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

Georgios Petropoulos
University of Southern California

November 3, 2026

Abstract

How a firm configures generative artificial intelligence (AI) determines whether the technology develops or hollows out its junior talent. We build a dynamic theory of early-career skill formation in which a worker can learn from work the AI completes by reviewing it. Deriving the learning technology from these events makes the premise that reviewing teaches a local, measurable condition: Reviewing out-teaches doing when workers have developed some basic verification ability. It miseducates, in the case of junior workers who lack such an ability. Seniority-blind deployments therefore face a dilemma: thin review of work the AI completes on its own hollows out junior careers, while heavy review traps entrants in a low-expertise basin. A seniority-contingent rule that withholds autonomy until a worker can verify never lowers learning below its no-AI level, whatever the size of the review premium and raises the expertise ceiling when the premium exists. However, a competitive market will not choose it: because expertise is general and review is hard to contract on, firms delegate to AI all capable work and review none of it. Mandating review without gating autonomy deepens the harm. AI deployment is 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, early-career workers, 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 adopting firms: junior employment falls while senior employment continues to rise, and 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 AI-era hiring (Lodefalk et al., 2026), a sharp slowdown in coder employment growth (Crane and Soto, 2026), and broader evidence on AI exposure (Klein Teeselink, 2025).

The worry this raises is dynamic. 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 while the future stock of expertise falls. We ask whether this harm is an intrinsic property of capable AI or a consequence of *how firms deploy it* and, if the latter, 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 system can be run with the human on the task and the machine assisting, or with the task handed to the machine and the human checking the result. We call the share of AI-capable tasks a firm hands to the machine in this second way its autonomy (the AI's autonomy), and the share of that output a worker checks in depth its review intensity. The prevailing intuition, and the structure of existing models, treats the second arrangement as pedagogically dead: a worker who does not perform a task does not learn from it. We consider the channel this ignores, that a worker can learn by reviewing what the machine produced. We build the learning technology from what happens when a worker checks AI output. She can confirm a correct step and absorb its reasoning, she can catch an error and learn where right and wrong part, or she can miss an error and absorb it as if it were right. How much reviewing teaches then depends on three measurable objects: the worker's ability to verify, which grows with seniority, the transparency of the AI's reasoning and how often, and how instructively, the AI errs.

Above a verification threshold, reviewing an expert-grade solution and its mistakes can teach more than performing a routine version of the task oneself, so handing capable work

to the machine builds skill. Below a lower threshold, reviewing teaches less than nothing: a junior who cannot yet tell right from plausible, asked to sign off on autonomous output, learns its errors. For such a worker more review is worse, not better, and delegation does not merely remove a rung; it corrupts it.

The dynamic consequences follow. A deployment that sets one configuration for every worker faces a dilemma. With thin review, delegation strips the tasks juniors learn from and their careers flatten. With heavy review, juniors absorb errors faster than they learn, and an entrant below a critical level of expertise is pulled into a low-expertise basin from which she does not recover. Whether any uniform review intensity can both keep entrants out of this trap and lift incumbents above their no-AI expertise depends on how costly mislearning is relative to the review premium. When it is costly, none can. The dilemma dissolves under a *seniority-contingent rule*: keep the worker on the task, with the AI assisting, until she can verify, and only then delegate with mandated review. The rule maximizes learning at every seniority among the configurations it chooses between, never lowers learning below its no-AI level, and makes the trap impossible. Its safety does not depend on whether reviewing can out-teach doing; its upside does. When it can, and when skills are durable enough that workers reach the verification threshold, the rule raises the long-run expertise ceiling above both its no-AI and its copilot-only levels.

A configuration that develops the workforce 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 bears its cost and captures only part of its return. Each firm therefore favors the configuration that maximizes current output: delegate every capable task and review none. Competition sharpens rather than disciplines the resulting under-provision of development, which disappears only when the deployment itself can be contracted upon. The missing rung, on this account, is a *deployment externality*, a hold-up in the formation of general human capital rather than a property of the technology. The same logic disciplines the remedy: a mandate that forces review onto autonomous output without gating autonomy by seniority can convert the market's myopic corner into the junior trap. The binding lever is how autonomy is allocated across a career, not how much of it is allowed.

The remainder proceeds as follows. Section 2 places the paper in the literature. Section 3 sets up the environment and derives the learning technology. Section 4 establishes when review and autonomy teach. Section 5 characterizes career dynamics under seniority-blind and seniority-contingent deployment. Section 6 closes the model with a labor market and derives the deployment externality. Section 7 develops the managerial framework and Section 8 concludes. Numerical illustrations under a common leading parametrization accompany the

results as figures. Appendix A contains proofs and Appendix B extensions.

2 Related Literature

The paper connects five strands literature. Our contribution has a common core: we treat the learning content of work as something AI configuration acts on, so that a delegated task can teach, fail to teach, or mislead, depending on who reviews it.

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 measures AI's task-level exposure and traces 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 (output and learning), so that a statically efficient deployment can be dynamically inefficient and the relevant policy unit is a task–seniority pair rather than an occupation. The augmentation–automation dichotomy is then incomplete: a delegated task is not pedagogically dead, and depending on who reviews it, it can have a learning value.

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), how heterogeneous agents sort into such teams and how communication technology reshapes the assignment (Antràs et al., 2006), how information and communication technologies push authority in opposite directions (Bloom et al., 2014), and which contracts implement the efficient hierarchy (Fuchs et al., 2015). Ide and Talamàs (2025) bring AI into this frame, with agents that operate autonomously or only as copilots, and Ide and Talamàs (2026) extend it to global knowledge work. Xu et al. (2025) study how generative AI reorganizes a knowledge hierarchy across automation and augmentation of the worker and expert layers, with fallible output that must be verified, and find that better AI lowers the knowledge requirement of the worker role. Friebel et al. (2026) model how productivity and learning shocks move a firm's span of control between pyramid and diamond structures. Both papers take AI's effect on skill as a primitive and ask how the organization adjusts around it. While we share the premise that AI is both autonomous and fallible and that verification is central, we follow a different approach. We endogenize the skill itself: reviewing AI output is the technology by which a worker's expertise accumulates,

so configuration acts on the *production* of the future expertise distribution rather than the *allocation* of the current one.

The economics of training. Training theory since Becker (1964) asks how investment in transferable skill is financed when it can walk out the door. 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 apprentice contracts and 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: careers remain viable when an expertise leverage ratio, the value of a trained graduate relative to AI's standalone output, exceeds a threshold. We depart from it in two ways. First, this tradition takes the learning technology as given and asks whether its returns can be financed. We ask a logically prior question that AI makes pressing: does the work environment still *deliver the learning* to be financed? Appendix B.1 shows that the learning and the financial viability conditions are logically independent. Second, we adopt the hold-up logic of this tradition. When the "training input" is a production configuration whose learning value changes sign over a career, the hold-up produces consequences that training theory does not: the market *misallocates* autonomy and review across seniority rather than simply under-investing, the natural remedy of mandating the input can backfire, and the size of the failure grows with AI capability.

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), internal wage and promotion structures (Baker et al., 1994), early-career scarring (Kahn, 2010), and the role of expertise in shielding workers from automation (Autor and Thompson, 2025). Deming (2023) shows that nearly all life-cycle wage growth occurs within jobs and that sorting into complex, learning-intensive work drives the steeper trajectories of educated workers, so a technology that alters the learning content of entry-level work bears directly on lifetime outcomes. Because that expertise is largely portable across employers (Gathmann and Schönberg, 2010), the firm cannot appropriate its return. We build on the career-dynamics tradition linking accumulation to wage and promotion paths (Farber and Gibbons, 1996; Gibbons and Waldman, 1999, 2006), adding the channel through which AI configuration governs the rate, and possibly the sign, of accumulation.

Human–AI collaboration and learning. A body of work studies how organizations combine human and machine effort: how decisions can be delegated to, split with, or aggregated across algorithms (Shrestha et al., 2019), how delegation is bounded by humans' knowledge of what they do not know (Fügener et al., 2022), when workers follow algorithmic recommendations (Kawaguchi, 2021), and why humans nominally kept in the loop fail to catch automation errors (Parasuraman and Manzey, 2010). Field experiments document large contemporaneous productivity gains from generative AI (Brynjolfsson et al., 2025b; Cui et al., 2026; Noy and Zhang, 2023; Peng et al., 2023), a jagged frontier beyond which AI degrades performance (Dell'Acqua et al., 2023) and an AI teammate that lets less experienced workers reach performance that previously required senior colleagues (Dell'Acqua et al., 2025). AI advice can also erode the unique knowledge that distinguishes workers (Fügener et al., 2021). A newer set of experiments measures learning directly. Unrestricted access to GPT-4 raised high-school students' practice performance and lowered their subsequent unassisted exam performance, while a tutor that withheld answers removed the harm (Bastani et al., 2025). Randomized AI assistance reduced mastery of an unfamiliar programming library (Shen and Tamkin, 2026) and persistence and independent performance in mathematics and reading (Liu et al., 2026). Consultants who used AI to perform data-science tasks did not retain the skill without it (Wiles et al., 2026). Delayed learning gains concentrated among students who used AI to explain rather than to produce (Contractor and Reyes, 2026) and evaluating flawed AI solutions taught as much as, but not more than, solving the problems (Dickey et al., 2026). This evidence establishes that AI's effect on learning depends on how the worker engages with it, which is the margin our model makes explicit. We add a learning technology in which engagement is decomposed into verification events, so that one configuration can teach, fail to teach, or mislead depending on seniority, and we derive the dynamic and market consequences: the configuration that optimizes current human–AI performance also sets the rate of accumulation, and with portable skill and non-contractible review the market selects too much autonomy and too little review.

3 The Model

3.1 Workers, tasks, and AI

Time is continuous, $t \geq 0$, and 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 receives a flow of tasks indexed by complexity $z \in [0, 1]$, drawn from a distribution $F(\cdot \mid h)$.

Assumption 1 (Seniority). *(i) $F(\cdot \mid h)$ is ordered by first-order stochastic dominance in h: for $h' > h$, $F(z \mid h') \leq F(z \mid h)$ for all z. (ii) The support of $F(\cdot \mid h)$ lies in $[h, 1]$: a worker is assigned tasks at or above her current level.*

Assumption 1 formalizes how work is allocated over a career. Part (i) says that more senior workers face harder tasks on average, as in knowledge hierarchies in which less experienced workers handle routine problems and pass harder ones up (Garicano, 2000), in internal job ladders (Baker et al., 1994), and in the sorting of experienced workers into more complex work (Deming, 2023). Part (ii) says that a worker's queue contains no tasks below her level.[1] It implies that the AI-capable share of a worker's queue shrink as she advances and vanish once her queue lies beyond the frontier.

Each task carries learning content $\ell(z) \geq 0$, the marginal, transferable increment to general human capital a worker can obtain by engaging with a task of complexity z. We take $\ell > 0$ on $[0, 1)$.

The AI has a *capability frontier* $\theta \in [0, 1)$: it can complete tasks with $z \leq \theta$ (*capable* tasks). It does so fallibly: each component of a capable task it completes is correct with probability $q \equiv A(\theta) \in (0, 1]$, with $A' > 0$. The frontier and the error rate are distinct objects, the "jagged frontier" of Dell'Acqua et al. (2023): an AI can attempt a task and still err on parts of it. The deployed system exposes its reasoning (a derivation, a chain of thought, citations) with probability $e \in [0, 1]$, its *explanation quality*. Tasks beyond the frontier ($z > \theta$) are always performed by the worker. On capable tasks the firm chooses a *deployment* $u = (a, r, m)$:

- *Autonomy* $a \in [0, 1]$: the share of capable tasks the AI completes without the worker performing them.
- *Review intensity* $r \in [0, 1]$: the share of delegated output the worker verifies in depth.
- *Copilot intensity* $m \geq 0$: the assistance the worker receives on capable tasks she performs herself.

The worker's *verification ability* is $p(h) \in [0, 1]$, the probability that she correctly classifies a component she checks, confirming a correct one or detecting an erroneous one, with $p' > 0$ and $p(0) = 0$: an entrant cannot tell a correct solution from a plausible one.

[1] In practice senior workers still perform some routine work, but such tasks carry little learning content for someone who has mastered them, so what part (ii) requires is that a worker's *learnable* tasks lie at or above her level.

3.2 Learning by doing, with assistance, and by reviewing

We decompose each task into a unit measure of knowledge components, a fraction $\ell(z)$ of which is learnable. We measure learning in *units of doing*: performing one learnable component unaided teaches one unit, and $\alpha > 0$ converts units into human capital per unit of time. Every learning value below is expressed in these units, so reviewing and doing are compared on a single scale.

When the worker reviews a delegated component, one of three events occurs. With probability qp the component is correct and she confirms it, absorbing $\omega(e) \equiv \bar{\omega}\,[e + (1-e)\gamma]$ units: $\bar{\omega} > 0$ when the reasoning is exposed, the worked-example value, and $\gamma\bar{\omega}$, with $\gamma \in [0, 1)$, when she sees only an opaque correct answer. With probability $(1-q)p$ the component is wrong and she catches it. Locating and correcting a mistake teaches $\xi \geq 0$ units. With probability $(1-q)(1-p)$ the component is wrong and she misses it, absorbing the error as if it were correct and losing $\chi \geq 0$ units: *mislearning*. A correct component she fails to confirm contributes nothing. The net value of reviewing a learnable component is

$$v_R(q, e, p) \;=\; \underbrace{p\,q\,\omega(e)}_{\text{confirmed}} \;+\; \underbrace{p\,(1-q)\,\xi}_{\text{caught}} \;-\; \underbrace{(1-p)(1-q)\,\chi}_{\text{missed}} \;=\; p\,\Sigma(q, e) - (1-q)\chi, \qquad (1)$$

with $\Sigma(q, e) \equiv q\,\omega(e) + (1-q)(\xi + \chi) > 0$. Delegated output the worker does not review, a fraction $1 - r$, does not engage her: she neither learns from it nor absorbs its errors.

