreciation δ. If instead $h^*(1) \leq h_0^*$, no interior r^* exists.

Where the static reversal fixes the sign of autonomy's effect at a given seniority, the watershed fixes the sign of its effect on the *whole workforce's* steady-state expertise and shows that the sign is not a property of the technology but a choice. The same capable, autonomous AI raises long-run expertise above its no-AI level when review is intense enough ($r > r^*$) and lowers it when review is too thin ($r < r^*$). Capability θ sets only how much is at stake, never the direction. The governable margin is review, not how good the AI is.

The watershed exists precisely when the review channel is strong enough that maximal review overturns the autonomy loss but weak enough that minimal review does not. Outside this window the long-run effect is single-signed: AI builds at every feasible review intensity (if even $r = 0$ clears h_0^*, which requires copilot-dominated, low-autonomy deployments) or erodes at every intensity (review exhausted).

Two features make this more than a comparative static. First, the same autonomous AI is expertise-building or expertise-eroding depending on a single governable margin, review. Second, the watershed does *not* produce a poverty trap: although the review channel makes learning locally increasing in h, the steady state remains unique (Lemma 1), because strengthening review simultaneously lowers the scaffolding threshold h_ψ that would form a basin's edge and the two effects offset. Deployment, not initial human capital, decides the long run. Figure 1 traces $h^*(r)$ crossing the no-AI ceiling at r^*. We see that holding au-

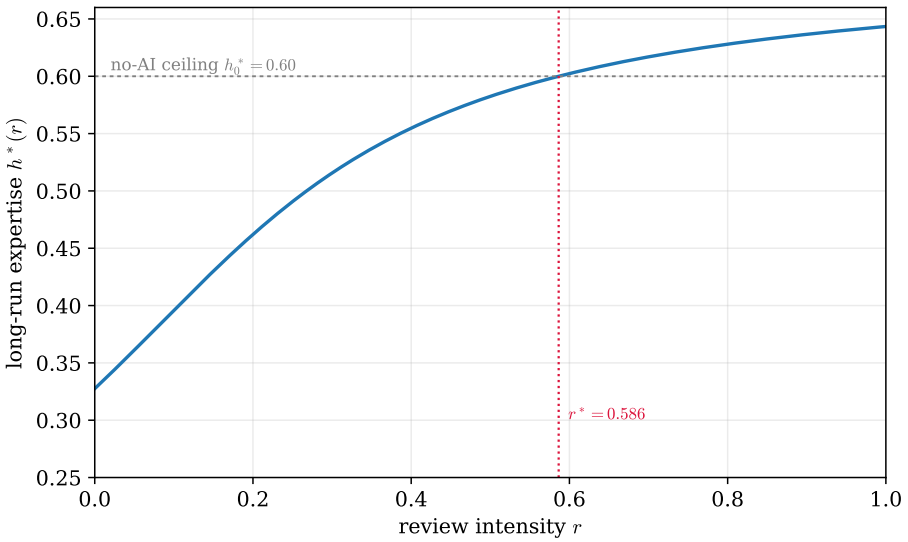


Figure 1: The review watershed. Holding autonomy at $a = 0.85$ and $\theta = 0.70$, the long-run expertise stock $h^*(r)$ rises with review intensity and crosses the no-AI ceiling $h_0^* = 0.60$ at $r^* \approx 0.586$. Below r^*, autonomous AI permanently lowers the workforce's expertise. Above it, the same autonomy permanently raises it. The steady state is unique throughout with a smooth sign reversal. Parameter values: $\alpha = 0.045$, $\beta = 0.18$, $\delta = 0.015$, $\theta = 0.70$, $a = 0.85$, $m = 0.20$, $A(\theta) = \theta = 0.70$, $p(h) = h$, $\ell(z) = 1 - z$, $z \mid h \sim \text{Uniform}[h, 1]$ and $h_0^* = \alpha/(\alpha + 2\delta) = 0.60$.

tonomy fixed and raising review across the watershed r^* flips the sign of the long-run effect, from AI *lowering* steady-state expertise below its no-AI level to *raising* it above.

6 Deployment Choice, Contractibility, and the Autonomy Externality

The wage block fixes pay given the deployment. It does not say which deployment is chosen. That depends on who internalizes the worker's future human capital, which in turn depends on the market structure and on whether the deployment is contractible. The relevant benchmark throughout is the planner, who values output and internalizes all future human capital (and any cross-firm spillover $\nu \geq 0$) at the social discount rate ρ:

$$\rho\, V_S(h) = \max_u \Big\{ Y(h;u) + \big(V_S'(h) + \nu\big)\big[\Phi(h;u) - \delta h\big] \Big\}, \tag{18}$$

with social shadow value $\mu_S \equiv V_S'$ and interior conditions $Y_u + (\mu_S + \nu)\Phi_u = 0$. We call the maximizer

$$u^S(h) \equiv \arg\max_u \left\{ Y(h;u) + \big(\mu_S(h) + \nu\big)\big[\Phi(h;u) - \delta h\big] \right\}$$

the *social optimum*: the deployment a planner would choose, weighting the future human capital that deployment builds by its full social shadow value $\mu_S + \nu$. At an interior optimum it satisfies $Y_u + (\mu_S + \nu)\Phi_u = 0$. The social optimum is the benchmark against which we measure the market's deployment and the gap between u^S and the firm's choice.

Because the deployment results turn on what is contracted and when, we state the environment explicitly. Time is again continuous. At each instant a worker of type h in a match experiences:

(t1) **Pricing.** A competitive spot market values the worker's *current* type at the zero-profit wage $w(h)$ of (12). Pay is set on the current type, not on a promised trajectory.

(t2) **Deployment.** The matched firm chooses $u = (a, r, m)$.

(t3) **Production and accumulation.** Output $Y(h;u)$ is produced, $w(h)$ paid, and h evolves by $\dot{h} = \Phi(h;u) - \delta h$.

(t4) **Separation.** With hazard s the match ends and the worker re-enters and is re-priced at the new type. Human capital is general, so, h is carried across employers.

A *deployment policy* is a map $h \mapsto u(h) = (a(h), r(h), m(h))$ specifying the firm's choice at each seniority. One policy will recur.

Definition 3 (Seniority-contingent rule)**.** *The* seniority-contingent rule *is the deployment policy that conditions autonomy on the worker's seniority relative to the scaffolding threshold h_ψ of Proposition 1: while $h < h_\psi$ the firm keeps the worker in* copilot*, withholding autonomy and assigning review; once $h \geq h_\psi$ it switches to* delegate-and-review*, granting autonomy paired with review intensity $r \geq r^*$. Autonomy is thus released exactly when reviewing begins to teach more than doing.*

A learning-poor deployment lowers the worker's *future* type, for which they are compensated later, at the future spot price, *regardless of which firm they are then at*, since human capital is general. Hence, even a fully *observed* deployment cannot be priced by the current employer: the firm cannot capture a return that the worker will collect elsewhere, and the worker cannot bill the current employer for a portable future gain. What disciplines deployment is therefore the ability to *contract and commit* on review intensity, or to *direct search* toward it, not mere observability. This separates two main regimes:[3]

[3] A third regime is the one in which review intensity is unobservable. For the purposes of this study, this is identical to the regime where review intensity is observable and non-contractible.

(i) *Observable but non-contractible*, with spot wages and no directed search: the spot wage prices current type only, so μ_F reflects retained rents alone.

(ii) *Contractible and committable*: firms post (w, u) and workers direct search, so the training environment is a priced term and efficiency is restored.

Review intensity is costly now, pays off only through portable future skill, and is a *practice* rather than a verifiable output: a court or an outside market cannot confirm how intensively AI output was checked. With deployment non-verifiable and interaction non-repeated, a firm can neither write an enforceable promise to review nor sustain an informal one. "Non-contractible" thus means not that AI adoption is hidden but that the learning-relevant *intensity of human engagement* cannot be verified or committed. We develop three distinct cases out of these two regimes.

Case 1: Competitive market, non-contractible deployment. This is the competitive benchmark case which is defined by the following four conditions:

(a) *Price-taking.* Free entry and atomistic firms, so a filled job has zero value, $J \equiv 0$, and the firm's shadow value of human capital is $\mu_F = 0$: a firm that competes away its rents and cannot retain the worker prices none of the worker's future skill.

(b) *Non-contractibility.* Deployment, in particular review intensity, is non-verifiable and cannot be written into an enforceable contract, so its sole return, the portable future skill it builds, is not something the firm can be paid for.

(c) *No commitment device.* There is no directed search or relational commitment on the training environment, so neither posted contracts nor reputation substitutes for the missing formal contract.

(d) *Spot wages.* Wages are set each instant at the competitive price for the worker's current type, which is what delivers $\mu_F = 0$ in (a).

Then, the following proposition applies:

Proposition 4 (The competitive equilibrium selects myopic deployment)**.** *Let Assumptions 1 and 2 hold. In the competitive, non-contractible Case 1, the firm maximizes current output alone. The unique equilibrium deployment is the myopic corner*

$$\bar{u}: \quad a = 1, \quad r = 0,$$

with the copilot margin vacuous at $a = 1$ and normalized to $m = 0$. Every capable task is then delegated and unreviewed, so capable tasks contribute zero learning and the induced steady state solves $\alpha\bar{\mathcal{L}}_\theta(h) = \delta h$. Capability acts as a trap:

(i) if $\theta < h_0^* = \alpha/(\alpha + 2\delta)$, the worker outgrows the frontier ($h^* > \theta$) and recovers the no-AI ceiling h_0^*; the loss is purely transitional;

(ii) if $\theta > h_0^*$, the worker is permanently below the frontier, capable tasks are delegated forever without review, and the steady state collapses far below h_0^*.

When AI is capable enough that the no-AI steady state lies below the frontier ($h_0^* < \theta$), the competitive steady state collapses far below it, and the more capable the AI the deeper the collapse. The calibration of Section 7 quantifies this. Below this threshold ($\theta < h_0^*$), by contrast, the worker eventually outgrows the AI frontier and the market recovers the no-AI ceiling: the myopic deployment costs only a transitional slowdown, not a permanent loss. The externality therefore bites precisely when AI is capable enough to matter and bites harder the more capable it is. *Endogenizing wages thus makes the externality starker, not weaker:* absent contractibility, competition for general human capital drives deployment to the corner that maximizes current output and minimizes review, exactly the configuration that hollows out learning.

The starkness of this corner is a modeling choice. The $r = 0$ corner reflects that in (4) review carries only a current cost and no current-output benefit: its sole return is the portable, non-contractible learning it generates, which a myopic firm ($\mu_F = 0$) does not value. Suppose instead that reviewing autonomous output also raises its current quality by catching AI errors, adding a term $a\,Q_\theta(h)\,B(r)$ to (4) with $B' > 0$, $B(0) = 0$, and B concave. Then even the myopic firm reviews: it sets the *output-quality optimum* r^Y solving $B'(r^Y) = C_r'(r^Y)$, interior and strictly positive whenever $B'(0) > 0$, and autonomy still corners at $a = 1$ (reviewed output is more valuable, so $Y_a > 0$ a fortiori). The equilibrium is then $(1, r^Y, \cdot)$ rather than $(1, 0, \cdot)$. The externality is nonetheless robust: the planner's review condition $Y_r + (\mu_S + \nu)\Phi_r = 0$ adds the learning value $\Phi_r = a\beta A(\theta)p(h)\mathcal{L}_\theta > 0$, giving $B'(r^S) - C_r'(r^S) = -(\mu_S + \nu)\beta A(\theta)p(h)\mathcal{L}_\theta/Q_\theta < 0$ and hence $r^S > r^Y$ by concavity of $B - C_r$. The firm under-provides review relative to the social optimum. Only the *starkness* of the corner softens, from $r = 0$ to the output-quality level r^Y and the $r = 0$ corner is the limit $B \equiv 0$. This sharpens rather than weakens the mechanism: what the market fails to supply is specifically the *learning* value of review, which is portable and non-contractible. Its current-output value, which the firm internalizes, is not the source of the failure and does not close it. Distinguishing the two is precisely the content of the externality.

Case 2: Firm rents, non-contractible deployment. With search frictions or imperfect competition the firm retains a share of match rents, so it partially internalizes future skill but discounts it at the retention-adjusted rate $\rho + s$. Writing the firm's value as V_F in the

rents case,

$$(\rho + s)\, V_F(h) = \max_u \Big\{ Y(h;u) - w(h) + V_F'(h)\big[\Phi(h;u) - \delta h\big] \Big\}, \tag{19}$$

where the wage schedule is the following.

Assumption 3 (Frictional wage from search)**.** *A worker of type h can relocate to a competitive firm, which prices general skill at $Y(h;\bar{u}(h))$. By searching, the worker meets such an offer at Poisson rate λ. The value of search is therefore $\rho U(h) = \frac{\lambda}{\rho+\lambda} Y(h;\bar{u}(h))$, and the current firm, holding within-match bargaining power, pays this outside option:*

$$w(h) = \sigma\, Y\big(h;\bar{u}(h)\big), \qquad \sigma \equiv \frac{\lambda}{\rho+\lambda} \in (0,1), \qquad w'(h) = \sigma\, Y_h\big(h;\bar{u}(h)\big). \tag{20}$$

Two properties follow. (i) *Price-taking*: the wage equals the worker's outside option because the competitive price of their *general* skill is evaluated at the market deployment $\bar{u}$, so it depends on the type h, not on the current firm's deployment: $\partial w/\partial u = 0$, and $w' = \sigma Y_h(\cdot;\bar{u})$ by the envelope theorem ($Y_u = 0$ at $\bar{u}$). (ii) *Endogenous pass-through*: σ is the equilibrium markdown, not a free parameter. As frictions vanish ($\lambda \to \infty$) the worker can instantly recontract at the competitive wage, $\sigma \to 1$, $w = Y$, and the firm's rent vanishes (this is the competitive benchmark, Case 1, with $\mu_F = 0$). With frictions (λ finite), $\sigma < 1$: the worker collects only a fraction of the return to general skill and the firm retains the complement, the wage-compression mechanism of Acemoglu and Pischke (1998, 1999). The firm's rent $J(h) = V_F(h) = \mathrm{PV}\big[Y(h;u) - w(h)\big] \geq 0$ is the present value of that retained share, and is also what the firm discounts at the retention-adjusted rate $\rho + s$.

Because $\partial w/\partial u = 0$, (19) leaves the deployment conditions unchanged from Case 1. In the leading *linear* case, $\pi = a\,\bar{\pi}(h,\theta)$, with constant per-task costs C_r, C_m, the maximand is multilinear in u and the optimum is bang-bang (each margin at a corner). The competitive corner $a = 1$, $r = 0$ of Proposition 4 is exactly this case.

To characterize an *interior* optimum, we impose the following regularity condition.

Assumption 4 (Regular deployment)**.** *In (4) the throughput $\pi(a;h,\theta)$ is C^2 and strictly concave in a ($\pi_{aa} < 0$), and the review and copilot cost schedules $C_r(r), C_m(m)$ are C^2 and strictly convex ($C'' > 0$, $C'(0) = 0$).*

As a result, the deployment objective $Y + \mu(\Phi - \delta h)$ is strictly concave in each control and admits a unique interior maximizer. The interior optimum is characterized by

$$\underbrace{Y_a + \mu\,\Phi_a = 0}_{\text{autonomy}}, \qquad \underbrace{Y_r + \mu\,\Phi_r = 0}_{\text{review}}, \qquad \underbrace{Y_m + \mu\,\Phi_m = 0}_{\text{copilot}}, \tag{21}$$

with $\mu = \mu_F$ for the firm and $\mu = \mu_S + \nu$ for the planner, where $\Phi_a = \mathcal{L}_\theta\, \alpha(1 + mAp)(\psi - 1)$, $\Phi_r = a\beta Ap\, \mathcal{L}_\theta > 0$, $\Phi_m = (1-a)\alpha Ap\, \mathcal{L}_\theta > 0$, $Y_a = \pi_a(a;\theta) - [C_r(r) - C_m(m)]Q_\theta$ (with $\pi_a > 0$, $\pi_{aa} < 0$), $Y_r = -C_r'(r)aQ_\theta < 0$, and $Y_m = -C_m'(m)(1-a)Q_\theta < 0$.

Lemma 2 (Undervaluation of human capital)**.** *In the environment of Case 2, if $Y_h \geq 0$ then at any common allocation $\mu_F(h) \leq \mu_S(h) + \nu$, with strict inequality whenever $s > 0$, $\sigma > 0$, or $\nu > 0$. The wedge has three sources: the worker's wage pass-through, the retention hazard s, and the spillover ν. The wedge is increasing in the pass-through σ,* decreasing *in the firm's retained share $1 - \sigma$ and increasing in s and in ν.*

Lemma 2 says the firm always values the worker's human capital by (weakly) less than the planner does, and isolates why. Each source is a distinct wedge between private and social returns to the same future skill. The pass-through σ measures how much of the worker's marginal product the worker collects: because general skill is portable and priced competitively, a fraction σ of any future gain accrues to the worker (and, when they separate, is carried to the next employer), so the firm capitalizes only the retained complement $(1 - \sigma)Y_h$. The hazard s compounds this: the firm discounts future skill at the retention-adjusted rate $\rho + s$ because it may lose the worker before the return arrives, whereas the worker and hence the planner discounts at ρ, the skill surviving the match. The spillover ν is the part of the social return that no single employer internalizes at all. The two limits make the range transparent: at $\sigma = 1$ (Case 1) the firm retains nothing, $\mu_F = 0$, and the wedge is maximal. As $\sigma, s, \nu \to 0$ the firm becomes a residual claimant on the full, undiscounted, un-spilled return and $\mu_F \to \mu_S$, closing the gap.

This wedge is the engine of the deployment externality. The firm and the planner choose the autonomy–review mix by the same first-order logic but with different valuations of the learning it produces: the planner weights net learning by $\mu_S + \nu$, the firm by the smaller μ_F. Wherever raising review (or lowering autonomy) builds future skill (i.e. wherever Φ responds to deployment), the firm, valuing that skill less, supplies less of it than the planner would. Proposition 5 turns this gap in valuations into a gap in chosen deployments, showing that the competitive firm sets autonomy too high and review too low relative to the social optimum, by an amount that scales with the wedge $\mu_S + \nu - \mu_F$.

Proposition 5 (Over-autonomy, under-review)**.** *Let Assumption 1 hold. Suppose the firm earns rents but retention is imperfect ($s > 0$) or there is a spillover ($\nu > 0$), so $0 < \mu_F < \mu_S + \nu$. Then at the privately optimal deployment, relative to the socially optimal one:*

(i) in the rung-removing region with $\psi < 1$, the firm sets autonomy a too high*: it underweights the learning loss from delegation;*

(ii) the firm sets review r and copilot intensity m too low*: both are costly today and pay off only through future human capital, which the firm undervalues;*

(iii) the distortion is larger the lower is retention (higher s), the higher the wage pass-through σ (less compression), and the larger the spillover ν, and is maximal in the competitive full-pass-through limit $\mu_F \to 0$ ($\sigma \to 1$) of the competitive equilibrium of Proposition 4.

