The Broken Job Ladder:
AI, Vanishing Apprenticeships, and the Role of Universities

Miao Ben Zhang*

March 25, 2026

Abstract

Mounting evidence suggests that artificial intelligence (AI) is eliminating entry-level cognitive jobs, which historically served as the apprenticeship through which college graduates develop expert judgment to climb the job ladder. I build a dynamic model in which firms profit from hiring junior workers at a markdown on wages while these workers learn by doing. The arrival of AI renders this apprenticeship uneconomical: firms can no longer capture enough surplus from junior labor to justify the hiring, particularly when trained workers may be poached by competitors. Managerial myopia compounds the underhiring, resulting in the economy to face a long-run scarcity of senior expertise as seasoned professionals retire without replacements. I show that this market failure cannot be corrected by standard policy instruments among firms. Instead, universities, by redesigning curricula around structured micro-residency programs that bundle real-case judgment development with immediate deliverable production, can address the failure even when firms are arbitrarily myopic. The paper identifies universities as the institutional actors best positioned to address what is not just a labor-market disruption but a structural crisis in human capital formation.

JEL: C72, D86, I23, J24, M53, O33

Keywords: AI, human capital, learning-by-doing, career ladders, mechanism design

*Miao Ben Zhang is from the University of Southern California Marshall School of Business and can be reached at miao.zhang@marshall.usc.edu.

1. Introduction

Most skilled professions share a somewhat common developmental architecture. Universities teach *knowledge*—the codifiable, transferable principles of a field, such as legal doctrine, accounting standards, and programming languages. Firms then develop *judgment*—the tacit, experience-based capacity to act decisively when the problem itself is undefined, the stakes are irreversible, and the right answer depends on reading people rather than analyzing documents.[1] Judgment is the product of thousands of supervised decisions accumulated over years of practice. A novice enters the workforce with knowledge, performs routine tasks under supervision, and gradually develops the judgment that distinguishes an expert from a beginner. For decades, production and judgment formation were bundled: the same activity that generated output for the firm also developed expertise in the worker.

Motivated by recent empirical evidence that generative AI is disproportionately displacing entry-level professional employment,[2] this paper argues that the most consequential effect of generative AI on the labor market is not only the elimination of jobs per se—a concern common to all waves of automation[3]—but rather the destruction of the *apprenticeship* through which entry-level workers develop judgment. When AI makes routine cognitive tasks cheaper to automate than to assign to juniors, firms rationally reassign those tasks. But in doing so, they inadvertently eliminate the substrate on which professional learning depended. The career ladder breaks not because the technology for learning has changed, but because the *economic conditions* for deploying it have vanished.

The core insight of our model is that this disruption triggers an endogenous regime transition in human capital formation. Before AI, firms had incentives to hire juniors, as firms can charge a wage markdown on entry-level employees—paying juniors less than their marginal product in exchange for on-the-job learning (Acemoglu and Pischke 1998, 1999).[4] In this regime, judgment developed as a costless byproduct of production, in the

[1]For example, judgment enables a senior M&A partner to walk away from a deal whose financials are clean because the management team's behavior signals undisclosed risk, or a senior litigator to advise settling a case the firm would likely win because opposing counsel's unusual confidence suggests a hidden witness.

[2]See, for example, Brynjolfsson, Chandar, and Chen (2025), Hosseini and Lichtinger (2025), Eloundou et al. (2024), Noy and Zhang (2023), among others. Section 2 details the evidence on AI disproportionately displacing entry-level jobs.

[3]The task-based framework of Acemoglu and Autor (2011), Acemoglu and Restrepo (2019), Acemoglu and Restrepo (2022), Eisfeldt, Schubert, and Zhang (2023) provides the canonical treatment of how automation or generative AI displaces workers from tasks in which they previously had comparative advantage. Autor et al. (2024) extend this framework to study the creation of new work. Our paper differs from this literature in that we focus not on the displacement of tasks but on the destruction of the *developmental function* that task performance served.

[4]This claim is supported by many empirical findings. E.g., Loewenstein and Spletzer (1998) document that employers often bear the explicit costs of such training; Barron, Berger, and Black (1999) provide evidence

tradition of Arrow (1962)'s learning-by-doing. Universities complemented this apprenticeship by providing the foundational knowledge that made on-the-job learning more efficient or by screening job candidates with greater learning ability (Cunha and Heckman 2007 and Spence (1973)). Figure 1 (left panel) illustrates this self-sustaining career ladder.

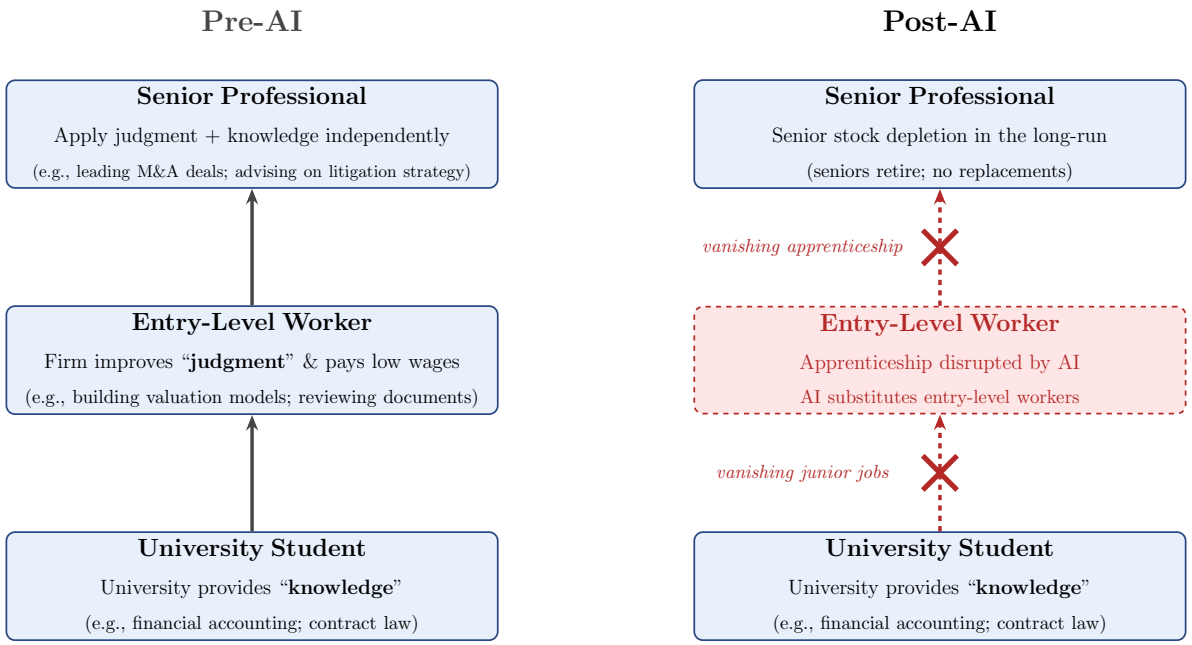


Figure 1: **AI Disrupting the Job Ladder.** The left panel illustrates the pre-AI career ladder in which university education (knowledge) and firm apprenticeship (judgment) combine to produce senior professionals. The right panel shows how generative AI, by cheaply substituting for entry-level cognitive tasks, breaks the apprenticeship rung and leads to long-run depletion of senior talent.

The introduction of generative AI tools fundamentally disrupts this architecture. By cheaply substituting for entry-level cognitive tasks, AI eliminates the junior surplus that sustained the implicit training subsidy, causing firms to stop assigning—and therefore stop developing—junior workers. The right panel of Figure 1 shows the result: the apprenticeship rung of the job ladder breaks, the flow of new seniors dries up, and the stock of senior professionals erodes through retirement with no replacements maturing. The long-run depletion of senior talent is a market failure: each firm's rational decision to substitute AI for juniors is individually optimal but collectively self-destructive, because each firm bears the full cost of training but captures only a fraction of the return — the remainder accruing to competitors who hire away the trained senior without having invested in development.

that training is associated with at least modestly lower starting wages and faster subsequent productivity growth. More recently, Gao, Wang, and Wu (2024) shows that finance workers sustain low compensation in initial jobs to harvest greater salary jumps in later careers.

This market failure is self-reinforcing due to two features of the post-AI economy. First, even firms that recognize the problem will not solve it unilaterally: developing a junior now requires costly, deliberate investment—forgoing the cheaper AI alternative—yet the resulting senior can be poached by other firms (Becker, 1964). Managerial myopia compounds the problem, as the consequences of a depleted senior pipeline materialize well beyond the typical executive's compensation horizon. Each firm therefore under-hires juniors, hoping others will train instead.[5] Second, firms cannot use AI to fully displace senior expertise. Unlike knowledge, judgment is forged in undocumented, real-time experience—the body language of a counterparty under pressure, the boardroom dynamics that never appear in any memo—signals that large language models cannot learn from because they were never written down. Put together, AI makes judgment that they cannot easily substitute harder to develop, resulting in the long-run depletion of expertise that no market participant has the incentive to prevent.