On a capable task she performs herself, the worker learns the unit from doing and, with intensity m, receives the AI's suggestion, which she absorbs when it is correct and she recognizes it. Copilot learning per learnable component is therefore $1 + m\,q\,p$. Because she authors the work, an incorrect suggestion is either rejected or exposed when her own work fails, and we normalize its learning effect to zero.[2]

Reviewing can out-teach doing only if some event in (1) is worth more than doing. Even with perfect verification, $v_R(q, e, 1) = q\,\omega(e) + (1-q)\,\xi$ is a weighted average of the value of a confirmed component and of a caught error, so the review channel can dominate only if $\omega(e) > 1$ or $\xi > 1$. Both have empirical counterparts. For a learner with the prerequisites to follow it, studying a fully worked expert solution yields better subsequent independent performance than solving the equivalent problem unaided, the *worked-example effect* (Sweller and Cooper, 1985; Cooper and Sweller, 1987; Sweller et al., 1998), confirmed for retention and transfer by a recent meta-analysis (Barbieri et al., 2023). It operates through the active processing of a model solution (Chi et al., 1989) and the observation of expert performance a

[2] If incorrect suggestions could also be absorbed, at cost χ_C each, the copilot gain mqp becomes $m[qp - (1-q)(1-p)\chi_C]$, which can be negative for raw entrants. The natural response is then $m = 0$ for them, and every result below holds with copilot learning replaced by hands-on learning on that range.

novice cannot yet produce, the mechanism of apprenticeship (Collins et al., 1989). Detecting and correcting errors raises transfer performance, as in error-management training (Keith and Frese, 2008) and instruction with incorrect worked examples (Durkin and Rittle-Johnson, 2012; Booth et al., 2013). The same literatures supply the limits the model builds in: the advantage of studying solutions requires that the learner can follow them and fades with expertise, the *expertise-reversal effect* (Kalyuga et al., 2003), and humans who monitor automation absorb its errors and grow complacent as it becomes more reliable (Parasuraman and Manzey, 2010), which grounds the mislearning cost χ and the role of q.

3.3 Aggregation, output, and accumulation

The learning masses of capable and beyond-frontier tasks and the measure of capable tasks are

$$\mathcal{L}_\theta(h) \equiv \int_{\{z\le\theta\}} \ell(z)\,\mathrm{d}F(z \mid h), \qquad \bar{\mathcal{L}}_\theta(h) \equiv \int_{\{z>\theta\}} \ell(z)\,\mathrm{d}F(z \mid h), \qquad Q_\theta(h) \equiv F(\theta \mid h). \tag{2}$$

We assume that F and ℓ are regular enough for these masses to be Lipschitz in h.[3] By Assumption 1(ii), $\mathcal{L}_\theta(h) = Q_\theta(h) = 0$ for $h \ge \theta$: AI cannot touch the learning of a worker whose entire queue lies beyond its frontier. Aggregating the three channels, the gross learning rate is

$$\Phi(h;u) = \alpha\Big[\bar{\mathcal{L}}_\theta(h) + (1-a)\big(1+m\,q\,p(h)\big)\mathcal{L}_\theta(h) + a\,r\,v_R\big(q,e,p(h)\big)\mathcal{L}_\theta(h)\Big]. \tag{3}$$

Without AI every task is performed hands-on and $\Phi_0(h) \equiv \alpha\int \ell(z)\,\mathrm{d}F(z \mid h) = \alpha\big[\mathcal{L}_\theta(h) + \bar{\mathcal{L}}_\theta(h)\big]$: capability relabels tasks between channels without changing the total learnable mass. Unlike standard learning technologies, Φ can be negative: a worker who reviews more errors than she catches loses expertise.

Flow output combines the worker and the AI:

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

where $y_0(h)$ is baseline output, increasing in h, $\pi(a;h,\theta)$ is the throughput gain from delegating capable tasks, with $\pi_a > 0$, so that autonomy raises current output and $C_r, C_m \ge 0$ are strictly increasing per-task costs of review and assistance with $C_r(0) = C_m(0) = 0$. Costs scale with the *measure* of capable tasks rather than their learning content: a task is not

[3] Lipschitz continuity of the task masses ensures that the law of motion has a unique solution from any initial condition. It rules out discrete jumps in task assignment but permits kinks.

cheaper to review because the worker learns little from it. In this baseline review carries a current cost and no current-output benefit. We later relax this. Human capital evolves as

$$\dot{h} = \Phi(h; u) - \delta h, \tag{5}$$

with depreciation rate $\delta > 0$, on the interval $[0, 1]$.[4]

4 When Do Review and Autonomy Teach?

Whether handing a task to AI helps or harms a worker's learning turns on two comparisons: does reviewing the AI's output teach anything, and does it teach more than performing the task would have? We hold h fixed here. We begin with the review value itself.

Lemma 1 (Review value)**.** *The review value* (1) *satisfies:*

(i) Verification. *v_R is affine and strictly increasing in p, with slope $\Sigma(q, e)$, and $v_R < 0$ if and only if $p < p_\chi \equiv (1 - q)\chi/\Sigma(q, e)$: below the* mislearning threshold *p_χ, reviewing teaches less than nothing.*

(ii) Transparency. *$\partial v_R/\partial e = p\, q\, \bar{\omega}(1 - \gamma) \geq 0$.*

(iii) Capability is double-edged. *$\partial v_R/\partial q = \chi - p\,(\xi + \chi - \omega(e))$. If $\xi > \omega(e)$, this is positive for $p < \bar{p} \equiv \chi/(\xi + \chi - \omega(e)) \in [0, 1)$ and negative for $p > \bar{p}$: more reliable AI raises the review value of weak verifiers, who absorb fewer errors, and lowers that of strong verifiers, who lose the instructive errors they learn most from. If $\xi \leq \omega(e)$, v_R is nondecreasing in q.*

(iv) Ceiling. *$v_R(q, e, p) \leq v_R(q, e, 1) = q\,\omega(e) + (1 - q)\,\xi \leq \max\{\omega(e), \xi\}$.*

Comparing reviewing with doing requires a per-task benchmark: the alternative to delegating a capable task is performing it with assistance.

Definition 1 (Scaffolding ratio)**.** *For a capable task, the* scaffolding ratio *is learning when the task is delegated and reviewed relative to learning when it is performed with assistance,*

$$\psi(h) \equiv \frac{r\, v_R\big(q, e, p(h)\big)}{1 + m\, q\, p(h)}. \tag{6}$$

When $\psi > 1$, delegating the task and having the worker check it builds more skill than having her do it. When $\psi < 1$, delegation strips learning she would otherwise have gained.

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

When $\psi < 0$, delegation with review actively miseducates. Substituting (3), the effect of AI on learning relative to the no-AI benchmark decomposes as

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

and the two margins the firm controls on delegated work are

$$\frac{\partial \Phi}{\partial a} = \alpha\, \mathcal{L}_\theta(h)\big(1 + m\, q\, p(h)\big)\big(\psi(h) - 1\big), \qquad \frac{\partial \Phi}{\partial r} = \alpha\, a\, \mathcal{L}_\theta(h)\, v_R\big(q, e, p(h)\big). \tag{8}$$

The autonomy margin switches sign where reviewing and doing teach the same. The review margin switches sign where reviewing starts to teach at all. Because v_R is affine in p, both switch points are available in closed form.

Proposition 1 (When review and autonomy teach)**.** *Fix capability $\theta < 1$, explanation quality e, and a deployment $u = (a, r, m)$ with $r > 0$, and consider a worker who faces capable tasks, $h < \theta$.*

(i) Review margin. *For $a > 0$,* $\operatorname{sign}(\partial\Phi/\partial r) = \operatorname{sign}\big(p(h) - p_\chi\big)$*: for a worker below the mislearning threshold, more review lowers learning.*

(ii) Autonomy margin. $\operatorname{sign}(\partial\Phi/\partial a) = \operatorname{sign}\big(\psi(h) - 1\big)$*. If $r\,\Sigma(q,e) > m\,q$, then $\psi(h) \gtrless 1$ as $p(h) \gtrless p_\psi$, with* scaffolding threshold

$$p_\psi \equiv \frac{1 + r(1-q)\chi}{r\,\Sigma(q,e) - m\,q}; \tag{9}$$

otherwise $\psi < 1$ for every p and autonomy always lowers learning. In either case $p_\chi < p_\psi$.

(iii) Scaffolding region. *Autonomy strictly builds learning on $\{h : p(h) > p_\psi,\ h < \theta\}$. This region is non-empty if and only if*

$$r\, v_R\big(q, e, p(\theta)\big) > 1 + m\, q\, p(\theta), \tag{10}$$

for which a necessary condition is $q\,\omega(e) + (1-q)\,\xi > 1/r \geq 1$: a confirmed component or a caught error must teach more than doing.

(iv) Comparative statics. *The threshold p_ψ is decreasing in r, e, $\bar{\omega}$, and ξ, increasing in m, and increasing in χ whenever $p_\psi < 1$. The mislearning threshold $p_\chi = \chi/\big[\xi + \chi + \omega(e)\, q/(1-q)\big]$ is decreasing in q.*

Proposition 1 sorts workers who face capable tasks into three bands. Writing h_χ and h_ψ for the seniorities at which $p(h)$ reaches p_χ and p_ψ:

Definition 2 (Seniority bands)**.** *Given* (θ, e, r, m)*, a worker is a* raw junior *if* $p(h) < p_\chi$*,* *an* intermediate *if* $p_\chi \leq p(h) < p_\psi$*, and* seniority-ready *if* $p(h) \geq p_\psi$*.*

For seniority-ready workers below the frontier, reviewing expert-grade output and its errors out-teaches hands-on practice, and restricting autonomy would slow their development. The heuristic "keep the human hands-on to preserve learning" is wrong for them. For intermediates, reviewing teaches, but less than doing. For raw juniors, the result is starker than a missing rung. Delegating their capable work removes learning, and asking them to review it subtracts more: a worker who cannot yet verify absorbs the errors she signs off on. Keeping a junior "in the loop" by routing autonomous output to her for approval is therefore exactly the wrong protection; the protective configuration keeps her on the task.

Capability shifts these thresholds, but not in a fixed direction. More reliable AI lowers the mislearning threshold, since there are fewer errors to absorb, while its effect on the scaffolding threshold is ambiguous, since it also removes errors that able reviewers learn from (Lemma 1(iii)). Capability changes how much is at stake, through the mass of capable tasks, but does not by itself decide whether delegation builds skill; the margins the firm controls do.

Figure 1 illustrates these properties. In panel (a), learning from a delegated and reviewed task rises with verification ability: it is negative for raw juniors, positive but below learning with assistance for intermediates, and above it only for seniority-ready workers, and more transparent output lowers both thresholds. In panel (b), more reliable AI raises the review value of weak verifiers and lowers that of strong ones.

The seniority profile follows from two opposing forces. Raw juniors have low verification ability and, by Assumption 1, the largest share of capable tasks, $Q_\theta(h)$ being non-increasing in h (in the leading parametrization $\mathcal{L}_\theta$ is strictly decreasing as well), so they bear the rung removal. Workers beyond the frontier face no capable tasks and AI leaves their learning unchanged. The scaffolding benefit is concentrated on seniority-ready workers whose queues still contain capable tasks. Because the margin along which these effects operate is verification ability, AI also reshapes the dispersion of careers:

Corollary 1 (Verification and dispersion)**.** *Consider two workers with verification profiles* $p_1(h) > p_2(h)$ *for all* $h \in (0, 1]$*. Under any common deployment with* $h < \theta$ *and either* $ar > 0$ *or* $(1-a)m > 0$*, the first worker has a strictly higher learning rate at every common seniority, and reaches the mislearning and scaffolding thresholds at lower seniority. Because the wage prices current type, the two can start at the same wage while their wage growth diverges.*

Within an AI-exposed occupation, the spread of early-career wage growth should there-

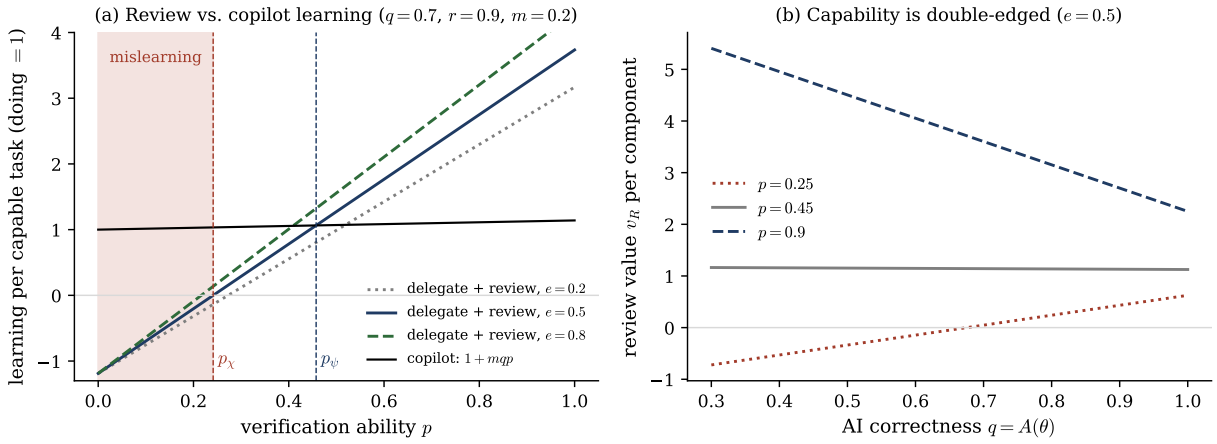


Figure 1: Review value and the two thresholds. *(a)* Learning per capable task when it is delegated and reviewed (rv_R) and when it is performed with assistance ($1 + mqp$), against verification ability, for three levels of explanation quality; the thresholds p_χ and p_ψ are marked for $e = 0.5$. *(b)* Review value per component against AI correctness, for three levels of verification ability. Unless stated otherwise, all figures use the following *leading parametrization*: tasks $z \mid h \sim \text{Uniform}[h, 1]$; learning content $\ell(z) = 1 - z$; AI correctness $A(\theta) = \theta$; verification ability $p(h) = h$; hands-on learning rate $\alpha = 0.045$; depreciation $\delta = 0.015$; capability $\theta = 0.70$, so $q = 0.70$; explanation quality $e = 0.5$; opacity discount $\gamma = 0.25$; worked-example value $\bar{\omega} = 4$; caught-error value $\xi = 8$; mislearning cost $\chi = 4.4$; and an entrant at $h_0 = 0.15$. Without AI, α and δ imply growth in expertise of about one half over the first five years and a depreciation rate in the range estimated for general human capital (Adda and Dustmann, 2023; Arrazola and de Hevia, 2004). Panel (a): $r = 0.9$, $m = 0.2$, $e \in \{0.2, 0.5, 0.8\}$. Panel (b): $e = 0.5$, $p \in \{0.25, 0.45, 0.9\}$, $q \in [0.3, 1]$.

fore widen with AI adoption and line up with measurable verification skill rather than with exposure as such, a divergence in slopes that entry-wage comparisons would miss.