The distortion in Cases 1 and 2 has a single root: deployment is non-contractible, so the firm cannot be paid for the portable skill that review builds, and prices it at $\mu_F < \mu_S + \nu$ rather than at its social value.

Case 3: Contractible deployment (competitive search). Suppose now firms can post and commit to contracts specifying both pay and deployment, $(w(\cdot), u(\cdot))$, and workers direct their search across posted contracts. Free entry then implies that the equilibrium contract maximizes the worker's value subject to the firm breaking even, i.e. it maximizes the *joint match surplus*.

Corollary 1 (Contractibility restores efficiency)**.** *If deployment is contractible and the market is competitive in (w, u) packages, the equilibrium deployment maximizes the joint match surplus, so the hold-up wedge of Lemma 2 closes ($\mu_F = \mu_S$) and the over-autonomy/under-review distortion of Proposition 5 vanishes. Absent a cross-match spillover ($\nu = 0$) the allocation coincides with the social planner's. The worker finances learning-intensive deployment by accepting lower current wages, recovering more than the concession through the higher future wage path that higher h commands.*

This is in fact the Becker logic for general training, extended to the choice of how AI is deployed. It also makes precise why the worker cannot simply self-finance the learning and dissolve the externality on their own. The natural objection to a hold-up in *general* skill is Beckerian: if review builds portable skill the worker values, the worker should review more intensively on their own effort and accept a lower wage during the learning phase, internalizing the return. That trade is exactly what Case 3 implements, and it is why the externality vanishes there. It fails in the baseline (Case 1) for two reasons that the objection elides. First, the learning-relevant input is not private study effort the worker can add at the margin. It is the deployment configuration u, and by the timing above the firm chooses u, including whether the worker is given the slack to review rather than ship. Second, even a worker willing to make the Becker concession cannot write the contract that would buy it: the priced object would have to be the *learning content* of the deployment, but, as above, whether review actually *taught* is non-verifiable to any third party, so no enforceable "deploy

me to learn and I will accept lower pay" term exists. The Becker wage discount therefore cannot operate through a spot contract. It requires the committable (w, u) package of Case 3. The directed-search environment of Appendix B.5 is precisely the Becker channel operating *imperfectly*, through workers sorting toward developmental postures on a noisy public signal rather than a contract, and Proposition 11 shows that even then it does not fully dissolve the externality.

We conclude that the autonomy externality is a hold-up problem, not a feature of AI adoption as such. It is absent only under full contractibility of deployment (Case 3). It is present whenever "how intensively AI output is reviewed" or "whether juniors are kept hands-on" cannot be written into an enforceable, portable contract. The distortion ranges from interior (firm rents, Case 2) to full autonomy and zero review (competition, Case 1).

7 Calibration

The per-task results of Section 4 and the dynamic results of Section 5 hold for an arbitrary FOSD-ordered task distribution $F(\cdot \mid h)$ and learning profile ℓ. To obtain explicit magnitudes and to calibrate, we now specialize the general environment of Section 3 to a transparent leading example.

Tasks are uniform on the worker's frontier and above,

$$z \mid h \sim \text{Uniform}[h, 1], \qquad f(z \mid h) = \frac{1}{1-h} \text{ on } [h, 1],$$

with learning content $\ell(z) = 1 - z$, output quality $A(\theta) = \theta$, and verification ability $p(h) = h$. A worker of seniority h is handed problems at or above their current level; as h rises the support shifts right and the most elementary, highest-content tasks drop out of the queue. The choices $A(\theta) = \theta$ and $p(h) = h$ are normalizations consistent with $A' > 0$, $A(0) = 0$ and $p' > 0$, $p(0) = 0$. That ℓ is decreasing while $F(\cdot \mid h)$ rises with seniority means juniors face the highest-content tasks.

The capable-task measure (2) becomes

$$Q_\theta(h) = \frac{\theta - h}{1-h} \quad (h < \theta), \qquad Q_\theta(h) = 0 \quad (h \geq \theta),$$

and, evaluating the learning masses over the support $[h, 1]$,

$$\mathcal{L}_\theta(h) = \int_h^\theta \frac{1-z}{1-h}\,\mathrm{d}z = \frac{(1-h)^2 - (1-\theta)^2}{2(1-h)}, \quad \bar{\mathcal{L}}_\theta(h) = \int_\theta^1 \frac{1-z}{1-h}\,\mathrm{d}z = \frac{(1-\theta)^2}{2(1-h)} \quad (h < \theta). \tag{22}$$

Their sum is $\mathcal{L}_\theta(h) + \bar{\mathcal{L}}_\theta(h) = \frac{1}{2}(1-h)$, independent of θ. At the top $\mathcal{L}_\theta(1) = \bar{\mathcal{L}}_\theta(1) = 0$, so the top-mass condition of Lemma 1 holds and $\Phi(1;u) = 0$. The no-AI rate is $\Phi_0(h) = \alpha \int_h^1 \frac{1-z}{1-h}\,\mathrm{d}z = \frac{\alpha}{2}(1-h)$. The feasibility condition for a non-empty scaffolding region (Proposition 1(ii)) specializes to $\beta/\alpha > (\theta^{-2}+m)/r$.

Under the no-AI baseline (and, more generally, for any worker with $h \geq \theta$ who faces only frontier tasks), $\Phi_0(h) = \alpha(1-h)/2$ and (5) is linear:

$$\dot{h} = \tfrac{\alpha}{2} - \big(\tfrac{\alpha}{2} + \delta\big)h, \qquad h_t = h^* + (h_0 - h^*)e^{-(\alpha/2+\delta)t}, \qquad h_0^* = \frac{\alpha}{\alpha + 2\delta}. \tag{23}$$

The career path is an exponential approach to the no-AI steady state h_0^* at rate $\alpha/2 + \delta$. With AI active ($h < \theta$), Φ is smooth and (5) is solved numerically. The path has a kink at $h = \theta$, above which it coincides with (23).

We calibrate this parametrization to *illustrate magnitudes*, not to discipline them: the calibration shows what the mechanism implies for plausible parameters, and is not an estimate. We set $\alpha = 0.045$ and $\beta = 0.18$, so $\beta/\alpha = 4$. Depreciation is $\delta = 0.015$.[4] These imply a no-AI human-capital steady state $h_0^* = \alpha/(\alpha+2\delta) = 0.60$ and, for an entrant at $h_0 = 0.15$, a five-year *human-capital* growth of 51.3% via (23), in the range of acceptable early-career wage slopes according to empirical evidence.[5] AI quality is $A(\theta) = \theta$ and verification ability $p(h) = h$, both normalizations rather than measurements. A final normalization concerns the mapping from human capital to pay. The model pins down the *slope* of the wage path (governed by Φ) but not its *level*. The unmodeled functions $y_0(h)$ and $\pi(a;h,\theta)$ enter the level of $w(h) = Y(h;\bar{u}(h))$ without affecting the law of motion (5). For the calibration we

[4] We set $\delta = 0.015$, a representative rate in the lower part of the empirical range for the depreciation of *general*, portable human capital. The appropriate object here, since h is general skill held by an *employed* worker. Wage-profile estimates for the general qualification stock centre near 1% per year (Arrazola and de Hevia, 2004), with rates below that for workers in active training and somewhat above for those exposed to skill-eroding interruptions; task-based estimates span roughly 1–6%, the higher rates attaching to *specific* skills and to occupations heavy in non-routine or rapidly-upgraded tasks, while general skills sit at the bottom of that range (Walter and Lee, 2022). A representative value modestly above the 1% floor is therefore the natural choice for durable but not permanent professional expertise. Note that rates of 4–6% typically identify depreciation from *time out of employment* (Dinerstein et al., 2022), a different and faster process than the steady-state erosion of an employed worker's general stock that δ captures. $\delta = 0.015$ is consistent with the general-skill evidence and sits safely below the obsolescence watershed of Appendix B.2, so the baseline remains in the expertise-building region.

[5] A cumulative five-year human-capital growth of 51.3% (about 8.6% per year) is a reasonable magnitude for the experience-driven component of early-career wage growth. Administrative and panel data place raw early-career wage growth at roughly 10–11% per year over the first five years, flattening sharply thereafter (Rubinstein and Weiss, 2006; Topel and Ward, 1992; Adda and Dustmann, 2023). Adda and Dustmann (2023) further show that the *experience* (human-capital) channel accounts for the large majority of this slope, with the remainder reflecting search, mobility, and match effects that the present model abstracts from. The figure here therefore sits below the raw compounded slope ($\approx$ 68% over five years) by design. It isolates the human-capital channel while remaining in its empirically relevant range. As with the other primitives, the calibration is illustrative rather than estimated.

Table 1: Calibrated deployment regimes (entrant $h_0 = 0.15$; $\alpha = 0.045$, $\beta = 0.18$, $\delta = 0.015$). The final column is five-year proportional growth in human capital.

Regime	θ	a	r	m	$\psi(0.2)$	$\psi(0.5)$	h^*	5-yr growth
Pre-AI baseline	0.00	0.00	0.00	0.00	–	–	0.600	51.3%
Basic autonomous AI	0.35	0.90	0.10	0.10	0.028	0.069	0.600	32.6%
Advanced copilot	0.70	0.15	0.20	0.70	0.102	0.225	0.614	48.6%
Advanced autonomous, low review	0.70	0.85	0.20	0.30	0.107	0.253	0.465	12.8%
Advanced autonomous + review	0.70	0.85	0.90	0.20	0.490	1.178	0.637	27.6%
Predictive + targeted review	0.70	0.40	0.85	0.50	0.445	1.013	0.626	42.2%
Competitive eq. (Prop. 4)	0.70	1.00	0.00	0.00	0.000	0.000	0.161	0.4%

therefore report proportional growth in human capital and read it as wage growth under the normalization that the market price of general skill is proportional to human capital, $w(h) \propto h$ (unit pay-elasticity, $\varepsilon_{w,h} \equiv Y_h h/Y = 1$). The reported five-year figures are then simultaneously human-capital growth and wage growth.[6]

We then solve (5) along the career for six deployment regimes spanning the policy space. Table 1 reports the scaffolding ratio at two seniority levels, the human-capital steady state, and five-year early-career growth (proportional growth in h, read as wage growth under the normalization above). The table entries are *outcomes* computed by integrating the law of motion given the parameters above. What is chosen is the parameter vector, above all β/α.

The first six rows are deployment *choices*. The last is the deployment the competitive market with non-contractible deployment *selects*. With capable AI ($\theta = 0.70 > h_0^*$), that equilibrium drives the steady state to 0.161 and early-career wage growth to 0.4%. This row is the *stark limiting case*: it is the $B \equiv 0$ corner in which review carries no current-output benefit, so the myopic firm sets $r = 0$ exactly (Section 6). The robust prediction is *under-provision* of review relative to the social optimum: once review also raises current quality, the competitive choice is an interior $r^Y > 0$ that still falls short of the planner's r^S, because the firm internalizes review's current-output value but not its portable learning value. The 0.4% row is the most dramatic illustration of that under-provision, not a claim that real firms set review exactly to zero; the gap between it and, say, the predictive regime is the

[6] In general $\dot{w}/w = (Y_h/Y)\dot{h} = \varepsilon_{w,h}(h)\,(\dot{h}/h)$, so a different, fully specified pay schedule would rescale each regime's five-year figure by its pay elasticity $\varepsilon_{w,h}$ rather than change what is computed. Two facts make the normalization innocuous here. First, the reported number is exactly five-year human-capital growth, independent of any pay map. Second, because all regimes share the entry type h_0 and the wage is a single increasing function of type, the market price of general skill, with $\partial w/\partial u = 0$ (Section 5.2), the cross-regime *ordering* of five-year growth is preserved under any such schedule. The qualitative comparisons are therefore properties of the h-paths, not of the normalization. We adopt $w \propto h$ only so the figures read directly as wage growth.

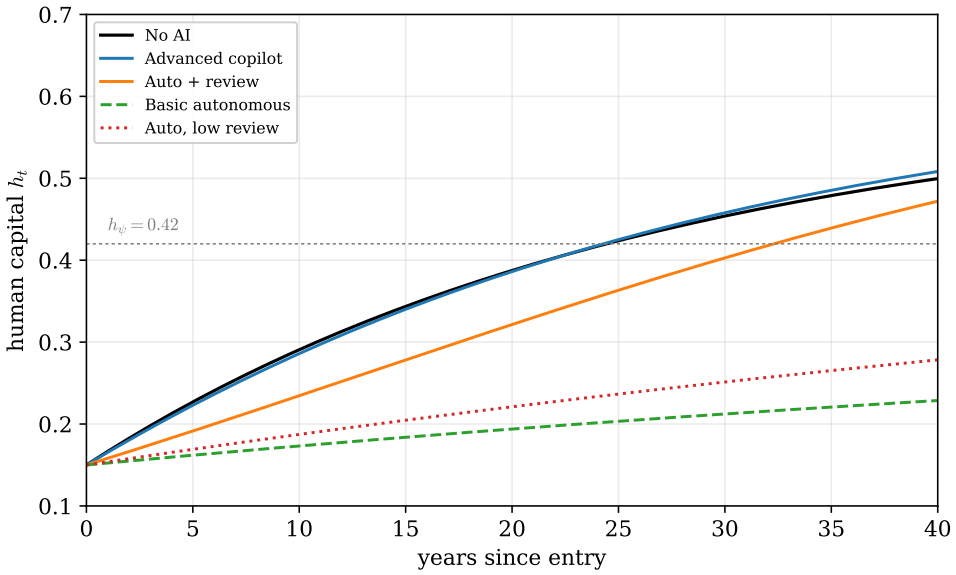


Figure 2: Career human-capital paths h_t from entry ($h_0 = 0.15$), by deployment regime. Automation regimes (dashed, dotted) flatten and, with thin review, depress the path. Over the early-career window shown, the review regime (solid) lies below copilot (the early-career loss of Proposition 2(i)), while raw entrants sit far below the scaffolding threshold $h_\psi = 0.42$ (the slow approach reflects $\alpha/2 + \delta = 0.0375$). Its advantage is asymptotic: it converges to the highest long-run ceiling among all regimes ($h^* = 0.637$, Table 1).

value the worker would pay to make deployment contractible.

Four patterns stand out for the given calibration. First, *automation hollows out the early career*: "basic autonomous AI" and especially "advanced autonomous, low review" collapse five-year wage growth (to 32.6% and 12.8%) and, in the low-review case, drag the steady state down to 0.465. This is the model's rendering of the Brynjolfsson et al. (2025a) and Hosseini and Lichtinger (2025) pattern that early-career losses concentrate in automation exposure. Second, *copilot preserves learning*: "advanced copilot" nearly matches the no-AI slope (48.6%), consistent with the large junior productivity gains from assistance in Brynjolfsson et al. (2025b) and Dell'Acqua et al. (2025). Third, *review rescues autonomy*: holding autonomy at 0.85, raising review from 0.20 to 0.90 turns a near-catastrophe (12.8%) into a survivable trajectory (27.6%) and lifts the steady state from 0.465 to 0.637. Fourth, *the review regime raises the ceiling above no-AI* ($h^* = 0.637 > 0.600$) even as it lowers early growth, because intermediate workers ($\psi(0.5) = 1.18 > 1$) learn faster from reviewing than from doing. The predictive regime, which keeps autonomy moderate while targeting review, comes closest to dominating: it nearly preserves the slope (42.2%) and raises the ceiling (0.626). These are illustrations at the maintained $\beta/\alpha = 4$, not validation of it: the first two patterns (automation hollows, copilot preserves) hold for any β/α, but the third and fourth

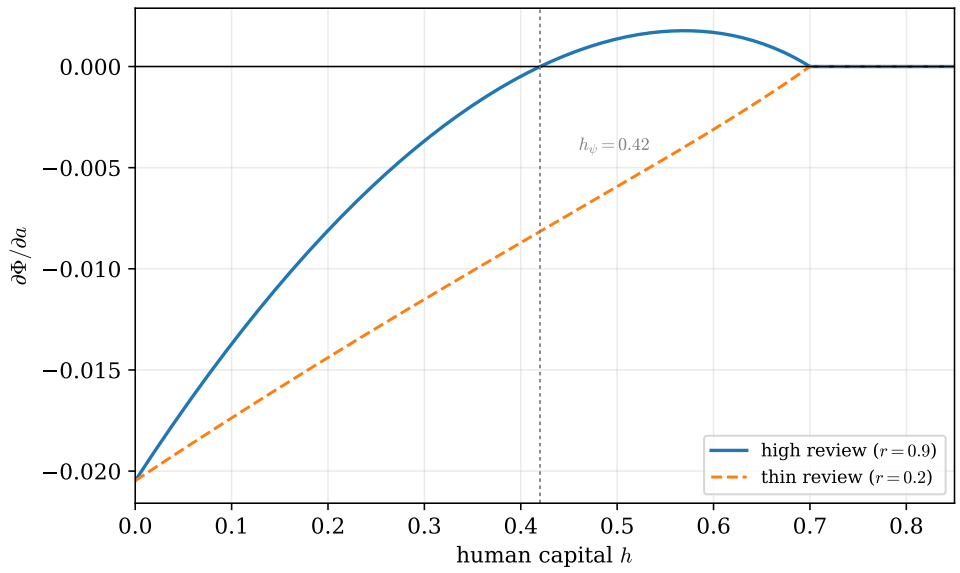


Figure 3: The marginal effect of autonomy on learning, $\partial\Phi/\partial a$, against human capital h (advanced AI, $\theta = 0.70$). Under high review the curve crosses zero at the scaffolding threshold $h_\psi = 0.42$: autonomy removes rungs for juniors ($h < h_\psi$) and builds scaffolds for intermediates and seniors ($h > h_\psi$). Under thin review it is negative throughout.Parameters: $r = 0.90, m = 0.20$ (high) and $r = 0.20, m = 0.30$ (thin).

(review rescues autonomy and the ceiling exceeds no-AI) are exactly the threshold-dependent claims, requiring β/α above the feasibility floor $(\theta^{-2} + m)/r \approx 2.5$ of Proposition 1(ii) (here comfortably met) and weakening or reversing below it.

Figure 2 shows the career paths: the review regime begins below copilot (raw entrants cannot verify) and remains below it over the early-career window, but converges to a higher long-run ceiling once the worker crosses h_ψ (Table 1). Figure 3 shows the sign reversal of $\partial\Phi/\partial a$ at $h_\psi = 0.42$. Figure 4 maps the scaffolding region in capability–seniority space. We treat these as illustrative of mechanisms, not as a causal estimate.

Because β/α is set rather than estimated, the calibration is only as informative as that choice, so we now report how the model's objects move with it and with the other primitives: the verification profile $p(h)$, AI quality $A(\theta)$, review intensity r, and the task distribution.