While firms cannot resolve this market failure on its own, our model implies a promising possibility that universities can. Because universities have dinstict incentives from firms and operate on longer planning horizons than firms, they are uniquely positioned to serve as *mechanism designers* rather than passive providers of knowledge and screening. The paper proposes a micro-residency mechanism in which the university designs structured professional experiences—modeled on the medical residency, the clinical legal education, and the Harvard Business School case method—that bundle judgment development with immediate deliverable production. The key design principle is that the deliverable must *require* the exercise of judgment, so that the developmental function operates as a byproduct of the transaction. In this way, the university effectively re-bundles production and learning—the very bundling that AI destroyed—through deliberate institutional design rather than organic market forces. This has concrete implications for university education reform: expanding clinical practica, live-client consulting projects, AI-augmented simulation environments, and cross-disciplinary capstone residencies, all structured so that participating firms receive immediate value while students develop judgment.

Our model basis on a dynamic game among N symmetric firms, a university, and a continuum of junior and senior workers. The university's role follows the education eco-

[5]The poaching concern is less acute for purely firm-specific skills that lose value when transferred across firms. However, the judgment developed through professional apprenticeship is largely *transferable*—particularly within the same industry. Neal (1995) provides direct evidence that human capital is predominantly industry-specific rather than firm-specific; Gathmann and Schönberg (2010) show that labor market skills are "more portable than previously considered," with workers moving systematically to jobs requiring similar task profiles; and Gao, Wang, and Wu (2024) document this pattern in finance, showing that M&A bankers' deal-making expertise transfers across advisory firms.

nomics literature: it certifies minimum ability in the sense of Spence (1973), and builds the knowledge base that Cunha and Heckman (2007) show is complementary with subsequent on-the-job learning. University quality enters as an exogenous amplifier of firm training—graduates with stronger knowledge develop judgment faster from the same work experience—reflecting the institutional reality that firms interact with universities through the labor market for graduates, not by directly controlling program quality.[6] Labor market frictions follow Becker (1964) and Acemoglu and Pischke (1999): trained seniors separate from their training firms with some probability each period, and destination firms acquire fully developed workers without bearing any of the training cost. We model managerial myopia as a shortened effective planning horizon, capturing the agency problem in which executives whose compensation is tied to near-term performance systematically undervalue training investments whose returns materialize over a decade or more.

The model yields sharp results on how AI disrupts the career ladder and why market forces alone cannot repair it. When firms are patient, the equilibrium still features underinvestment in training relative to the social optimum—because each firm anticipates losing some of its trained seniors to competitors—but the pipeline functions. When firms are myopic, the poaching distortion and the myopia distortion compound multiplicatively: training collapses to a small fraction of the optimal level, and in the extreme, no firm develops juniors at all. The senior stock then erodes to zero through retirement with no replacements. A standard corrective training levy—which taxes firms and reimburses those that successfully develop and retain juniors—can restore the optimum when firms are patient. But it fails when firms are myopic: the agents whose short-termism creates the need for intervention are precisely the agents who refuse to participate in it. This fundamental asymmetry means the problem cannot be resolved within the firm sector alone. An external, long-horizon institution must intervene.

The paper's central result establishes that a university-designed micro-residency mechanism achieves *dominant-strategy* incentive compatibility for arbitrarily myopic firms. The mechanism works by providing simulated learning-by-doing substrate—structured professional experiences in which students produce tangible deliverables for participating firms while developing judgment under faculty supervision. The participation constraint reduces to a static inequality (i.e., the deliverable value exceeds the participation fee), which involves no discount factor. This means firms participate not because they value the long-run talent pipeline, but because the immediate transaction is profitable—a la the economic logic of the

[6] Allowing university quality to depend endogenously on firms' donations does not resolve the problem, as each firm has an incentive to free-ride on others' contributions: the university cannot credibly commit to channel its graduates exclusively to the firms that funded it.

teaching hospital, where hospitals accept medical residents not only because the long-run growth value but also because residents produce clinical services that exceed the cost of supervising them.

We rank the equilibrium regimes by welfare: the social planner's optimum (the first-best, achievable via a mandatory training levy when firms are patient) dominates the patient-firm Nash equilibrium (which features poaching-driven underinvestment), which dominates the micro-residency equilibrium (which mitigates but does not eliminate the friction), which strictly dominates the myopic-Nash collapse (in which the pipeline shuts down entirely). The welfare ordering $W^* > W^{\mathrm{NE}} > W^R > W^{\mathrm{MN}}$ highlights that the micro-residency does not achieve the first-best—but it is the *only* mechanism that works when firms are too myopic to accept the levy. University-led curriculum reform is thus not merely a complement to market solutions; it is the solution of last resort when standard corrective mechanisms fail.

Our model and analysis also rely on some limited assumptions. First, we assume that firms continue to need senior professionals whose judgment cannot be fully replicated by AI. If future AI systems develop the capacity to learn from the undocumented, real-time, interpersonal signals, then there is no senior role needed to fill. Second, our model treats workers as passive recipients of firm training decisions: juniors do not invest in their own judgment development by, for example, accepting lower wages in exchange for richer apprenticeship experiences, self-funding simulations, or seeking out informal mentorship, which may attenuate the severity of the market failure.[7] Third, we model firms as symmetric and atomistic, whereas in practice firms differ substantially in size, market power, and training capacity, with the dominant firm may have a stronger incentive to invest in training to maintain longer term dominance.

The paper contributes to two strands of literature. First, we build on the classical theory of human capital accumulation on the job. The idea that workers develop skill through supervised practice has deep roots: Becker (1964) formalized the distinction between general and specific training; Ben-Porath (1967) modeled human capital investment as a life-cycle optimization; Arrow (1962) showed that learning-by-doing generates productivity growth as a byproduct of production; and Acemoglu and Pischke (1998, 1999) explained why firms bear the cost of general training when labor markets are imperfect. On-the-job skill formation has been among the most robust empirical regularities in labor economics.[8] Our paper

[7] Allowing workers to co-invest in training would introduce a two-sided investment problem in the spirit of Acemoglu and Pischke (1999), in which the split of training costs between workers and firms depends on the structure of labor market frictions. Our single-sided formulation is a simplification that likely overstates the severity of underinvestment but does not change the qualitative predictions.

[8] See Mincer (1974) for the foundational empirical treatment of on-the-job training and wage profiles, and Heckman (1998) for a review of the returns to human capital investment across the life cycle.

differs from this literature in two respects. We use the classical frameworks to study a fundamentally new challenge—the AI-driven destruction of the apprenticeship channel—rather than to analyze the traditional questions of training incidence and wage dynamics. And we derive a fundamentally new policy prescription: that the solution lies not in asking firms to internalize the training externality (the standard approach), but in redesigning university education to substitute for the apprenticeship function that firms can no longer profitably provide.

Second, we contribute to the rapidly growing empirical literature on how generative AI is reshaping labor markets. Brynjolfsson, Chandar, and Chen (2025) document disproportionate employment declines among young workers in AI-exposed occupations using ADP payroll data. Hosseini and Lichtinger (2025) corroborate these findings at the firm level, showing that AI-adopting firms reduce junior hiring by 9–10% while senior employment remains stable. Eloundou et al. (2024) estimate that over 46% of U.S. jobs could be substantially affected by LLM-complementary software. Noy and Zhang (2023) provide experimental evidence that generative AI substantially raises productivity on mid-level writing tasks, with the largest gains accruing to lower-ability workers. Acemoglu (2024) provides a macroeconomic framework for assessing AI's impact on growth and wages. Fuller, Sigelman, and Fenlon (2025) and Edmondson and Chamorro-Premuzic (2025) articulate the organizational concern that AI is eliminating the developmental function of entry-level roles. This literature has been predominantly empirical—documenting the facts of displacement—or focused on aggregate labor-market outcomes at the task level (Acemoglu and Restrepo, 2019, 2022; Autor et al., 2024). Our paper contributes a theoretical framework that explains *why* the displacement of entry-level jobs is structurally different from previous waves of automation (because it destroys the apprenticeship channel, not just the tasks), and provides *actionable guidance* on institutional reform—specifically, how universities can redesign their curricula to address the resulting market failure.

The paper proceeds as follows. Section 2 reviews the empirical evidence in greater detail. Section 3 develops the model, beginning with the pre-AI career ladder and tracing the endogenous consequences of the AI shock. Section 4 presents the equilibrium analysis and mechanism design results. Section 5 discusses policy implications for universities, governments, and professional education. Section 6 concludes.

2. Motivating Evidence on AI Disrupting Entry-Level Jobs

A rapidly growing body of research documents the impact of generative AI on professional work. We organize the evidence in three categories.

Labor-market displacement. Recent studies using large-scale administrative data reveal that the productivity gains documented above are translating into measurable declines in entry-level employment. Brynjolfsson, Chandar, and Chen (2025), using ADP payroll records covering millions of U.S. workers, document a 13% relative decline in employment for workers aged 22–25 in AI-exposed occupations since late 2022, while employment for experienced workers in the same occupations remained stable. The decline is concentrated in occupations where AI *automates* rather than *augments* existing tasks—precisely the routine cognitive work that constitutes the apprenticeship substrate in our model. Hosseini and Lichtinger (2025), analyzing 62 million résumés across 285,000 U.S. firms, corroborate these findings with a different identification strategy. They report that firms adopting generative AI reduced junior hiring by 9–10% within six quarters, driven entirely by reduced hiring rather than increased separations. Critically, senior hiring was unaffected, suggesting that AI is displacing the *entry point* of the career ladder rather than the ladder as a whole. Eloundou et al. (2024) provide a complementary perspective, estimating that over 46% of U.S. jobs could be substantially affected by LLM-complementary software, with the highest exposure concentrated in precisely the white-collar, cognitively routine occupations where entry-level apprenticeship has historically been most important.