5 Career Dynamics

We now let human capital evolve under (5). A *deployment policy* is a map $h \mapsto u(h)$ specifying the configuration a worker of seniority h faces. We compare two classes: *seniority-blind* policies, which apply one configuration to every worker, the natural reading of a firm-wide automation target, and *seniority-contingent* policies, which condition autonomy on verification ability.

5.1 Steady states and traps

Lemma 2 (Steady states and traps)**.** *Let $u(\cdot)$ be a deployment policy for which $G(h) \equiv \Phi(h; u(h)) - \delta h$ is Lipschitz on $[0,1]$, and suppose the learnable content of the most senior tasks cannot sustain full expertise, $\Phi(1; u(1)) < \delta$.*

(i) If $\Phi(0; u(0)) > 0$, (5) has a steady state in $(0,1)$, and paths from $h = 0$ converge to the smallest one. If $\Phi(h; u(h))/h$ is strictly decreasing on $(0,1]$, equivalently the elasticity of gross learning in human capital is below one, the steady state is unique and globally asymptotically stable.

(ii) If $\Phi(0; u(0)) < 0$, $h = 0$ is absorbing and locally stable. Let $h_{\text{trap}} \equiv \inf\{h \in (0,1] : G(h) \geq 0\}$, with $h_{\text{trap}} \equiv 1$ if $G < 0$ on $(0,1]$. Every worker who starts below h_{trap} converges to $h = 0$. If G is single-peaked with a positive maximum, it has exactly two zeros in $(0,1)$, the trap edge *h_{trap} and a stable steady state $h^H > h_{\text{trap}}$, and every worker who starts above h_{trap} converges to h^H: the long run is history-dependent.*

Case (ii) is new relative to models in which learning is a nonnegative flow. Since $p(0) = 0$, an entrant confirms nothing and catches nothing, so under a deployment u her learning rate is

$$\Phi(0; u) = \alpha \Big[\underbrace{\bar{\mathcal{L}}_\theta(0)}_{\text{frontier}} + \underbrace{(1-a)\mathcal{L}_\theta(0)}_{\text{copilot}} - \underbrace{a\, r\, (1-q)\chi\, \mathcal{L}_\theta(0)}_{\text{mislearning}} \Big], \tag{11}$$

which is negative when enough of her capable work is delegated and routed back to her for approval. The single-crossing condition in (i) accommodates a feature special to the review channel: because review value rises with seniority, learning can be locally increasing in h, a virtuous circle in which becoming more senior makes review more productive. Uniqueness requires only that this force not overturn the eventual decline driven by the shrinking stock of learnable content.

5.2 Seniority-blind deployment

Proposition 2 (Seniority-blind deployment)**.** *Fix a deployment $u = (a, r, m)$ with $a \in (0,1]$ and $r \in (0,1]$, applied at every seniority, and let $\theta < 1$.*

(i) Early loss. *For $h < \theta$, $\Phi(h; u) < \Phi_0(h)$ if and only if $p(h) < p_c$, where*

$$p_c \equiv \frac{a\,[1 + r(1-q)\chi]}{(1-a)\, m\, q + a\, r\, \Sigma(q, e)} \;\leq\; p_\psi. \tag{12}$$

(ii) Junior trap. $\Phi(0;u) < 0$ *if and only if*

$$r > \hat{r}(a) \equiv \frac{\bar{\mathcal{L}}_\theta(0) + (1-a)\mathcal{L}_\theta(0)}{a\,(1-q)\,\chi\,\mathcal{L}_\theta(0)}. \tag{13}$$

For given task masses, $\hat{r}(a)$ is decreasing in a and χ, increasing in q, and independent of e, $\bar{\omega}$, and ξ. If the trap edge lies among raw juniors who face capable tasks, $p(h_{\text{trap}}) < p_\chi$ and $h_{\text{trap}} < \theta$, and G crosses zero transversally there, then h_{trap} is strictly increasing in r.

(iii) Review watershed for incumbents. *Suppose $h_0^* < \theta$, where h_0^* is the no-AI steady state, and $p(h_0^*) > p_\chi$, and let*

$$r^* \equiv \frac{a - (1-a)\,m\,q\,p(h_0^*)}{a\,v_R\big(q,e,p(h_0^*)\big)}. \tag{14}$$

A worker at the no-AI ceiling h_0^* who is exposed to u converges to a long-run level above h_0^* if $r > r^*$ and below it if $r < r^*$. The watershed r^* is increasing in δ whenever $\Phi_0(h)/h$ is strictly decreasing; when $r^* > 1$, no feasible review intensity makes u expertise-building for incumbents.

(iv) Entrants versus incumbents. *Let an entrant start at $h_0 < \theta$ with $p(h_0) < p_\chi$, and let*

$$r_e \equiv \max\Big\{0,\ \frac{\Phi\big(h_0;(a,0,m)\big) - \delta h_0}{\alpha\,a\,\mathcal{L}_\theta(h_0)\,\big|v_R\big(q,e,p(h_0)\big)\big|}\Big\}. \tag{15}$$

The entrant ends with less expertise than she started with for every $r > r_e$, and with more for every $r < r_e$. The threshold r_e is nondecreasing and r^* nonincreasing in $\bar{\omega}$, ξ, and e, and the reverse holds for χ. Hence, under the conditions of (iii), if $r_e < r^*$, every review intensity that lifts incumbents above h_0^* leaves entrants worse off than at entry.

Part (i) is the dynamic counterpart of the static reversal: wherever $p(h) < p_c$ the early-career learning rate, and with it the wage slope, is flatter than without AI. Part (ii) says that a seniority-blind deployment with thin review hollows out early careers without trapping anyone. The same deployment with heavy review traps them. The trap condition (13) depends on autonomy, the error rate, the mislearning cost, and the entry task masses, and not at all on how much review could teach an able reviewer, because an entrant confirms nothing and catches nothing. Raising review intensity, the natural corrective for hollowed-out junior work, is precisely what creates the trap.

Part (iii) is the long-run counterpart for workers who can verify. For an incumbent at the no-AI ceiling, review sets the *sign* of AI's long-run effect: above r^* the same autonomous deployment raises her expertise, below it lowers it. The watershed has a closed form because

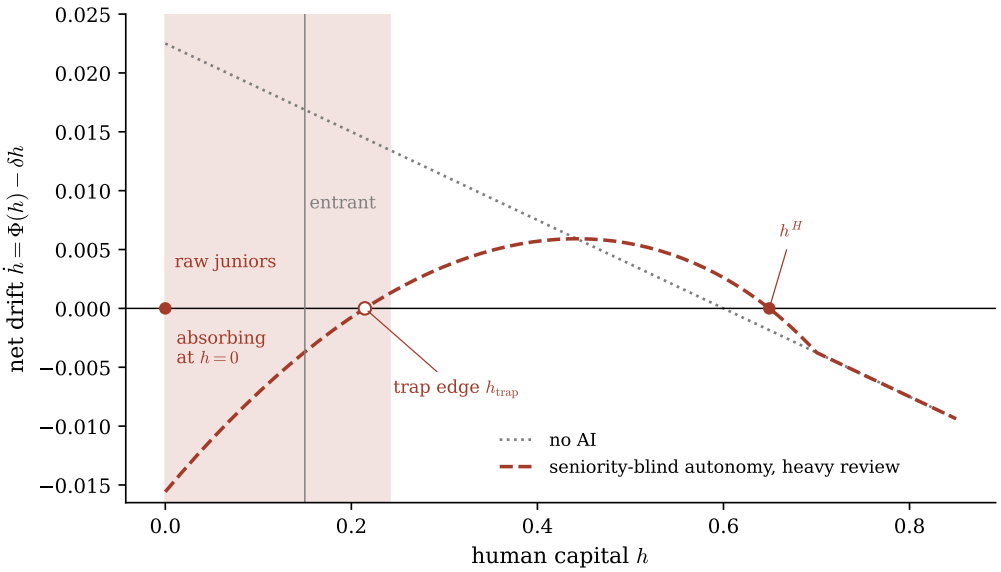


Figure 2: Net drift of human capital, $\dot{h} = \Phi(h) - \delta h$, under seniority-blind autonomy with heavy review and without AI. Filled and open circles mark stable and unstable steady states. The shaded band marks raw juniors. The leading parametrization is used with seniority-blind deployment $(a, r, m) = (0.85, 0.9, 0.2)$. No AI corresponds to $\theta = 0$.

the sign of the drift at h_0^* equals the sign of $\Phi - \Phi_0$ there, so r^* is the review intensity at which the loss threshold p_c of part (i) equals $p(h_0^*)$. It is therefore independent of the task distribution and the learning-content profile, which enter only through h_0^*. Durability enters the same way: faster obsolescence lowers h_0^*, raising r^*, until review is exhausted at $r^* = 1$. This yields a cross-occupation prediction. In fields where expertise is durable (steep tenure–earnings profiles, long skill half-lives), AI paired with review can build incumbents' expertise. In fast-churning fields, the same configuration erodes it because workers never settle at a level where they verify well enough.

Parts (ii)–(iv) together define the dilemma of seniority-blind deployment. One review intensity governs both groups, and it pushes them in opposite directions: entrants need little review of autonomous output, incumbents need a lot. Part (iv) makes the conflict precise and says when it binds. A larger review premium lowers the intensity incumbents need and raises the intensity entrants can bear, while a larger mislearning cost does the reverse, so the dilemma arises when mislearning is costly relative to the premium.

Figure 2 shows the mechanism of the trap: under seniority-blind autonomy with heavy review, raw juniors absorb the errors they sign off on, so the drift is negative at entry, turns positive once workers verify well enough, and turns negative again as learnable content runs out. Workers on either side of the trap edge end in opposite places. Figure 3 traces the

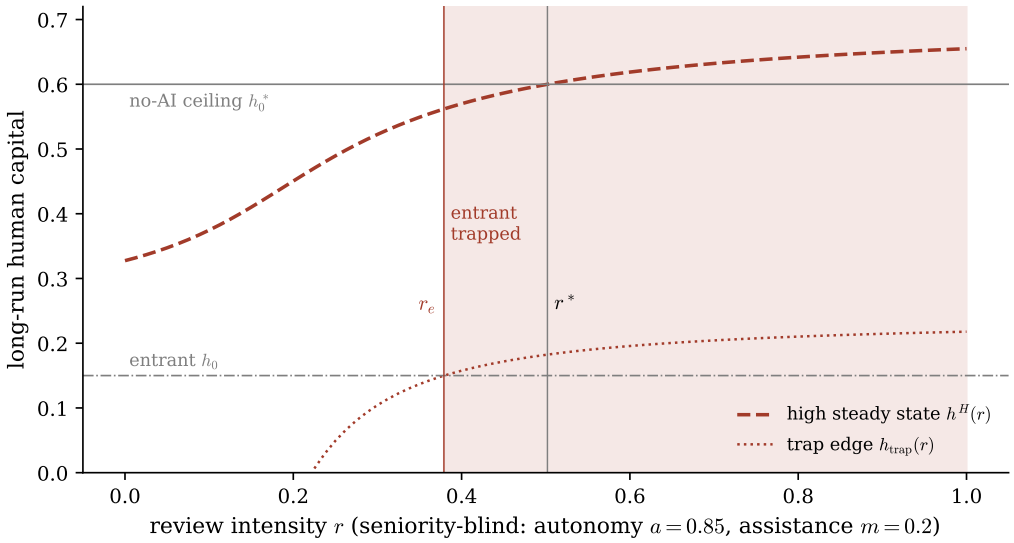


Figure 3: Long-run human capital against review intensity under seniority-blind autonomy: the high steady state $h^H(r)$ and the trap edge $h_{\text{trap}}(r)$, which appears at $\hat{r}$ of (13), relative to the no-AI ceiling and the entrant's initial level; r_e and r^* are given by (15) and (14). Same parameter values as in Figure 2, with review intensity r varying over $[0, 1]$.

dilemma as review rises: the trap edge appears and reaches the entrant before the high steady state clears the no-AI ceiling, so no uniform review intensity both protects the entrant and lifts incumbents. Figure 4 shows that this is a property of a region of parameters rather than of a point: the dilemma holds when mislearning is costly relative to the review premium, gives way to a window of uniform intensities that do both when the premium is strong or mislearning mild, and gives way to review exhaustion when the premium is weak.

5.3 Seniority-contingent deployment

The dilemma of Section 5.2 arises because a seniority-blind deployment applies one configuration to workers who differ in their ability to learn from it: the review intensity that builds incumbents' expertise is the one that traps entrants. The natural alternative conditions the configuration on verification ability.

Definition 3 (Seniority-contingent rule)**.** *Given delegation settings (a, r, m) with $a > 0$, the* seniority-contingent rule $u^g(\cdot)$ *assigns* copilot $u^C \equiv (0, \cdot, m)$*, with autonomy withheld and review therefore moot, to workers with $p(h) < p_\psi$, and* delegate-and-review $u^D \equiv (a, r, m)$ *to workers with $p(h) \geq p_\psi$, where p_ψ is the scaffolding threshold (9) at (r, m). If the scaffolding region is empty, the rule assigns u^C throughout. We write $\Phi^g(h) \equiv \Phi(h; u^g(h))$ and $\Phi^C(h) \equiv \Phi(h; u^C)$.*

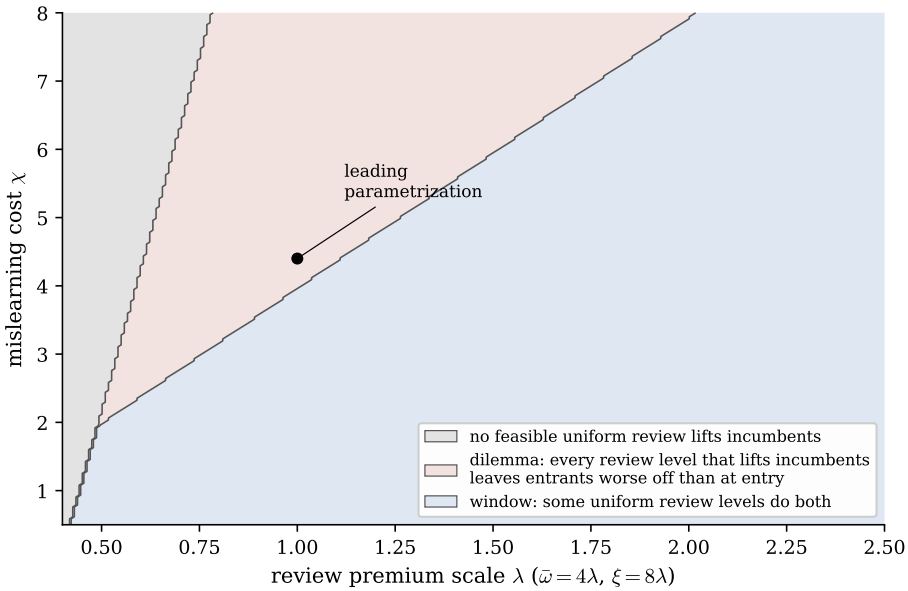


Figure 4: Where the dilemma arises: regions in which no feasible uniform review intensity lifts incumbents above the no-AI ceiling ($r^* > 1$), in which every intensity that does so leaves entrants worse off than at entry ($r_e < r^* \leq 1$), and in which some uniform intensities do both. The leading parametrization is used, except that the worked-example and caught-error values are scaled by a common factor, $\bar{\omega} = 4\lambda$ and $\xi = 8\lambda$, and the mislearning cost χ varies. $(a, m) = (0.85, 0.2)$. The dot marks the leading parametrization.