At the baseline (θ, r, m), a scaffolding region exists and the central reversal is operative, only if $\beta/\alpha > (\theta^{-2} + m)/r \approx 2.5$ (the feasibility condition of Proposition 1(ii)). This floor is the paper's pivotal premise and we treat it as such rather than burying it: everything distinctive, that review builds expertise above the no-AI level, that an interior watershed r^* exists, that autonomy is rescued rather than merely neutralized, holds above it and *vanishes* below it. Below the floor, $\psi < 1$ for every $h < \theta$, so autonomy only removes rungs where capable tasks remain and is neutral beyond the frontier; the safe prescription (gate juniors,

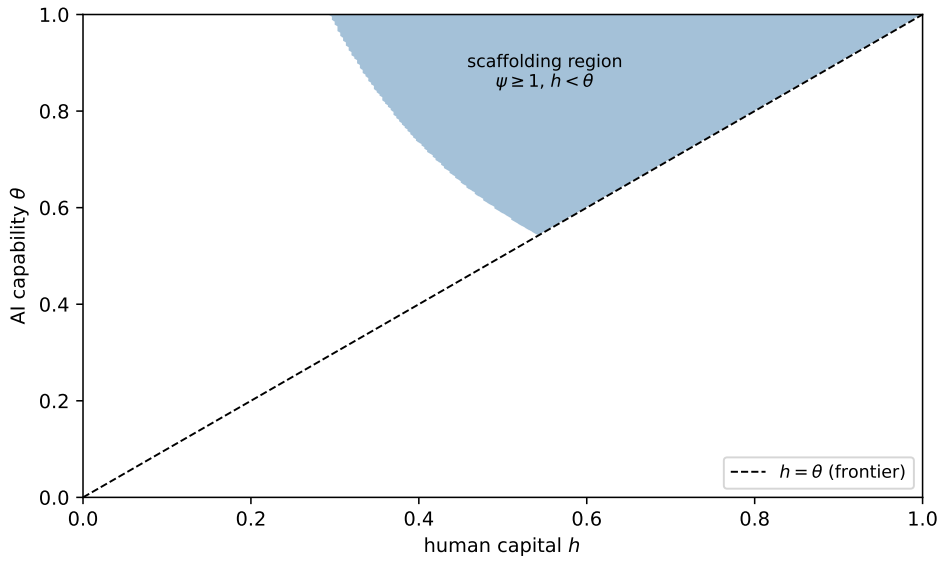


Figure 4: The scaffolding region (shaded) in (h, θ) space, where capable tasks exist ($h < \theta$, below the dashed frontier) and $\psi \geq 1$. Above the frontier the worker faces no AI-capable tasks. Inside the region, delegation-with-review out-teaches copilot mode. Parameters: $r = 0.90, m = 0.20, \beta/\alpha = 4, A(\theta) = \theta, p = h$.

Table 2: Sensitivity of the scaffolding threshold h_ψ, the watershed review intensity r^*, and the auto-plus-review steady state and five-year growth, to the central premise parameters. Baseline (boxed logic): $\beta/\alpha = 4$, $p(h) = h$, $A(\theta) = \theta$, $\theta = 0.70$, $a = 0.85$, $r = 0.90$, $m = 0.20$, $\delta = 0.015$. The *qualitative* pattern (review rescues; a watershed r^* exists; the ceiling can exceed no-AI) holds throughout; the *magnitudes* move substantially.

Parameter	value	h_ψ	r^*	5-yr growth	ceiling h^*
β/α	2.5	0.697	0.938	20.2%	0.596
	3.0	0.571	0.782	22.7%	0.614
	4.0 (baseline)	0.420	0.586	27.6%	0.637
	5.0	0.332	0.469	32.8%	0.650
	6.0	0.275	0.391	38.0%	0.659
	8.0	0.204	0.293	49.0%	0.670
$p(h) = h^{\phi_p}$	$\phi_p = 0.5$ (early verify)	0.177	0.452	55.9%	0.650
	$\phi_p = 1.0$ (baseline)	0.420	0.586	27.6%	0.637
	$\phi_p = 2.0$ (late verify)	0.648	0.983	11.4%	0.588
$A(\theta) = \theta^{\phi_A}$	$\phi_A = 0.5$ ($A = 0.84$)	0.352	0.489	—	0.648
	$\phi_A = 1.0$ ($A = 0.70$)	0.420	0.586	—	0.637
	$\phi_A = 2.0$ ($A = 0.49$)	0.600	0.842	—	0.607

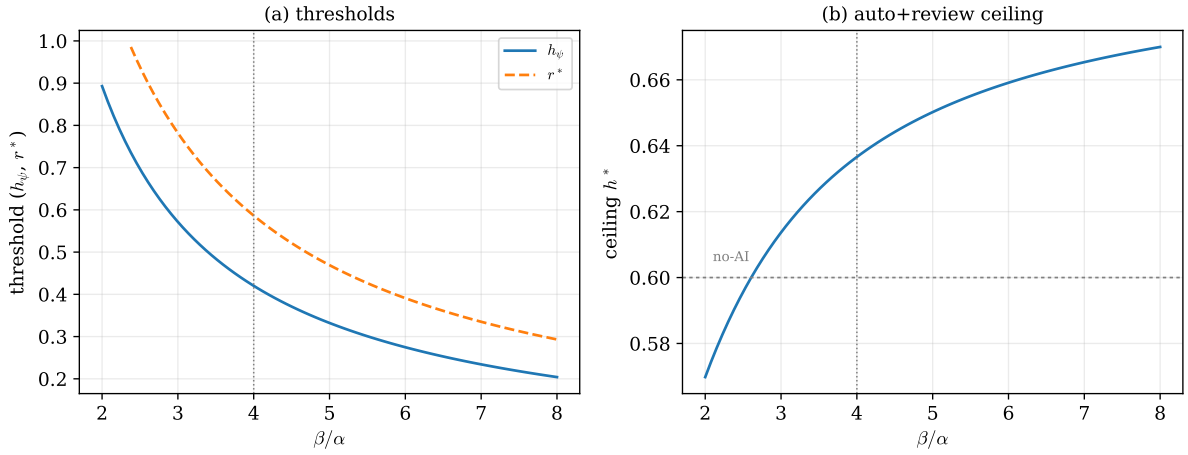


Figure 5: Sensitivity to the review-to-doing learning ratio β/α (baseline $\theta = 0.70$, $a = 0.85$, $r = 0.90$, $m = 0.20$, $A(\theta) = \theta$, p(h)=h). *(a)* The scaffolding threshold h_ψ and the watershed review intensity r^* both fall as the premise strengthens. *(b)* The auto-plus-review long-run ceiling rises with β/α and crosses the no-AI level near $\beta/\alpha \approx 3$. The qualitative structure is invariant; the thresholds and the policy intensity required are not.

keep them in copilot) still protects the worker, but it delivers no more than the no-AI ceiling, because there is nothing for review to build on. The chosen value 4 sits above this floor but not far above it, and at $\beta/\alpha = 2.5$ the auto-plus-review ceiling (0.596) lies marginally *below* no-AI while requiring near-total review ($r^* = 0.94$): the knife-edge is where the upside disappears, not where the policy becomes unsafe.

Across the entire range of Table 2, the *signs* are invariant: the autonomy reversal exists (Proposition 1), ψ is increasing in seniority and, in the main model, in capability, the watershed r^* exists (Proposition 3), and review rescues autonomy. What moves, often by a lot, are the *thresholds and policy magnitudes*. As β/α rises from 2.5 to 8, the scaffolding threshold h_ψ falls from 0.70 to 0.20 and the watershed review r^* from 0.94 to 0.29: the same managerial prescription ("mandate review") requires near-total review under a weak premise but only modest review under a strong one. The verification profile matters comparably: if verification ability ramps late ($\phi_p = 2$), h_ψ rises to 0.65, r^* to 0.98, and the ceiling falls below no-AI, so, review cannot rescue when juniors verify too slowly to reach the scaffolding region. Lower AI quality ($\phi_A = 2$) pushes the same way. Figure 5 traces the β/α dependence.

The scaffolding ratio ψ, the reversal, and the seniority threshold are invariant to the learning-content profile ℓ and to the task distribution F, which enter only the task *masses* and hence the magnitude of the effects. Table 3 confirms numerically that the regime ordering survives across four ℓ-profiles. We therefore do not re-vary F here.

The dynamic magnitudes above, the steady-state level, the five-year growth, the conver-

Table 3: Robustness across learning-content profiles ($\alpha = 0.045$, $\beta = 4\alpha$; δ set so the no-AI ceiling is $h_0^* = 0.60$ in every row). Five-year early-career growth in human capital from $h_0 = 0.15$ (read as wage growth under $w \propto h$, text), and the auto-plus-review ceiling. In every profile: copilot preserves growth, automation-with-thin-review depresses it, mandated review rescues it, and high-autonomy-with-review lifts the long-run ceiling above the no-AI level 0.60.

		five-year early-career growth (h)				ceiling
Profile $\ell(z,h)$	δ	no AI	copilot	auto, low rev.	auto + rev.	$h^*_{\text{auto+rev}}$
decreasing $1-z$ (baseline)	0.015	51.3%	48.6%	12.8%	27.6%	0.637
constant 1	0.075	93.8%	90.9%	33.4%	60.9%	0.649
increasing z	0.060	51.3%	49.2%	22.3%	33.4%	0.658
frontier $1-(z-h)$ (ZPD)	0.060	51.3%	47.7%	1.7%	19.7%	0.668

gence speed, run through the level and curvature of the aggregate learning rate $\Phi(h)$, and so, unlike the per-task reversal and threshold (which Section 4 and Appendix A establish for arbitrary F and ℓ), they are *not* invariant to the learning-content profile. We therefore verify numerically that the calibration's *qualitative* conclusions do not depend on the decreasing profile $\ell = 1 - z$ maintained above. Table 3 re-solves the model under four profiles: decreasing ($\ell = 1 - z$), constant ($\ell = 1$), increasing ($\ell = z$), and a state-dependent "frontier" profile $\ell(z,h) = 1 - (z - h)$ peaked at the worker's own frontier, recalibrating δ in each so that the no-AI ceiling is held at the common baseline $h_0^* = 0.60$, and holding the premise $\beta/\alpha = 4$ fixed. Across all four profiles the three signatures survive: automation with thin review depresses early growth relative to copilot, mandated review rescues autonomy, and high-autonomy-with-review lifts the long-run ceiling above the no-AI level, the latter two conditional, on β/α clearing the feasibility floor $(\theta^{-2}+m)/r \approx 2.5$ of Proposition 1(ii), which the maintained value does.

The frontier profile is the sharpest test. Under $\ell(z,h) = 1 - (z - h)$, learning content is highest for tasks at the worker's own frontier ($z = h$) regardless of seniority, so a senior learns most from their hardest tasks and total learnable mass *rises* with h. The "seniors face lower-content tasks" feature of the baseline disappears entirely. Yet the AI results are unchanged, and the long-run ceiling gain from autonomy-with-review is in fact largest (0.668). The decreasing-content profile thus shapes the baseline life cycle and the dynamic magnitudes but leaves the calibration's qualitative conclusions intact.

8 A Microfoundation of Learning by Reviewing

The analysis so far has taken the review-learning rate $\beta\, r\, A(\theta)\, p(h)$ as a primitive: reviewing transfers learning that increases in AI quality and in verification ability. On that primitive rest the autonomy reversal and scaffolding ratio of Section 4, the dynamic career path, watershed, and obsolescence results of Section 5, and the deployment externality of Section 6. This section derives that review-learning technology from a model of *verification, error detection, and explanation*, so that the conversion rate is an *output* rather than an assumption. Doing so does three things: it grounds the comparative statics that the earlier sections derived. This makes the primitives measurable, and most importantly it shows the central reversal is *conditional*, identifying the settings in which reviewing teaches more than doing and the settings in which it teaches less or even produces negative learning. The main results of Sections 4–6 survive this grounding. What the microfoundation adds is a precise account of *when* they hold and one change to the long-run dynamic structure that we isolate at the end.

Decompose a capable task into a unit measure of *knowledge components* (sub-steps), of which the learnable fraction is $\ell(z)$. The AI completes the task. The worker reviews it. For a representative learnable component:

- the AI's component is *correct* with probability $q = A(\theta)$ and *erroneous* with probability $1 - q$ (higher-capability AI errs less).
- the worker *verifies correctly* (she confirms a correct component or detects an erroneous one) with probability $p = p(h)$, the verification ability.
- the AI's reasoning is *exposed* (chain-of-thought, derivation, citations) with probability $e \in [0, 1]$, the *explanation quality* of the deployment, and otherwise the worker sees a correct-or-incorrect answer without the reasoning.

Learning accrues per component as follows. A *correct, verified* component (prob. qp) transfers knowledge worth $\omega(e) \equiv e + (1 - e)\gamma$: the full unit if the reasoning is exposed, and only $\gamma \in [0, 1)$ if the worker must learn from an opaque correct answer (confirming an answer whose derivation is hidden). A *detected error* (prob. $(1 - q)p$) is a high-value diagnostic event worth $\nu \geq 0$: locating and correcting a mistake teaches the boundary between right and wrong. An *undetected error* (prob. $(1 - q)(1 - p)$) is absorbed as if correct, corrupting the worker's knowledge at cost $\chi \geq 0$: *mislearning*. Correct but unverified components contribute negligibly and are normalized to zero. The net review value per component is

therefore

$$v_R(q,e,p) \;=\; \underbrace{p\,q\,\omega(e)}_{\text{learn from correct output}} \;+\; \underbrace{p\,(1-q)\,\nu}_{\text{learn from detected errors}} \;-\; \underbrace{(1-p)\,(1-q)\,\chi}_{\text{mislearning}}, \tag{24}$$

while performing the component hands-on teaches a constant $d > 0$ (learning by doing, independent of the AI). Reviewing a delegated component therefore out-teaches performing it precisely when $v_R(q,e,p) > d$.

Proposition 6 (Comparative statics of review-learning)**.** *The review value* (24) *satisfies:*

(i) ***Transparency.*** $\partial v_R/\partial e = p\,q\,(1-\gamma) > 0$*: explanation quality raises review-learning.*

(ii) ***Verification ability.*** $\partial v_R/\partial p = q\,\omega(e) + (1-q)\nu + (1-q)\chi > 0$*: ability raises it through more confirmation, more error detection, and less mislearning. Review-learning is monotone in seniority through* $p(h)$*.*

(iii) ***AI quality is non-monotone.*** $\partial v_R/\partial q = p\,[\,\omega(e) - \nu\,] + (1-p)\chi$*. When error diagnosis is the more valuable channel* ($\nu > \omega(e)$)*, this is* positive for low-ability workers *(mislearning reduction dominates) and* negative for high-ability workers *(the loss of instructive errors dominates): better AI helps novices review-learn but can hurt experts.*

(iv) ***Mislearning region.*** $v_R < 0$*, reviewing produces net-negative learning when verification ability and explanation quality are low and the error rate is non-trivial:* $(1-p)(1-q)\chi > p\,[\,q\,\omega(e) + (1-q)\nu\,]$.

The distinctive results of the main linear model require the learning ratio β/α to clear the feasibility floor $(\theta^{-2}+m)/r$ of Proposition 1(ii). That aggregate is calibrated, not estimated, and we do not claim otherwise. The microfoundation does not estimate it either, but it does two things that bear on how much weight the floor must carry. First, it recasts the premise as a *local* condition on objects the firm measures: at a (task, worker) cell, reviewing out-teaches doing exactly when $v_R(A(\theta), e, p(h)) > d$, with $q = A(\theta)$ from internal evaluation, $p(h)$ from seeded-error audits, e from the exposure of AI reasoning, and the benchmark d from a controlled review-versus-doing comparison (Section 9). Under the reduced-form identification that matches $\kappa\, v_R$ to the linear review rate $\beta A(\theta)p(h)$ at the reference cell (and d to the doing rate), this measurable condition is *no more permissive* than the aggregate floor: because v_R sits below its linear image by the mislearning gap $(1-p)(1-q)\chi$, one has $v_R > d \implies \beta/\alpha > (\theta^{-2}+m)/r$, with equality only in the degenerate full-verification, full-capability limit. Conditioning the rule on $v_R > d$ therefore cannot manufacture feasibility the asserted aggregate lacks; if anything it is the stricter test.[7]

[7]The constant $\kappa > 0$ is a unit-conversion factor, not a structural parameter: the main-model rate

Second, and more importantly for the policy reading, the optimal deployment is *safe for any value of* β/α. The seniority-contingent rule grants autonomy with review only where $v_R > d$ and withholds it elsewhere, so it never deploys autonomy into the mislearning region whatever the learning ratio turns out to be. When β/α clears the floor, the rule does this by *gating*, routing the worker through copilot into a nonempty scaffolding region. When β/α lies below the floor, the scaffolding region is empty and the same rule does it by *abstaining*, keeping the worker in copilot throughout. In neither case is autonomy granted where review would teach less than doing. This separates two claims that the calibration should not be asked to bear jointly. The *safety* of the prescription, that configuring AI at the task–seniority cell never converts it into a rung-remover, is unconditional. The *upside*, that review builds expertise strictly above the no-AI level, that an interior watershed r^* exists, and that autonomy is rescued rather than merely neutralized, is conditional on $\beta/\alpha > (\theta^{-2} + m)/r$ and is the genuine empirical premise of the paper. We do not infer that premise from the optimal rule, doing so would be circular, since the rule avoids the mislearning region in both regimes. We state it as conditional and identify, in $v_R > d$, the measurable object on which it turns.

One implication of Proposition 6 is worth drawing out separately, because it is not visible in the linear model. The monotonicity of ψ in AI *capability* is *not* robust: by Proposition 6(iii), more capable AI can *reduce* expert review-learning by removing the errors experts learn most from. Figure 6 plots how the review value v_R depends on verification ability, explanation quality, and AI capability.

The microfounded learning rate v_R can be negative and non-monotone in capability, where the main model's review rate $\beta A(\theta)p(h)$ is linear and everywhere positive. The two therefore cannot be globally identical: replacing $\beta A(\theta)p(h)$ by $\kappa\, v_R(A(\theta), e, p(h))$ in Φ makes $\Phi(h) < 0$ for raw juniors under a fixed high-autonomy deployment, so $h = 0$ becomes an absorbing low-expertise trap that the linear model, with $\Phi \geq 0$, can never produce.

The divergence is, however, confined to a region that the model's own prescription tells the firm to avoid, and this yields a precise sense in which the two models agree. The trap is reachable only under a deployment that holds autonomy high on workers who cannot yet verify, that is, on the segment $h < h_\psi$ where $v_R < 0$. The optimal, seniority-contingent rule (Section 4) withholds autonomy exactly there. Along the trajectory that rule induces, the worker never enters the mislearning region, the microfounded drift coincides in sign and steady state with the main model, and the two systems are equivalent. The next proposition

$\beta A(\theta)p(h)$ is in human-capital-per-unit-time units (in which β is calibrated), whereas v_R in (24) is built from per-component primitives (γ, ν, χ, d) on a scale where a fully transparent, correctly verified component is worth one unit. Matching $\kappa\, v_R(A(\theta), e, p(h)) = \beta A(\theta)p(h)$ at the reference point $h = h_\psi$ fixes κ. Being a positive rescaling, it leaves every sign and monotonicity property of v_R unchanged.

