

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

Miao Ben Zhang\*

April 2026

## Abstract

Generative AI disproportionately displaces entry-level cognitive jobs that historically served as apprenticeships through which college graduates developed expert judgment. I build a dynamic model in which AI breaks that ladder by making junior work uneconomical and eliminating learning-by-doing—previously a byproduct of junior production. Employee poaching and managerial myopia then lead firms to underhire juniors, resulting in long-term depletion of senior expertise. Standard firm-side corrections fail under myopia. Universities, by designing micro-residency programs delivering immediate value for firms while embedding judgment development, achieve dominant-strategy incentive compatibility—identifying universities as best positioned to address this structural crisis in human capital formation.

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

**Keywords:** AI, human capital, job ladders, mechanism design, university education

---

\*Miao Ben Zhang is from the University of Southern California Marshall School of Business and can be reached at [miao.zhang@marshall.usc.edu](mailto:miao.zhang@marshall.usc.edu).

Most skilled professions share a common developmental architecture. Universities teach *knowledge*—codifiable, transferable principles such as legal doctrine, accounting standards, and programming languages. Firms then develop *judgment*—the tacit, experience-based capacity to act decisively when the problem is undefined, the stakes are irreversible, and the right answer depends on reading people rather than analyzing documents.<sup>1</sup> 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 evidence that generative AI is disproportionately displacing entry-level professional employment,<sup>2</sup> 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<sup>3</sup>—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 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).<sup>4</sup> In this regime, judgment developed as a costless byproduct of production, in the tradition of Arrow (1962)’s learning-by-doing. Figure 1 (left panel) illustrates this self-sustaining career ladder.

---

<sup>1</sup>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.

<sup>2</sup>See, for example, Brynjolfsson, Chandar, and Chen (2025), Hosseini and Lichtinger (2025), Eloundou et al. (2024), Noy and Zhang (2023), among others. Section 1 details the evidence.

<sup>3</sup>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 displaces workers from tasks. Autor et al. (2024) extend this framework to study the creation of new work. Our paper differs in that we focus not on the displacement of tasks but on the destruction of the *developmental function* that task performance served.

<sup>4</sup>E.g., Loewenstein and Spletzer (1998) document that employers bear explicit costs of training; Barro, Berger, and Black (1999) provide evidence that training is associated with lower starting wages and faster subsequent productivity growth. Gao, Wang, and Wu (2024) shows that finance workers sustain low compensation in initial jobs to harvest greater salary jumps later.