The rule withholds autonomy exactly until reviewing begins to teach more than doing, and in particular never routes autonomous output to a raw junior.

Proposition 3 (The seniority-contingent rule)**.** *Let $\theta < 1$ and $\Phi_0(1) < \delta$.*

(i) Learning maximization and safety. $\Phi^g(h) = \max\{\Phi^C(h), \Phi(h; u^D)\} \geq \Phi^C(h) \geq \Phi_0(h)$ *for every h, so the rule weakly dominates both the no-AI benchmark and the seniority-blind deployment u^D at every seniority. Φ^g is continuous; $\Phi^g > \Phi^C$ on the scaffolding region; and $\Phi^C > \Phi_0$ wherever $h < \theta$, $p(h) > 0$, and $m > 0$. The rule never assigns autonomy, and hence never assigns review, to a raw junior.*

(ii) No trap and no early loss. $\Phi^g(0) = \Phi_0(0) > 0$. *From any initial condition, the worker's path under the rule lies weakly above her paths without AI, under copilot alone, and under u^D, at every date.*

(iii) Ceiling. *From any initial condition, the long-run level under the rule is at least that under copilot alone, which is at least that without AI. If $\Phi^g(h)/h$ is strictly decreasing, the rule has a unique, globally stable steady state h^*_g.*

(iv) Activation and durability. *Suppose $\Phi^C(h)/h$ is strictly decreasing, with copilot-only ceiling h^*_C, that $h_\psi < \theta$, and that an entrant starts at $h_0 < \min\{h_\psi, h^*_C\}$. She reaches*

the delegation phase if and only if

$$\delta < \delta^{**} \equiv \Phi^{C}(h_\psi)/h_\psi, \tag{16}$$

in which case her long-run expertise is at least h_C^*, and strictly above it when $h_C^* < \theta$; otherwise it equals h_C^* and the review channel never operates along her path.

(v) Review. *Φ^g is pointwise nondecreasing in r and in a; hence so are career paths and long-run expertise under the rule.*

Proposition 3 separates two claims. The *safety* of the rule is unconditional: whatever the review premium, configuring AI by seniority never places a worker below her no-AI learning rate at any date, and the junior trap cannot form. When the scaffolding region is non-empty the rule achieves this by gating; when it is empty, by abstaining. The *upside*, expertise strictly above what copilot alone delivers, is conditional on two things: that reviewing can out-teach doing for some workers below the frontier, condition (10), and that skills are durable enough for workers to reach the threshold, condition (16). The second condition gives the durability prediction its sharpest form: under seniority-contingent deployment, the review channel matters in fields where workers reach verification competence, and is idle elsewhere. The rule's cost is not learning but current output, since withholding autonomy from juniors forgoes throughput. Remark 1 shows that it is the limit of the planner's policy as the social value of skill grows; who bears its cost, and why a market will not, is the subject of Section 6.

Figure 5 compares the rule with the seniority-blind deployment it builds on. Panel (a) shows why the trap cannot form. Below h_ψ the rule follows copilot, whose drift is positive at entry and lies above the no-AI drift, and above h_ψ it coincides with the seniority-blind deployment, so it reaches the same high steady state without passing through a trap edge (Proposition 3(i)–(ii)). Panel (b) shows the same comparison as review intensity varies. The rule's long-run expertise never falls below the seniority-blind high state or the no-AI ceiling; it equals the copilot ceiling until review is intensive enough for workers to reach the delegation phase, and rises with review thereafter (Proposition 3(iv)–(v)). Across the range of review intensities at which seniority-blind deployment leaves the entrant worse off than at entry, the rule keeps her on a rising path.

In Figure 6 panel (a), the rule coincides with copilot while the worker is below h_ψ, so her early career tracks the copilot path, above the no-AI path and far above the seniority-blind and competitive paths, and it pulls ahead once she can verify; the seniority-blind deployment with heavy review sends the entrant toward zero while a worker who starts above the trap edge climbs. In the leading parametrization the rule begins to delegate late in the career,

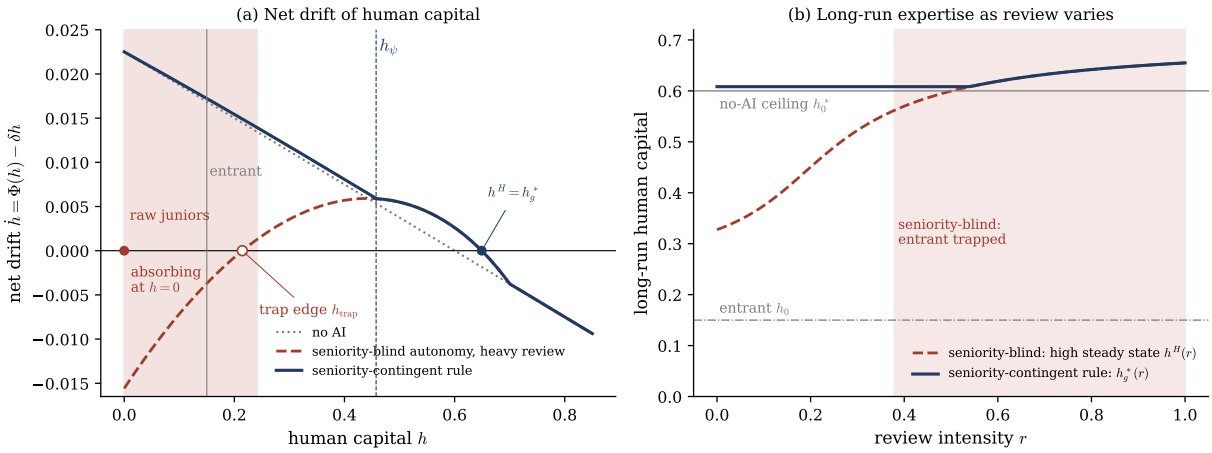


Figure 5: The seniority-contingent rule compared with seniority-blind deployment. *(a)* Net drift of human capital under the rule, under seniority-blind autonomy with heavy review, and without AI. *(b)* Long-run human capital against review intensity: the rule's steady state $h_g^*(r)$ and the seniority-blind high steady state $h^H(r)$, relative to the no-AI ceiling and the entrant's initial level. The shaded range marks the review intensities at which seniority-blind deployment leaves the entrant worse off than at entry ($r > r_e$). Same parameter values as in Figure 2. The rule assigns copilot $(0, \cdot, 0.2)$ below h_ψ and $(0.85, 0.9, 0.2)$ above, and in panel (b) review intensity r varies over $[0, 1]$ in both deployments.

so its gain there is a higher ceiling rather than faster early growth. Panel (b) shows that this timing is governed by how fast verification ability develops relative to general expertise: when verification develops early, the rule begins to delegate within the first years of a career and pulls ahead of copilot within the first decade; when it develops late, the rule abstains along the path, as in Proposition 3(iv).

6 Wages, Deployment Choice, and the Autonomy Externality

6.1 Wages

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 u, produces $Y(h; u)$, pays $w(h)$, and loses the worker at hazard s, after which she re-enters the market and is valued at her type. The firm's match value $J(h)$

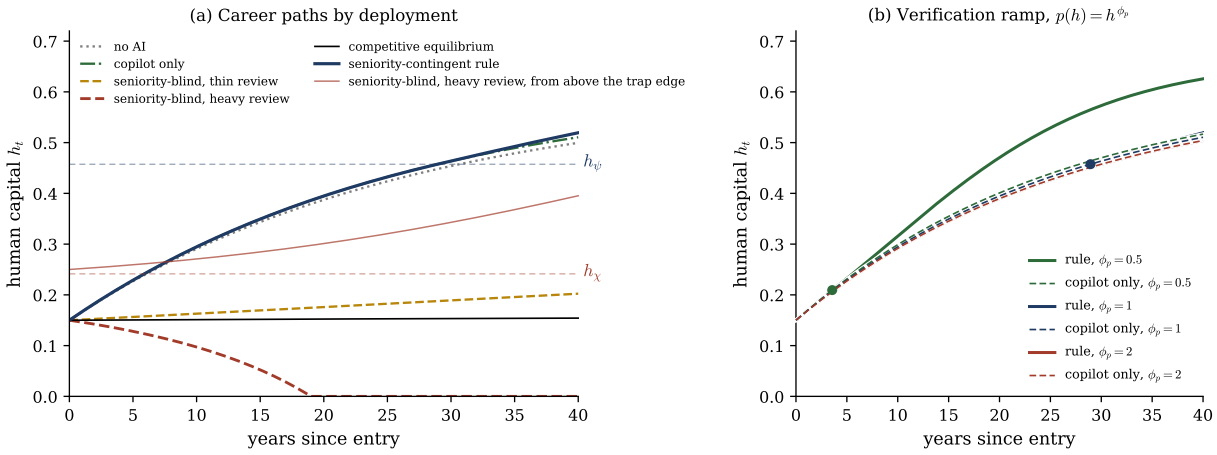


Figure 6: Career paths from entry. *(a)* By deployment. One seniority-blind heavy-review path starts above the trap edge. *(b)* Under the seniority-contingent rule (solid) and copilot alone (dashed) for three verification profiles $p(h) = h^{\phi_p}$; dots mark the date at which the rule begins to delegate. The leading parametrization is used. Panel (a): copilot only $(0, \cdot, 0.2)$, seniority-blind thin review $(0.85, 0.2, 0.2)$, seniority-blind heavy review $(0.85, 0.9, 0.2)$, from h_0 and from $h = 0.25$, competitive equilibrium $(1, 0, \cdot)$ of Proposition 4, seniority-contingent rule as in Figure 5. Panel (b): the rule as in Figure 5 and copilot $(0, \cdot, 0.2)$, with $\phi_p \in \{0.5, 1, 2\}$.

and the worker's value $W(h)$ solve

$$(\rho+s)J(h) = Y(h;u) - w(h) + J'(h)\big[\Phi(h;u) - \delta h\big], \qquad \rho W(h) = w(h) + W'(h)\big[\Phi(h;u) - \delta h\big]. \tag{17}$$

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

$$w(h) = Y\big(h; \bar{u}(h)\big), \tag{18}$$

current output at the deployment $\bar{u}(\cdot)$ that firms play. Wage growth is then

$$\frac{\dot{w}}{w} = \frac{Y_h(h;\bar{u})}{Y(h;\bar{u})}\big(\Phi(h;\bar{u}) - \delta h\big), \tag{19}$$

proportional to net learning. The *level* of pay includes the throughput gain from autonomy; the *slope* is governed by Φ. Entry pay can therefore rise while the trajectory flattens, and under mislearning the trajectory turns down. Equation (19) holds exactly within any range on which $\bar{u}$ is constant, and by the envelope theorem wherever $\bar{u}(h)$ is an interior maximizer

of Y; where a seniority-contingent program switches deployment, the pay path has a kink or a step, which Appendix B.2 signs.

6.2 Who chooses the deployment

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. The benchmark is the planner, who values output and internalizes all future human capital, together with any cross-firm spillover $\nu \geq 0$, at the social discount rate ρ:

$$\rho\, V_S(h) = \max_u \Big\{ Y(h;u) + \big(V_S'(h) + \nu\big)\big[\Phi(h;u) - \delta h\big] \Big\}, \tag{20}$$

with social shadow value $\mu_S \equiv V_S'$. Its maximizer $u^S(h)$ is the *social optimum*.

Because the results turn on what is contracted and when, we state the timing. At each instant a worker of type h in a match experiences: (t1) *pricing*, a competitive spot market values her current type at $w(h)$; (t2) *deployment*, the matched firm chooses u; (t3) *production and accumulation*, $Y(h;u)$ is produced, $w(h)$ paid, and h evolves by (5); (t4) *separation*, with hazard s the match ends and she is re-priced at her new type elsewhere.

A learning-poor deployment lowers the worker's future type, for which she is compensated later, at the future spot price, at whichever firm she is then employed. Even a fully observed deployment therefore cannot be priced by the current employer: the firm cannot capture a return the worker collects elsewhere, and the worker cannot bill the current employer for a portable gain. What disciplines deployment is the ability to *contract and commit* on it, not observability. 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, let alone whether the check taught anything. "Non-contractible" means that the learning-relevant intensity of human engagement cannot be verified or committed, not that AI adoption is hidden.[5] We consider three cases.

Case 1: Competitive market, non-contractible deployment. Firms are price-takers under free entry, so $J \equiv 0$ and $\mu_F = 0$; deployment cannot be written into an enforceable contract; there is no directed search or relational commitment on the training environment; and wages are spot prices for the worker's current type.

Proposition 4 (The competitive equilibrium selects myopic deployment)**.** *In Case 1 the firm maximizes current output, and the unique equilibrium deployment is the myopic corner*

[5] A regime in which review intensity is unobservable is equivalent, for our purposes, to the observable but non-contractible regime.