Task-level productivity. Controlled studies consistently show that generative AI substantially raises productivity on the routine cognitive tasks that constitute entry-level work. Noy and Zhang (2023), in a randomized experiment with 453 college-educated professionals, find that access to ChatGPT reduces task completion time by 40% and raises output quality by 18% on mid-level writing tasks, with the largest gains accruing to lower-ability workers. Brynjolfsson, Li, and Raymond (2025) examine the staggered deployment of a generative AI assistant across 5,172 customer-support agents and find a 15% average productivity increase, with a 30% improvement for novice and low-skilled workers but minimal impact on experienced agents. Peng et al. (2023) report that software engineers using GitHub Copilot complete coding tasks 55% faster, again with larger gains for less experienced developers. The consistent finding across these studies is that generative AI compresses the productivity distribution by disproportionately augmenting the output of junior workers—precisely the

workers whose tasks our model identifies as the substrate for learning-by-doing.[9]

Organizational consequences. A parallel literature documents how firms are reorganizing in response to these shifts. Fuller, Sigelman, and Fenlon (2025) show that firms are reshaping from pyramidal to diamond structures as entry-level tasks are absorbed by AI. Edmondson and Chamorro-Premuzic (2025) articulate the developmental concern at the heart of our paper—that removing junior roles eliminates the "apprenticeship dividend" through which novices develop expert judgment. Walther (2025) highlights the convergence of AI adoption with senior retirements, creating a pipeline crisis from both ends simultaneously.

These pieces of evidence together suggest a pattern that our model formalizes: generative AI disproportionately displaces the routine cognitive work that historically served as the apprenticeship through which juniors develop judgment, and firms respond by reducing junior hiring rather than by investing in alternative developmental pathways.

3. The Model

We develop the model in four steps. We first describe the environment—the features of the economy that are invariant across technological regimes. We then characterize the pre-AI economy, in which professional development was sustained through learning-by-doing. Third and key, we introduce the AI shock and track its endogenous consequences. Finally, we formalize the post-AI training game that is the paper's primary object of analysis.

3.1 Environment

Consider an infinite-horizon discrete-time economy with three types of agents: $N \geq 2$ identical firms, one university, and a continuum of workers. Workers are of two types: *juniors* (J) and *seniors* (S).[10] A senior produces per-period output $y^S > 0$ at wage w^S, yielding net surplus $\sigma \equiv y^S - w^S > 0$ to the employing firm. Seniors retire each period with probability

[9] Dell'Acqua et al. (2023), in a field experiment with Boston Consulting Group consultants, document a related pattern: AI substantially improved performance on tasks within the "technological frontier" but led to worse outcomes on tasks requiring judgment beyond AI's capabilities—what the authors term the "jagged technological frontier." This finding is consistent with our model's distinction between knowledge (which AI can replicate) and judgment (which it cannot).

[10] The binary skill structure is a simplification. In practice, professional development is continuous—a second-year associate is more productive than a first-year, even though neither is a partner. We adopt the two-type framework for tractability, following the standard approach in the training literature (Acemoglu and Pischke, 1999). The qualitative results—underinvestment due to poaching, amplification by myopia, and the university mechanism's dominance—extend to a model with continuous skill accumulation, provided the externality structure is preserved.

$\delta \in (0,1)$, independently. The social discount factor is $\beta \in (0,1)$; we introduce myopic discounting ($\beta^M < \beta$, equivalently a higher discount *rate* than the social optimum) when analyzing managerial short-termism in Section 4.3.[11]

Throughout, we denote by

$$v \equiv \frac{\sigma}{1-\beta(1-\delta)} \tag{1}$$

the present discounted value of the net surplus from one senior, accounting for the geometric probability of retirement.[12]

The university. The university provides two services to the economy. First, it *screens*: a university degree certifies that the graduate possesses minimum cognitive ability, discipline, and domain knowledge, resolving part of the information asymmetry between workers and employers in the tradition of Spence (1973). Second, it *builds foundational human capital*: graduates arrive at firms with codifiable knowledge—legal doctrine, accounting principles, programming skills, statistical methods—that determines how prepared they are to absorb further development on the job.

We represent both functions through a scalar quality parameter $q_0 > 0$, which indexes the university's effectiveness at producing well-prepared, well-screened graduates. The parameter q_0 is *exogenous to the firms*: it is determined by tuition revenue, public funding, endowment income, accreditation standards, and faculty quality—none of which are choice variables in the firms' optimization.[13] The role of q_0 in the model is to serve as a *complementary input* to on-the-job development: graduates with stronger foundational preparation learn faster from the same work experience, consistent with the dynamic complementarity of investments in human capital established by Cunha and Heckman (2007).

[11] For simplicity, we model myopia in reduced form rather than deriving β^M from an explicit principal-agent problem. This formulation is standard in the human capital literature (Acemoglu and Pischke, 1998) and can be microfounded in several ways: an executive who is replaced with probability p each period and receives no continuation value upon departure has an effective discount factor $\beta^M = \beta(1-p)$; alternatively, a manager whose bonus depends only on the current year's earnings rationally applies $\beta^M \approx 0$ to cash flows beyond the bonus horizon. Our results require only that $\beta^M < \beta$, not a specific mechanism generating the wedge.

[12] This is the standard asset-pricing formula for a perpetuity with stochastic termination. Under myopia, the corresponding value is $v^M = \sigma/[1-\beta^M(1-\delta)] < v$, since $\beta^M < \beta$ implies a larger denominator and hence a lower valuation of the senior.

[13] This reflects the institutional reality that firms interact with universities primarily through the labor market for graduates. While firms may sponsor chairs, fund research, or participate in advisory boards, these activities are not the primary determinants of program quality at the margin, and modeling them as such would distort the analysis. We return to the university's *endogenous* role as a mechanism designer in Section 4.5.

Labor market frictions. Each newly matured senior separates from the training firm with probability $\lambda \in [0, 1)$, independently, driven by non-wage factors: relocation, career preferences, industry transitions, or outside-option shocks. Departing seniors are absorbed by other firms at the common equilibrium wage w^S.[14] In symmetric equilibrium, each firm's outflow of departing seniors equals its expected inflow from other firms' departures, so the *net* senior stock is unaffected by turnover. The distortion operates not through wages but through *training costs*: the origin firm invested t_i per junior to develop the departing senior, while the destination firm acquires a fully trained senior at zero training cost. This is the classical externality of Becker (1964), analyzed in the imperfect-labor-markets setting of Acemoglu and Pischke (1999): each firm captures only a fraction $(1 - \lambda)$ of the return on its training investment, leading to systematic underinvestment.

3.2 The Pre-AI Economy: Learning-by-Doing

Before the arrival of generative AI, routine cognitive tasks—document review, financial modeling, research memoranda, basic coding, data assembly—were performed by junior professionals. These tasks constituted the bulk of entry-level knowledge work and had two properties that are central to our analysis.

First, juniors performing these tasks were *productive*. Their output y^J exceeded their wage w^J, generating a positive net surplus $\tau^0 \equiv y^J - w^J > 0$ for the firm. This is the *wage compression* condition identified by Acemoglu and Pischke (1999) as the key mechanism enabling firm-sponsored general training in imperfect labor markets: because frictions prevent trained workers from capturing their full marginal product, firms earn a rent on junior employment that implicitly finances on-the-job development.[15]

Second, performing these tasks *developed the junior's expertise*. The act of reviewing hundreds of contracts built the pattern recognition to spot risk. The act of building financial models developed the intuition to question assumptions. This is *learning-by-doing* in the precise sense of Arrow (1962): human capital accumulated as a costless byproduct of production, without requiring any deliberate investment in training.

We formalize the pre-AI developmental process as follows. Let y^J denote the volume of

[14] In symmetric equilibrium, all firms offer the same wage w^S and earn the same surplus σ per senior. Turnover is not driven by wage differentials—it reflects the empirical reality that a fraction of professionals leave their training firm each period regardless of compensation, due to geographic, personal, or career-stage factors.

[15] Acemoglu and Pischke show that labor market imperfections—minimum wages, unionization, search frictions, asymmetric information—compress the wage structure relative to the productivity distribution. This compression turns "technologically general skills into de facto specific skills," making firms willing to bear training costs that they would avoid under perfect competition.

routine tasks a junior performs per period. The junior's per-period probability of maturing into a senior is:

$$m^{\text{pre}} \;=\; \left(\gamma \cdot y^{J}\right) \cdot q_0^{b}, \tag{2}$$

where $\gamma > 0$ is the Arrow learning rate—the efficiency with which productive experience translates into expertise—and q_0^b with $b \in (0,1)$ captures the Cunha-Heckman complementarity: graduates with stronger foundational preparation (higher q_0) extract more developmental value from the same work experience.[16]

Under this regime, the senior stock at firm i evolves according to:

$$S_i^{t+1} \;=\; (1-\delta)\, S_i^{t} \;+\; (1-\lambda) \cdot m^{\text{pre}} \cdot J_i\,, \tag{3}$$

and the firm's per-period profit from its junior workforce is:

$$\pi_i^{\text{pre}} \;=\; \underbrace{\tau^0 \cdot J_i}_{\text{immediate surplus}} \;+\; \underbrace{(1-\lambda) \cdot m^{\text{pre}} \cdot v \cdot J_i}_{\text{maturation option value}}\,. \tag{4}$$

Both terms are positive. The first is the Acemoglu-Pischke rent: the firm earns $\tau^0 > 0$ on each junior employed. The second is the option value of developing a future senior. The firm hires juniors until the marginal benefit (the sum of these two terms) equals the marginal cost.