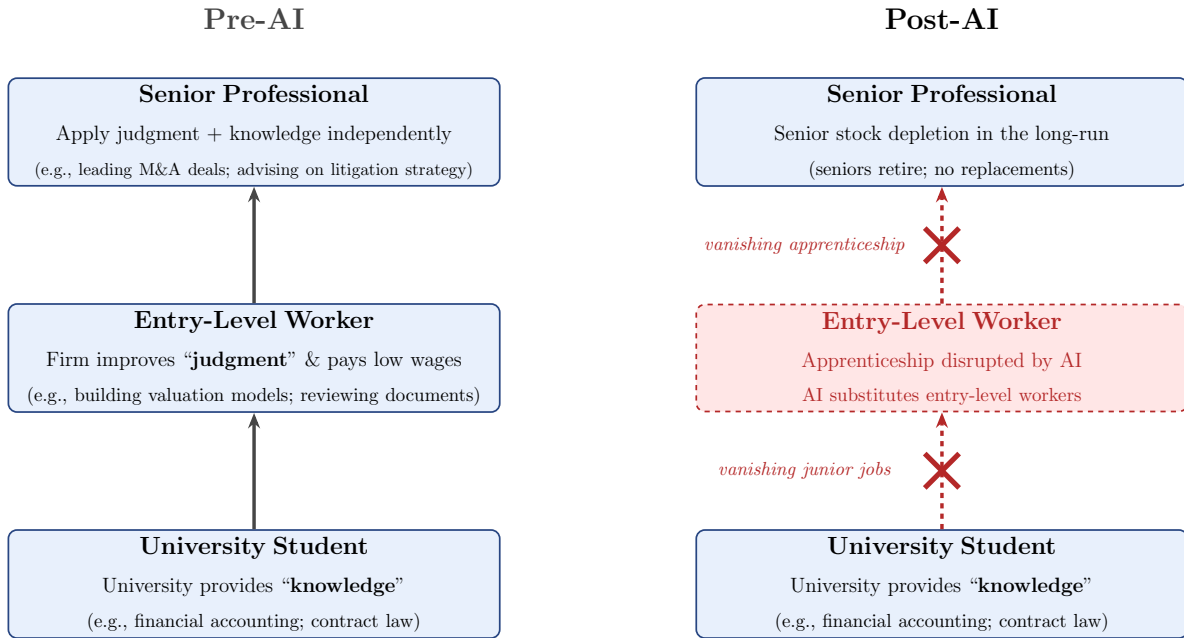


Figure 1: **AI Disrupting the Job Ladder.** The left panel illustrates the pre-AI career ladder in which university education (knowledge) and firm apprenticeship (judgment) combine to produce senior professionals. The right panel shows how generative AI, by cheaply substituting for entry-level cognitive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

The introduction of generative AI 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 breaks, the flow of new seniors dries up, and the stock of senior professionals erodes through retirement with no replacements. The long-run depletion of senior talent is a market failure: 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.

This market failure is self-reinforcing. 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 consequences of a depleted pipeline materialize well beyond the typical executive’s compensation horizon.<sup>5</sup> Moreover, firms cannot use AI to fully displace senior expertise. Unlike knowledge, judgment

<sup>5</sup>The poaching concern is less acute for purely firm-specific skills. However, the judgment developed through professional apprenticeship is largely *transferable*: Neal (1995) provides evidence that human capital is predominantly industry-specific; Gathmann and Schönberg (2010) show that skills are “more portable than previously considered”; Gao, Wang, and Wu (2024) document this pattern in finance.

is forged in undocumented, real-time experience—body language under pressure, boardroom dynamics that never appear in any memo—signals that large language models cannot learn from because they were never written down.

While firms cannot resolve this failure alone, our model implies that universities can. Because universities operate on longer planning horizons than firms, they are uniquely positioned to serve as *mechanism designers*. The paper proposes a micro-residency mechanism in which the university designs structured professional experiences—modeled on the medical residency, clinical legal education, and the 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. The university effectively re-bundles production and learning—the very bundling that AI destroyed—through deliberate institutional design.

Our model features a dynamic game among  $N$  symmetric firms, a university, and a continuum of junior and senior workers. The university 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. Labor market frictions follow Becker (1964) and Acemoglu and Pischke (1999): trained seniors separate from their training firms with some probability, and destination firms acquire fully developed workers without bearing training cost. Managerial myopia captures executives whose compensation is tied to near-term performance systematically undervaluing training investments whose returns materialize over a decade or more.

The model yields sharp results. When firms are patient, the equilibrium features underinvestment relative to the social optimum—because each firm anticipates losing some trained seniors to competitors—but the pipeline functions. When firms are myopic, the poaching distortion and myopia compound multiplicatively: training collapses, and in the extreme, the senior stock erodes to zero. A standard corrective training levy can restore the optimum when firms are patient but fails when firms are myopic: the agents whose short-termism creates the need for intervention are precisely those who refuse to participate. 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 participation constraint reduces to a static inequality (deliverable value exceeds participation fee), independent of the discount factor. Firms participate because the immediate transaction is profitable—the economic logic of the teaching hospital, where hospitals accept

residents because residents produce clinical services exceeding supervision costs.

We rank the equilibrium regimes: the social planner’s optimum dominates the patient-firm Nash equilibrium (poaching-driven underinvestment), which dominates the micro-residency equilibrium (second-best), which strictly dominates the myopic-Nash collapse. The welfare ordering  $W^* > W^{\text{NE}} > W^R > W^{\text{MN}}$  highlights that the micro-residency is the *only* mechanism that works when firms are too myopic to accept the levy.

Our model relies on some assumptions. First, we assume firms continue to need senior professionals whose judgment cannot be replicated by AI. Second, workers are passive recipients of firm training decisions, which may overstate the severity of the market failure.<sup>6</sup> Third, we model firms as symmetric, whereas in practice firms differ in size and training capacity.