$\bar{u}$: $a = 1$, $r = 0$, *with the copilot margin vacuous and normalized to* $m = 0$. *Capable tasks then contribute no learning, and no mislearning, so* $\Phi(h; \bar{u}) = \alpha \bar{\mathcal{L}}_\theta(h)$ *and entry learning is positive.*

(i) If $\theta < h_0^*$ *and* $\alpha \bar{\mathcal{L}}_\theta(h) > \delta h$ *for all* $h \in [0, \theta)$, *every entrant outgrows the frontier and converges to* h_0^*: *the loss is transitional.*

(ii) If $\theta > h_0^*$, *every worker who starts below* h_0^* *converges to a steady state strictly below* h_0^*, *and the lowest such steady state is decreasing in* θ: *the more capable the AI, the deeper the collapse.*

The externality bites precisely when AI is capable enough to matter; Figure 6(a) illustrates the permanent case, in which the competitive path barely rises above its entry level. Endogenizing wages makes the externality starker, not weaker: absent contractibility, competition for general human capital drives deployment to the configuration that maximizes current output and removes review, the configuration that hollows out learning.

The starkness of the corner reflects the baseline assumption that review has no current-output value. Suppose instead that reviewing autonomous output also raises its quality by catching errors, adding $a\, Q_\theta(h)\, B(r)$ to (4) with $B' > 0$, $B(0) = 0$, and $B - C_r$ concave. Then even the myopic firm reviews, at the output-quality optimum r^Y solving $B'(r^Y) = C_r'(r^Y)$, and autonomy still corners at $a = 1$. The planner's review condition adds the learning value of review, $B'(r^S) - C_r'(r^S) = -(\mu_S + \nu)\, \alpha\, v_R\, \mathcal{L}_\theta / Q_\theta$. Wherever review teaches ($v_R > 0$), this gives $r^S > r^Y$: the market under-reviews. Where it miseducates ($v_R < 0$, raw juniors), it gives $r^S < r^Y$: the market's quality-driven review of juniors' delegated output is itself excessive, and the planner would rather withhold autonomy. What the market fails to price is specifically the *learning* value of review, positive or negative, which is portable and non-contractible; its current-output value, which the firm internalizes, neither causes nor closes the failure.

Case 2: Firm rents, non-contractible deployment. With search frictions the firm retains a share of match rents, so it partially internalizes future skill, discounted at the retention-adjusted rate $\rho + s$:

$$(\rho + s)\, V_F(h) = \max_u \Big\{ Y(h; u) - w(h) + V_F'(h) \big[\Phi(h; u) - \delta h \big] \Big\}. \tag{21}$$

Assumption 2 (Frictional wage)**.** *A worker of type* h *can relocate to a competitive firm, which prices general skill at* $Y(h; \bar{u}(h))$, *and meets such an offer at Poisson rate* λ. *The*

current firm, holding within-match bargaining power, pays this outside option:

$$w(h) = \sigma\, Y\big(h; \bar{u}(h)\big), \qquad \sigma \equiv \frac{\lambda}{\rho + \lambda} \in (0,1), \qquad w'(h) = \sigma\, Y_h\big(h; \bar{u}(h)\big). \tag{22}$$

The wage depends on type, not on the current firm's deployment, so $\partial w/\partial u = 0$, and the pass-through σ is the equilibrium markdown: as $\lambda \to \infty$, $\sigma \to 1$ and Case 1 obtains; with frictions the firm retains the complement of the return to general skill, the wage-compression mechanism of Acemoglu and Pischke (1998, 1999). In the linear case $\pi = a\,\bar{\pi}(h,\theta)$ with linear costs, the deployment problem is multilinear in u and the optimum is bang-bang; the corner of Proposition 4 is this case. To characterize interior optima we impose:

Assumption 3 (Regular deployment)**.** *$\pi(a; h, \theta)$ is C^2 and strictly concave in a, and C_r, C_m are C^2 and strictly convex with $C_r'(0) = C_m'(0) = 0$.*

The deployment objective $Y + \mu(\Phi - \delta h)$ is then strictly concave in each control, and an interior optimum satisfies

$$Y_a + \mu\,\Phi_a = 0, \qquad Y_r + \mu\,\Phi_r = 0, \qquad Y_m + \mu\,\Phi_m = 0, \tag{23}$$

with $\mu = \mu_F$ for the firm and $\mu = \mu_S + \nu$ for the planner, where $\Phi_a = \alpha \mathcal{L}_\theta (1 + mqp)(\psi - 1)$, $\Phi_r = a\,\alpha\, v_R\, \mathcal{L}_\theta$, $\Phi_m = (1-a)\,\alpha\, q\, p\, \mathcal{L}_\theta$, $Y_a = \pi_a - [C_r(r) - C_m(m)]Q_\theta$, $Y_r = -C_r'(r)\, a\, Q_\theta$, and $Y_m = -C_m'(m)(1-a)Q_\theta$.

Lemma 3 (Undervaluation of human capital)**.** *In Case 2, if $Y_h > 0$, then along any common allocation $\mu_F(h) < \mu_S(h) + \nu$. The wedge $\mu_S + \nu - \mu_F$ is increasing in the pass-through σ, in the hazard s, and in the spillover ν.*

Each source is a distinct wedge between private and social returns to the same future skill: a fraction σ of any future gain accrues to the worker, and is carried to her next employer; the firm discounts at $\rho + s$ because it may lose her before the return arrives, while the skill survives the match; and the spillover is internalized by no single employer. At $\sigma = 1$ the firm retains nothing and the wedge is maximal; as $\sigma, s, \nu \to 0$ the firm becomes residual claimant on the full return and the gap closes. This wedge is the engine of the externality.

Proposition 5 (Over-autonomy, under-review)**.** *Let Assumptions 1 and 3 hold and suppose $0 < \mu_F < \mu_S + \nu$. Relative to the social optimum, the privately optimal deployment has:*

(i) too much autonomy *wherever reviewing teaches less than doing, $\psi < 1$, at the deployments chosen for every $\mu \in [\mu_F, \mu_S + \nu]$, in particular for every raw junior, for whom $\psi < 0$;*

(ii) too little review *wherever review teaches, $p(h) > p_\chi$; for raw juniors both firm and planner set $r = 0$, so the distortion is entirely in autonomy;*

(iii) too little assistance *wherever capable tasks are retained and $p(h) > 0$;*

and the distortions are larger the higher the hazard s, the pass-through σ, and the spillover ν, maximal in the competitive limit $\mu_F \to 0$.

Remark 1 (The planner's rule is a seniority threshold)**.** *In the linear case, for given (r, m) the deployment objective is linear in a with slope $\bar{\pi} - [C_r(r) - C_m(m)]Q_\theta + \mu\,\alpha\,\mathcal{L}_\theta(1+mqp)(\psi - 1)$. Where the net throughput gain from delegation is positive, the planner ($\mu = \mu_S + \nu$) therefore delegates wherever ψ exceeds a cutoff below one, that is, at seniorities below h_ψ, trading some learning for output, and the cutoff converges to one, and the threshold to h_ψ, as the social value of skill grows relative to that gain. The competitive firm ($\mu = 0$) delegates at every seniority. The seniority-contingent rule of Definition 3 is thus the learning-maximizing limit of the planner's policy, and the market's policy is its opposite.*

Case 3: Contractible deployment. Suppose 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 joint match surplus.

Corollary 2 (Contractibility restores efficiency)**.** *If deployment is contractible and the market is competitive in (w, u) packages, the equilibrium deployment maximizes joint match surplus, the wedge of Lemma 3 closes ($\mu_F = \mu_S$), and the distortions of Proposition 5 vanish. Absent a cross-match spillover ($\nu = 0$) the allocation is the planner's. The worker finances learning-intensive deployment by accepting lower current wages, recovering more than the concession through the higher wage path that higher h commands.*

This is the Becker logic for general training, extended to how AI is deployed, and it clarifies why the worker cannot dissolve the externality on her own. The learning-relevant input is not private study effort she can add at the margin but the configuration u, which the firm chooses at (t2); and even a worker willing to accept lower pay cannot buy a learning-intensive configuration through a spot contract, because whether review *taught* is not verifiable to any third party. The Becker discount requires the committable package of Case 3. Appendix B.3 shows that workers sorting toward developmental firms on a noisy public signal, the Becker channel operating imperfectly, does not overturn the myopic corner unless the signal is precise.

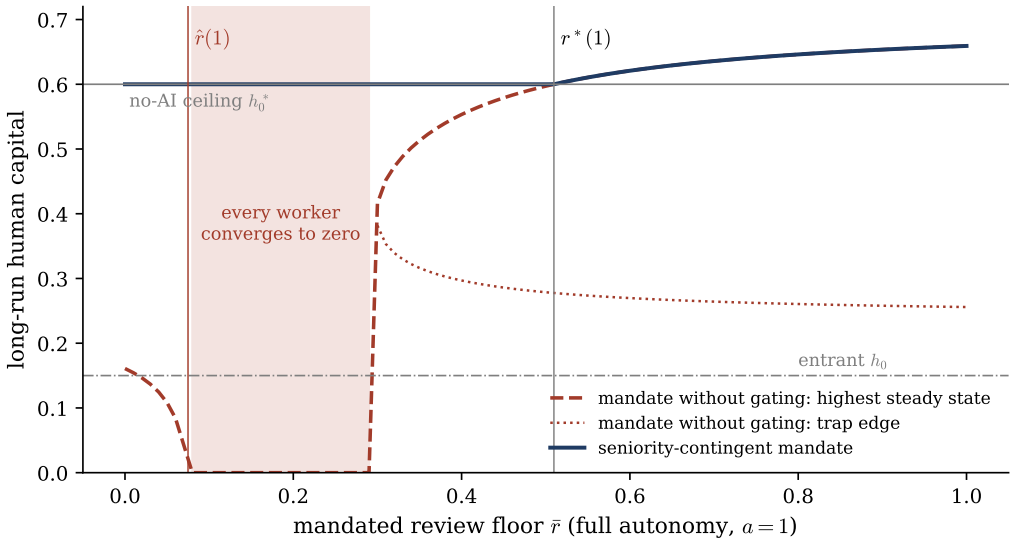


Figure 7: Review mandates under full autonomy: the highest steady state and the trap edge when a uniform review floor $\bar{r}$ is imposed on the competitive deployment, and the steady state under a seniority-contingent mandate; in the shaded range no positive steady state exists. The thresholds $\hat{r}(1)$ and $r^*(1)$ are those of Corollary 3. The leading parametrization is used. The mandate without gating sets $(a, r, m) = (1, \bar{r}, 0)$. The seniority-contingent mandate assigns $(0, \cdot, 0)$ below h_ψ and $(1, \bar{r}, 0)$ above.

6.3 Review mandates

The externality invites a regulatory response: require that autonomous output be reviewed. The learning technology shows why such a mandate can backfire if it is not conditioned on seniority.

Corollary 3 (Review mandates without gating)**.** *In Case 1, suppose a regulator imposes a uniform review floor $\bar{r} \in (0, 1]$ on delegated output and that the throughput gain still exceeds the review cost, $\pi_a > C_r(\bar{r})Q_\theta$, so that firms continue to delegate every capable task ($a = 1$). Then:*

(i) entry learning is negative, and entrants below the trap edge converge to $h = 0$, if and only if $\bar{r} > \hat{r}(1) = \bar{\mathcal{L}}_\theta(0)/[(1-q)\chi\mathcal{L}_\theta(0)]$;

(ii) under the conditions of Proposition 2(iii), an incumbent at the no-AI ceiling gains in the long run if and only if $\bar{r} > r^*(1) = 1/v_R\big(q, e, p(h_0^*)\big)$;

(iii) a seniority-contingent mandate, which withholds autonomy below h_ψ and requires review above it, never lowers any worker's learning below its no-AI level.

Figure 7 illustrates. The trap appears at a low review floor, well before the floor is high enough to lift incumbents above the no-AI ceiling. In between lies a range in which no

positive steady state exists at all: the floor is too thin to teach able reviewers much but thick enough to miseducate juniors, and every worker's expertise decays. The seniority-contingent mandate never falls below the no-AI ceiling and rises above it once the floor exceeds the incumbents' watershed. Mandating review onto the market's myopic corner can thus convert hollowed-out junior work into miseducated junior work. The appropriate instrument conditions autonomy on demonstrated verification ability, and requires review only of those who can learn from it.

7 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 what a manager observes, how the primitives are measured, and the rule they imply.

Most primitives sit in existing telemetry or cheap evaluations. Verification ability, the pivotal input because it places each worker relative to the thresholds, is recovered by seeded-error audits. The learning values are not in telemetry, but the decomposition (1) identifies them. Matched workers at a common seniority either review AI-produced solutions to a task set or perform the same tasks unaided, holding the task set and time on task fixed. In the review arm the experimenter seeds known errors and varies whether the reasoning is exposed, so the numbers of confirmed, caught, and missed components are observed for each worker. Learning is read not from performance during the episode but from subsequent independent performance on held-out tasks, so the contemporaneous boost from AI does not contaminate it. Regressing the learning gain, relative to the doing arm, on these counts identifies $\omega(e)$, ξ, and χ. Variation in e separates $\bar{\omega}$ from γ and repeating the design across seniority traces the thresholds.

This measurability sits in no tension with the non-contractibility that drives the externality. Telemetry is *internally observable*: a noisy signal the firm reads for its own deployment decisions, which need not be gaming-proof because the firm has no incentive to fool itself. Contractibility requires that a third party, a court enforcing a contract or a regulator disbursing a transfer, can verify the variable. Edit distance, check time, and detection rates fail that bar where it bites, because they can be inflated by box-checking that logs the act of review without the learning it is meant to proxy. The object on which the externality turns is whether review *taught*, not whether it occurred. As review becomes more verifiable to outsiders the setting moves toward Case 3 and the wedge shrinks, so the size of the problem and the feasibility of contracting on review co-move rather than conflict.

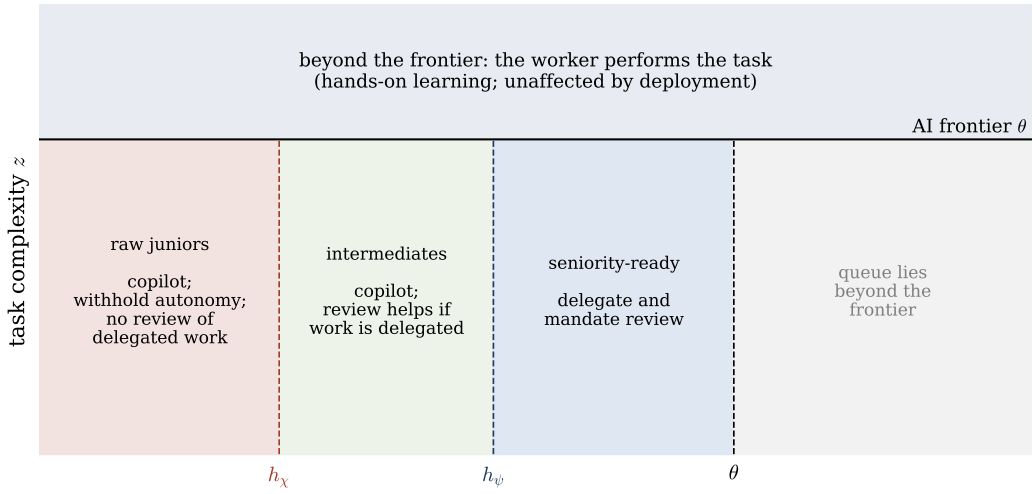


Figure 8: The deployment rule by task–seniority cell. The AI frontier θ separates tasks the worker performs (top) from capable tasks. Among capable tasks, the mislearning threshold h_χ and the scaffolding threshold h_ψ separate raw juniors, intermediates, and seniority-ready workers. Thresholds drawn at the same parameter values as in Figure 5.