Three features of this regime merit emphasis.

No explicit training decision. The maturation rate m^{pre} depends on y^J (the volume of productive tasks) and q_0 (the quality of the graduate's preparation), neither of which is a firm choice variable in the standard sense. The firm does not "choose to train"—it simply assigns juniors to productive work, and learning occurs as a byproduct. The variable t_i (explicit training investment) does not appear.

No coordination problem. Because the developmental function is bundled with the productive function at zero marginal cost, the poaching externality (λ) imposes no training-cost distortion. A departing senior takes human capital that the origin firm developed, but that development cost the origin firm *nothing beyond the productive surplus it already captured.* The Beckerian externality is present in principle but dormant in practice.

A self-sustaining pipeline. The pre-AI career ladder maintained itself automatically.

[16]The specification (2) embeds Arrow's insight that "learning is the product of experience" and that "experience" is itself a joint product of the act of producing. The multiplicative interaction with q_0 reflects Cunha and Heckman (2007)'s finding that early investments in human capital raise the marginal return to later investments.

Firms hired juniors because they were productive ($\tau^0 > 0$); juniors developed expertise because they were doing productive work ($m^{\text{pre}} > 0$); seniors were replaced as they retired because new juniors were continually entering and maturing. No deliberate institutional design was required. The ladder was, in the language of mechanism design, *self-enforcing*.

3.3 The AI Shock

Generative AI enters the economy as a technology that can perform routine cognitive tasks—the same tasks previously assigned to juniors—at per-unit cost c^{AI} that is strictly below the junior's wage: $c^{\text{AI}} < w^J$. Moreover, AI performs these tasks at quality $y^{\text{AI}} \geq y^J$: as well as or better than the typical junior.[17]

The firm's response is immediate and follows from elementary cost minimization: it reassigns routine cognitive tasks from juniors to AI. This is not a strategic decision about training; it is the same logic by which firms adopted word processors, spreadsheets, and databases. The firm does what is cheapest.

The consequences, however, are far-reaching, and they all follow *endogenously* from this single reallocation decision:

Consequence 1: Junior surplus collapses. With routine tasks reassigned to AI, the junior's productive output on those tasks falls to zero: $y^J \to 0$. The net junior surplus becomes $\tau = y^J - w^J \leq 0$. A junior who is not assigned routine cognitive work produces little billable output but still draws a wage. The Acemoglu-Pischke wage-compression mechanism that sustained implicit training investment is destroyed.[18] In contrast, senior tasks—advising under ambiguity, reading counterparties, exercising irreversible judgment—depend on undocumented experiential knowledge that does not appear in AI training data; hence $\sigma = y^S - w^S$ is not immediately affected by the AI shock.[19]

Consequence 2: Arrow's learning-by-doing channel goes dormant. From equation (2), the maturation rate was $m^{\text{pre}} = \left(\gamma \cdot y^J\right) \cdot q_0^b$. With $y^J \to 0$ (no routine tasks

[17] We do not require AI to be superior in quality—only cheaper at comparable quality. The assumption $c^{\text{AI}} < w^J$ is the operative one. In practice, current large language models often produce higher-quality first drafts of routine documents than entry-level professionals, making the assumption conservative.

[18] The wage w^J may also adjust downward, but institutional frictions—minimum wages, professional norms, recruitment competition for talent—prevent it from falling as fast as y^J, so τ turns negative or approaches zero. Even if $\tau = 0$ exactly, the implicit training subsidy vanishes.

[19] In the long run, as the pipeline depletes and senior talent becomes scarcer, the senior labor market may tighten, raising equilibrium wages toward marginal product and compressing σ. This reinforcing feedback is outside our steady-state analysis but would strengthen the case for intervention.

assigned to the junior), the maturation rate converges to zero: $m^{\text{pre}} \to 0$. The learning technology itself has not changed—a junior who reviews hundreds of contracts *would* still develop pattern recognition. But no firm will let the junior do it, because AI does it cheaper. The Arrow channel shuts down not for technological reasons but for *economic* ones: the substrate on which learning-by-doing depends has been reallocated to a non-human agent.

Consequence 3: Developing juniors now requires explicit, costly investment. If a firm still wishes to develop a junior into a senior, it must deliberately assign the junior to tasks that AI could handle more cheaply, purely for developmental purposes. This is the regime transition from Arrow to Ben-Porath.

Let $t_i \geq 0$ denote firm i's per-junior training investment, measured as the opportunity cost of assigning routine tasks to juniors rather than to AI. Concretely, each dollar of t_i represents output the firm forgoes by deploying a junior instead of AI on a unit of routine work. Pre-AI, this opportunity cost was zero or negative because the junior was the cost-effective option. Post-AI, it is strictly positive. The maturation rate under this regime is now a function of the firm's chosen training investment t_i:

$$m^{\text{post}}(t_i) \;=\; (\alpha \cdot t_i^{a}) \cdot q_0^{b}\,, \tag{5}$$

where $\alpha > 0$ is a productivity parameter and $a \in (0, 1)$ captures diminishing returns to training intensity: the first hours of hands-on experience are most valuable; subsequent hours add progressively less.[20] The critical point is that equation (5) embeds the *same* learning-by-doing process as equation (2): the junior still learns from performing tasks under supervision. What has changed is not the technology of learning but its *price*—every hour of developmental experience now carries an opportunity cost that did not exist when juniors were the cheapest way to perform the work.

Remark. The Arrow-to-Ben-Porath transition is driven entirely by a change in *prices*, not in *technology*. The human brain's capacity for learning-by-doing is identical before and after AI. What changes is the opportunity cost of deploying it. Pre-AI, the junior was the cheapest way to perform routine tasks, so learning was a costless byproduct of production: $t_i = 0$. Post-AI, AI is cheaper, so every hour of developmental experience carries a positive opportunity cost: $t_i > 0$. The broken ladder is fundamentally a problem of prices, not of capabilities.

[20] The exponent $a < 1$ also ensures that the firm's problem is strictly concave in t_i, guaranteeing interior solutions.

Consequence 4: The poaching externality activates. Pre-AI, the cost of losing a trained senior to a competitor was low because the training had been costless (embedded in productive work). Post-AI, the firm has invested $t_i \cdot J_i$ in developing its juniors. A fraction λ of newly matured seniors depart, representing a direct loss of the firm's explicit investment. The Beckerian externality, dormant when training was bundled with production, is now first-order.

Consequence 5: Junior hiring declines. With $\tau \leq 0$ (juniors are cost centers) and $t_i > 0$ (training is costly), the marginal benefit of hiring a junior falls sharply relative to the pre-AI regime. Firms rationally reduce junior hiring—precisely the pattern documented by Brynjolfsson, Chandar, and Chen (2025) and Hosseini and Lichtinger (2025).

3.4 The Post-AI Training Game

We now formalize the post-AI economy as a dynamic game among the N firms. The university operates in the background, producing graduates of exogenous quality q_0. Each firm's strategic choices are (i) how many juniors to hire (J_i) and (ii) how intensively to train them (t_i).

Transition. Firm i's senior stock evolves as:

$$S_i^{t+1} \;=\; (1-\delta)\, S_i^t + (1-\lambda)\, m^{\text{post}}(t_i)\, J_i\,, \tag{6}$$

where $m^{\text{post}}(t_i) = \alpha\, t_i^a\, q_0^b$ from (5) and we have used the symmetric-equilibrium result that net spot-market flows are zero.

Payoffs. Firm i's per-period profit is:

$$\pi_i^t \;=\; \sigma \cdot S_i^t + \tau \cdot J_i - t_i \cdot J_i\,. \tag{7}$$

The first term is the surplus from the senior workforce; the second is the (non-positive) surplus from juniors; the third is the training cost.

Value function. Each firm maximizes $\Pi_i = \sum_{t=0}^{\infty} \beta^t \pi_i^t$. Because the per-period payoff is linear in the state S_i and the transition is linear, the value function takes the affine form

$V(S_i) = v \cdot S_i + K$, where v is given by (1) and K collects terms independent of the state.[21] The optimal policy is state-independent, and the system converges to a steady state in which

$$\delta \cdot S_i = (1 - \lambda) \cdot m^{\text{post}}(t_i) \cdot J_i \,. \tag{8}$$

Substituting and simplifying, the firm's steady-state present value reduces to:

$$\Pi_i \;=\; v \cdot (1 - \lambda) \cdot \alpha \, t_i^a \, q_0^b \cdot J_i + \frac{(\tau - t_i) \cdot J_i}{1 - \beta} \,. \tag{9}$$

This is the firm's objective in the analysis that follows.