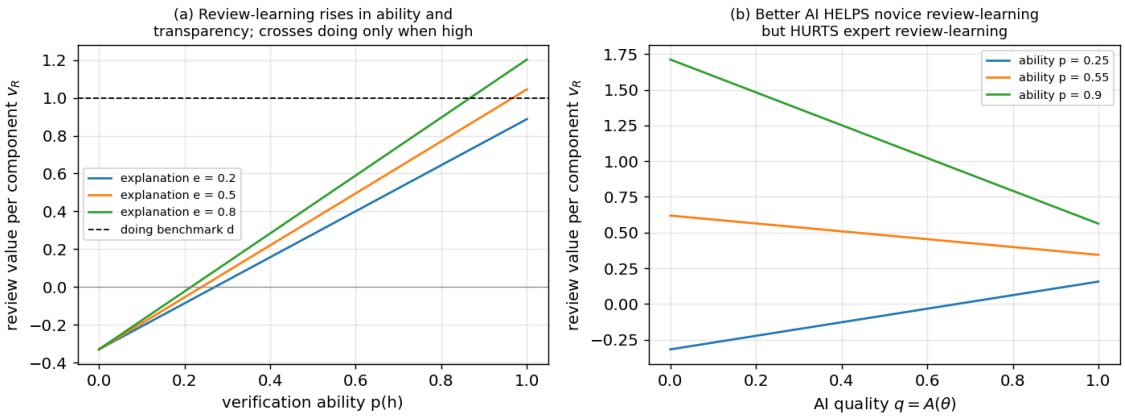


Figure 6: Microfounded review value v_R. *(a)* v_R rises in verification ability $p(h)$ and in explanation quality e, crossing the learning by doing benchmark d only for sufficiently able reviewers (the conditional reversal). *(b)* The effect of AI quality $q = A(\theta)$ on review-learning is positive for low-ability reviewers (mislearning reduction) and negative for high-ability reviewers (loss of instructive errors). Parameters (per-component illustrative units): $\gamma = 0.25$, $\nu = 2.025$, $\chi = 1.10$, $d = 1.00$, $\omega(e) = e + (1-e)\gamma$. Panel (a) fixes $q = 0.70$, varies $e \in \{0.2, 0.5, 0.8\}$; panel (b) fixes $e = 0.50$, varies $p \in \{0.25, 0.55, 0.9\}$.

states this.

Proposition 7 (Conditional equivalence under optimal deployment)**.** *Let* Φ^{lin} *use the linear review rate* $\beta r A(\theta) p(h)$ *and* Φ^{mic} *the microfounded rate* $r\kappa v_R(A(\theta), e, p(h))$*, with all other channels identical.*

(i) Global non-equivalence. *Under a deployment that holds autonomy high from entry,* $\Phi^{\text{mic}}(0) < 0$ *while* $\Phi^{\text{lin}}(0) \geq 0$*, so the microfounded dynamics have an absorbing boundary at* $h = 0$ *and an unstable interior edge* h_{trap} *that the linear dynamics lack: the long run is history-dependent under* Φ^{mic} *and unique under* Φ^{lin}*.*

(ii) Conditional equivalence. *Under the seniority-contingent rule that withholds autonomy for* $h < h_\psi$*,* $\Phi^{\text{mic}}(0) \geq 0$*: the lower boundary is repelling,* h_{trap} *vanishes, and* Φ^{mic} *admits a unique, globally asymptotically stable interior steady state. On the induced trajectory* Φ^{mic} *and* Φ^{lin} *agree in sign throughout and share the same reversal, the same review-controlled long-run sign, and the same comparative statics. The linear model is thus the reduced form of the microfoundation along the optimal-deployment path.*

Two points complete the reconciliation. First, the equivalence is not an artifact of the linear form's tractability: it holds because the one property the main results actually use, that review-learning rises with seniority, holds *globally* in the microfoundation ($\partial v_R / \partial p > 0$ for all (q, e, p)), so the linear model gets the sign of the seniority gradient right everywhere, and exactly right on the optimal path. Second, the divergence that remains is not a defect but a

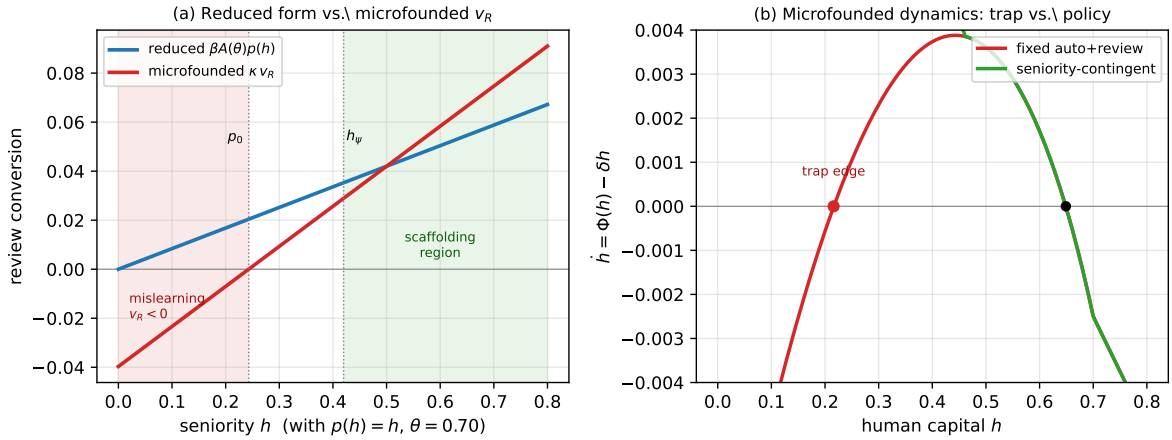


Figure 7: Microfoundation versus linear model. *(a)* The main-model review conversion $\beta A(\theta)p(h)$ and the microfounded $\kappa\, v_R$ (at $\theta = 0.70$, $p(h) = h$): they diverge in *sign* only for raw juniors ($h < p_0 = 0.24$, mislearning), and agree in sign on the scaffolding region ($h > h_\psi = 0.42$). *(b)* Microfounded dynamics $\dot{h} = \Phi(h) - \delta h$: a fixed high-autonomy-with-review deployment makes $h = 0$ an absorbing low-expertise boundary, with an unstable interior edge at $h_{\text{trap}} = 0.22$ and a high stable ceiling ($h = 0.65$); the seniority-contingent deployment makes $\Phi(0) > 0$, so $h = 0$ is repelling, leaving the single high steady state on the optimal path. Parameter values: $\alpha = 0.045$, $\beta = 0.18$, $\delta = 0.015$, $\theta = 0.70$, $A(\theta) = \theta$, $p(h) = h$, $\ell(z) = 1 - z$, $z \mid h \sim \text{Uniform}[h, 1]$. Microfoundation primitives: $\gamma = 0.25$, $\nu = 2.025$, $\chi = 1.10$, explanation quality $e = 0.50$; the unit-conversion constant $\kappa = 0.214$ is fixed by matching $\kappa\, v_R = \beta A(\theta)p(h)$ at $h = h_\psi$. Deployments: fixed high-autonomy-with-review $(a, r, m) = (0.85, 0.90, 0.20)$; seniority-contingent deployment uses copilot $(0.15, 0.20, 0.70)$ for $h < h_\psi$ and $(0.85, 0.90, 0.20)$ for $h \geq h_\psi$. $h_0 = 0.15$.

sharper statement of the paper's thesis: the microfoundation says that deploying autonomy on workers who cannot verify does not merely slow learning but can trap it, so setting deployment at the task–seniority level is, under the microfoundation, not merely optimal but *necessary* to avoid a low-expertise basin. Re-solving with $\kappa\, v_R$ confirms the magnitudes (Figure 7): along the seniority-contingent path the autonomous-plus-review ceiling is 0.659, above both the no-AI level 0.60 and the linear-model ceiling 0.637, while a fixed high-autonomy deployment from entry sends a raw entrant ($h_0 = 0.15$) into the trap at $h = 0$.

9 Managerial and Organizational Implications

The model is, at bottom, a theory of an operating decision: how to configure AI across a workforce. We translate it into an operational framework: what a manager observes, how the levers are measured, and the rule they imply and then draw four implications.

Each primitive maps to an object a manager can read from existing telemetry or elicit

cheaply (Table 4). The deployment levers (a, r, m) are already logged by most AI-in-the-loop workflows. Capability θ is an internal evaluation. The one primitive not in standard telemetry, verification ability $p(h)$, is recovered by *seeded-error audits*, planting known errors in AI output and recording the detection rate by seniority. It is also the pivotal input, since it sets where each worker sits relative to the thresholds.

From these, the manager computes the scaffolding ratio $\psi(z, h)$ per cell, the threshold h_ψ, the loss threshold h_c, and the watershed review intensity r^* for each role.

This measurability sits in no tension with the non-contractibility that drives the externality of Section 6. The two refer to different standards of proof. The telemetry above is *internally observable*: a noisy signal that the firm reads for its own deployment decisions. For that use it need not be gaming-proof, because the firm sets the configuration to raise its own match value and has no incentive to fool its own signal. Even a coarse, biased proxy for review intensity supports a sensible cell-level rule. Contractibility is a higher bar: a variable is contractible only if a third party enforcing a private contract or a regulator disbursing a transfer, can verify it. Edit distance, check time, and seeded-error detection rates fail that bar precisely where it would bite, because they can be inflated by box-checking that logs the act of review without producing the learning it is meant to proxy. The object the externality turns on is therefore not whether review *occurred* but whether it *taught*: a firm can log the former for itself, yet neither it nor an outside party can verify the latter in an enforceable, gaming-proof way. Observability for the firm's own use and verifiability to a third party are thus distinct, and the externality requires only the absence of the second. The two are endpoints of a continuum: as review becomes more verifiable to outsiders the setting moves toward the contractible benchmark (Case 3 of Section 6), the wedge shrinks, and instruments that condition on review become more feasible, so the size of the problem and the feasibility of contracting on review co-move rather than conflict.

The policy is a rule over *(task type, worker) cells*, not a firm-wide automation target (Figure 8). For each cell, classify the task by whether it lies beyond the AI frontier and the worker by where h sits relative to the thresholds:

(1) *Beyond the frontier* ($z > \theta$): the worker performs the task. This is where frontier learning and the bottom rung of the ladder live. Protect it from premature automation.

(2) *Capable task, worker below* h_c (cannot yet verify, $\psi < 1$ and copilot cannot compensate): *copilot mode*, low autonomy, high assist and withhold autonomy. Delegating here removes rungs and (under the microfoundation) risks the junior trap.

(3) *Capable task, worker in the scaffolding region* ($h > h_\psi$, $\psi > 1$): *delegate aggressively and mandate review at* $r \geq r^*$. Reviewing expert-grade output out-teaches doing, this is where autonomy builds expertise.

Table 4: Mapping the model to measurable managerial objects.

Model object	What it is	How a firm measures it
θ (capability)	expert-grade complexity frontier of the AI	internal evaluation: grade a task sample by difficulty, record the AI's expert-verified pass rate. θ is where it falls off
a (autonomy)	share of capable tasks shipped without human authorship	workflow logs: fraction of outputs (commits, drafts, tickets, analyses) AI-authored and accepted without substantive edit
r (review)	fraction and depth of AI outputs verified	review logs and version diffs: share of AI outputs checked, time per check, edit distance, corrections logged
m (copilot)	assist intensity on retained tasks	tool telemetry: suggestion-acceptance rate under human authorship
$p(h)$ (verification)	probability the worker detects an AI error	*seeded-error audits*: inject known errors, record detection rate by level
h (seniority)	human capital / verification competence	level, tenure, or skill assessment
β/α (learning ratio)	review-to-doing learning premium	controlled review-vs-doing comparison (Section 8)

(4) *Intermediate band* ($h_c < h < h_\psi$): copilot still dominates, but raise review as $p(h)$ rises to move the worker into the scaffolding region.

Setting one mix per occupation or firm is the wrong instrument: a uniform "automate X%" over-automates juniors (regions 1–2) and under-uses the scaffolding region for seniors (region 3).

The one entry that is not read directly from telemetry is the learning ratio β/α itself, the premise on which the distinctive results turn. It is identified, in principle, by a controlled comparison that isolates the learning channel from the output channel: matched workers at a common seniority are assigned either to *review* AI-produced solutions to a set of tasks or to *perform* those same tasks unaided, holding fixed the task set, the time on task, and the AI output reviewed. Learning is then read not from performance during the episode but from *subsequent independent* performance on held-out tasks of the same type, so that the contemporaneous productivity boost from AI assistance does not contaminate the estimate of what was learned. The ratio of the two subsequent-performance gains, review-induced to doing-induced, estimates β/α at that seniority, repeating across seniorities traces its profile. We emphasize that the existing field evidence on generative AI measures the output channel, not this learning premium, and so does not identify β/α. The design above is what

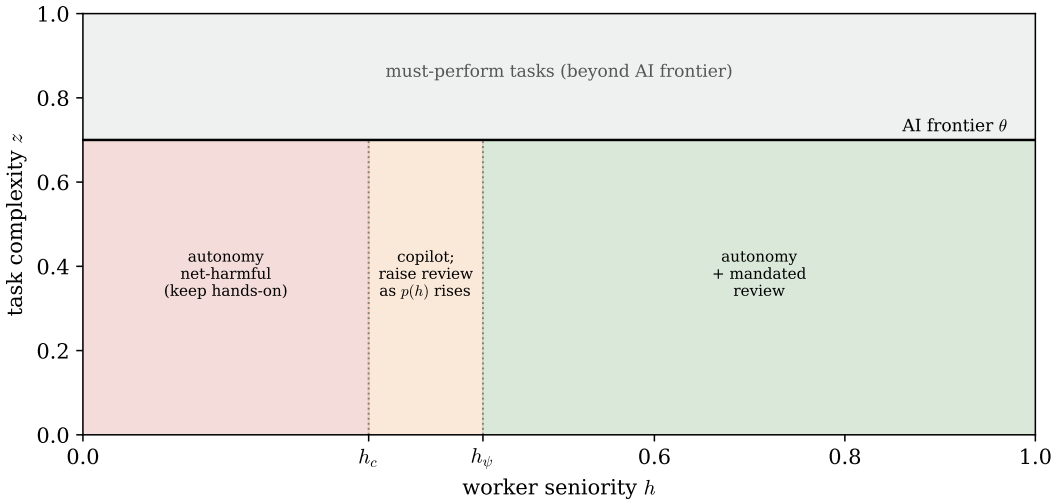


Figure 8: The deployment rule by task–seniority cell. The AI capability frontier θ separates tasks the worker must perform (top) from capable tasks; among capable tasks, the scaffolding threshold h_ψ separates juniors who should be kept in copilot mode with autonomy withheld (left) from seniors for whom autonomy should be paired with mandated review (right). The loss threshold $h_c \leq h_\psi$ marks where pure autonomy turns net learning negative.

the parameter would require. Until it is run, β/α is calibrated rather than estimated, which is why we carry the feasibility floor as an explicit conditional premise rather than a settled magnitude.

The manager monitors ψ per cell and watches for three failure modes the framework names precisely: the *hollow job* is the one with high a, low r, $\psi < 1$, where output is high but learning collapses. The *junior trap* is realized when raw entrants held at high autonomy mislearn into a low-expertise basin. *Review exhaustion* is another possibility with fast-churn skills where $r^* > 1$. Understanding the causes of a failure is a necessary first step for making the proper decisions. Those can include raising r where $p(h)$ allows it or else holding in copilot (*hollow job*), making the amount of autonomy a worker is given conditional on their seniority (*junior trap*) or redesigning AI configuration rather than mandating more review (*review exhaustion*).

Four implications follow. They divide into a part that holds whatever the learning ratio β/α turns out to be and a part contingent on it: the rule below is *safe* for any β/α, never deploying autonomy where review would teach less than doing, while the *gains* it can deliver, expertise above the no-AI level and a review intensity that rescues autonomy, materialize only where β/α clears the feasibility floor $(\theta^{-2} + m)/r$ (the conditional premise made precise in Section 8). *Treat AI configuration as a talent-development decision, not only a productivity decision.* The autonomy, review, and assist intensities (a, r, m) that maximize current

human–AI output are generally not those that develop the junior pipeline. A firm that optimizes only the first objective (Proposition 4) drives its workforce toward a low-expertise steady state. The configuration choice should be owned jointly by operations and talent functions, not by either alone.

Set the configuration at the task–seniority cell, not the occupation or the firm. The optimal mix is contingent on where the worker sits relative to the scaffolding threshold h_ψ (Proposition 1). For raw juniors (below h_ψ, low verification ability), keep humans in copilot mode and reserve autonomy for tasks beyond their reach. For intermediate and senior staff (above h_ψ), delegate aggressively *and mandate review*, because reviewing expert-grade output now out-teaches hands-on practice. A single firm-wide "automate X%" target is the wrong instrument: it over-automates juniors and under-uses the scaffolding region for seniors.

Make review a managed, measured input, not an afterthought. Mandated review is what converts autonomy from learning-destroying to learning-building (Table 1: raising r from 0.20 to 0.90 at fixed autonomy lifts five-year wage growth from 12.8% to 27.6% and the expertise ceiling from 0.465 to 0.637). Review intensity, review quality, and who reviews should be tracked like any other operating metric. Human mentorship and AI-assisted review are substitutes in one congestible channel: managers who let AI raise nominal review capacity while thinning senior ranks can leave juniors below the viability threshold.

Anticipate that competition will under-provide development. Because general human capital is portable and "how the AI is run" is hard to contract on, the labor market does not reward a firm for developing skills a departing worker takes elsewhere (Section 6). Firms that nonetheless want the higher-expertise steady state must substitute internal commitment devices, explicit development tracks, retention-linked investment, rotation through review-intensive roles for the missing market contract. This is also a recruiting proposition: the worker would pay (in accepted wage) for a credibly development-intensive deployment, so credibly offering one is a way to attract talent at lower wage cost.

10 Conclusion

This paper builds around an overlooked channel in the economics of AI and early careers: workers can learn by reviewing what AI does on its own, and reviewing expert-grade output can teach more than performing routine work oneself. This turns the worry that AI removes the early rungs of the career ladder into a sharper, conditional claim. Letting AI work autonomously removes rungs when workers cannot yet judge its output and the AI is not yet good enough to be worth learning from. It builds new rungs when workers can judge the output and the AI is good enough. Which of the two occurs shifts with the worker's

experience and with how capable the AI is: the same arrangement that slows a novice's early growth can lift the expertise a seasoned worker ultimately reaches. And firms that expect workers to leave soon tend to configure AI shortsightedly, leaning on automation and skimping on the human review that builds lasting skill.