The rule. The policy is a rule over task–worker cells, not a firm-wide automation target (Figure 8), built on the seniority bands of Definition 2:

(1) *Beyond the frontier* ($z > \theta$): the worker performs the task, and deployment does not touch her learning there. As the frontier moves, protect these tasks from premature automation.

(2) *Capable task, raw junior* ($p(h) < p_\chi$): copilot, with autonomy withheld. If output pressure forces delegation, do not route the review of delegated output to the junior, assign it to someone who can verify.

(3) *Capable task, intermediate* ($p_\chi \leq p(h) < p_\psi$): copilot. Where work is delegated anyway, pair it with review, which now teaches, and track the worker's verification ability toward the threshold.

(4) *Capable task, seniority-ready* ($p(h) \geq p_\psi$): delegate and mandate review, preferably on transparent deployments.

A uniform "automate $X\%$" target is the wrong instrument: it over-automates juniors and under-uses the scaffolding region. A uniform "keep a human in the loop" requirement is wrong for the opposite reason: it routes autonomous output to raw juniors for approval, the configuration that traps them.

The framework names three failures modes. The *hollow job* combines high autonomy with

thin review: output is high and learning collapses, the market's own corner. The *junior trap* arises when autonomous output is routed for review to workers who cannot verify: juniors absorb errors and fall into a low-expertise basin. *Idle scaffolding* arises when skills churn too fast for workers to reach verification competence ($\delta \geq \delta^{**}$): the review channel never operates and mandating more review cannot help, so the response is to act on durability (continuous retraining, deeper hands-on exposure) rather than on review.

Four implications follow. First, *treat AI configuration as a talent-development decision, not only a productivity decision.* The intensities that maximize current output are not those that develop the junior pipeline, and a firm that optimizes only the first (Proposition 4) drives its workforce toward a low-expertise steady state; configuration should be owned jointly by operations and talent functions. Second, *gate autonomy on measured verification ability, cell by cell.* The rule's safety does not depend on the size of the review premium, so it is the robust default. Third, *manage review as an input and route it to those who can learn from it.* Review intensity, review quality, and who reviews should be tracked like any other operating metric; transparent deployments raise the value of review and lower both thresholds (Lemma 1(ii), Proposition 1(iv)) and human mentorship and AI review compete for the same congestible channel, so thinning senior ranks while raising nominal review capacity can leave juniors without anyone who can verify on their behalf. Fourth, *anticipate that competition will under-provide development.* Firms that want the high-expertise steady state must substitute internal commitment devices for the missing contract: explicit development tracks, retention-linked investment, and rotation through review-intensive roles. Because workers would accept lower pay for a credibly developmental deployment (Corollary 2), offering one is also a way to attract talent at lower wage cost.

8 Conclusion

This paper builds a theory of early-career skill formation around a channel the economics of AI has largely ignored: workers can learn by reviewing what AI does on its own. Deriving that channel from what happens when a worker checks AI output, confirming correct steps, catching errors, and absorbing the ones she misses, turns the worry that AI removes the early rungs of the career ladder into a sharper claim. Delegating capable work removes rungs for workers who cannot yet verify, and asking them to review it corrupts the rungs that remain. For workers who can verify, the same delegation paired with review can build skill faster than doing the work. A deployment that treats all workers alike must choose between hollowing out entrants' careers and trapping them; a deployment that conditions autonomy on verification ability avoids both, never lowers learning below its no-AI level, and

raises the expertise ceiling when the review premium exists. A competitive market will not choose it, because the skill is portable and review cannot be contracted on.

The policy implications concern how AI is configured, not whether it is allowed. The remedy for hollowed-out junior work is to gate autonomy by demonstrated verification ability and to require review only of those who can learn from it. Capping adoption sacrifices the expertise that review-paired autonomy can build for seasoned workers; mandating review without gating can do worse, converting hollowed-out junior work into miseducated junior work. Because the gains and losses from deployment fall on different parties, closing the gap also requires giving firms a larger stake in development they cannot otherwise capture: contractible development tracks, apprenticeship-style commitments, and retention-linked investment.

Among jobs exposed to AI, what happens to early-career workers depends on how automation and human review are combined and on whether workers can check the AI's work, not on exposure 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. and N. Thompson (2025). Expertise. *Journal of the European Economic Association* 23(4), 1203–1271.

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.

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

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.

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.

Bastani, H., O. Bastani, A. Sungu, H. Ge, Ö. Kabakçı, and R. Mariman (2025). Generative AI Without Guardrails Can Harm Learning: Evidence from High School Mathematics. *Proceedings of the National Academy of Sciences* 122(26), e2422633122.

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.

Booth, J. L., K. E. Lange, K. R. Koedinger, and K. J. Newton (2013). Using Example Problems to Improve Student Learning in Algebra: Differentiating between Correct and Incorrect Examples. *Learning and Instruction* 25, 24–34.

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.

Contractor, Z. and G. Reyes (2026). Experimental Evidence on the Learning Impact of Generative AI. IZA Discussion Paper No. 18792.

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.

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.

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).

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.

Dickey, E., M. Mertzanidis, and A. Psomas (2026). Is Solving Better Than Evaluating GenAI Solutions? Working paper, arXiv:2607.27586.

Durkin, K. and B. Rittle-Johnson (2012). The Effectiveness of Using Incorrect Examples to Support Learning about Decimal Magnitude. *Learning and Instruction* 22(3), 206–214.

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.

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.

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.

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.

Kalyuga, S., P. Ayres, P. Chandler, and J. Sweller (2003). The Expertise Reversal Effect. *Educational Psychologist* 38(1), 23–31.

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.

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

Liu, G., B. Christian, T. Dumbalska, M. A. Bakker, and R. Dubey (2026). AI Assistance Reduces Persistence and Hurts Independent Performance. Working paper, arXiv:2604.04721.

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.

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.

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.

Shen, J. H. and A. Tamkin (2026). How AI Impacts Skill Formation. Working paper, arXiv:2601.20245.

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.

Wiles, E., L. Krayer, M. Abbadi, et al. (2026). Generative AI and the Temporary Upskilling of Knowledge Workers. *Nature Human Behaviour*, doi:10.1038/s41562-026-02566-6.

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

Proof of Lemma 1. (i) Collecting terms in (1), $v_R = pq\omega + p(1-q)\xi - (1-q)\chi + p(1-q)\chi = p\,\Sigma(q,e) - (1-q)\chi$, where $\Sigma = q\omega(e) + (1-q)(\xi+\chi)$ is positive whenever $\omega(e) > 0$, or $q < 1$ and $\xi + \chi > 0$; we exclude the degenerate case $\Sigma = 0$. Hence v_R is affine and strictly increasing in p and vanishes at $p_\chi = (1-q)\chi/\Sigma$. (ii) Only the first term depends on e, and $\omega'(e) = \bar{\omega}(1-\gamma)$. (iii) $\partial v_R/\partial q = p\,\omega - p\,\xi + (1-p)\chi = \chi - p(\xi + \chi - \omega)$, which is affine in p, equal to $\chi \geq 0$ at $p = 0$ and to $\omega - \xi$ at $p = 1$. If $\xi > \omega$ it is strictly decreasing in p and changes sign at $\bar{p} = \chi/(\xi + \chi - \omega) \in [0,1)$; if $\xi \leq \omega$ it is nonnegative at both endpoints and hence on $[0,1]$. (iv) By (i), $v_R \leq v_R(q,e,1) = q\omega(e) + (1-q)\xi$, a convex combination of $\omega(e)$ and ξ. □

Derivation of (7) and (8). Since $\Phi_0 = \alpha(\mathcal{L}_\theta + \bar{\mathcal{L}}_\theta)$, (3) gives $\Phi - \Phi_0 = \alpha\mathcal{L}_\theta\big[(1-a)(1+mqp) + arv_R - 1\big] = \alpha\mathcal{L}_\theta\big[mqp + a\big(rv_R - (1+mqp)\big)\big]$. Substituting $rv_R = \psi\,(1+mqp)$ from (6) yields (7). Differentiating (3) in a and in r yields (8). □

Proof of Proposition 1. For $h < \theta$, Assumption 1 and $\ell > 0$ on $[0,1)$ give $\mathcal{L}_\theta(h) > 0$. (i) By (8), $\partial\Phi/\partial r = \alpha a \mathcal{L}_\theta v_R$, whose sign for $a > 0$ is that of v_R, hence of $p(h) - p_\chi$ by Lemma 1(i). (ii) By (8) the sign of $\partial\Phi/\partial a$ is that of $\psi - 1$. Using Lemma 1(i), $\psi \geq 1 \iff rv_R \geq 1 + mqp \iff p\,(r\Sigma - mq) \geq 1 + r(1-q)\chi$. If $r\Sigma > mq$ this is $p \geq p_\psi$; if $r\Sigma \leq mq$ the left side is nonpositive and the right side positive, so $\psi < 1$ for every p. Since $1 + r(1-q)\chi > r(1-q)\chi$ and $0 < r\Sigma - mq \leq r\Sigma$, $p_\psi > r(1-q)\chi/(r\Sigma) = p_\chi$. (iii) On $\{h < \theta : p(h) > p_\psi\}$, $\psi > 1$ and $\mathcal{L}_\theta > 0$, so $\partial\Phi/\partial a > 0$. Because p is continuous and increasing, this set is non-empty if and only if $p_\psi < p(\theta)$, that is, if and only if $\psi > 1$ at $p = p(\theta)$, which is (10); when $r\Sigma \leq mq$ both the region and (10) fail. For necessity, (10) implies $v_R(q,e,p(\theta)) > (1 + mqp(\theta))/r \geq 1/r$, while $v_R(q,e,p(\theta)) \leq q\omega(e) + (1-q)\xi$ by Lemma 1(iv). (iv) Write $p_\psi = N/D$ with $N = 1 + r(1-q)\chi$ and $D = r\Sigma - mq > 0$. Then $\partial p_\psi/\partial r = [(1-q)\chi D - N\Sigma]/D^2 = -[(1-q)\chi mq + \Sigma]/D^2 < 0$. The threshold depends on e, $\bar{\omega}$, and ξ only through Σ, which is increasing in each, and D is increasing in Σ; it is increasing in m through D. For χ, $\partial N/\partial\chi = \partial D/\partial\chi = r(1-q)$, so $\partial p_\psi/\partial\chi = r(1-q)(D-N)/D^2$, positive if and only if $p_\psi < 1$. Finally, dividing numerator and denominator of p_χ by $1-q$ gives $p_\chi = \chi/[\xi + \chi + \omega(e)q/(1-q)]$, decreasing in q. □

Proof of Corollary 1. In (3), p enters only through the copilot term $\alpha(1-a)mqp\mathcal{L}_\theta$ and the review term $\alpha arv_R(q,e,p)\mathcal{L}_\theta$, with $\partial v_R/\partial p = \Sigma > 0$. For $h < \theta$, $\mathcal{L}_\theta > 0$, so Φ is strictly increasing in p when $ar > 0$ or $(1-a)m > 0$, and $p_1(h) > p_2(h)$ gives the first worker a strictly higher learning rate. The thresholds p_χ and p_ψ do not depend on the worker, so the seniorities at which they are reached are lower under the higher profile. The wage (18)

depends on h and not on p, so the two workers can share an entry wage, while by (19) wage growth is proportional to net learning. □

Proof of Lemma 2. G is continuous with $G(1) = \Phi(1; u(1)) - \delta < 0$. (i) If $G(0) > 0$, let h_1 be the smallest zero of G, which exists by the intermediate value theorem. Since $G > 0$ on $[0, h_1)$, paths starting there increase monotonically and converge to h_1. Zeros of G in $(0, 1]$ are the solutions of $\Phi(h; u(h))/h = \delta$; if this ratio is strictly decreasing it crosses δ once, so $G > 0$ below the unique zero and $G < 0$ above it, which is global asymptotic stability on $(0, 1]$. The ratio is decreasing if and only if $h\,\Phi' - \Phi < 0$, the elasticity condition. (ii) If $G(0) < 0$, the projected drift at $h = 0$ is zero, so $h = 0$ is absorbing. By continuity $h_{\text{trap}} > 0$, and $G < 0$ on $(0, h_{\text{trap}})$ by definition of h_{trap}. A path starting there decreases monotonically; it cannot converge to an interior point, which would have to be a zero of G, so it converges to 0. This also gives local stability of 0. If G is strictly increasing on $[0, \hat{h}]$ and strictly decreasing on $[\hat{h}, 1]$ with $G(\hat{h}) > 0$, it has exactly one zero on $(0, \hat{h})$, which is h_{trap}, and one on $(\hat{h}, 1)$, call it h^H, since $G(0) < 0$ and $G(1) < 0$. Paths starting in (h_{trap}, h^H) rise to h^H, and paths starting above h^H fall to it. □