4. Equilibrium Analysis

4.1 The Social Planner's Benchmark

A social planner who internalizes all externalities sets $\lambda = 0$ (poaching is zero-sum in aggregate) and chooses (t, J) for each of N symmetric firms to maximize total per-period welfare:

$$W(t, J) \;=\; N \left[\frac{\sigma \alpha \, t^a \, q_0^b \, J}{\delta} + \tau J - tJ \right]. \tag{10}$$

The first-order condition with respect to t yields the planner's optimum:

$$t^* \;=\; \left[\frac{\sigma \alpha a \, q_0^b}{\delta} \right]^{\frac{1}{1-a}} . \tag{t^*}$$

Proposition 1 (Cooperative Optimum)**.** *Under the stated assumptions, the planner's problem has a unique interior solution $t^* > 0$ given by (t^*). The optimal training level is increasing in the senior surplus σ, the maturation productivity α, and the university quality q_0, and decreasing in the retirement rate δ. Intuitively, the planner equates the marginal cost of training (one dollar per junior) with the marginal benefit (the increment to the discounted value of the senior pipeline).*

The proof is in Appendix A. The key comparative static for our purposes is $\partial t^* / \partial q_0 > 0$: better university preparation raises the socially optimal level of firm training. This is the Cunha-Heckman dynamic complementarity operating between institutions—it implies that

[21] The proof follows the standard guess-and-verify approach for linear-quadratic control problems. Conjecture $V(S_i) = vS_i + K$, substitute into the Bellman equation, and verify that $v = \sigma/[1 - \beta(1 - \delta)]$. The optimization over (J_i, t_i) enters only through K, confirming the conjecture. See Appendix A.

investments in university quality and investments in firm training are *strategic complements* at the social level. A society that improves its universities but neglects firm-side training, or vice versa, captures less than the full return from either investment.

4.2 Decentralized Equilibrium: The Poaching Externality

We now characterize the symmetric Nash equilibrium when each firm independently chooses (t_i, J_i) to maximize (9), taking other firms' actions as given.

The firm's first-order condition with respect to t_i is:

$$(1-\lambda) \cdot \frac{\sigma \alpha a \, t_i^{a-1} \, q_0^b}{\delta} \;=\; 1\,. \tag{11}$$

Comparing this to the planner's condition (which lacks the $(1-\lambda)$ factor), the firm's private marginal return to training is scaled down by $(1-\lambda)$: it captures the return from a matured senior only if that senior stays, which occurs with probability $(1-\lambda)$. Solving:

$$t^{\mathrm{NE}} \;=\; (1-\lambda)^{\frac{1}{1-a}} \cdot t^*\,. \tag{t^{NE}}$$

Since $(1-\lambda)^{1/(1-a)} < 1$ for $\lambda > 0$, firms underinvest relative to the social optimum. The gap is increasing in λ (more poaching $\to$ more underinvestment) and in a (higher returns to training $\to$ larger absolute distortion).

Proposition 2 (Nash Equilibrium and Inefficiency)**.** *The post-AI training game has a unique symmetric Nash equilibrium with training level $t^{NE} < t^*$. The welfare loss $W^* - W^{NE}$ is:*

(i) strictly positive for any $\lambda > 0$;

(ii) increasing in λ (poaching intensity) and a (returns to training);

(iii) decreasing in q_0: higher university quality partially compensates for firm-side underinvestment by raising the maturation rate at any given t.

Result (iii) is worth highlighting. It implies that university quality and firm training are *partial substitutes from a welfare perspective*: a society with excellent universities can tolerate more firm-side underinvestment without catastrophic talent-pipeline consequences than a society with weak universities. This has direct policy implications, which we discuss in Section 5.

4.3 Managerial Myopia

We now introduce the paper's second friction. Each firm's decision-maker uses a myopic discount factor $\beta^M < \beta$, reflecting executive turnover, quarterly earnings pressure, and compensation structures tied to near-term performance. Because β is a discount *factor* (not a rate), a smaller β^M means the manager discounts future cash flows *more* heavily: β^M close to zero corresponds to extreme myopia in which only current-period payoffs matter, while $\beta^M = \beta$ recovers the socially optimal horizon.[22] This is not irrationality; it is an *agency problem*: the executive making the training decision will not bear the consequences of talent-pipeline depletion that materializes 10–15 years hence.

Under myopia, the firm values a marginal senior at $v^M = \sigma/[1-\beta^M(1-\delta)] < v$, reducing the perceived return to training.

Proposition 3 (Myopic Equilibrium)**.** *Under managerial myopia, the unique symmetric Nash equilibrium has training level*

$$t^{MN} \;=\; \left(\frac{v^M}{v}\right)^{\frac{1}{1-a}} \cdot t^{NE} \;=\; (1-\lambda)^{\frac{1}{1-a}} \cdot \left(\frac{v^M}{v}\right)^{\frac{1}{1-a}} \cdot t^{*}\,. \tag{12}$$

The two distortions—poaching and myopia—compound ***multiplicatively****:*

$$\frac{t^{MN}}{t^*} = \underbrace{(1-\lambda)^{\frac{1}{1-a}}}_{\textit{poaching distortion}} \;\times\; \underbrace{\left(\frac{v^M}{v}\right)^{\frac{1}{1-a}}}_{\textit{myopia distortion}}\,.$$

The multiplicative structure has a stark implication: even moderate levels of poaching and myopia can produce severe underinvestment. If $\lambda = 0.3$ (30% of trained seniors depart) and $v^M/v = 0.5$ (the manager's effective horizon is half the social optimum), then $t^{\mathrm{MN}}/t^* = 0.7^{1/(1-a)} \times 0.5^{1/(1-a)}$, which for $a = 0.5$ gives $t^{\mathrm{MN}}/t^* = 0.49 \times 0.25 = 0.12$—the myopic equilibrium provides barely one-eighth of the socially optimal training.

Corollary 1 (Talent Extinction)**.** *As $\beta^M \to 0$ (extreme myopia): $t^{MN} \to 0$, the maturation rate converges to zero, and the senior stock declines geometrically to extinction: $S_i^t = (1-\delta)^t S_i^0 \to 0$.*

Corollary 2 (Pipeline Shutdown)**.** *There exists a critical threshold $\bar{\tau}$ such that for $\tau < \bar{\tau}$, no firm hires juniors ($J_i = 0$) and the talent pipeline shuts down entirely, regardless of training intensity or university quality.*

[22] The relationship between the discount factor β and the discount rate r is $\beta = 1/(1+r)$. A myopic manager with $\beta^M < \beta$ equivalently applies a higher discount rate $r^M > r$ to future cash flows.

The proof derives $\bar{\tau}$ from the firm's FOC for J_i (Appendix A). The threshold $\bar{\tau}$ is more negative when university quality q_0 is higher—better universities extend the range of AI shock severity that the economy can absorb before the pipeline collapses.

4.4 Mechanism Design: The Training Levy

Can the coordination failure be corrected? We first consider a Pigouvian intervention.

A *training levy mechanism* $\mathcal{M}^L = (\tau^L, r)$ imposes a per-senior levy τ^L on each firm and uses the revenue to reimburse firms at rate r for each junior that matures *and is retained*. The reimbursement offsets the poaching externality by compensating the training firm for the expected loss from departing seniors.

Proposition 4 (Implementation via Training Levy)**.** *Setting $r^* = \lambda \cdot v \cdot m^{post}(t^*)$, the training levy implements the cooperative optimum t^* as the unique symmetric Nash equilibrium.*

The proof (Appendix A) shows that the reimbursement r^* restores the firm's effective marginal return to training from $(1-\lambda)\sigma\alpha a\, t^{a-1} q_0^b/\delta$ to $\sigma\alpha a\, t^{a-1} q_0^b/\delta$—exactly offsetting the poaching distortion.

However, the levy requires firms to accept an immediate cost (the levy payment) in exchange for a long-run benefit (the restored training equilibrium). When firms are myopic, this exchange may not be attractive:

Proposition 5 (Participation Failure Under Myopia)**.** *If the training levy is offered on a voluntary basis, a firm with discount factor β^M participates if and only if $\beta^M \geq \bar{\beta}$, where $\bar{\beta}$ is a critical threshold defined in Appendix A. For $\beta^M < \bar{\beta}$, the firm's myopic valuation of future senior services is insufficient to justify the immediate levy cost.*

This result reveals a fundamental asymmetry: the agents whose myopia necessitates intervention are the same agents who refuse to participate in it. The mechanism designed for patient agents is rejected by impatient ones. Solving this requires a different approach.

4.5 The University as Mechanism Designer

The university has a structural advantage: its planning horizon far exceeds the typical firm's. Endowments, tenured faculty, and institutional missions create effective discount factors close to β, while corporate decision-makers operate at $\beta^M \ll \beta$. We now show that this asymmetry allows the university to design a mechanism that succeeds where the Pigouvian levy fails.

The key insight from Section 3.3 is that AI destroyed the *substrate* for learning-by-doing—the routine tasks that simultaneously produced output and developed expertise. The university's mechanism provides a *synthetic substitute* for this substrate: structured professional experiences in which juniors work on real problems under senior supervision, generating tangible deliverables for participating firms while developing judgment as a byproduct. This is not a hypothetical construct—it is the logic of the medical residency, the clinical legal education, and the teaching hospital, now generalized to knowledge-work professions.