The implications for policy follow directly from this mechanism, and they concern how AI is configured, not whether it is allowed. First, the remedy for hollowed-out junior work is to pair automation with review rather than to cap automation. Handing work to AI erodes learning only when too little of its output is reviewed. The fix is to raise the intensity of review wherever workers are able to check the AI's work, not to hold automation back, which would sacrifice the greater expertise that review-paired automation can ultimately build. A blunt "automate this share of the work" rule is the wrong instrument, because it automates too much of what juniors should still be doing while leaving senior workers under-served. The better rule grants autonomy gradually, as workers become able to supervise it. Second, the gains and losses from how AI is deployed fall on different parties, and that gap needs closing. Because firms that cannot retain their workers skimp on review and over-rely on automation, a competitive labor market provides too little of the review that builds the kind of skill a worker carries to their next employer. What the market leaves in short supply is precisely the learning value of review, skill that is portable across employers and that no single firm can be paid to provide; closing the gap therefore requires giving firms a larger stake in development they cannot otherwise capture. The common thread is that the decision that matters is how automation and human review are combined, task by task and at each stage of a career.

Among jobs exposed to AI, what happens to early-career workers depends on how automation and human review are combined and on whether workers are able to check the AI's work, not on mere exposure to AI itself.

References

Acemoglu, D. and D. Autor (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics*, Vol. 4B, 1043–1171.

Acemoglu, D. and J.-S. Pischke (1998). Why Do Firms Train? Theory and Evidence. *Quarterly Journal of Economics* 113(1), 79–119.

Acemoglu, D. and J.-S. Pischke (1999). The Structure of Wages and Investment in General Training. *Journal of Political Economy* 107(3), 539–572.

Acemoglu, D. and P. Restrepo (2018). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review* 108(6), 1488–1542.

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

Acemoglu, D. and P. Restrepo (2020). Robots and jobs: Evidence from U.S. labor markets. *Journal of Political Economy*, 128(6), 2188–2244.

Acemoglu, D. and P. Restrepo (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica* 90(5), 1973–2016.

Adda, J. and C. Dustmann (2023). Sources of wage growth. *Journal of Political Economy* 131(2), 456–503.

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

Arrazola, M. and J. de Hevia (2004). More on the estimation of the human capital depreciation rate. *Applied Economics Letters* 11(3), 145–148.

Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the U.S. labor market. *American Economic Review*, 103(5), 1553–1597.

Autor, D., F. Levy, and R. Murnane (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *Quarterly Journal of Economics* 118(4), 1279–1333.

Autor, D. and N. Thompson (2025). Expertise. *Journal of the European Economic Association* 23(4), 1203–1271.

Baker, G., M. Gibbs, and B. Holmstrom (1994). The Internal Economics of the Firm: Evidence from Personnel Data. *Quarterly Journal of Economics* 109(4), 881–919.

Becker, G. S. (1964). *Human Capital: A Theoretical and Empirical Analysis*. University of Chicago Press.

Ben-Porath, Y. (1967). The Production of Human Capital and the Life Cycle of Earnings. *Journal of Political Economy* 75(4), 352–365.

Bloom, N., L. Garicano, R. Sadun, and J. Van Reenen (2014). The Distinct Effects of Information Technology and Communication Technology on Firm Organization. *Management Science* 60(12), 2859–2885.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024). Artificial Intelligence, Firm Growth, and Product Innovation. *Journal of Financial Economics* 151, 103745.

Barbieri, C. A., D. Miller-Cotto, S. N. Clerjuste, and K. Chawla (2023). A Meta-Analysis of the Worked Examples Effect on Mathematics Performance. *Educational Psychology Review* 35(1), Article 11.

Brynjolfsson, E., B. Chandar, and R. Chen (2025). 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, and L. R. Raymond (2025). Generative AI at Work. *Quarterly Journal of Economics* 140(2), 889–942.

Chi, M. T. H., M. Bassok, M. W. Lewis, P. Reimann, and R. Glaser (1989). Self-Explanations: How Students Study and Use Examples in Learning to Solve Problems. *Cognitive Science* 13(2), 145–182.

Collins, A., J. S. Brown, and S. E. Newman (1989). Cognitive Apprenticeship: Teaching the Crafts of Reading, Writing, and Mathematics. In L. B. Resnick (Ed.), *Knowing, Learning, and Instruction: Essays in Honor of Robert Glaser*, pp. 453–494. Hillsdale, NJ: Lawrence Erlbaum.

Cooper, G. and J. Sweller (1987). Effects of Schema Acquisition and Rule Automation on Mathematical Problem-Solving Transfer. *Journal of Educational Psychology* 79(4), 347–362.

Cui, K. Z., M. Demirer, S. Jaffe, L. Musolff, S. Peng, and T. Salz (2026). The effects of generative AI on high-skilled work: Evidence from three field experiments with software developers. *Management Science*, 0(0).

Crane, L. D. and P. E. Soto (2026). AI and Coder Employment: Compiling the Evidence. Finance and Economics Discussion Series 2026-018, Board of Governors of the Federal Reserve System.

Dell'Acqua, F., E. McFowland III, E. R. Mollick, H. Lifshitz-Assaf, K. Kellogg, S. Rajendran, L. Krayer, F. Candelon, and K. R. Lakhani (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organization Science* 37(2), 403–423.

Dell'Acqua, F., C. Ayoubi, H. Lifshitz, R. Sadun, E. Mollick, L. Mollick, Y. Han, J. Goldman, H. Nair, S. Taub, and K. R. Lakhani (2025). The cybernetic teammate: A field experiment on generative AI reshaping teamwork and expertise. NBER Working Paper 33641.

Deming, D. J. (2023). Why do wages grow faster for educated workers? NBER Working Paper 31373.

Dinerstein, M., R. Megalokonomou, and C. Yannelis (2022). Human capital depreciation and returns to experience. *American Economic Review* 112(11), 3725–3762.

Dustmann, C. and U. Schönberg (2012). What Makes Firm-Based Vocational Training Schemes Successful? The Role of Commitment. *American Economic Journal: Applied Economics* 4(2), 36–61.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science* 384(6702), 1306–1308.

Farber, H. S. and R. Gibbons (1996). Learning and Wage Dynamics. *Quarterly Journal of Economics* 111(4), 1007–1047.

Fuchs, W., L. Garicano, and L. Rayo (2015). Optimal Contracting and the Organization of Knowledge. *Review of Economic Studies* 82(2), 632–658.

Fügener, A., J. Grahl, A. Gupta, and W. Ketter (2021). Will Humans-in-the-Loop Become Borgs? Merits and Pitfalls of Working with AI. *MIS Quarterly* 45(3), 1527–1556.

Fügener, A., J. Grahl, A. Gupta, and W. Ketter (2022). Cognitive Challenges in Human–Artificial Intelligence Collaboration: Investigating the Path Toward Productive Delegation. *Information Systems Research* 33(2), 678–696.

Friebel, G., Y. Huang, J. Li, S. Shukla, and A. Zhang (2026). Pyramids, Diamonds, and Oscillations: AI and the Structure of Internal Labor Markets. Working paper.

Garicano, L. (2000). Hierarchies and the Organization of Knowledge in Production. *Journal of Political Economy* 108(5), 874–904.

Garicano, L. and L. Rayo (2026). Training in the Age of AI: A Theory of Career Viability. CEPR Discussion Paper 20634 (updated).

Garicano, L. and E. Rossi-Hansberg (2006). Organization and Inequality in a Knowledge Economy. *Quarterly Journal of Economics* 121(4), 1383–1435.

Gathmann, C. and U. Schönberg (2010). How General Is Human Capital? A Task-Based Approach. *Journal of Labor Economics* 28(1), 1–49.

Gibbons, R. and M. Waldman (1999). A Theory of Wage and Promotion Dynamics Inside Firms. *Quarterly Journal of Economics* 114(4), 1321–1358.

Gibbons, R. and M. Waldman (2006). Enriching a Theory of Wage and Promotion Dynamics Inside Firms. *Journal of Labor Economics* 24(1), 59–108.

Goos, M., A. Manning, and A. Salomons (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8), 2509–2526.

Hampole, M., D. Papanikolaou, L. D. W. Schmidt, and B. Seegmiller (2025). Artificial Intelligence and the Labor Market. NBER Working Paper No. 33509, National Bureau of Economic Research.

Hosseini, S. M. and G. Lichtinger (2025). Generative AI as seniority-biased technological change: Evidence from U.S. résumé and job posting data. Working paper, SSRN 5425555.

Ide, E. and E. Talamàs (2025). Artificial Intelligence in the Knowledge Economy. *Journal of Political Economy* 133(12), 3762–3800.

Ide, E. and E. Talamàs (2026). The Impact of AI on Global Knowledge Work. *Journal of Monetary Economics* 157, 103876.

Kahn, L. B. (2010). The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy. *Labour Economics* 17(2), 303–316.

Katz, E. and A. Ziderman (1990). Investment in General Training: The Role of Information and Labour Mobility. *Economic Journal* 100(403), 1147–1158.

Kawaguchi, K. (2021). When Will Workers Follow an Algorithm? A Field Experiment with a Retail Business. *Management Science* 67(3), 1670–1695.

Keith, N. and M. Frese (2008). Effectiveness of Error Management Training: A Meta-Analysis. *Journal of Applied Psychology* 93(1), 59–69.

Loewenstein, M. A. and J. R. Spletzer (1998). Dividing the Costs and Returns to General Training. *Journal of Labor Economics* 16(1), 142–171.

Malcomson, J. M., J. W. Maw, and B. McCormick (2003). General Training by Firms, Apprentice Contracts, and Public Policy. *European Economic Review* 47(2), 197–227.

Klein Teeselink, B. (2025). Generative AI and Labor Market Outcomes: Evidence from the United Kingdom. SSRN Working Paper 5516798.

Lodefalk, M., L. Löthman, M. Koch, and E. Engberg (2026). Same Storm, Different Boats: Generative AI and the Age Gradient in Hiring. Working Paper 2026:2, Örebro University School of Business.

Mincer, J. (1974). *Schooling, Experience, and Earnings*. New York: Columbia University Press.

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

Parasuraman, R. and D. H. Manzey (2010). Complacency and Bias in Human Use of Automation: An Attentional Integration. *Human Factors* 52(3), 381–410.

Peng, S., E. Kalliamvakou, P. Cihon, and M. Demirer (2023). The Impact of AI on Developer Productivity: Evidence from GitHub Copilot. arXiv preprint arXiv:2302.06590.

Rubinstein, Y. and Y. Weiss (2006). Post-Schooling Wage Growth: Investment, Search and Learning. In E. A. Hanushek and F. Welch (Eds.), *Handbook of the Economics of Education*, Volume 1, pp. 1–67. Amsterdam: Elsevier.

Shrestha, Y. R., S. M. Ben-Menahem, and G. von Krogh (2019). Organizational Decision-Making Structures in the Age of Artificial Intelligence. *California Management Review* 61(4), 66–83.

Stevens, M. (1994). A Theoretical Model of On-the-Job Training with Imperfect Competition. *Oxford Economic Papers* 46(4), 537–562.

Sweller, J. and G. A. Cooper (1985). The Use of Worked Examples as a Substitute for Problem Solving in Learning Algebra. *Cognition and Instruction* 2(1), 59–89.

Sweller, J., J. J. G. van Merriënboer, and F. G. W. C. Paas (1998). Cognitive Architecture and Instructional Design. *Educational Psychology Review* 10(3), 251–296.

Topel, R. and M. Ward (1992). Job Mobility and the Careers of Young Men. *Quarterly Journal of Economics* 107(2), 439–479.

Walter, S. and J.-D. Lee (2022). How susceptible are skills to obsolescence? A task-based perspective of human capital depreciation. *Foresight and STI Governance* 16(2), 32–41.

Xu, F., J. Hou, W. Chen, and K. Xie (2025). Generative AI and Organizational Structure in the Knowledge Economy. Working paper, arXiv:2506.00532.

A Proofs

Derivation of the decomposition (7)–(8). Since the capable region $\{z \le \theta\}$ and the beyond-frontier region $\{z > \theta\}$ partition the support of $F(\cdot \mid h)$, $\mathcal{L}_\theta(h) + \bar{\mathcal{L}}_\theta(h) = \int \ell \, \mathrm{d}F = \Phi_0(h)/\alpha$, i.e. $\Phi_0 = \alpha(\mathcal{L}_\theta + \bar{\mathcal{L}}_\theta)$ (so capability only relabels tasks between the two channels and leaves the total learnable mass unchanged). Substituting this and (3),

$$\Phi - \Phi_0 = (1-a)\alpha(1+mAp)\mathcal{L}_\theta + a\beta r Ap\,\mathcal{L}_\theta - \alpha\mathcal{L}_\theta = \mathcal{L}_\theta\big[(1-a)\alpha(1+mAp) + a\beta r Ap - \alpha\big].$$

Add and subtract α: the bracket equals $\alpha(1+mAp) - \alpha - a\alpha(1+mAp) + a\beta r Ap = \alpha m Ap + a[\beta r Ap - \alpha(1+mAp)]$. Factor the second term using $\psi = \beta r Ap/[\alpha(1+mAp)]$: $\beta r Ap - \alpha(1+mAp) = \alpha(1+mAp)(\psi - 1)$, giving (7). Differentiating in a (the copilot term is independent of a) yields (8); since $\mathcal{L}_\theta \ge 0$ and $\alpha(1+mAp) > 0$, the sign is that of $\psi - 1$. □

Proof of Proposition 1. Let $x \equiv A(\theta)p(h) \ge 0$ and write $\psi = \dfrac{\beta r x}{\alpha(1+mx)}$. Then

$$\frac{\mathrm{d}\psi}{\mathrm{d}x} = \frac{\beta r\,\alpha(1+mx) - \beta r x\,\alpha m}{\alpha^2(1+mx)^2} = \frac{\beta r}{\alpha(1+mx)^2} > 0.$$

Since $A' > 0$ and $p' > 0$, x is strictly increasing in θ and in h, so ψ is strictly increasing in each. Solving $\psi = 1$ gives $\beta r x = \alpha(1+mx)$, i.e. $x^* = \alpha/(\beta r - \alpha m)$, positive and finite iff $\beta r > \alpha m$; $\psi \gtrless 1 \iff x \gtrless x^*$. By (8), $\partial\Phi/\partial a = \mathcal{L}_\theta(h)\,\alpha(1+mx)(\psi-1)$, and $\mathcal{L}_\theta(h) > 0$ for $h < \theta$ while $\mathcal{L}_\theta(h) = 0$ for $h \ge \theta$ (no capable tasks remain in the queue). Hence for $h < \theta$ the sign of $\partial\Phi/\partial a$ is that of $\psi - 1$, and for $h \ge \theta$ it is zero, proving (i). For fixed θ, the seniority threshold solves $A(\theta)p(h_\psi) = x^*$, i.e. $p(h_\psi) = x^*/A(\theta)$; on $\{h_\psi < h < \theta\}$ both $\psi > 1$ and $\mathcal{L}_\theta > 0$, so autonomy strictly builds learning. Since p is increasing, this region is non-empty iff $h_\psi < \theta \iff p(h_\psi) < p(\theta) \iff x^* < A(\theta)p(\theta) \iff A(\theta)p(\theta)(\beta r - \alpha m) > \alpha$. This proves (ii). Finally $x^* = \alpha/(\beta r - \alpha m)$ is decreasing in r and increasing in m, and h_ψ inherits these signs (and decreases in $A(\theta)$); the feasibility inequality $A(\theta)p(\theta)(\beta r - \alpha m) > \alpha$ slackens as β/α, r, or θ rises and as m falls, since its left side is increasing in each of βr, θ and decreasing in m. This is (iii). □

Proof of Lemma 1. $\Phi(\cdot; u)$ is continuous on $[0,1]$. Since $G(0) \equiv \Phi(0;u) > 0$ (Assumption 2) and $G(1) = \Phi(1;u) - \delta = -\delta < 0$ (top-mass exhaustion), by the intermediate value theorem there is $h^* \in (0,1)$ with $G(h^*) = 0$. For uniqueness, note $G(h) = 0 \iff \Phi(h;u)/h = \delta$. If $h \mapsto \Phi(h;u)/h$ is strictly decreasing it meets the constant δ at most once. This is equivalent to $\frac{\mathrm{d}}{\mathrm{d}h}(\Phi/h) = (\Phi' h - \Phi)/h^2 < 0$, i.e. $\varepsilon_\Phi = h\Phi'/\Phi < 1$, which is (10). Under it G is positive then negative, so h^* is unique and $\dot{h} = G(h)$ has $G > 0$ below and $G < 0$ above h^*: the

steady state is globally asymptotically stable on $(0,1)$. □

Proof of Proposition 2. (i) By (7), $\Phi - \Phi_0 = \mathcal{L}_\theta(h)\, D(h)$ with $D(h) = \alpha m A p + a\alpha(1 + mAp)(\psi - 1)$. Substituting $\psi = \beta r A p/[\alpha(1+mAp)]$ and cancelling the common $(1+mAp)$,

$$D(h) = A(\theta)p(h)\left[\alpha m(1-a) + a\beta r\right] - a\alpha,$$

which is *strictly increasing* in $A(\theta)p(h)$ (the bracketed coefficient is positive), and hence strictly increasing in h, since p is increasing and $A(\theta) > 0$. It therefore crosses zero exactly once, at the loss threshold h_c where $A(\theta)p(h_c) = a\alpha/[\alpha m(1-a) + a\beta r]$. Equivalently $\psi(h_c, \theta) = 1 - \frac{mAp}{a(1+mAp)} \leq 1$, which is (15). The bound is ≤ 1 with equality iff $m = 0$, so $h_c \leq h_\psi$. Since $\mathcal{L}_\theta(h) > 0$ on $\{h < \theta\}$, $\Phi - \Phi_0$ inherits the sign of D: negative on $\{h < h_c\}$ and positive on (h_c, θ). For $h_0 < h_c$, net learning $\Phi - \delta h$ is below baseline over $[h_0, h_c)$, and by (13) the wage slope is flatter. On (h_c, h_ψ) the copilot gain dominates and AI raises learning even though $\psi < 1$. (ii) Compare steady states. The baseline solves $\Phi_0(h) = \delta h$ at the no-AI steady state h_0^*. Under the deployment, if $h^* \in (h_\psi, \theta)$ then at h^* we have $\psi(h^*, \theta) > 1$, so (7) gives $\Phi(h^*; u) > \Phi_0(h^*)$ (both the copilot and the autonomy terms are nonnegative, the latter strictly positive since $\psi > 1$ and $\mathcal{L}_\theta(h^*) > 0$ because $h^* < \theta$). Evaluate $G_0(h) \equiv \Phi_0(h) - \delta h$ at h^*: $G_0(h^*) = \Phi_0(h^*) - \delta h^* < \Phi(h^*; u) - \delta h^* = 0$. Since G_0 is strictly decreasing (baseline satisfies (10)) and $G_0(h_0^*) = 0$, $G_0(h^*) < 0$ implies $h^* > h_0^*$. □

Proof of Proposition 3. Write $h^*(r)$ for the unique stable steady state (Lemma 1 under Assumption 2), solving $\Phi(h^*; r) = \delta h^*$. By the implicit function theorem, $\mathrm{d}h^*/\mathrm{d}r = \Phi_r(h^*)/(\delta - \Phi_h(h^*))$. The numerator $\Phi_r = a\beta A(\theta)\, p(h)\, \mathcal{L}_\theta(h) \geq 0$, strictly positive when $a > 0$, $p(h^*) > 0$, and $\mathcal{L}_\theta(h^*) > 0$ (i.e. $h^* < \theta$), and zero when $h^* \geq \theta$. The denominator is strictly positive at a stable steady state, where $G' = \Phi_h - \delta < 0$. Hence $h^*(r)$ is continuous and nondecreasing, and strictly increasing on the range where $h^*(r) < \theta$ and $p(h^*) > 0$. On that range $r \mapsto h^*(r)$ is a strictly increasing bijection onto its image, so a solution of $h^*(r) = h_0^*$ exists and is unique *if and only if* $h^*(0) < h_0^* < h^*(1)$, which is (16) (with $h_0^* < \theta$ placing the crossing in the capable region), the sign equivalence is immediate from monotonicity. It remains to sign the endpoints. At $r = 0$, $\Phi(\cdot; 0) - \Phi_0 = \alpha\mathcal{L}_\theta[(1-a)(1+mAp) - 1] = \alpha\mathcal{L}_\theta[(1-a)mAp - a] < 0$ iff $a > mAp/(1+mAp)$, in which case $h^*(0) < h_0^*$. At $r = 1$, $\Phi(\cdot; 1) - \Phi_0 = \mathcal{L}_\theta[\alpha m A p + a\alpha(1 + mAp)(\psi|_{r=1} - 1)]$, which is positive at $h^*(1)$ and hence $h^*(1) > h_0^*$ by the comparison in the proof of Proposition 2(ii), if and only if $h^*(1)$ lies above the loss threshold h_c of (15) evaluated at $r = 1$. This requires the review channel to be strong enough, failing which $h^*(1) \leq h_0^*$ and no interior r^* exists. □

Proof of Proposition 4. Under conditions of Case 1, the firm is a price taker with $J \equiv 0$, so its shadow value of human capital is $\mu_F = 0$ and the deployment problem reduces to maximizing current output $Y(h;u)$ alone. From the output derivatives in (21) (equivalently (4)), $Y_r = -C_r'(r)\, a\, Q_\theta \leq 0$ with strict inequality for $r > 0$, so the firm sets $r = 0$ and $Y_m = -C_m'(m)(1-a)\, Q_\theta \leq 0$, so $m = 0$ for every $a < 1$. With $r = m = 0$ and using $C_r(0) = C_m(0) = 0$, the autonomy derivative is $Y_a = \pi_a(a;\theta) > 0$ (since $\pi_a > 0$), so Y is strictly increasing in a and the firm sets $a = 1$ at $a = 1$ neither current output nor learning depends on m, so $m = 0$ is a normalization (the limit of the optimal m as $a \uparrow 1$). Hence the unique equilibrium deployment is $\bar{u} = (a, r, m) = (1, 0, \cdot)$, the myopic corner.