Proof of Proposition 2. (i) From the derivation of (7), $\Phi - \Phi_0 = \alpha\mathcal{L}_\theta\big[(1-a)mqp + a(rv_R - 1)\big] = \alpha\mathcal{L}_\theta\big[p\big((1-a)mq + ar\Sigma\big) - a\big(1 + r(1-q)\chi\big)\big]$, using Lemma 1(i). For $h < \theta$, $\mathcal{L}_\theta > 0$ and the bracket is affine and strictly increasing in p, vanishing at p_c. At $p = p_\psi$, $\psi = 1$ and (7) gives $\Phi - \Phi_0 = \alpha\mathcal{L}_\theta mqp \geq 0$, so $p_c \leq p_\psi$; if $r\Sigma \leq mq$, $p_\psi = \infty$ and the inequality is trivial. (ii) Since $p(0) = 0$, $v_R(q, e, 0) = -(1-q)\chi$ and the copilot gain vanishes at entry, so $\Phi(0; u)$ is given by (11). Because $\theta > 0$ and $\ell > 0$, $\mathcal{L}_\theta(0) > 0$, and $\Phi(0; u) < 0$ if and only if $r > \hat{r}(a)$. The comparative statics follow from (13): its numerator is decreasing and its denominator increasing in a, its denominator is increasing in χ and decreasing in q, and e, $\bar{\omega}$, ξ do not appear. At a transversal up-crossing, $G(h_{\text{trap}}) = 0$ and $G_h(h_{\text{trap}}) > 0$, while $\partial G/\partial r = \alpha a\mathcal{L}_\theta(h_{\text{trap}})v_R(q, e, p(h_{\text{trap}})) < 0$ because $h_{\text{trap}} < \theta$ and $p(h_{\text{trap}}) < p_\chi$. By the implicit function theorem, $\mathrm{d}h_{\text{trap}}/\mathrm{d}r = -G_r/G_h > 0$. (iii) Since $\Phi_0(h_0^*) = \delta h_0^*$, the drift at h_0^* is $G(h_0^*) = \Phi(h_0^*; u) - \Phi_0(h_0^*) = \alpha\mathcal{L}_\theta(h_0^*)\big[(1-a)mqp^* + a(rv_R(q, e, p^*) - 1)\big]$, with $p^* = p(h_0^*)$ and $\mathcal{L}_\theta(h_0^*) > 0$ because $h_0^* < \theta$. Since $v_R(q, e, p^*) > 0$, the bracket is strictly increasing in r and vanishes at r^*, so $\operatorname{sign} G(h_0^*) = \operatorname{sign}(r - r^*)$. If $G(h_0^*) > 0$, the path from h_0^* increases monotonically and converges to the smallest zero of G above h_0^*, which exists because $G(1) < 0$; if $G(h_0^*) < 0$, it decreases to the largest zero below h_0^*, or to 0. In both cases the limit lies strictly on the side given by the sign. If $\Phi_0(h)/h$ is strictly decreasing, h_0^* solves $\Phi_0(h)/h = \delta$ and is decreasing in δ; then p^* falls with δ, raising the numerator of (14) and lowering its positive denominator, so r^* increases. (iv) The drift at the entrant's initial level is $G(h_0; r) = \Phi(h_0; (a, 0, m)) - \delta h_0 + r\,\alpha a\mathcal{L}_\theta(h_0)v_R(q, e, p(h_0))$, affine in r with

a strictly negative slope, because $\mathcal{L}_\theta(h_0) > 0$ for $h_0 < \theta$ and $v_R < 0$ for $p(h_0) < p_\chi$. It is therefore negative for $r > r_e$ and positive for $r < r_e$. Since paths are monotone, the entrant's long-run level lies below h_0 exactly when $G(h_0; r) < 0$, and above it when $G(h_0; r) > 0$. The numerator of (15) does not depend on $\bar{\omega}$, ξ, e, or χ, and $|v_R(q, e, p(h_0))|$ is decreasing in $\bar{\omega}$, ξ, and e and increasing in χ by Lemma 1, since $\partial v_R/\partial\chi = -(1-q)(1-p) < 0$; this gives the comparative statics of r_e. The numerator of (14) does not depend on these primitives either, and $v_R(q, e, p(h_0^*)) > 0$ is increasing in $\bar{\omega}$, ξ, and e and decreasing in χ, which gives those of r^*. If $r_e < r^*$, any $r > r^*$ also exceeds r_e, so the entrant ends below h_0 whenever incumbents end above h_0^*. □

Proof of Proposition 3. (i) From (3), $\Phi(h; u^D) - \Phi^C(h) = \alpha\mathcal{L}_\theta\, a\big[rv_R - (1+mqp)\big] = \alpha\mathcal{L}_\theta\, a\,(1+mqp)(\psi - 1)$, which is nonnegative if and only if $\psi \geq 1$, that is, by Proposition 1(ii), if and only if $p(h) \geq p_\psi$; when the scaffolding region is empty it is negative for every $h < \theta$, and for $h \geq \theta$ both sides equal Φ_0. Hence $\Phi^g = \max\{\Phi^C, \Phi(\cdot; u^D)\}$, with strict dominance of Φ^C on the scaffolding region. By (7) with $a = 0$, $\Phi^C - \Phi_0 = \alpha\mathcal{L}_\theta mqp \geq 0$, strictly when $h < \theta$, $p(h) > 0$, and $m > 0$. Both branches are continuous and coincide at $p(h) = p_\psi$, where $\psi = 1$, so Φ^g is continuous, and as the maximum of two Lipschitz functions it is Lipschitz. A raw junior has $p(h) < p_\chi < p_\psi$ and is assigned u^C. (ii) At $h = 0 < h_\psi$ the rule assigns u^C, and $p(0) = 0$ gives $\Phi^g(0) = \alpha[\bar{\mathcal{L}}_\theta(0) + \mathcal{L}_\theta(0)] = \Phi_0(0) > 0$. For scalar autonomous dynamics with Lipschitz drifts $G_1 \geq G_2$ on $[0, 1]$, the comparison theorem gives $h_1(t) \geq h_2(t)$ for all t whenever $h_1(0) \geq h_2(0)$, and the projection onto $[0, 1]$ preserves this order. Applying it with $G_1 = \Phi^g - \delta h$ and $G_2 \in \{\Phi_0 - \delta h,\ \Phi^C - \delta h,\ \Phi(\cdot; u^D) - \delta h\}$ gives the claim. (iii) Paths are monotone and bounded, so they converge, and the order in (ii) carries over to the limits. Uniqueness and global stability under single crossing follow from Lemma 2(i), since $\Phi^g(0) > 0$. (iv) By Lemma 2(i), $G^C \equiv \Phi^C - \delta h$ is positive below h_C^* and negative above it. The rule coincides with u^C on $[0, h_\psi)$. If $\delta < \delta^{**}$, then $G^C(h_\psi) > 0$, so $h_C^* > h_\psi$ and $G^C > 0$ on $[h_0, h_\psi]$; the entrant reaches h_ψ in finite time. Thereafter $G^g \geq G^C > 0$ on $[h_\psi, h_C^*)$, and $G^g(h_C^*) \geq 0$, with strict inequality when $h_C^* < \theta$ because h_C^* then lies in the scaffolding region, so the path converges to a zero of G^g at or above h_C^*, strictly above it when $h_C^* < \theta$. If $\delta \geq \delta^{**}$, then $G^C(h_\psi) \leq 0$, so $h_C^* \leq h_\psi$; starting below h_C^*, the entrant rises toward h_C^* under $G^g = G^C$ without entering the delegation phase. (v) Fix h. $\Phi(h; u^D) = \Phi^C(h) + \alpha\mathcal{L}_\theta a\big[rv_R - (1+mqp)\big]$ is affine in r with slope $\alpha\mathcal{L}_\theta a v_R$. If $v_R \geq 0$ it is nondecreasing in r, so the maximum is; if $v_R < 0$, $\Phi(h; u^D) < \Phi^C(h)$ for every $r \geq 0$ and the maximum equals $\Phi^C(h)$, independent of r. Likewise $\Phi(h; u^D) - \Phi^C(h) = a\,\alpha\mathcal{L}_\theta(1 + mqp)(\psi - 1)$ is affine in a, with ψ independent of a, so the maximum is nondecreasing in a. The comparison theorem transfers pointwise monotonicity of the drift to paths and limits. □

Proof of Proposition 4. In Case 1, $J \equiv 0$ and $\mu_F = 0$, so the firm maximizes $Y(h;u)$. From (4), $Y_r = -C_r'(r)aQ_\theta < 0$ for $r > 0$ and $aQ_\theta > 0$, so $r = 0$; $Y_m = -C_m'(m)(1-a)Q_\theta < 0$ for $m > 0$ and $a < 1$, so $m = 0$; and with $r = m = 0$, $Y_a = \pi_a > 0$, so $a = 1$, at which m affects neither output nor learning and is normalized to zero. At $\bar{u}$ every capable task is delegated and unreviewed, so it contributes neither learning nor mislearning, and $\Phi(h;\bar{u}) = \alpha\bar{\mathcal{L}}_\theta(h)$, with $\Phi(0;\bar{u}) = \alpha\bar{\mathcal{L}}_\theta(0) > 0$ because $\theta < 1$. Write $G^{\bar{u}}(h) = \alpha\bar{\mathcal{L}}_\theta(h) - \delta h$ and $G_0(h) = \Phi_0(h) - \delta h$, and recall that $G^{\bar{u}} = G_0$ for $h \geq \theta$. (i) If $G^{\bar{u}} > 0$ on $[0,\theta)$, a path from $h_0 < \theta$ rises past θ, after which it follows the no-AI drift, which (under single crossing of Φ_0, as in the leading parametrization) is positive on $[\theta, h_0^*)$; it converges to h_0^*. (ii) If $\theta > h_0^*$, then for $h < \theta$, $\mathcal{L}_\theta(h) > 0$ and $G^{\bar{u}}(h) < G_0(h)$; in particular $G^{\bar{u}}(h_0^*) < G_0(h_0^*) = 0$. A path starting below h_0^* cannot cross it and converges to a zero of $G^{\bar{u}}$ in $(0, h_0^*)$, which exists because $G^{\bar{u}}(0) > 0$. The smallest such zero is the first crossing of $\alpha\bar{\mathcal{L}}_\theta(h)$ and δh; since $\bar{\mathcal{L}}_\theta(h) = \int_{\{z>\theta\}} \ell \, \mathrm{d}F$ is nonincreasing in θ for every h, raising θ lowers $G^{\bar{u}}$ pointwise and moves this crossing down. □

Proof of Lemma 3. Fix a common allocation $\bar{u}(\cdot)$ and its convergent path h_t, along which $\Phi_h - \delta < \rho$, as holds near a stable steady state. Differentiating (21) in h, using $\partial w/\partial u = 0$ and the envelope property, and writing $\dot{\mu} = \mu'(h)\dot{h}$, the firm's costate satisfies $(\rho + s)\mu_F - \dot{\mu}_F = (1-\sigma)Y_h + \mu_F(\Phi_h - \delta)$, since $w' = \sigma Y_h$ by (22). Its bounded solution is

$$\mu_{F,t} = \int_t^\infty e^{-\int_t^\tau [\rho + s - (\Phi_h - \delta)]\,\mathrm{d}v} (1-\sigma)\, Y_h \, \mathrm{d}\tau.$$

Differentiating (20) likewise and writing $M \equiv \mu_S + \nu$, $\rho M - \dot{M} = Y_h + \rho\nu + M(\Phi_h - \delta)$, so

$$M_t = \int_t^\infty e^{-\int_t^\tau [\rho - (\Phi_h - \delta)]\,\mathrm{d}v} \left(Y_h + \rho\nu\right) \mathrm{d}\tau.$$

The planner's integrand weakly exceeds the firm's, $Y_h + \rho\nu \geq (1-\sigma)Y_h$, strictly when $\sigma Y_h > 0$, as in Case 2, where $\sigma \in (0,1)$ and $Y_h > 0$; and its discount rate is weakly lower, strictly when $s > 0$. Hence $M_t > \mu_{F,t}$. The firm's costate is decreasing in s and in σ, while M does not depend on either and is increasing in ν, so the wedge $M - \mu_F$ is increasing in σ, s, and ν. □

Proof of Proposition 5. Write the deployment objective $\Pi(u;\mu) = Y(h;u) + \mu[\Phi(h;u) - \delta h]$ with the derivatives listed after (23). Because Φ is linear in each control and costs are additively separable, $\Pi_{aa} = \pi_{aa} < 0$, $\Pi_{rr} = -C_r''(r)aQ_\theta < 0$, $\Pi_{mm} = -C_m''(m)(1-a)Q_\theta < 0$, $\Pi_{rm} = 0$, $\Pi_{ar} = -C_r'(r)Q_\theta + \mu\alpha v_R\mathcal{L}_\theta$, and $\Pi_{am} = C_m'(m)Q_\theta - \mu\alpha qp\mathcal{L}_\theta$. Moreover r enters only through the factor a and m only through $1-a$, so $\Pi_r = a\,\Pi_{ar}$ and $\Pi_m = -(1-a)\,\Pi_{am}$.

Review. If $v_R \leq 0$, which holds exactly for raw juniors, $\Pi_r = a[-C_r'(r)Q_\theta + \mu\alpha v_R \mathcal{L}_\theta] < 0$ for every $r > 0$ and every $\mu \geq 0$, so firm and planner both set $r = 0$. If $v_R > 0$, the interior review condition $C_r'(r) = \mu\alpha v_R \mathcal{L}_\theta / Q_\theta$ has a solution increasing in μ by convexity of C_r, so $r(\mu_F) < r(\mu_S + \nu)$. *Assistance.* Likewise $C_m'(m) = \mu\alpha qp\mathcal{L}_\theta / Q_\theta$, increasing in μ when $p > 0$. *Autonomy.* At an interior optimum $\Pi_r = \Pi_m = 0$ with $a \in (0, 1)$ implies $\Pi_{ar} = \Pi_{am} = 0$ wherever r and m are interior; where r is at the corner $r = 0$, $\mathrm{d}r/\mathrm{d}\mu = 0$. In either case the indirect effects of μ on Π_a through r and m vanish, and $\mathrm{d}a/\mathrm{d}\mu = -\Phi_a/\Pi_{aa}$, which is negative wherever $\Phi_a = \alpha\mathcal{L}_\theta(1 + mqp)(\psi - 1) < 0$. If $\psi < 1$ at the chosen deployments for every $\mu \in [\mu_F, \mu_S + \nu]$, integrating gives $a(\mu_F) > a(\mu_S + \nu)$; for raw juniors $v_R < 0$ implies $\psi < 0$ for every r. In the bang-bang case the factorization implies that the sign of each control's marginal value moves in μ in the same directions, so the corners shift likewise. Finally, by Lemma 3 the gap $\mu_S + \nu - \mu_F$ is increasing in s, σ, and ν and approaches $\mu_S + \nu$ as $\mu_F \to 0$, so the divergence in deployments widens accordingly. □

Proof of Corollary 3. With $r = \bar{r}$ imposed and m irrelevant at $a = 1$, $Y_a = \pi_a - C_r(\bar{r})Q_\theta > 0$, so firms set $a = 1$. (i) is Proposition 2(ii) at $a = 1$, for which the copilot term in (13) vanishes. (ii) is Proposition 2(iii) at $a = 1$, for which the numerator of (14) equals one. (iii) is Proposition 3(i). □

B Extensions

B.1 Two Viability Conditions: Learning versus 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 that Garicano and Rayo (2026) analyze through an expertise leverage ratio. The second is also the margin on which juniors are hired. We show that the two conditions are logically independent and that, in autonomy, they move in opposite directions. At the entry margin define

$$\mathcal{V}_L(u) \equiv \frac{\Phi(h_0; u)}{\Phi_0(h_0)}, \qquad \mathcal{V}_F(u) \equiv \frac{Y(h_0; u) - w(h_0)}{\kappa}, \tag{24}$$

where $\kappa > 0$ is the per-period cost of sustaining the apprenticeship (training, supervision, vacancy) and $w(h_0)$ is the entrant's market wage under Assumption 2, so that $\partial w/\partial u = 0$. Learning is viable if $\mathcal{V}_L \geq 1$, so that AI does not reduce the entrant's learning below its no-AI level; the arrangement is financed if $\mathcal{V}_F \geq 1$. Under mislearning $\mathcal{V}_L$ can be negative.