Definition 1. A *micro-residency mechanism* $\mathcal{M}^R = (f, v^R, t^R)$ specifies: (i) a per-slot fee $f > 0$ paid by each participating firm; (ii) a guaranteed deliverable of value $v^R > 0$ produced by each resident within the current period (a consulting report, AI audit, prototype, or analysis); and (iii) an embedded training intensity t^R structured so that each resident matures with probability $m^R = m^{\text{post}}(t^R)$.

The university designs the program so that producing the deliverable *requires* exercising professional judgment under supervision—the same developmental mechanism that routine work provided pre-AI, but now deliberately engineered rather than organically emergent. The condition $v^R \geq f$ requires that the university can design projects in which supervised students produce deliverables whose value to the firm covers the participation fee. This is an empirical condition, not a theoretical guarantee: it is satisfied in settings where faculty supervision converts student effort into professional-quality output—as in teaching hospitals, law clinics, and consulting practica—but may fail in fields where the gap between student and professional output is too large to bridge with supervision alone.

Proposition 6 (Dominant-Strategy Incentive Compatibility)**.** *The micro-residency mechanism is incentive-compatible in* ***dominant strategies*** *if $v^R \geq f$. That is, for any discount factor $\beta^M \in (0, 1)$, any beliefs about other firms' actions, and any realization of the poaching shock, each firm weakly prefers participation to non-participation.*

Proof. The firm's per-period payoff from one residency slot is:

$$\pi_i^R = \underbrace{(v^R - f)}_{\geq 0} + \underbrace{\beta^M (1 - \lambda)\, m^R\, v^M}_{\geq 0}.$$

The first term is the immediate net value of the deliverable. The second is the discounted option value of the resident maturing into a retained senior. Both terms are non-negative, so $\pi_i^R \geq 0$ whenever $v^R \geq f$. This condition involves no discount factor, no beliefs about other firms, and no realization of λ. Participation is a dominant strategy. □

The result works because the mechanism *recreates the economic structure of the pre-AI regime at the transaction level*: the firm receives immediate productive value (v^R, analogous to $\tau^0 \cdot J_i$) from the resident's work, while the developmental function operates in the background (analogous to Arrow's learning-by-doing). The university has effectively *re-bundled production and development*—the very bundling that AI destroyed—by engineering a program in which the deliverable cannot be produced without exercising and developing professional judgment.

Remark. Proposition 6 is the paper's central result. It identifies the university—not the firm, not the government—as the institutional actor best positioned to solve the broken-ladder problem. The university succeeds where the Pigouvian levy fails because it offers *immediate value* rather than *deferred benefit*. The mechanism works for arbitrarily myopic firms because its participation constraint is a *static* inequality, independent of the discount factor. This is the economic logic of the teaching hospital: hospitals accept medical residents not because they are altruistic about medical education, but because residents produce clinical services that exceed the cost of supervising them. The university's micro-residency generalizes this logic to knowledge-work professions disrupted by AI.

4.6 Welfare Ranking

Proposition 7 (Welfare Ordering)**.** *The equilibrium regimes are strictly ranked:*

$$W^* > W^{NE} > W^R > W^{MN}.$$

The micro-residency achieves a second-best outcome that strictly dominates the myopic-Nash collapse.

Table 1: Equilibrium regimes and welfare ranking.

Regime	**Training**	**Senior Stock**	**Welfare**	**Key Friction**
Social planner	t^*	S^*	W^*	None
Nash (patient firms)	$t^{\mathrm{NE}} < t^*$	$S^{\mathrm{NE}} < S^*$	W^{NE}	Poaching (λ)
Micro-residency	t^R (embedded)	$S^R \in (0, S^*)$	W^R	Poaching (mitigated)
Myopic Nash	$t^{\mathrm{MN}} \ll t^*$	$S^{\mathrm{MN}} \to 0$	W^{MN}	Poaching $\times$ myopia

To make the welfare ranking concrete, we provide a numerical illustration of the ranking by parameterizing the model as follows: $\lambda = 0.30$, $a = 0.5$, $\delta = 0.10$, $\beta = 0.95$, $\beta^M = 0.60$, $\sigma = 1$, $\alpha = 0.20$, $q_0 = 1$, $\tau = -0.10$, and $J = 1$. The value $\beta^M = 0.60$ corresponds roughly to a manager who discounts anything beyond a three-year horizon; $\lambda = 0.30$ means that

three in ten trained seniors depart per period; and $\tau = -0.10$ captures the mild negative junior surplus that results when AI makes routine cognitive work cheaper to automate than to assign to entry-level employees. Appendix B derives the training levels and welfare values for each regime under this parameterization; Table 2 summarizes the results.

Table 2: Numerical Illustration of the Welfare Ranking

Regime	Training t	Maturation $m(t)$	Senior Stock S	Welfare W	W as % of W^*
Social planner	1.00	0.200	2.00	0.90	100%
Nash (patient)	0.49	0.140	0.98	0.81	90%
Micro-residency	0.25	0.100	0.70	0.65	72%
Myopic Nash	0.05	0.044	0.31	0.30	33%

Three features of these numbers deserve emphasis. First, the poaching externality alone—holding myopia aside—already destroys ten percent of first-best welfare and cuts the training level roughly in half, a significant distortion even before managerial short-termism enters the picture. Second, poaching and myopia compound *multiplicatively*: combined, they reduce training to five percent of the social optimum and leave the senior stock at one-sixth of its efficient level, eroding toward zero as seniors retire with no replacements maturing. Third, the micro-residency more than doubles the senior stock relative to myopic Nash (0.70 versus 0.31) and recovers seventy-two percent of first-best welfare—the only improvement achievable when β^M is low. The residual gap between the micro-residency and the social optimum reflects the persistent poaching externality and the university's imperfect ability to replicate organic learning-by-doing; nonetheless, it strictly dominates every outcome the market can produce unaided.

5. Implications for University Education

5.1 What Universities Should Do

The model's central prescription is that universities must act as mechanism designers rather than passive credential providers. Proposition 6 shows that the micro-residency succeeds where market forces and Pigouvian intervention both fail, because icro-residency delivers immediate value to myopic firms while embedding developmental experiences that only a long-horizon institution would design.

In practice, this means designing programs that bundle professional judgment development with tangible deliverables: embedded consulting projects in which students solve real firm problems under faculty supervision; AI-audit residencies in which students evaluate and

improve firms' AI implementations; cross-functional simulations that replicate the complexity of senior-level decisions. The key design principle is that the deliverable must *require* the exercise of judgment—otherwise the developmental function is not embedded.

University quality q_0 plays a dual role in the post-AI economy. Its pre-AI function—screening and foundational human capital—persists. But the comparative static $\partial t^*/\partial q_0 > 0$ from Proposition 1 implies an additional role: higher q_0 raises the marginal return to whatever firm-side training does occur, partially compensating for the coordination failure. Investments in university quality are thus more socially valuable post-AI than pre-AI, because they amplify the scarce resource (firm training) rather than merely supplementing an abundant one (organic learning-by-doing).

5.2 The Role of Government

Table 3 shows that U.S. university revenues derive primarily from state, federal and local government funding. When firms are patient ($\beta^M \approx \beta$), the training levy of Proposition 4 implements the first-best. When firms are myopic, the levy must be mandatory to overcome the participation constraint of Proposition 5. The optimal levy rate $r^* = \lambda \cdot v \cdot m^{\text{post}}(t^*)$ provides a concrete calibration target: it equals the expected present value of the poaching externality, and it can be estimated from observable quantities (senior wages, turnover rates, maturation durations).

Government can also address myopia directly by mandating minimum supervised-practice requirements for professions in which the broken ladder threatens public welfare—the logic of the medical residency mandate, now extended to AI-disrupted knowledge professions.

Table 3: U.S. University Funding Sources

This table reports the sources of funding for the U.S. public universities in 2023. Source: Visual Capitalist.

Source of Revenue	**Total 2023 (Billions)**	**Share**
State Government	$109.5	28%
Tuition and Fees	$80.8	21%
Federal Government	$68.9	18%
Private Sources	$51.2	13%
Other	$36.5	9%
Local Government	$20.7	5%
Sales and Services	$15.5	4%
Government Capital Appropriations	$9.0	2%
Total	**$392**	**100%**

5.3 Path of University Education: Existing Models and New Possibilities

Medicine solved a version of this problem long ago: structured supervision, progressive responsibility, real patients as training substrates, and public subsidization via Medicare GME funding. Our model identifies why this institutional form was necessary and why analogous institutions may be needed for knowledge work: the medical model does not rely on market forces to produce the socially optimal level of training. It mandates and subsidizes it. In a world where AI has broken the market-based training mechanism in law, finance, consulting, and technology, the medical model provides a blueprint.

The micro-residency mechanism of Proposition 6 is not without precedent. Perhaps the most influential existing model is the *case method* pioneered by Harvard Business School, in which students are presented with detailed, real-world business situations and required to analyze ambiguous information, weigh competing considerations, and defend a course of action under time pressure and peer scrutiny. The case method does not teach codifiable knowledge in the traditional sense—students do not memorize formulas or doctrine. Rather, it develops judgment: the ability to identify what matters in a complex situation, to reason under uncertainty, and to make defensible decisions when the data are incomplete. This is precisely the type of human capital that our model's maturation function $m^{\text{post}}(t_i)$ captures, and that AI cannot substitute for because it requires repeated, supervised exposure to situations where context, stakeholder dynamics, and ambiguity dominate.