The paper contributes to two strands of literature. First, we build on classical human capital theory—Becker (1964), Ben-Porath (1967), Arrow (1962), Acemoglu and Pischke (1998, 1999)—to study the AI-driven destruction of apprenticeship, deriving a new prescription: the solution lies not in firms internalizing the training externality but in redesigning university education to substitute for the apprenticeship function.<sup>7</sup> Second, we contribute to the fast expanding literature on AI and labor markets, including empirical discovery of AI’s biased impact on different jobs (Brynjolfsson, Chandar, and Chen, 2025; Hosseini and Lichtinger, 2025; Eloundou et al., 2024), task-level productivity effects of generative AI (Brynjolfsson, Li, and Raymond, 2025; Noy and Zhang, 2023; Peng et al., 2023; Dell’Acqua et al., 2023), macroeconomic frameworks for assessing AI’s aggregate impact (Acemoglu, 2024; Jones, 2026; Aghion, Jones, and Jones, 2018), the microeconomics of AI as prediction technology (Gans, 2025a,b), AI’s erosion of human learning incentives and long-run knowledge accumulation (Acemoglu, Kong, and Ozdaglar, 2026), and the case for steering AI development toward pro-worker applications (Acemoglu, Autor, and Johnson, 2026; Acemoglu and Restrepo, 2019, 2022; Autor et al., 2024). Our paper provides a theoretical framework explaining *why* entry-level displacement is structurally different from previous automation waves and offers *actionable guidance* on institutional reform. Our work thus complement recent studies

---

<sup>6</sup>Allowing workers to co-invest would introduce a two-sided investment problem (Acemoglu and Pischke, 1999). Our single-sided formulation likely overstates underinvestment severity but does not change qualitative predictions.

<sup>7</sup>See Mincer (1974) for the foundational empirical treatment of on-the-job training and Heckman (1998) for a review of returns to human capital investment across the life cycle.

## 1. 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.** 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 remained stable. The decline is concentrated in occupations where AI *automates* rather than *augments* tasks—precisely the routine cognitive work constituting the apprenticeship substrate in our model. Hosseini and Lichtinger (2025), analyzing 62 million résumés across 285,000 U.S. firms, report that AI-adopting firms reduced junior hiring by 9–10% within six quarters, driven entirely by reduced hiring rather than increased separations. Critically, senior hiring was unaffected, suggesting AI displaces the *entry point* of the career ladder. Eloundou et al. (2024) estimate that over 46% of U.S. jobs could be substantially affected by LLM-complementary software, with highest exposure in white-collar, cognitively routine occupations.

**Task-level productivity.** Controlled studies consistently show that generative AI substantially raises productivity on routine cognitive tasks. Noy and Zhang (2023) find that ChatGPT access reduces task completion time by 40% and raises output quality by 18% on mid-level writing tasks, with largest gains for lower-ability workers. Brynjolfsson, Li, and Raymond (2025) find a 15% average productivity increase from a generative AI assistant across 5,172 customer-support agents, with 30% improvement for novices but minimal impact on experienced agents. Peng et al. (2023) report that software engineers using GitHub Copilot complete coding tasks 55% faster. The consistent finding is that AI compresses the productivity distribution by disproportionately augmenting junior workers—the workers whose tasks our model identifies as the learning-by-doing substrate.<sup>8</sup>

**Organizational consequences.** 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 through which

---

<sup>8</sup>Dell’Acqua et al. (2023) document that AI improved performance on tasks within the “technological frontier” but worsened outcomes on tasks requiring judgment beyond AI’s capabilities, consistent with our knowledge–judgment distinction.

novices develop expert judgment. [Walther \(2025\)](#) highlights the convergence of AI adoption with senior retirements, creating a pipeline crisis from both ends.

These pieces of evidence 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 investing in alternative developmental pathways.

## 2. The Model

We develop the model in four steps: the environment, the pre-AI economy, the AI shock, and the post-AI training game.

### 2.1 Environment

Consider an infinite-horizon discrete-time economy with  $N \geq 2$  identical firms, one university, and a continuum of workers of two types: *juniors* ( $J$ ) and *seniors* ( $S$ ).<sup>9</sup> A senior produces output  $y^S > 0$  at wage  $w^S$ , yielding net surplus  $\sigma \equiv y^S - w^S > 0$ . Seniors retire each period with probability  $\delta \in (0, 1)$ . The social discount factor is  $\beta \in (0, 1)$ ; we introduce myopic discounting ( $\beta^M < \beta$ ) when analyzing managerial short-termism in Section 3.3.<sup>10</sup>

We denote by  $v \equiv \sigma/[1 - \beta(1 - \delta)]$  the present discounted value of the net surplus from one senior, accounting for stochastic retirement.