Proposition 6 (Divergence of the two conditions)**.** *For an entrant with $h_0 < \theta$, $\partial \mathcal{V}_F/\partial a = Y_a/\kappa$ and $\operatorname{sign}(\partial \mathcal{V}_L/\partial a) = \operatorname{sign}(\psi(h_0) - 1)$. Wherever $Y_a > 0$ and $\psi(h_0) < 1$, raising autonomy moves the arrangement toward the financed column and away from the learning-viable row, and all four configurations are attainable:*

	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*

Proof. Since $\partial w/\partial u = 0$, $\partial \mathcal{V}_F/\partial a = Y_a/\kappa$, with $Y_a = \pi_a - [C_r(r) - C_m(m)]Q_\theta$ by (4); at any interior optimum with $\mu > 0$ in the region $\psi < 1$, (23) gives $Y_a = -\mu \Phi_a > 0$. Since $\Phi_0(h_0)$ does not depend on a, $\partial \mathcal{V}_L/\partial a = \Phi_a(h_0)/\Phi_0(h_0)$, whose sign is that of $\psi(h_0) - 1$ by (8). The output primitives (π, w, κ) move $\mathcal{V}_F$ without affecting Φ, and $(r, m, e, \bar{\omega}, \xi, \chi)$ move $\mathcal{V}_L$ without affecting Y_a beyond the cost terms, so each cell is attainable. □

For a raw-junior entrant the review margin adds a further divergence: $\partial \mathcal{V}_L/\partial r$ has the sign of $v_R(q, e, p(h_0)) < 0$, while $\partial \mathcal{V}_F/\partial r = -C_r'(r) a Q_\theta/\kappa \leq 0$, so mandated review lowers both indices. The *level* of $\mathcal{V}_F$ depends on how wages are set. With frictions ($\sigma < 1$) the firm retains a surplus and the financed column is attainable. In the competitive 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 $\mathcal{V}_F < 1$ for every deployment: the apprenticeship is unfunded for exactly the hold-up reason of Proposition 4. The leverage ratio of Garicano and Rayo (2026) governs the financed column and can be comfortably satisfied by exactly the high-autonomy deployments that hollow out learning; the two conditions are complementary halves of a single requirement.

B.2 The Wage Path under the Seniority-Contingent Rule

Section 6.1 gives wage growth within a fixed regime. Under the seniority-contingent rule the deployment switches at h_ψ, and the pay path bends there.

Proposition 7 (The seniority-contingent wage path)**.** *Let the firm follow the rule of Definition 3, with copilot $u^C = (0, \cdot, m)$ below h_ψ and delegate-and-review $u^D = (a, r, m)$ above it, and let pay track the deployment played, $w(h) = \sigma Y(h; u^g(h))$ with $\sigma \in (0, 1]$. For an entrant with $h_0 < h_\psi < h_g^*$:*

(i) Within phases *the path is C^1 with $w'(h) = \sigma Y_h(h; u)$ and growth given by (19) at the phase's deployment. At every seniority both phases produce less current output than the myopic corner $(1, 0, \cdot)$; this shortfall is the development concession of Corollary 2.*

(ii) At the switch *pay jumps by*

$$\Delta w(h_\psi) = \sigma\Big[\pi(a; h_\psi, \theta) - \pi(0; h_\psi, \theta) - a\, Q_\theta(h_\psi)\big(C_r(r) - C_m(m)\big)\Big],$$

which is positive whenever the throughput gain from delegation exceeds the net change in per-task operating costs, and the slope changes by $\sigma[Y_h(h_\psi; u^D) - Y_h(h_\psi; u^C)]$. *When the jump is positive, the concession is larger in the copilot phase.*

Proof. (i) On each phase the deployment is constant, so $w'(h) = \sigma Y_h(h; u)$ and (19) applies. By (4), output is increasing in a at $r = m = 0$ and decreasing in r and m, so $(1, 0, \cdot)$ maximizes current output and both u^C and u^D produce less. (ii) $\Delta w = \sigma[Y(h_\psi; u^D) - Y(h_\psi; u^C)]$; substituting (4) and cancelling $y_0(h_\psi)$ and the common assistance cost on the $1 - a$ retained share gives the expression. The slope change follows from $w' = \sigma Y_h$ on each side. If $\Delta w > 0$, $Y(h; u^C) < Y(h; u^D)$ near h_ψ, so the copilot shortfall from the myopic corner is the larger. □

The development-oriented pay path is thus a flatter junior segment, an upward step at h_ψ, and a steeper senior segment, the mirror image of the myopic path, whose pay is high and flat from entry and stalls with the low steady state of Proposition 4. Empirically, the rule 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.3 Worker Sorting under a Noisy Development Signal

The competitive benchmark imposes spot wages and no directed search. The natural market correction is worker sorting: early-career workers value development and would accept lower pay at firms that provide it. We ask whether directed search overturns the myopic corner, 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. Signal precision indexes how much development is observable *ex ante*; Cases 1 and 3 are its two extremes.

Environment. A unit mass of risk-neutral early-career workers faces free entry of firms. (i) Each entering firm chooses a posture, develop (D) or not (N); D costs $c_R > 0$, the current cost of mandated review, and raises the worker's portable continuation value by $\Delta_D > 0$, while N costs and adds nothing. (ii) A public signal $\varsigma \in \{g, b\}$ realizes with precision $\zeta \in [\frac{1}{2}, 1]$, $\Pr(g \mid \mathsf{D}) = \Pr(b \mid \mathsf{N}) = \zeta$. (iii) Firms post a wage and open a vacancy at cost

c_V. (iv) Workers observe wage and signal, form the posterior $b(\varsigma) = \Pr(\mathsf{D} \mid \varsigma)$, and direct their search. (v) Matching and production follow. Posture is neither postable nor verifiable. With a fraction n_D of developers, Bayes' rule gives

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

A submarket is a pair (w, ς) with tightness τ, vacancies per worker. Matching is constant-returns, increasing, and strictly concave in each argument; the worker's job-finding rate $\eta(\tau)$ satisfies $\eta(0) = 0$, $\eta' > 0$, $\eta'' < 0$, $\eta(\infty) = 1$, and $\lim_{\tau\to\infty} \tau\eta'(\tau) = 0$, and the vacancy-filling rate $\eta(\tau)/\tau$ is strictly decreasing (urn-ball and Cobb–Douglas matching qualify). A match yields the firm flow output $y > 0$, net of AI operating cost, and the worker the wage plus, in expectation, $b(\varsigma)\Delta_D$.

Competitive-search equilibrium. Workers attain a common market value $\bar{U}$, so an active submarket satisfies

$$\eta\big(\tau(w,\varsigma)\big)\,\big[w + b(\varsigma)\Delta_D\big] = \bar{U}. \tag{25}$$

A firm in signal class ς chooses w and τ to maximize $\frac{\eta(\tau)}{\tau}(y-w)$ subject to (25); substituting w, its problem becomes

$$\mathcal{G}(\varsigma) \equiv \max_{\tau>0}\Big\{\frac{\eta(\tau)}{\tau}\,P(\varsigma) - \frac{\bar{U}}{\tau}\Big\}, \qquad P(\varsigma) \equiv y + b(\varsigma)\Delta_D, \tag{26}$$

the value of a post whose total match value is $P(\varsigma)$. Free entry pins $\bar{U}$. Posts that workers believe more likely to develop fill more easily at any wage, equivalently command a wage discount, and that discount is the only channel through which a firm is compensated for the portable return.

Lemma 4 (Within-submarket existence and uniqueness)**.** *For any post value $P > 0$ and market utility $\bar{U} \in (0, P)$, problem (26) has a unique maximizer $\tau^* \in (0, \infty)$, and its value $\mathcal{G}(P, \bar{U})$ is continuous, strictly increasing in P, and strictly decreasing in $\bar{U}$.*

Proof. Write the objective as $\Lambda(\tau) = [\eta(\tau)P - \bar{U}]/\tau$, so $\Lambda'(\tau) = N(\tau)/\tau^2$ with $N(\tau) = P[\tau\eta'(\tau) - \eta(\tau)] + \bar{U}$. Since $\frac{\mathrm{d}}{\mathrm{d}\tau}[\tau\eta' - \eta] = \tau\eta'' < 0$ and $\tau\eta' - \eta$ starts at zero, N is strictly decreasing, with $N(0^+) = \bar{U} > 0$ and $N(\infty) = \bar{U} - P < 0$ by the boundary condition. Hence N has a unique zero τ^*, which is the global maximizer. The envelope theorem gives $\partial\mathcal{G}/\partial P = \eta(\tau^*)/\tau^* > 0$ and $\partial\mathcal{G}/\partial\bar{U} = -1/\tau^* < 0$; continuity follows from the theorem of the maximum. □

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

Proof. By Lemma 4 the expected value of an entering vacancy is continuous and strictly decreasing in $\bar{U}$, exceeds c_V as $\bar{U} \to 0$, and falls below c_V as $\bar{U}$ approaches $\min_\varsigma P(\varsigma)$. The intermediate value theorem gives exactly one $\bar{U}$ at which it equals c_V. □

Posture choice. Taking beliefs as given, D lands in g with probability ζ and N with probability $1-\zeta$, so the gain from developing is

$$\Gamma(\zeta, n_D) = (2\zeta - 1)\big[\mathcal{G}(g) - \mathcal{G}(b)\big] - c_R, \tag{27}$$

evaluated at the beliefs prevailing when a fraction n_D develops. Both factors of the benefit vanish when the signal is uninformative ($\zeta = \frac{1}{2}$) and when the population is homogeneous ($n_D \in \{0, 1\}$, so $b(g) = b(b)$).

Proposition 8 (Sorting does not overturn the myopic corner unless the signal is precise)**.** *In the competitive-search economy with signal precision* $\zeta \in [\frac{1}{2}, 1]$*:*

(i) The corner always exists. $n_D = 0$ *is an equilibrium for every* ζ*: worker sorting never eliminates the externality.*

(ii) Full development never exists. $n_D = 1$ *is never an equilibrium.*

(iii) Threshold. *There is* $\zeta^* \in (\frac{1}{2}, 1]$*, possibly exceeding one, such that for* $\zeta < \zeta^*$ *the corner is the unique equilibrium, and for* $\zeta \geq \zeta^*$ *interior development equilibria* $n_D \in (0, 1)$ *exist alongside the corner. The threshold is increasing in* c_R *and decreasing in* Δ_D*.*

(iv) Poles. *As* $\zeta \to \frac{1}{2}$ *the corner is the unique equilibrium, reproducing Case 1; as* $\zeta \to 1$ *the development equilibrium implements the contractible allocation of Case 3.*

Proof. By Lemmas 4 and 5, Γ is well defined and continuous on $[\frac{1}{2}, 1] \times [0, 1]$. (i) At $n_D = 0$, $b(g) = b(b) = 0$, so $\mathcal{G}(g) = \mathcal{G}(b)$ and $\Gamma = -c_R < 0$: no firm deviates to D. (ii) At $n_D = 1$, $b(g) = b(b) = 1$ and again $\Gamma = -c_R < 0$: a hidden deviation to N saves c_R at no wage penalty. (iii) For $\zeta > \frac{1}{2}$ the posterior spread $b(g) - b(b)$ is zero at $n_D \in \{0, 1\}$ and positive in between, and $P(g) - P(b) = [b(g) - b(b)]\Delta_D$, so by Lemma 4 $\mathcal{G}(g) - \mathcal{G}(b)$ is positive on $(0, 1)$ and zero at the endpoints. Let $\Gamma^*(\zeta) = \max_{n_D} \Gamma(\zeta, n_D)$. At any interior n_D the benefit is zero at $\zeta = \frac{1}{2}$ and strictly increasing in ζ, so Γ^* is continuous and strictly increasing with $\Gamma^*(\frac{1}{2}) = -c_R$. Define ζ^* by $\Gamma^*(\zeta^*) = 0$ if $\Gamma^*(1) \geq 0$, and $\zeta^* > 1$ otherwise. For $\zeta < \zeta^*$, $\Gamma < 0$ for all n_D, so only the corner survives; for $\zeta \geq \zeta^*$, $\Gamma(\zeta, \cdot)$ is non-negative somewhere in the interior and equals $-c_R$ at the endpoints, so it has interior zeros, at which firms are indifferent. The benefit is decreasing in c_R and increasing in Δ_D, which moves ζ^* accordingly. (iv) At $\zeta = \frac{1}{2}$,

$\Gamma = -c_R$ everywhere. As $\zeta \to 1$, supporting beliefs give $b(g) \to 1$ and $b(b) \to 0$, so the wage discount approaches the full competitive-search value of Δ_D and the firm internalizes the portable return as if it had posted a binding development contract, which is Case 3. □

The externality is therefore not an artifact of the spot-wage benchmark: the myopic corner is an equilibrium for every signal precision, and sorting can at most add a high-development equilibrium when the signal is precise. The relevant precision is low precisely where the mechanism bites. The signal workers can read *ex ante*, a firm's development reputation or stated programs, tracks whether review *occurs*, not whether it *teaches*, the same observability–verifiability gap developed in Section 7. Sorting on coarse signals moves wages on the appearance of development without disciplining its substance; allowing the signal to be manipulated could only raise ζ^* and strengthen the conclusion.