The case method illustrates a general principle: universities can build judgment when they design curricula around *authentic decision-making under uncertainty*, rather than around the transmission of codifiable knowledge alone. Several existing and emerging pedagogical models embody this principle and could be expanded in a post-AI environment:

Clinical legal education. Law school clinics place students in real client-facing roles—representing tenants in eviction proceedings, advising small businesses on regulatory compliance, drafting briefs for appellate cases—under the supervision of faculty attorneys. Unlike traditional doctrinal courses that teach what the law *is*, clinics develop the judgment to determine what the law *means* in a specific factual context: which arguments are strongest, which risks are tolerable, and how to counsel a client whose interests are not purely legal. The clinic model is a direct analog of the micro-residency: the student produces a real deliverable (legal representation) while developing judgment as a byproduct.

Consulting practica and live-client projects. Business schools increasingly offer courses in which student teams solve real strategic problems for partner firms—market entry analyses, operational restructurings, product launch assessments—and present recommendations to senior management. These engagements develop the judgment to scope an ambiguous

problem, to distinguish signal from noise in qualitative data, and to tailor recommendations to organizational politics and resource constraints. From the firm's perspective, the engagement delivers a tangible consulting output (v^R in our model) at below-market cost (the fee f), while the developmental externality accrues to the student.

AI-augmented simulation environments. Generative AI itself may paradoxically provide the infrastructure for rebuilding the judgment pipeline it disrupted. Universities can use AI to create high-fidelity simulations of professional environments: simulated client negotiations in which AI plays the counterparty with realistic behavioral patterns; simulated deal processes in which students must evaluate AI-generated financial models and identify the assumptions that require human judgment; simulated litigation scenarios in which AI generates document sets that students must review for relevance and risk. These simulations recreate the learning-by-doing substrate at a fraction of the cost of real client work, allowing students to accumulate the thousands of repetitions that the Dreyfus model identifies as necessary for developing proficiency.

Cross-disciplinary capstone residencies. Universities can design capstone programs that place students at the intersection of multiple professional domains—for example, a joint finance-legal-technology residency in which students work on AI governance problems that require financial analysis, regulatory interpretation, and technical assessment simultaneously. These programs develop the integrative judgment that is most resistant to AI substitution: the ability to synthesize across domains, to recognize when a problem requires expertise beyond one's own, and to coordinate multi-disciplinary teams toward a coherent recommendation.

Structured mentorship networks with alumni practitioners. Universities can formalize the developmental function that informal professional networks once provided by creating structured mentorship programs in which senior practitioners—partners, managing directors, senior engineers—supervise cohorts of students on real professional challenges. The university provides the institutional infrastructure (matching, scheduling, quality control, liability coverage) while the practitioners provide the authentic senior supervision that is the scarce input in the post-AI economy. This model leverages the university's convening power and long-horizon planning to solve a coordination problem that individual firms cannot: no single firm will invest in training juniors who may leave, but a university-coordinated program pools the developmental function across firms while distributing the supervision burden.

Each of these models shares the economic structure identified in Proposition 6: the student produces a deliverable whose immediate value to the participating firm exceeds the fee, while the judgment-building function operates as a byproduct that the long-horizon

university internalizes. The common design principle is that the deliverable must *require* the exercise of judgment—mere knowledge application, which AI can replicate, is insufficient. The university's comparative advantage is its ability to engineer this bundling deliberately, at scale, and across cohorts, in a way that no individual firm would find privately optimal.

5.4 Challenges and Limitations of Micro-Residency Implementation

The micro-residency mechanism provides a theoretical blueprint, but its practical implementation faces several challenges that may weaken the condition $v^R \geq f$ on which incentive compatibility depends.

The fidelity gap between simulation and practice. Case studies and simulations, however well designed, lack the psychological pressure of irreversible, high-stakes decision-making. A student analyzing a merger case knows that no real capital is at risk. The Dreyfus model of skill acquisition emphasizes that expert judgment is forged through repetition *under consequential uncertainty*—simulations that remove this pressure may develop analytical skill but not the full spectrum of judgment. Immersive technologies—including virtual reality environments with realistic time constraints, adversarial counterparties, and immediate consequence feedback—offer a potential avenue for closing this gap, drawing on early evidence from simulation-based medical education that high-fidelity stress conditions improve skill transfer to practice.[23]

Corporate reluctance to share real-time problems. The strongest form of the micro-residency requires firms to share sensitive information: pending deals, internal reorganizations, competitive vulnerabilities. Many firms will resist this even under nondisclosure agreements. Where real-time engagement is infeasible, universities can design *retrospective case residencies* built from past corporate decisions, reconstructed with sufficient detail to replicate the decision-making environment. Firms are more willing to share problems from several years ago, when competitive sensitivity has diminished. The trade-off is that retrospective cases weaken the immediate-deliverable channel (v^R declines), potentially violating $v^R \geq f$. This scenario requires either public subsidization (analogous to Medicare GME funding) or a hybrid design in which some engagements are live and others are retrospective.

The quality gap between student and professional output. The condition $v^R \geq f$ assumes that supervised students can produce deliverables of sufficient quality to justify the

[23] See Dreyfus and Dreyfus (1986) on the role of consequential decision-making in expert development.

firm's participation fee. This is plausible where faculty supervision can bridge the gap—a law clinic supervised by an experienced litigator, a consulting practicum led by a former partner—but may fail in fields requiring deep proprietary expertise. Universities can mitigate this by carefully scoping projects: assigning well-defined, modular problems rather than open-ended strategic mandates, and structuring multi-cohort engagements in which returning students build on prior work, accumulating institutional knowledge that raises deliverable quality over time.

Despite these challenges, the core economic logic suggests that universities are the relatively more suitable institutions with both the planning horizon and the convening power to re-bundle production and judgment development. The medical profession's century-long experience with residency programs suggests that the mechanism can work at scale, but that it requires sustained institutional commitment, public subsidization, and iterative refinement over decades.

6. Conclusion

Generative AI is dismantling the developmental infrastructure through which novices become experts. This paper has identified the mechanism: AI makes junior labor uncompetitive on routine cognitive tasks, causing firms to reassign that work to AI, which endogenously destroys the Arrow learning-by-doing channel that sustained professional career ladders at zero marginal training cost. The economy is forced from a regime of costless, bundled human capital formation onto a regime of costly, explicit Ben-Porath investment, under the labor-market frictions that activate the classical Beckerian poaching externality.

The coordination failure that results is amplified by managerial myopia and compounds multiplicatively with the poaching distortion. Standard Pigouvian mechanisms fail when the agents whose myopia creates the problem are the same agents who refuse to participate in the solution. The paper's main result shows that universities can break this impasse by designing micro-residency mechanisms that provide simulated learning-by-doing substrate, achieving dominant-strategy incentive compatibility for arbitrarily myopic firms by reducing the participation constraint to a static inequality.

The broken ladder is not a temporary disruption. It is a structural consequence of a price change: the cost of routine cognitive work, which was formerly anchored to the junior wage, has fallen to near zero. Until institutions are built to provide the developmental experiences that the old price structure sustained organically, the pipeline through which novices become experts will continue to erode. Building those institutions is the defining challenge for higher

education in the age of generative AI.

A. Proofs and Technical Details

A.1 Proof of Proposition 1 (Cooperative Optimum)

The planner's objective (10) is strictly concave in t for fixed $J > 0$: $W_{tt} = N\sigma\alpha a(a-1)t^{a-2}q_0^b J/\delta < 0$ since $a < 1$. The FOC $\partial W/\partial t = 0$ yields $N\sigma\alpha a\, t^{a-1}q_0^b J/\delta = NJ$, which simplifies to ($t^*$). The comparative statics follow from implicit differentiation. □

A.2 Proof of Proposition 2 (Nash Equilibrium)

From (9), firm i's FOC with respect to t_i is (11). Comparing to the planner's FOC, the factor $(1-\lambda)$ scales down the private return. Solving gives (t^{NE}). Uniqueness follows from strict concavity of Π_i in t_i ($a < 1$). Strict inefficiency: $t^{\mathrm{NE}}/t^* = (1-\lambda)^{1/(1-a)} < 1$ for $\lambda > 0$. For (iii), $W^{\mathrm{NE}} = N\sigma\alpha(t^{\mathrm{NE}})^a q_0^b J/\delta + N(\tau - t^{\mathrm{NE}})J$; since $\partial W^{\mathrm{NE}}/\partial q_0 > 0$ via q_0^b, higher q_0 raises welfare at any t^{NE}. □

A.3 Proof of Proposition 3 (Myopic Equilibrium)

Under myopia, v is replaced by $v^M < v$ in the firm's Bellman equation. The FOC becomes $(1-\lambda)\sigma^M\alpha a\, t^{a-1}q_0^b/\delta = 1$ where $\sigma^M = \sigma v^M/v$. Solving and chaining with (t^{NE}) yields (12). The multiplicative structure follows from the power-function form of the solution. □

A.4 Proof of Corollary 2 (Pipeline Shutdown)