**The university.** The university screens graduates and builds foundational human capital. We represent both functions through a scalar quality parameter  $q_0 > 0$ , exogenous to firms.<sup>11</sup> The role of  $q_0$  is to serve as a *complementary input* to on-the-job development: graduates with stronger preparation learn faster from the same work experience, consistent with the dynamic complementarity of [Cunha and Heckman \(2007\)](#).

**Labor market frictions.** Each newly matured senior separates from the training firm with probability  $\lambda \in [0, 1)$ , driven by non-wage factors. In symmetric equilibrium, each

---

<sup>9</sup>The binary skill structure follows the standard approach in the training literature ([Acemoglu and Pischke, 1999](#)). The qualitative results extend to continuous skill accumulation, provided the externality structure is preserved.

<sup>10</sup>We model myopia in reduced form. This can be microfounded as an executive replaced with probability  $p$  per period having effective discount factor  $\beta^M = \beta(1 - p)$ , or as a manager whose bonus depends only on current-year earnings applying  $\beta^M \approx 0$  beyond the bonus horizon.

<sup>11</sup>This reflects that firms interact with universities primarily through the labor market for graduates. We return to the university's *endogenous* role as mechanism designer in Section 3.5.

firm’s outflow equals its expected inflow, so the *net* senior stock is unaffected by turnover. The distortion operates through *training costs*: the origin firm invested  $t_i$  per junior to develop the departing senior, while the destination firm acquires a trained senior at zero cost. This is the classical externality of [Becker \(1964\)](#), analyzed in [Acemoglu and Pischke \(1999\)](#): each firm captures only fraction  $(1-\lambda)$  of the return on training, leading to systematic underinvestment.

## 2.2 The Pre-AI Economy: Learning-by-Doing

Before AI, routine cognitive tasks were performed by junior professionals. These tasks had two properties central to our analysis. First, juniors were *productive*: their output  $y^J$  exceeded their wage  $w^J$ , generating positive net surplus  $\tau^0 \equiv y^J - w^J > 0$  for the firm—the *wage compression* condition of [Acemoglu and Pischke \(1999\)](#). Second, performing these tasks *developed expertise*: the act of reviewing hundreds of contracts built pattern recognition; building financial models developed intuition to question assumptions. This is *learning-by-doing* in the sense of [Arrow \(1962\)](#).

The junior’s per-period probability of maturing into a senior is:

$$m^{\text{pre}} = (\gamma \cdot y^J) \cdot q_0^b, \tag{1}$$

where  $\gamma > 0$  is the learning rate and  $b \in (0, 1)$  captures the Cunha-Heckman complementarity. Three features of this regime merit emphasis. *No explicit training decision*: maturation depends on productive task volume and graduate quality, neither of which is a firm choice variable. *No coordination problem*: because development was bundled with production at zero marginal cost, the poaching externality imposed no training-cost distortion. *A self-sustaining pipeline*: the pre-AI career ladder maintained itself automatically, requiring no deliberate institutional design.

## 2.3 The AI Shock

Generative AI performs routine cognitive tasks at cost  $c^{\text{AI}} < w^J$  and quality  $y^{\text{AI}} \geq y^J$ . Firms reassign these tasks from juniors to AI by elementary cost minimization. The consequences follow endogenously:

**Consequence 1: Junior surplus collapses.** With routine tasks reassigned,  $y^J \rightarrow 0$  and  $\tau = y^J - w^J \leq 0$ . The Acemoglu-Pischke mechanism is destroyed.<sup>12</sup> Senior tasks—advising under ambiguity, reading counterparties, exercising irreversible judgment—depend on undocumented experiential knowledge not in AI training data; hence  $\sigma$  is unaffected.

**Consequence 2: Learning-by-doing goes dormant.** From (1), with  $y^J \rightarrow 0$ , the maturation rate  $m^{\text{pre}} \rightarrow 0$ . The learning technology is unchanged—a junior who reviews contracts *would* still develop pattern recognition—but no firm will assign the work because AI is cheaper.