At $\bar{u}$ every capable task ($z \leq \theta$) is delegated and unreviewed, so it carries zero learning (ψ is irrelevant because the review channel is off and autonomy is full), and only beyond-frontier tasks contribute: $\Phi(h;\bar{u}) = \alpha \bar{\mathcal{L}}_\theta(h)$. The induced law of motion is $\dot{h} = \alpha \bar{\mathcal{L}}_\theta(h) - \delta h$, whose stable steady state solves $\alpha \bar{\mathcal{L}}_\theta(h) = \delta h$ (existence, uniqueness, and stability by Lemma 1 under Assumption 2, since $\bar{\mathcal{L}}_\theta$ inherits the boundary and single-crossing properties). Denote it h^*.

For the dichotomy, compare h^* to the no-AI steady state h_0^*, which solves $\Phi_0(h) = \delta h$ with $\Phi_0 = \alpha(\mathcal{L}_\theta + \bar{\mathcal{L}}_\theta)$. (i) If $\theta < h_0^*$, then at every $h \geq \theta$ the capable region is empty ($\mathcal{L}_\theta(h) = 0$), so $\Phi(h;\bar{u}) = \alpha \bar{\mathcal{L}}_\theta(h) = \Phi_0(h)$ there. The two drifts coincide above θ, the worker outgrows the frontier, and $h^* = h_0^*$, the loss being purely transitional (slower approach over $[h_0, \theta)$ where capable tasks are delegated unreviewed). (ii) If $\theta > h_0^*$, then $h^* < \theta$: throughout the convergent path capable tasks carry positive mass ($\mathcal{L}_\theta(h) > 0$) and are delegated without review, so $\Phi(h;\bar{u}) = \alpha \bar{\mathcal{L}}_\theta(h) < \alpha(\mathcal{L}_\theta(h) + \bar{\mathcal{L}}_\theta(h)) = \Phi_0(h)$ for all $h < \theta$. Hence $G(h) = \Phi(h;\bar{u}) - \delta h < \Phi_0(h) - \delta h = G_0(h)$ on $(0, \theta)$, since $G_0(h_0^*) = 0$ and G_0 is strictly decreasing (single crossing), $G(h_0^*) < 0$ forces the stable zero $h^* < h_0^*$. Moreover $\bar{\mathcal{L}}_\theta(h) = \int_{\{z > \theta\}} \ell \, \mathrm{d}F(z \mid h)$ is decreasing in θ, so h^* is strictly decreasing in θ: the more capable the AI, the deeper the collapse. □

Proof of Lemma 2. Both shadow values are costates of the same accumulation technology $\dot{h} = \Phi(h;u) - \delta h$, evaluated along the common allocation $\bar{u}(\cdot)$. They differ only in *whose* marginal return is capitalized, at *what* discount rate, and *with* or *without* the spillover. Differentiating the firm's value equation (19) in h and using the envelope property $\partial w / \partial u = 0$ (so the deployment terms vanish at the optimum), the firm's costate $\mu_F = V_F'$ satisfies

$$(\rho + s)\, \mu_F(h) = \big[Y_h(h;\bar{u}) - w'(h)\big] + \mu_F(h)\, \big[\Phi_h(h;\bar{u}) - \delta\big].$$

The forcing term is the marginal flow the firm *retains*. By the wage pass-through of (20), $w'(h) = \sigma\, Y_h(h;\bar{u})$, so $Y_h - w' = (1-\sigma)\, Y_h \geq 0$ under $Y_h \geq 0$. Differentiating the planner's

equation (18) likewise, the social costate $\mu_S = V_S'$ satisfies

$$\rho\,\mu_S(h) = Y_h(h;\bar{u}) + \big(\mu_S(h)+\nu\big)\big[\Phi_h(h;\bar{u}) - \delta\big],$$

whose forcing is the *full* marginal product Y_h, capitalized at the social rate ρ and augmented by the spillover ν through the accumulation term. Form the gap $D(h) \equiv \mu_S(h) + \nu - \mu_F(h)$. Subtracting the two costate equations and collecting terms, D solves a linear equation of the form

$$\big[\rho - (\Phi_h - \delta)\big]D(h) = \underbrace{\sigma\,Y_h}_{\text{(I) pass-through}} + \underbrace{s\,\mu_F(h)}_{\text{(II) retention hazard}} + \underbrace{\nu\big[\rho - (\Phi_h - \delta)\big] + \nu(\Phi_h - \delta)}_{\text{(III) spillover}},$$

where the three forcing components are: (I) the share σY_h of the marginal product that accrues to the worker rather than the firm, since the firm capitalizes only $(1-\sigma)Y_h$ while the planner capitalizes Y_h; (II) the term $s\,\mu_F$ by which the firm's higher effective discount rate $\rho + s$ shrinks its valuation relative to the planner's rate ρ, the worker carrying h across the separation while the firm does not; and (III) the spillover $\nu \geq 0$, which only the planner internalizes. At the relevant root (the stable costate, where the bracketed coefficient is positive along the convergent path), each component enters D with a nonnegative sign, so $D(h) \geq 0$, i.e. $\mu_F(h) \leq \mu_S(h) + \nu$. Each source contributes strictly: (I) is positive iff $\sigma > 0$ and $Y_h > 0$; (II) is positive iff $s > 0$ (and $\mu_F > 0$); (III) is positive iff $\nu > 0$. Hence the inequality is strict whenever $s > 0$, $\sigma > 0$, or $\nu > 0$, and holds with equality only in the knife-edge $s = \sigma = \nu = 0$.

For the comparative statics, the wedge is increasing in each source. Component (I) is increasing in the worker's pass-through σ: a larger σ transfers more of the marginal product to the worker, leaving the firm capitalizing the smaller share $(1-\sigma)Y_h$, so D rises. Equivalently, D is *decreasing* in the firm's retained share $1-\sigma$, which is exactly the force a firm would internalize. Component (II) is increasing in the retention hazard s: directly, through the forcing term $s\,\mu_F$, a higher hazard raises the firm's effective discount rate $\rho + s$ and widens the gap between the firm's and the planner's capitalization of the same future return. The induced effect reinforces it, since a higher hazard also lowers the firm's own valuation μ_F, enlarging the wedge further. Component (III) is increasing in the spillover ν, which the planner values and the firm does not. The three are additive and independent, so the wedge increases in σ, s, and ν and decreases in $1-\sigma$. □

Proof of Proposition 5. Write the deployment objective $\Pi(u;\mu) = Y(h;u) + \mu[\Phi(h;u) - \delta h]$, with $\Phi_a = \alpha(1+mAp)(\psi-1)\mathcal{L}_\theta$ (negative in the rung region $\psi < 1$), $\Phi_r = a\beta Ap\,\mathcal{L}_\theta > 0$,

and $\Phi_m = (1-a)\alpha Ap\,\mathcal{L}_\theta > 0$. We give a direct derivation rather than invoking a vector monotone-comparative-statics theorem: the model's separability makes the Hessian of Π in u *diagonal* at any interior optimum, so no joint supermodularity (lattice) condition on the control vector is required.

Separability and a diagonal Hessian. Because Φ is linear in each control and the costs are additively separable, the second derivatives of Π in (a, r, m) are

$$\Pi_{aa} = \pi_{aa} < 0, \qquad \Pi_{rr} = -C_r''(r)\,aQ_\theta < 0, \qquad \Pi_{mm} = -C_m''(m)(1-a)Q_\theta < 0,$$

$$\Pi_{ar} = -C_r'(r)Q_\theta + \mu\beta Ap\,\mathcal{L}_\theta, \qquad \Pi_{am} = C_m'(m)Q_\theta - \mu\alpha Ap\,\mathcal{L}_\theta, \qquad \Pi_{rm} = 0$$

Crucially, r enters Π only through the factor a and m only through $(1-a)$, so the first-order conditions factor as

$$\Pi_r = a\,\Pi_{ar}, \qquad \Pi_m = -(1-a)\,\Pi_{am}.$$

Hence at any interior optimum, where $\Pi_r = \Pi_m = 0$ with $a \in (0,1)$, the cross-partials vanish: $\Pi_{ar} = \Pi_{am} = \Pi_{rm} = 0$, and the Hessian $H = \Pi_{uu} = \mathrm{diag}(\Pi_{aa}, \Pi_{rr}, \Pi_{mm})$ is diagonal and negative definite.

Comparative statics. Differentiating the first-order conditions $Y_u + \mu\Phi_u = 0$ totally in μ gives $H\,(\mathrm{d}u/\mathrm{d}\mu) = -\Phi_u$, with H diagonal this decouples into three scalar comparative statics,

$$\frac{\mathrm{d}a}{\mathrm{d}\mu} = -\frac{\Phi_a}{\Pi_{aa}} < 0, \qquad \frac{\mathrm{d}r}{\mathrm{d}\mu} = -\frac{\Phi_r}{\Pi_{rr}} > 0, \qquad \frac{\mathrm{d}m}{\mathrm{d}\mu} = -\frac{\Phi_m}{\Pi_{mm}} > 0,$$

using $\Phi_a < 0$, $\Phi_r, \Phi_m > 0$, and $\Pi_{aa}, \Pi_{rr}, \Pi_{mm} < 0$. Equivalently and without matrix inversion: the r- and m-problems each separate, with r^* solving $C_r'(r^*) = \mu\beta Ap\,\mathcal{L}_\theta/Q_\theta$ and m^* solving $C_m'(m^*) = \mu\alpha Ap\,\mathcal{L}_\theta/Q_\theta$ (both independent of the remaining controls and increasing in μ by convexity of C_r, C_m) and, substituting them, the scalar a-problem has decreasing differences in $(a;\mu)$, since the indirect effects of μ through r^*, m^* enter $\mathrm{d}\Pi_a/\mathrm{d}\mu$ with the coefficients Π_{ar}, Π_{am}, which are zero at the optimum, leaving only the direct term $\Phi_a < 0$. Since $\mu_F < \mu_S + \nu$ (Lemma 2) and these signs hold at every interior optimum along the path, integrating over $[\mu_F, \mu_S + \nu]$ yields $a(\mu_F) > a(\mu_S + \nu)$, $r(\mu_F) < r(\mu_S + \nu)$, and $m(\mu_F) < m(\mu_S + \nu)$, establishing (i) and (ii). In the linear/bang-bang case the same factorization $\Pi_r = a\Pi_{ar}$, $\Pi_m = -(1-a)\Pi_{am}$ shows the sign of each control's marginal value flips in μ independently of the others, so the corners shift in the same directions. For (iii), the wedge $\mu_S + \nu - \mu_F$ is increasing in the separation hazard s, in the wage pass-through σ (equivalently decreasing in wage compression $1-\sigma$), and in the spillover ν (Lemma 2) and approaches $\mu_S + \nu$ as $\mu_F \to 0$ (the full-pass-through, competitive limit $\sigma \to 1$ of Proposition 4), so the divergence in optimal deployments widens accordingly. □

Proof of Proposition 6. Write $\omega(e) = e + (1-e)\gamma$, so $\omega'(e) = 1-\gamma > 0$ since $\gamma \in [0,1)$, and recall from (24)

$$v_R(q,e,p) = p\,q\,\omega(e) + p(1-q)\nu - (1-p)(1-q)\chi.$$

(i) Differentiating in e, only the first term depends on e, so $\partial v_R/\partial e = p\,q\,\omega'(e) = p\,q\,(1-\gamma) > 0$ for $p, q > 0$. *(ii)* Collecting the terms linear in p, $v_R = p\big[q\,\omega(e)+(1-q)\nu+(1-q)\chi\big]-(1-q)\chi$, whence $\partial v_R/\partial p = q\,\omega(e)+(1-q)\nu+(1-q)\chi > 0$, the three summands being the contributions of confirmation, error detection, and reduced mislearning, each nonnegative and positive in sum. Since $p = p(h)$ is increasing in h, v_R is increasing in seniority. *(iii)* Differentiating in q, $\partial v_R/\partial q = p\,\omega(e) - p\,\nu + (1-p)\chi = p\,[\omega(e) - \nu] + (1-p)\chi$. When $\nu > \omega(e)$ the first term is negative and the second positive; the expression is therefore positive at $p = 0$ (value $\chi \geq 0$) and decreasing in p (slope $\omega(e) - \nu - \chi < 0$), so it is positive for small p and negative for large p. It vanishes at $\bar{p} = \chi/(\nu - \omega(e) + \chi)$, giving $\partial v_R/\partial q > 0$ for $p < \bar{p}$ (low-ability reviewers, where reduced mislearning dominates) and $\partial v_R/\partial q < 0$ for $p > \bar{p}$ (high-ability reviewers, where the loss of instructive errors dominates). *(iv)* Directly from the expression, $v_R < 0 \iff (1-p)(1-q)\chi > p\big[q\,\omega(e) + (1-q)\nu\big]$, which holds when verification ability p and transparency e are low and the error rate $1-q$ is non-trivial: the mislearning loss then exceeds the combined value of confirmed-correct and detected-error components. □

Proof of Proposition 7. We start with the proof of *(i)* By Proposition 6(iv), $v_R < 0$ when verification ability lies below the threshold p_0. Holding autonomy high from entry, the nonnegative frontier and copilot contributions to $\Phi^{\text{mic}}(0)$ are dominated by the negative review contribution $a\,r\kappa\,v_R(A(\theta), e, p(0))\mathcal{L}_\theta(0)$, so $\Phi^{\text{mic}}(0) < 0$ and, by the boundary behavior of (5), $h = 0$ is absorbing. Above p_0, $v_R > 0$, so Φ^{mic} is positive on the scaffolding region; with $\Phi^{\text{mic}}(1) = 0$ this gives an up-crossing h_{trap} and a down-crossing $h^* > h_{\text{trap}}$, with $\Phi^{\text{mic}} - \delta h < 0$ throughout $(0, h_{\text{trap}})$, so h_{trap} is the unique unstable edge and the long run is history-dependent. By contrast $\Phi^{\text{lin}} \geq 0$ everywhere, so $\Phi^{\text{lin}}(0) \geq 0$ and (when $\Phi^{\text{lin}}(h)/h$ is decreasing) the steady state is unique and globally stable, with no trap. *(ii)* Withholding autonomy for $h < h_\psi$ inactivates the review channel at entry, so $\Phi^{\text{mic}}(0) = \alpha\bar{\mathcal{L}}_\theta(0) + (1-a)\alpha(1 + mA(\theta)p(0))\mathcal{L}_\theta(0) \geq 0$, strictly positive when a frontier or copilot channel is open. Hence the drift is inward at $h = 0$ and the absorbing state of (i) is removed. For $h \geq h_\psi$, where the worker has crossed into $p > p_0$, $v_R > 0$ and $\Phi^{\text{mic}}(h)/h$ is strictly decreasing; with $\Phi^{\text{mic}}(0) \geq 0$ and $\Phi^{\text{mic}}(1) = 0$ this yields a unique interior steady state, globally asymptotically stable on $(0, 1]$. On this trajectory h rises monotonically from h_0 through h_ψ to h^* without entering $\{p < p_0\}$, so $v_R > 0$ along the path and Φ^{mic} and Φ^{lin} share the sign of the seniority gradient (Proposition 6(ii), $\partial v_R/\partial p > 0$ globally), hence the same reversal, the same dependence of the long-run sign on r, and the same comparative statics. □

B Extensions and Robustness Checks

This appendix develops the extensions and robustness checks referenced in the main text.

B.1 Two Viability Conditions: Learning vs. Finance

A robust career ladder requires two distinct things: that the junior receive enough learning exposure to accumulate expertise (a *learning* condition), and that someone be willing to finance the arrangement given AI's outside productivity (a *financial* condition). Garicano and Rayo (2026) analyze the second through an expertise leverage ratio. We show the two conditions are logically independent and, in autonomy, move in opposite directions.

Define, at the entry margin, the *learning-viability* index and the *financial-viability* index

$$\mathcal{V}_L(\theta, u) \equiv \frac{\Phi(h_0; \theta, u)}{\Phi_0(h_0)}, \qquad \mathcal{V}_F(\theta, u) \equiv \frac{Y(h_0; \theta, u) - w(h_0)}{\kappa}, \tag{25}$$

where $\kappa > 0$ is the per-period cost of sustaining the apprenticeship (training, supervision, vacancy) and $w(h_0)$ is the market wage of the entrant's type under Assumption 3, so that $\partial w/\partial u = 0$: the wage prices the worker's general human capital, not the firm's deployment. Thus $\mathcal{V}_F$ is the firm's per-period *retained surplus* from the match net of the apprenticeship cost. Learning is viable if $\mathcal{V}_L \geq 1$ (AI does not reduce the entrant's learning below its no-AI level). The arrangement is financed if $\mathcal{V}_F \geq 1$.