From (9), the FOC for J_i is $v(1-\lambda)m^{\mathrm{post}}(t_i) + (\tau - t_i)/(1-\beta) = 0$. The first term is positive; the second is negative for $\tau < t_i$. At $\bar{\tau} = t^{\mathrm{MN}} - v^M(1-\beta^M)(1-\lambda)m^{\mathrm{post}}(t^{\mathrm{MN}})$, the two terms cancel and $J_i = 0$ is optimal. □

A.5 Proof of Proposition 4 (Training Levy)

Under the levy, firm i's modified payoff includes a reimbursement $r(1-\lambda)m^{\mathrm{post}}(t_i)J_i$ for retained matured seniors. The FOC becomes $[(\sigma - \tau^L)(1-\lambda)/\delta + r(1-\lambda)]\alpha a\, t_i^{a-1}q_0^b = 1$. Setting $r = r^*$ makes the bracketed term equal to σ/δ, replicating the planner's FOC. □

A.6 Proof of Proposition 5 (Participation Failure)

The firm participates iff its payoff under the levy exceeds autarky. The net benefit is positive iff $v^M(1-\lambda)m^{\mathrm{post}}(t^*)J/\delta \geq \tau^{L*}S^*/(\sigma - \tau^{L*})$. Solving for the critical β^M yields $\bar{\beta}$; for $\beta^M < \bar{\beta}$,

autarky dominates. □

A.7 Supermodular Game Structure

Define the reduced game $\hat{\Gamma}$ in which each firm chooses $t_i \in \mathbb{R}_+$. The payoff Π_i depends on rivals' actions only through the aggregate maturation rate (which affects the spot-market pool). In symmetric equilibrium, this dependence vanishes, and the game reduces to N identical single-agent problems. However, we can verify supermodularity of the extended game (with J_i as a second choice variable): $\partial^2\Pi_i/\partial t_i \partial J_i \propto \alpha a\, t_i^{a-1} q_0^b > 0$, confirming that training and hiring are complements. Existence and uniqueness of the symmetric equilibrium then follow from Tarski's fixed-point theorem (Topkis, 1998) and the contraction property of the best-response mapping. □

B. Numerical Illustration: Derivations

We adopt the parameterization of Section 4.7: $\lambda = 0.30$, $a = 0.5$, $\delta = 0.10$, $\beta = 0.95$, $\beta^M = 0.60$, $\sigma = 1$, $\alpha = 0.20$, $q_0 = 1$, $\tau = -0.10$, $J = 1$.

Discount-adjusted senior values. The present discounted value of a senior under the social and patient-firm discount factor β is

$$v = \frac{\sigma}{1 - \beta(1-\delta)} = \frac{1}{1 - 0.95 \times 0.90} = \frac{1}{0.145} \approx 6.90.$$

Under managerial myopia,

$$v^M = \frac{\sigma}{1 - \beta^M(1-\delta)} = \frac{1}{1 - 0.60 \times 0.90} = \frac{1}{0.46} \approx 2.17.$$

Social planner's optimum. From (t^*), the planner equates the marginal cost of training to its full social return:

$$t^* = \left(\frac{\sigma \alpha a\, q_0^b}{\delta}\right)^{\frac{1}{1-a}} = \left(\frac{1 \times 0.20 \times 0.5}{0.10}\right)^2 = (1.00)^2 = 1.00.$$

Nash equilibrium (patient firms). From (t^{NE}), the poaching distortion scales down the planner's optimum by $(1-\lambda)^{1/(1-a)}$:

$$t^{NE} = (1-\lambda)^{\frac{1}{1-a}} \cdot t^* = (0.70)^2 \times 1.00 = 0.49.$$

Myopic Nash equilibrium. From (t^{MN}), the poaching and myopia distortions compound multiplicatively:

$$t^{MN} = \left[(1-\lambda)\cdot\frac{v^M}{v}\right]^{\frac{1}{1-a}}\cdot t^* = \left[0.70\times\frac{2.17}{6.90}\right]^2\times 1.00 = (0.220)^2 \approx 0.05.$$

To appreciate the severity, note that myopia alone would reduce training to $(v^M/v)^{1/(1-a)} = (0.314)^2 \approx 0.10$ of the optimum; combined with poaching, it falls to five percent—the two frictions are more than additive.

Micro-residency. The university sets the embedded training intensity at $t^R = 0.25$, a level that myopic firms would never choose voluntarily but participate in because the deliverable value satisfies $v^R \geq f$. The value $t^R = 0.25$ is chosen for illustration; any $t^R \in (t^{MN}, t^{NE})$ would demonstrate the same qualitative improvement.

Maturation rates. The maturation function is $m(t) = \alpha t^a q_0^b = 0.20\cdot t^{0.5}$, so:

$$\begin{aligned} m(t^*) &= 0.20\times 1.00^{0.5} = 0.200,\\ m(t^{NE}) &= 0.20\times 0.49^{0.5} = 0.140,\\ m(t^R) &= 0.20\times 0.25^{0.5} = 0.100,\\ m(t^{MN}) &= 0.20\times 0.05^{0.5} \approx 0.044. \end{aligned}$$

Steady-state senior stocks. From the steady-state condition (**??**), $\delta S = (1-\lambda)\, m(t)\, J$, so $S = (1-\lambda)\, m(t)/\delta$:

$$\begin{aligned} S^* &= 0.70\times 0.200/0.10 = 2.00,\\ S^{NE} &= 0.70\times 0.140/0.10 = 0.98,\\ S^R &= 0.70\times 0.100/0.10 = 0.70,\\ S^{MN} &= 0.70\times 0.044/0.10 = 0.31. \end{aligned}$$

Welfare. We evaluate welfare using the planner's objective (**??**), treating poaching as zero-sum in aggregate:

$$W(t) = \frac{\sigma\alpha t^a q_0^b}{\delta} + \tau - t = 2\, t^{0.5} - 0.10 - t.$$

Substituting each training level:

$$\begin{aligned} W^{*} &= 2(1.00)^{0.5} - 0.10 - 1.00 = 0.90, \\ W^{NE} &= 2(0.49)^{0.5} - 0.10 - 0.49 = 0.81, \\ W^{R} &= 2(0.25)^{0.5} - 0.10 - 0.25 = 0.65, \\ W^{MN} &= 2(0.05)^{0.5} - 0.10 - 0.05 \approx 0.30. \end{aligned}$$

The ordering $W^{*} > W^{NE} > W^{R} > W^{MN}$ is confirmed, as required by Proposition 7.

References

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

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). Beyond Becker: Training in imperfect labour markets. *Economic Journal*, 109(453):112–142.

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 (2022). Tasks, automation, and the rise in U.S. wage inequality. *Econometrica*, 90(5):1973–2016.

Acemoglu, D. (2024). The simple macroeconomics of AI. NBER Working Paper No. 32487.

Arrow, K. J. (1962). The economic implications of learning by doing. *Review of Economic Studies*, 29(3):155–173.

Autor, D., C. Chin, A. Salomons, and B. Seegmiller (2024). New frontiers: The origins and content of new work, 1940–2018. *Quarterly Journal of Economics*, 139(3):1399–1465.

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

Barron, J. M., M. C. Berger, and D. A. Black (1999). Do workers pay for on-the-job training? *Journal of Human Resources*, 34(2):235–252.

Becker, G. S. (1964). *Human Capital*. 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.

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

Cunha, F. and J. J. Heckman (2007). The technology of skill formation. *American Economic Review*, 97(2):31–47.

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 AI on knowledge worker productivity and quality. Harvard Business School Working Paper No. 24-013.

Dreyfus, H. L. and S. E. Dreyfus (1986). *Mind Over Machine*. Free Press.

Edmondson, A. C. and T. Chamorro-Premuzic (2025). The perils of using AI to replace entry-level jobs. *Harvard Business Review*, September 2025.

Eisfeldt, A. L., G. Schubert, and M. B. Zhang (2023). Generative AI and firm values. NBER Working Paper No. 31222.

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

Eraut, M. (1994). *Developing Professional Knowledge and Competence*. Falmer Press.

Fuller, J., M. Sigelman, and M. Fenlon (2025). How gen AI could change the value of expertise. *Harvard Business Review*, March 2025.

Gao, J., W. Wang, and Y. Wu (2024). Human capital portability and careers in finance. *Review of Financial Studies*, 37(9):2732–2778.

Gathmann, C. and U. Schönberg (2010). How general is human capital? A task-based approach. *Journal of Labor Economics*, 28(1):1–49.

Heckman, J. J. (1998). What should be our human capital investment policy? *Fiscal Studies*, 19(2):103–119.

Hosseini, S. M. and G. Lichtinger (2025). Generative AI as seniority-biased technological change. SSRN Working Paper No. 5425555.

Jackson, M. O. and T. R. Palfrey (2001). Voluntary implementation. *Journal of Economic Theory*, 98(1):1–25.

Loewenstein, M. A. and J. R. Spletzer (1998). Dividing the costs and returns to general training. *Journal of Labor Economics*, 16(1):142–171.

Milgrom, P. and J. Roberts (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277.

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

Neal, D. (1995). Industry-specific human capital: Evidence from displaced workers. *Journal of Labor Economics*, 13(4):653–677.

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

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.

Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3):355–374.

Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton University Press.

Walther, C. C. (2025). Is AI pushing us to break the talent pipeline? Knowledge@Wharton, August 2025.