Proposition 8 (Divergence of the two conditions)**.** *The two indices respond to autonomy with opposite signs in the rung-removing region. Because $\partial w/\partial u = 0$ (Assumption 3), the wage does not move with the firm's deployment, so*

$$\frac{\partial \mathcal{V}_F}{\partial a} = \frac{Y_a}{\kappa} > 0, \qquad \operatorname{sign}\left(\frac{\partial \mathcal{V}_L}{\partial a}\right) = \operatorname{sign}\big(\psi(h_0, \theta) - 1\big) < 0 \quad \textit{when } \psi(h_0, \theta) < 1. \tag{26}$$

Hence all four configurations are possible:

	Financed *($\mathcal{V}_F \geq 1$)*	***Not financed*** *($\mathcal{V}_F < 1$)*
Learning-viable *($\mathcal{V}_L \geq 1$)*	*robust ladder*	*"unfunded apprenticeship"*
Learning-destroying *($\mathcal{V}_L < 1$)*	*"hollow job"*	*collapse*

In particular, raising autonomy with capable AI and thin review (high a, high θ, low r) moves an arrangement into *the financed column (output rises) while moving it* down *into the learning-destroying row that pays and produces but does not teach. Robust ladders require satisfying both conditions. Each detects a failure the other misses.*

Proof of Proposition 8. From (25), since $\partial w/\partial u = 0$ (Assumption 3), $\partial \mathcal{V}_F/\partial a = Y_a/\kappa$ with $Y_a = \pi_a - [C_r(r) - C_m(m)]Q_\theta$ by (4). This is positive whenever the throughput gain π_a exceeds the net per-task assistance-cost change, and is necessarily positive at any interior rung-region optimum, where the autonomy condition (21) gives $Y_a = \mu\,|\Phi_a| > 0$ (since $\Phi_a < 0$ when $\psi < 1$). For learning, $\mathcal{V}_L = \Phi(h_0; \theta, u)/\Phi_0(h_0)$ and $\Phi_0(h_0)$ is independent of a, so $\partial \mathcal{V}_L/\partial a = \Phi_a(h_0)/\Phi_0(h_0)$, whose sign is that of $\psi(h_0, \theta) - 1$ by (8). When $\psi(h_0, \theta) < 1$ the two derivatives have opposite signs. Independence of the two indices then follows because $\mathcal{V}_F$ can be made arbitrarily large or small through the output primitives (π, w, κ) without affecting Φ, while $\mathcal{V}_L$ can be moved through (r, m, β) without affecting Y_a; hence each cell of the 2×2 is attainable. □

The derivative result is the substance and it does not depend on how the *level* of the wage is pinned down, only on $\partial w/\partial u = 0$. The level does: under frictional wages (Assumption 3 with $\sigma < 1$) the firm retains a surplus, $\mathcal{V}_F$ can exceed one, and the financed column is attainable. In the competitive full-pass-through limit ($\sigma = 1$, $w = Y(h_0; \bar{u})$) the firm's retained surplus is $Y(h_0; u) - Y(h_0; \bar{u}) \leq 0$, so $\mathcal{V}_F \leq 0 < 1$ for every deployment: the apprenticeship is unfunded for exactly the hold-up reason of Proposition 4, no one will pay κ to build portable skill they cannot capture. Financing thus requires rents ($\sigma < 1$), fixed or bargained wages, or contractibility (Case 3). The table is read in that case, where autonomy still raises $\mathcal{V}_F$ through Y_a while lowering $\mathcal{V}_L$. The competitive limit is not a counterexample to the divergence but its sharpest instance: there the financial column collapses to "unfunded" precisely because the firm internalizes none of the learning it could finance.

This locates the present paper relative to Garicano and Rayo (2026). Their leverage ratio governs the financed/not-financed margin. It can be comfortably satisfied by exactly the high-autonomy deployments that hollow out learning. The learning condition is therefore not implied by the financial one, the two approaches are complementary.

B.2 When durability decides: the obsolescence watershed

Proposition 3 crosses the long-run threshold by changing the *deployment*. The same threshold can be crossed by a single fundamental: *how durable expertise is*, the obsolescence rate δ. Whether a worker ever reaches the verification competence the review channel requires depends on how high their human capital settles, and the no-AI steady state $h_0^*(\delta) = \alpha/(\alpha + 2\delta)$ is strictly decreasing in δ. If skills churn fast enough, the steady state lies below h_ψ no matter how the deployment is configured, and the worker spends their whole career in the rung-removal region. We isolate δ here because, unlike the other primitives, it maps to the long-run sign through a single monotone channel and so admits a closed-form watershed;

the learning-curve primitives (the verification profile p, the content profile ℓ) act through the same steady-state mechanism but functionally rather than through one threshold, and are explored numerically in the sensitivity analysis of Section 7.

Proposition 9 (Obsolescence watershed and review exhaustion)**.** *Fix a deployment (θ, a, r, m) with $a > 0$, and let $h_c = h_c(r)$ be the loss threshold of (15) (where $\Phi(h; u) = \Phi_0(h)$); write $h_0^*(\delta)$ for the no-AI steady state, strictly decreasing in δ. On the range where $h_0^*(\delta) < \theta$, the long-run effect $\Delta(\delta) \equiv h^*(\delta) - h_0^*(\delta)$ satisfies*

$$\Delta(\delta) > 0 \iff h_0^*(\delta) > h_c \iff \delta < \delta^{**}, \tag{27}$$

where δ^{**} is the unique rate at which $h_0^*(\delta^{**}) = h_c$. Hence Δ changes sign exactly once, from positive to negative, at δ^{**} (no monotonicity of Δ is needed or claimed): (i) *AI raises the long-run expertise stock iff $\delta < \delta^{**}$. Moreover, since $h_c(r)$ is strictly decreasing in r;* (ii) *the review intensity required to keep AI expertise-building, defined by $h_c(r^*(\delta)) = h_0^*(\delta)$, is strictly increasing in δ, and reaches $r^* = 1$ at a finite rate $\delta_{\max}$, for $\delta > \delta_{\max}$, $h_0^*(\delta) < h_c(r)$ for every $r \leq 1$, so* no feasible review intensity *makes AI expertise-building. Review can rescue the early-career slope, but in fast-obsolescence environments it cannot change the sign of the long-run effect.*

Proof of Proposition 9. The key is that the sign of $\Delta(\delta)$ is governed by the position of the *no-AI* steady state relative to the loss threshold, not by monotonicity of Δ. By the decomposition (7), $\Phi(h; u) - \Phi_0(h) = \mathcal{L}_\theta(h)\, D(h)$ with $D(h) = A(\theta)p(h)\,[\alpha m(1-a) + a\beta r] - a\alpha$, which is strictly increasing in h (through p) and vanishes exactly at the loss threshold h_c of (15): $D(h) \gtrless 0 \iff h \gtrless h_c$. Evaluate the AI drift $G(h) = \Phi(h; u) - \delta h$ at the no-AI steady state $h_0^*(\delta)$, where $\Phi_0(h_0^*) = \delta h_0^*$:

$$G\big(h_0^*(\delta)\big) = \Phi\big(h_0^*; u\big) - \Phi_0\big(h_0^*\big) = \mathcal{L}_\theta\big(h_0^*\big)\, D\big(h_0^*\big),$$

which (for $h_0^* < \theta$, so $\mathcal{L}_\theta > 0$) is positive iff $h_0^*(\delta) > h_c$. The AI steady state $h^*(\delta)$ is the unique stable zero of G (Lemma 1, condition (10)), so G is positive to its left and negative to its right; hence $h^*(\delta) > h_0^*(\delta)$ iff $G(h_0^*(\delta)) > 0$ iff $h_0^*(\delta) > h_c$. Since $h_0^*(\delta)$ is strictly decreasing in δ, $h_0^*(\delta) > h_c \iff \delta < \delta^{**}$ with δ^{**} solving $h_0^*(\delta^{**}) = h_c$. This is (27): Δ is positive below δ^{**} and negative above it, a single crossing, with no claim on the shape of Δ in between. For (ii), the threshold $h_c = h_c(r)$ is strictly decreasing in r (raising r raises the bracket $[\alpha m(1-a) + a\beta r]$, lowering the $p(h_c)$ that solves $D = 0$). The crossing locus is $h_c(r) = h_0^*(\delta)$; differentiating, $h_c'(r)\,\mathrm{d}r^* = \mathrm{d}h_0^*(\delta)$, and since $h_c' < 0$ and $\mathrm{d}h_0^*/\mathrm{d}\delta < 0$, $\mathrm{d}r^*/\mathrm{d}\delta > 0$. At $r = 1$ the locus defines $\delta_{\max}$. For $\delta > \delta_{\max}$, $h_0^*(\delta) < h_c(1) \leq h_c(r)$ for all

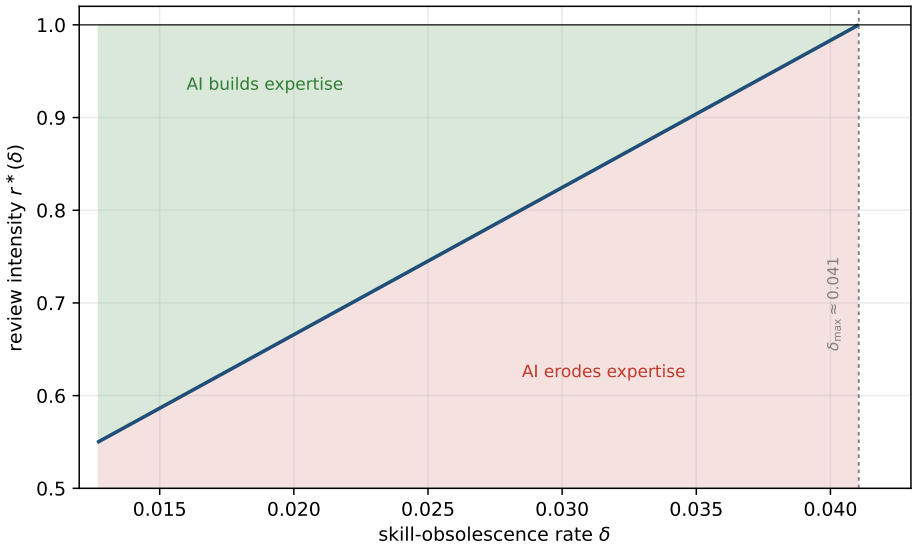


Figure 9: The obsolescence watershed. For each skill-obsolescence rate δ, the curve is the review intensity $r^*(\delta)$ needed to keep AI expertise-building. Above it AI raises the long-run stock, below it AI erodes it. As skills churn faster the required review rises, until at $\delta_{\max}$ even full review ($r = 1$) is insufficient: for $\delta > \delta_{\max}$ AI erodes expertise regardless of review. Parameters: $\alpha = 0.045$, $\beta = 0.18$, $a = 0.85$, $m = 0.20$, $\theta = 0.70$, with $A(\theta) = \theta$ and $p(h) = h$.

$r \leq 1$, so $\Delta(\delta) < 0$ for every feasible r. □

The mechanism is the interaction of skill durability with the verification gate. The review channel pays off only above the loss threshold h_c, where the worker can verify enough that reviewing out-teaches doing; below it autonomy net-removes rungs. Skill durability determines which side of the threshold the worker settles on. When expertise is durable (δ low), the steady state climbs into the scaffolding region and AI compounds expertise. When skills churn fast (δ high), the worker plateaus below h_c, autonomy keeps removing rungs, and the expertise base erodes. Mandating review steepens the climb but cannot move its destination past the threshold once δ is large enough. The required review is exhausted at $r = 1$. Figure 9 maps the two-margin frontier in the (δ, r) plane. In the calibration the long-run sign flips at $\delta^{**} \approx 0.035$ (at $r = 0.9$) and review is exhausted at $\delta_{\max} \approx 0.041$, with the required review rising from $r^* = 0.586$ at $\delta = 0.015$ to 0.825 at $\delta = 0.03$ before becoming infeasible.

Two implications follow. First, a *cross-occupation prediction*: AI's long-run effect on expertise should track skill durability. In durable-expertise fields, where tenure-earnings profiles are steep and skill half-lives long (as in medicine, law, or engineering fundamentals), autonomous AI paired with review builds expertise. In fast-churning fields, short skill half-lives, flat senior premia, the same configuration erodes it, because workers never reach

verification competence. The obsolescence rate, proxied by the steepness of the experience-earnings profile or by occupational skill half-lives, is the sign-determining covariate. The occupation- and exposure-level heterogeneity documented by Brynjolfsson et al. (2025a) and Hosseini and Lichtinger (2025), who report that junior task removal is concentrated in the most AI-exposed occupations, is the natural margin against which this prediction can be tested. Second, a *policy contrast*: review mandates, the natural corrective suggested by Proposition 3, are sufficient only where skills are durable enough. In fast-obsolescence occupations they preserve the early slope but not the long-run base. Preserving expertise there requires acting on the steady state itself: continuous retraining that lowers effective δ, deepening the learning curve, or holding juniors hands-on until they cross h_ψ, rather than review alone. The right instrument is field-specific, and a uniform "mandate review" rule will succeed in some occupations and fail in others for reasons the watershed makes precise.

B.3 Seniority Bias and Heterogeneity

The seniority profile of AI's learning impact follows from two opposing comparative statics in the model's primitives: the scaffolding ratio $\psi(h, \theta)$ rises in seniority (Proposition 1), while the capable-task mass $\mathcal{L}_\theta(h)$ falls in seniority and vanishes at the frontier (by the FOSD ordering of Assumption 1).

Lemma 3 (Task masses in seniority)**.** *For an arbitrary task distribution $F(\cdot \mid h)$ ordered by FOSD and an arbitrary nonnegative learning profile $\ell(z)$, the capable-task mass $\mathcal{L}_\theta(h) = \int_{\{z \le \theta\}} \ell(z)\, \mathrm{d}F(z \mid h)$ is non-increasing in h on $[0, \theta)$ and equal to zero for $h \ge \theta$. It is strictly decreasing wherever the capable region carries positive mass and the profile is bounded on it. The scaffolding ratio $\psi(h, \theta)$ is strictly increasing in h. Under the leading parametrization both hold in closed form, with $\mathcal{L}_\theta$ strictly decreasing on all of $[0, \theta)$.*

Proof. Capable-task mass. For $h \ge \theta$ the capable region $\{z \le \theta\}$ is disjoint from the support $[h, \infty)$ of $F(\cdot \mid h)$, so $\mathcal{L}_\theta(h) = 0$. For $h < \theta$, $\mathcal{L}_\theta(h) = \int_{\{z \le \theta\}} \ell(z)\, \mathrm{d}F(z \mid h)$. As h rises, $F(\cdot \mid h)$ shifts up in the FOSD order, so the capable-region probability $F(\theta \mid h)$ is non-increasing in h, since $\ell \ge 0$, integrating a nonnegative profile over a (weakly) shrinking capable region delivers a non-increasing $\mathcal{L}_\theta$. The inequality is strict wherever the capable region $\{z \le \theta\} \cap \operatorname{supp} F(\cdot \mid h)$ has positive probability and ℓ is bounded above zero on a subset of it, since then the FOSD shift strictly removes capable mass. *Scaffolding ratio.* $\psi(h, \theta) = \beta r A(\theta) p(h) / [\alpha(1 + m A(\theta) p(h))]$ is strictly increasing in $A(\theta)p(h)$, hence in h since $p' > 0$. □

These pull in opposite directions and partition the career into three regions:

- **Juniors** (h small): $\psi < 1$ *and* $\mathcal{L}_\theta$ large. Much of the queue is automatable, yet the worker cannot yet verify, so delegation removes rungs. Copilot mode is the learning-optimal deployment; mandated review helps little because $p(h)$ is too low.
- **Intermediates** ($h \approx h_\psi$): ψ crosses one while $\mathcal{L}_\theta$ is still positive. The review channel becomes the better teacher; autonomy paired with review now *builds* skill.
- **Seniors** ($h \geq \theta$): $\mathcal{L}_\theta = 0$. The worker's queue lies beyond the AI frontier; AI complements their output (through π) but is pedagogically neutral, $\Phi = \Phi_0$.

Corollary 2 (Heterogeneity and wage-slope dispersion)**.** *As a direct consequence of Proposition 1, suppose workers differ in verification ability, in AI access, or in copilot access. Each such parameter raises the scaffolding ratio $\psi(h, \theta)$, lowers the threshold h_ψ, and raises the net learning rate Φ wherever the review channel is active, hence steepens the early-career wage slope through (13). The same AI deployment therefore widens the dispersion of early-career wage growth across workers, increasing between-group inequality even with identical entry wages.*

Proof. By Proposition 1, $\psi(h, \theta)$ is increasing in any parameter scaling verification ability $p(\cdot)$, AI quality $A(\theta)$, or available review r, so the threshold h_ψ (defined by $\psi(h_\psi, \theta) = 1$) is decreasing in each. Net learning $\Phi(h; u)$ is increasing in those same parameters wherever the review channel is active, through the term $a\beta r A(\theta) p(h) \mathcal{L}_\theta$ in (3). By (13) the early-career wage slope is proportional to Φ, so a worker with higher verification ability or better review access has a strictly steeper wage profile under the same deployment. Two workers with identical entry wages but different verification ability thus diverge in wage growth, the gap monotone in the verification differential; dispersion widens accordingly. □

This result matters for three reasons. First, it identifies a *distributional* consequence of AI that the level-and-sign analysis of the preceding sections does not reach: the same deployment that determines whether a representative worker's trajectory bends up or down also determines how widely trajectories *fan out* across workers who face it. AI thus reshapes not only the average early career but its variance. Second, the margin along which it fans them out is *verification ability* (precisely the skill the review channel rewards) so the dispersion is not arbitrary inequality but a sorting on who can learn from supervising AI. Workers who can already verify pull away; those who cannot fall behind, even starting from the same wage. This gives the model a sharp, cross-sectional prediction distinct from its time-series ones: within an AI-exposed occupation, the spread of early-career wage growth should widen with AI adoption, and the spread should line up with measurable verification skill rather than with AI exposure as such. Third, the effect operates *without any change in entry pay*:

because the wage prices the worker's current type and the divergence is in the *slope*, standard entry-wage comparisons would miss it entirely.

B.4 The Seniority-Contingent Wage Path

Section 5.2 gives wage growth within a fixed deployment regime. We now draw out the pay path a worker experiences when the firm follows the seniority-contingent prescription of Figure 8, where the deployment switches with seniority and the path therefore bends at the scaffolding transition.

Proposition 10 (The seniority-contingent wage path)**.** *Let the firm follow the bang-bang seniority-contingent program of Figure 8: copilot $\bar{u}^C = (a_C, r_C, m_C)$ with autonomy withheld (a_C low, $m_C > 0$) for $h < h_\psi$, and delegate-and-review $\bar{u}^D = (a_D, r_D, m_D)$ with $a_D > a_C$ and $r_D \geq r^*$ for $h \geq h_\psi$; let pay track the deployment played, $w(h) = \sigma\, Y(h; \bar{u}(h))$, $\sigma \in (0, 1]$. For an entrant with $h_0 < h_\psi < h^*$:*

(i) Within phases. *On each of $[h_0, h_\psi)$ and $(h_\psi, h^*]$ the path is C^1 with slope $w'(h) = \sigma Y_h(h; \bar{u})$ and proportional growth governed by Φ through the within-regime form of (13). At a given seniority both phases produce less current output than the myopic deployment $Y(h; 1, 0, \cdot)$. The shortfall is the Becker development concession (Corollary 1) and is larger in the copilot phase, which withholds the autonomy throughput entirely.*

(ii) Kink at h_ψ (deployment switch). *Pay jumps by*

$$\Delta w(h_\psi) \equiv w(h_\psi^+) - w(h_\psi^-) = \sigma \Big[\underbrace{\pi(a_D; h_\psi, \theta) - \pi(a_C; h_\psi, \theta)}_{\text{throughput gain} > 0} - Q_\theta(h_\psi)\, \Delta C \Big],$$

with net per-capable-task operating-cost change $\Delta C \equiv [C_r(r_D)a_D - C_r(r_C)a_C] + [C_m(m_D)(1 - a_D) - C_m(m_C)(1 - a_C)]$. The jump is positive *whenever the throughput gain exceeds the net cost change, $\pi(a_D) - \pi(a_C) > Q_\theta(h_\psi)\, \Delta C$, which holds in the leading case of modest operating costs. The slope rises across the switch by $\sigma[Y_h(h_\psi; \bar{u}^D) - Y_h(h_\psi; \bar{u}^C)]$.*

(iii) No kink at h_c under the two-phase rule. *The loss threshold $h_c \leq h_\psi$ is* not *a deployment switch here, so the path is smooth at h_c. It instead marks where, under the* alternative *fixed high-autonomy deployment, the AI early-career wage slope crosses the no-AI slope. Were the firm to refine the program by switching review on at h_c (regions 2→4 of the rule), a second, smaller pay kink would appear there, signed by the formula of (ii) at the adjacent deployments. Under the two-phase rule the only pay discontinuity is the upward step at h_ψ.*

Proof. (i) On each phase $\bar{u}$ is constant, so $\bar{u}'(h) = 0$ and (14) reduces to $w'(h) = \sigma Y_h(h; \bar{u})$.

The growth statement is (13). By (4) current output rises in autonomy ($Y_a = \pi_a - [C_r(r) - C_m(m)]Q_\theta$, with $\pi_a > 0$ the leading term) and falls in r and m ($Y_r, Y_m < 0$), so at any h the myopic deployment $(1, 0, \cdot)$ attains the highest current output. Both $\bar{u}^C$ and $\bar{u}^D$ produce strictly less, and the copilot shortfall is the larger, since it forgoes the autonomy throughput π entirely and pays the assist cost $C_m(m_C)(1 - a_C)Q_\theta$. (ii) The jump is $w(h_\psi^+) - w(h_\psi^-) = \sigma[Y(h_\psi; \bar{u}^D) - Y(h_\psi; \bar{u}^C)]$. Substituting (4) and cancelling the common baseline $y_0(h_\psi)$ gives the displayed expression, with throughput difference $\pi(a_D; h_\psi, \theta) - \pi(a_C; h_\psi, \theta)$ and operating-cost difference $Q_\theta(h_\psi)\Delta C$. Since $a_D > a_C$ and π is increasing in a, the throughput term is positive, so $\Delta w(h_\psi) > 0 \iff \pi(a_D) - \pi(a_C) > Q_\theta(h_\psi)\Delta C$. The slope change follows by evaluating $w'(h) = \sigma Y_h(h; \bar{u})$ on each side of h_ψ. (iii) Under the two-phase program $\bar{u}$ is constant on a neighbourhood of h_c (which lies in the copilot phase, $h_c \leq h_\psi$), so w is C^1 at h_c. The slope-versus-baseline statement is Proposition 2(i) applied to the fixed high-autonomy deployment. A program that instead switches deployment at h_c produces a level jump there by the computation in (ii) at the adjacent deployments. □

The development-oriented pay path is thus a flatter, partly depressed junior segment (the cohort paying the concession), an upward step at h_ψ, and a steeper senior segment: the mirror image of the myopic path of Proposition 4, whose pay is high and flat from entry. The step at h_ψ is the pay-path signature of the scaffolding transition, where withheld autonomy is released and the worker's product jumps. Along the career the development-oriented worker, though paid less at every common seniority, reaches higher h and ultimately higher pay, because the myopic path stalls at the low steady state of Proposition 4. Empirically the result predicts a discontinuity in pay, or in measured autonomy and review intensity, at the seniority at which firms begin delegating-with-review, a feature absent under uniform automation.

B.5 Worker Sorting under a Noisy Development Signal

The competitive benchmark of Section 6 (Cases 1–2) imposes spot wages and no directed search, which is the configuration that maximizes the hold-up: the firm captures none of the portable return and under-provides review maximally. The natural market correction is worker sorting: early-career workers value development and would accept lower pay at firms that provide it (Section 9). This appendix asks whether directed search overturns the myopic corner. We model it in the canonical competitive-search form (posted wages, free entry, endogenous tightness) with one departure dictated by the paper's premise: the developmental content of a deployment is *not* contractible and cannot be posted, so workers direct search on a *noisy public signal* of a firm's development posture rather than on the

posture itself. Signal precision indexes how much development is observable *ex ante*; the contractible benchmark (Case 3) and the spot benchmark (Case 1) are the two precision extremes.

Environment. A unit mass of risk-neutral early-career workers and free entry of firms. Timing: (i) each entering firm chooses a *posture*, develop (D) or not (N); the developer posture costs $c_R > 0$ (the current cost of mandated review, whose return is portable) and raises the worker's portable continuation value by $\Delta > 0$, while N costs nothing and adds nothing; (ii) a public signal $s \in \{g, b\}$ realizes with precision $\pi \in [\frac{1}{2}, 1]$, $\Pr(g \mid \mathsf{D}) = \Pr(b \mid \mathsf{N}) = \pi$; (iii) firms post a wage w and open a vacancy; (iv) observing (w, s), workers form a posterior $b(s) = \Pr(\mathsf{D} \mid s)$ and direct their search; (v) matching, production. Posture is the non-contractible action: it is neither postable nor verifiable, and the signal is informative about it only to the degree $\pi > \frac{1}{2}$. Let n_D be the equilibrium mass fraction of developers; Bayes' rule gives

$$b(g) = \frac{n_D \pi}{n_D \pi + (1 - n_D)(1 - \pi)}, \qquad b(b) = \frac{n_D (1 - \pi)}{n_D (1 - \pi) + (1 - n_D)\pi}.$$

A submarket is a pair (w, s) with tightness $\tau(w, s)$ (vacancies per worker). The matching technology is standard: aggregate matches $M(u, v)$ are constant-returns, twice differentiable, increasing and strictly concave in each argument, with $M(0, v) = M(u, 0) = 0$. Writing the worker matching rate $f(\tau) \equiv M(1, \tau)$, this gives $f(0) = 0$, $f' > 0$, $f'' < 0$, $f(\infty) = 1$, and a vacancy-fill rate $g(\tau) = f(\tau)/\tau$ that is strictly decreasing; we also impose the mild boundary condition $\lim_{\tau \to \infty} \tau f'(\tau) = 0$ (equivalently $f(\tau) - \tau f'(\tau) \to 1$), satisfied by the urn-ball $f(\tau) = 1 - e^{-\tau}$ and by Cobb-Douglas matching. A match yields the firm flow output $y > 0$ (net of AI operating cost) and the worker the wage plus, in expectation, the portable development value $b(s)\Delta$.

Competitive-search equilibrium. Workers attain a common market value $\bar{U}$ (free mobility across submarkets); a submarket (w, s) is active only if it delivers $\bar{U}$:

$$f\big(\tau(w, s)\big)\,\big[w + b(s)\Delta\big] = \bar{U}. \tag{28}$$

A firm in signal class s posts w and faces tightness τ to maximize per-vacancy profit $g(\tau)(y - w)$ subject to (28). Substituting $w = \bar{U}/f(\tau) - b(s)\Delta$ and using $g = f/\tau$, its problem reduces to

$$G(s) \;\equiv\; \max_{\tau > 0} \Big\{ g(\tau) \underbrace{\big[y + b(s)\Delta\big]}_{\text{post value } Y(s)} - \frac{\bar{U}}{\tau} \Big\}, \tag{29}$$

the standard competitive-search value of a post whose total match value is $Y(s) = y + b(s)\Delta$. Free entry pins $\bar{U}$ so that the marginal active post earns the vacancy cost. Since $Y(s)$ is increasing in $b(s)$ and G is increasing in Y, $G(g) \geq G(b)$ with strict inequality whenever $\pi > \frac{1}{2}$: posts that workers believe more likely to develop fill more easily at any wage, equivalently command a wage discount. That discount is the only channel through which a firm is compensated for the portable return. The next two lemmas make the within-submarket and market-utility objects well defined; they are the standard competitive-search building blocks under the regularity above.

Lemma 4 (Within-submarket existence and uniqueness)**.** *For any post value $Y > 0$ and market utility $\bar{U} \in (0, Y)$, the firm's problem (29) has a unique maximizer $\tau^*(Y, \bar{U}) \in (0, \infty)$. The value $G(Y, \bar{U})$ is continuous, strictly increasing in Y, and strictly decreasing in $\bar{U}$, with $\partial G/\partial \bar{U} = -1/\tau^* < 0$.*

Proof. Write the objective as $\Pi(\tau) = g(\tau)Y - \bar{U}/\tau = [f(\tau)Y - \bar{U}]/\tau$. Then $\Pi'(\tau) = N(\tau)/\tau^2$ with $N(\tau) \equiv Y[\tau f'(\tau) - f(\tau)] + \bar{U}$. Differentiating, $\frac{d}{d\tau}[\tau f'(\tau) - f(\tau)] = \tau f''(\tau) < 0$, so $\tau f'(\tau) - f(\tau)$ is strictly decreasing; starting from 0 at $\tau = 0$ it is negative for $\tau > 0$, hence N is strictly decreasing. By the boundary condition $\tau f'(\tau) \to 0$, $N(0^+) = \bar{U} > 0$ and $N(\infty) = \bar{U} - Y < 0$ since $\bar{U} < Y$. Thus N has a unique zero τ^*, with $\Pi' > 0$ on $(0, \tau^*)$ and $\Pi' < 0$ beyond, so τ^* is the unique global maximizer. The envelope theorem gives $\partial G/\partial \bar{U} = -1/\tau^* < 0$ and $\partial G/\partial Y = g(\tau^*) > 0$; continuity follows from the theorem of the maximum. □

Lemma 5 (Market utility)**.** *Given beliefs $(b(g), b(b))$, for any vacancy cost k below the entry value evaluated at $\bar{U} \to 0$, free entry determines a unique market utility $\bar{U} > 0$.*

Proof. Free entry equates the expected value of an entering vacancy to k. By Lemma 4 the entry value is continuous and strictly decreasing in $\bar{U}$ (each $G(s)$ is), tends to a positive bound as $\bar{U} \to 0$ (posts earn the full match value) and falls below k as $\bar{U}$ rises toward $\min_s Y(s)$. By the intermediate value theorem there is exactly one $\bar{U}$ at which the entry value equals k. □

Posture choice. A firm taking the prevailing beliefs $b(g), b(b)$ as given compares postures. D lands in g with probability π, N with probability $1 - \pi$, so the gain from developing is

$$\Gamma(\pi, n_D) \;=\; \underbrace{(2\pi - 1)}_{\text{signal shift}} \underbrace{\big[G(g) - G(b)\big]}_{\text{wage discount}} - c_R, \tag{30}$$

evaluated at the beliefs $b(g), b(b)$ that prevail when a fraction n_D develops, and the firm develops iff $\Gamma \geq 0$. Both factors of the benefit vanish as $\pi \to \frac{1}{2}$ (an uninformative signal:

$2\pi - 1 \to 0$ and $b(g) \to b(b)$, so $G(g) - G(b) \to 0$), and both at the endpoints $n_D \in \{0, 1\}$, where the population is homogeneous and the signal again carries no information ($b(g) = b(b)$). The benefit is therefore zero on the boundary of $[0, 1]$ in n_D and single-peaked in the interior.

Proposition 11 (Equilibrium set: sorting does not overturn the myopic corner unless the signal is precise)**.** *Under the matching regularity above, in the competitive-search economy with development-signal precision $\pi \in [\frac{1}{2}, 1]$ the set of posture equilibria is as follows.*

(i) The corner always exists. *The myopic corner $n_D = 0$ (no firm provides developmental review) is an equilibrium for* every *π: with $n_D = 0$ the signal carries no information ($b(g) = b(b) = 0$), the wage discount in (30) is zero, and a unilateral deviation to* D *loses c_R. Worker sorting therefore never eliminates the externality of Proposition 5.*

(ii) Full development never exists. *$n_D = 1$ is never an equilibrium: with a homogeneous developer population the signal is again uninformative ($b(g) = b(b) = 1$), so a hidden deviation to* N *saves c_R at no wage penalty. Even where sorting helps, it never sustains universal development.*

(iii) Threshold and uniqueness below it. *There is a threshold $\pi^* \in (\frac{1}{2}, 1)$, defined by $\max_{n_D \in [0,1]} \Gamma(\pi^*, n_D) = 0$, such that: for $\pi < \pi^*$ the corner is the* unique *equilibrium ($\Gamma(\pi, n_D) < 0$ for all n_D, so no developer can be supported); for $\pi \geq \pi^*$ there exist in addition one or more interior development equilibria $n_D \in (0, 1)$ at which firms are indifferent, $\Gamma(\pi, n_D) = 0$. The corner persists alongside them. The threshold π^* is increasing in the review cost c_R and decreasing in the portable value Δ.*

(iv) Poles. *As $\pi \to \frac{1}{2}$ the corner is the unique equilibrium, reproducing the spot benchmark of Case 1; as $\pi \to 1$ the signal becomes fully revealing, the wage discount approaches the entire portable return, and the development equilibrium implements the contractible allocation of Case 3. The realistic case lies strictly between, and the externality is the unique outcome throughout $\pi < \pi^*$.*

Proof. By Lemmas 4–5, for each n_D the beliefs, market utility, and submarket values $G(g), G(b)$ are well defined, so $\Gamma(\pi, n_D)$ in (30) is a well-defined continuous function on $[\frac{1}{2}, 1] \times [0, 1]$.

(i) At $n_D = 0$, Bayes' rule gives $b(g) = b(b) = 0$, so $Y(g) = Y(b) = y$ and $G(g) = G(b)$ by Lemma 4; the bracket in (30) is zero and $\Gamma(\pi, 0) = -c_R < 0$ for all π. No firm deviates: the corner is an equilibrium at every precision.

(ii) At $n_D = 1$, Bayes' rule gives $b(g) = b(b) = 1$, so again $G(g) = G(b)$ and $\Gamma(\pi, 1) = -c_R < 0$: a firm strictly prefers to deviate to N, so $n_D = 1$ is never an equilibrium.

(iii) Fix $\pi > \frac{1}{2}$. The posterior spread $b(g) - b(b)$ is zero at $n_D \in \{0, 1\}$ and strictly positive in between, and G is strictly increasing in its post value (Lemma 4) with $Y(g) - Y(b) =$

$[b(g) - b(b)]\Delta$; hence $G(g) - G(b) > 0$ on $(0,1)$ and $= 0$ at the endpoints, so $\Gamma(\pi, \cdot)$ is continuous, negative at both endpoints (equal to $-c_R$), and single-peaked on $(0,1)$. Define $\Gamma^*(\pi) \equiv \max_{n_D \in [0,1]} \Gamma(\pi, n_D)$. The benefit $(2\pi - 1)[G(g) - G(b)]$ is, at any interior n_D, zero at $\pi = \frac{1}{2}$ and strictly increasing in π (the factor $2\pi - 1$ rises and the posterior spread $b(g) - b(b)$ rises), so $\Gamma^*(\pi)$ is continuous and strictly increasing, with $\Gamma^*(\frac{1}{2}) = -c_R < 0$. If $\Gamma^*(1) \geq 0$ there is a unique $\pi^* \in (\frac{1}{2}, 1]$ with $\Gamma^*(\pi^*) = 0$; otherwise set $\pi^* > 1$ (no development equilibrium at any precision). For $\pi < \pi^*$, $\Gamma(\pi, n_D) \leq \Gamma^*(\pi) < 0$ for all n_D, so no firm ever strictly prefers D and, with (ii) excluding $n_D = 1$, the corner of (i) is the *unique* equilibrium. For $\pi \geq \pi^*$, $\Gamma(\pi, \cdot)$ is positive on a subinterval and $= -c_R < 0$ at the endpoints, so by continuity it has interior zeros; at any such n_D firms are indifferent and a mixed-posture profile with that n_D is an equilibrium, coexisting with the corner. Monotone comparative statics: the benefit is decreasing in c_R and increasing in Δ (through $Y(s)$ in Lemma 4), so Γ^* shifts down with c_R and up with Δ, moving the crossing π^* up with c_R and down with Δ.

(iv) At $\pi = \frac{1}{2}$, $\Gamma(\frac{1}{2}, n_D) = -c_R < 0$ for all n_D, so by (iii) the corner is unique: the spot allocation of Case 1, with the maximal wedge of Lemma 2. As $\pi \to 1$, a developer lands in g with certainty and the supporting interior beliefs give $b(g) \to 1$, $b(b) \to 0$, so the wage discount $G(g) - G(b) \to G(y + \Delta) - G(y)$, the full competitive-search value of the portable return; the firm internalizes Δ through the accepted wage exactly as if it had posted a binding development contract, which is Case 3, and $\mu_F = \mu_S$ by Corollary 1. □

Two readings follow, both confirming rather than overturning the analysis. First, the externality is not an artifact of the spot-wage benchmark: the myopic corner is an equilibrium of the directed-search economy for every signal precision, so allowing worker sorting cannot by itself discipline configuration. What sorting can do, when the public signal is precise enough ($\pi \geq \pi^*$), is sustain an *additional* high-development equilibrium. It never removes the bad one. Second, the threshold is demanding precisely where the paper's mechanism bites. The signal that workers can read *ex ante*, a firm's development reputation or stated programs, tracks whether review *occurs*, not whether it *teaches*, the same observability-versus-verifiability gap developed in Section 9. The learning value of review is the least observable component of a deployment, so the relevant precision is low, plausibly below π^*, where the development equilibrium does not exist and the externality is the unique outcome. Sorting on coarse signals (brand, glassdoor reputation, advertised training) moves wages on the appearance of development without disciplining its substance, leaving the wedge open. A fuller treatment would let the signal be manipulable, so that it tracks review-occurrence while the firm games the gap to review-learning; this can only raise π^* and enlarge the region in which the corner is the unique equilibrium, strengthening the conclusion.