**Consequence 3: Developing juniors requires costly investment.** If a firm wishes to develop a junior, it must assign tasks AI could handle more cheaply, purely for developmental purposes. Let  $t_i \geq 0$  denote firm  $i$ 's per-junior training investment (the opportunity cost of deploying a junior instead of AI). The maturation rate becomes:

$$m^{\text{post}}(t_i) = (\alpha \cdot t_i^a) \cdot q_0^b, \quad (2)$$

where  $\alpha > 0$  is a productivity parameter and  $a \in (0, 1)$  captures diminishing returns. The learning process is identical; what has changed is its *price*—every hour of developmental experience now carries an opportunity cost that did not exist pre-AI. The broken ladder is fundamentally a problem of prices, not capabilities.

**Consequence 4: The poaching externality activates.** Post-AI, the firm has invested  $t_i \cdot J_i$  in developing juniors. A fraction  $\lambda$  of newly matured seniors depart, representing a direct loss. The Beckerian externality, dormant when training was bundled with production, is now first-order.

## 2.4 The Post-AI Training Game

Each firm chooses how many juniors to hire ( $J_i$ ) and how intensively to train them ( $t_i$ ). The senior stock evolves as:

$$S_i^{t+1} = (1 - \delta) S_i^t + (1 - \lambda) m^{\text{post}}(t_i) J_i. \quad (3)$$

---

<sup>12</sup>The wage  $w^J$  may adjust downward, but institutional frictions prevent it from falling as fast as  $y^J$ , so  $\tau$  turns negative or approaches zero.

Per-period profit is  $\pi_i^t = \sigma \cdot S_i^t + \tau \cdot J_i - t_i \cdot J_i$ . The value function takes the affine form  $V(S_i) = v \cdot S_i + K$  (see Supplemental Appendix A for verification). The firm's steady-state present value reduces to:

$$\Pi_i = v \cdot (1 - \lambda) \cdot \alpha t_i^a q_0^b \cdot J_i + \frac{(\tau - t_i) \cdot J_i}{1 - \beta}. \quad (4)$$

### 3. Equilibrium Analysis

#### 3.1 The Social Planner's Benchmark

A social planner sets  $\lambda = 0$  (poaching is zero-sum in aggregate) and maximizes total welfare. The first-order condition yields the planner's optimum:

$$t^* = \left[ \frac{\sigma \alpha a q_0^b}{\delta} \right]^{\frac{1}{1-a}}. \quad (t^*)$$

**Proposition 1** (Cooperative Optimum). *The planner's problem has a unique interior solution  $t^* > 0$ . The optimal training level is increasing in  $\sigma$ ,  $\alpha$ , and  $q_0$ , and decreasing in  $\delta$ . The comparative static  $\partial t^* / \partial q_0 > 0$  implies that university quality and firm training are strategic complements at the social level.*

See proof in the Supplemental Appendix A.

#### 3.2 Decentralized Equilibrium: The Poaching Externality

Each firm's first-order condition with respect to  $t_i$  includes a  $(1 - \lambda)$  factor reflecting that the firm captures the return only if the senior stays. Solving:

$$t^{\text{NE}} = (1 - \lambda)^{\frac{1}{1-a}} \cdot t^*. \quad (t^{\text{NE}})$$

**Proposition 2** (Nash Equilibrium and Inefficiency). *The post-AI training game has a unique symmetric Nash equilibrium with  $t^{\text{NE}} < t^*$ . The welfare loss is strictly positive for  $\lambda > 0$ , increasing in  $\lambda$  and  $a$ , and decreasing in  $q_0$ : higher university quality partially compensates for firm-side underinvestment.*

See proof in the Supplemental Appendix A. Result (iii) 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 consequences.

### 3.3 Managerial Myopia

Each firm’s decision-maker uses  $\beta^M < \beta$ , reflecting executive turnover and near-term compensation structures. Under myopia, the firm values a marginal senior at  $v^M = \sigma/[1 - \beta^M(1 - \delta)] < v$ .

**Proposition 3** (Myopic Equilibrium). *The unique symmetric Nash equilibrium has training level*

$$t^{MN} = (1 - \lambda)^{\frac{1}{1-a}} \cdot \left(\frac{v^M}{v}\right)^{\frac{1}{1-a}} \cdot t^*. \quad (5)$$

*The two distortions compound **multiplicatively**: even moderate poaching ( $\lambda = 0.3$ ) and myopia ( $v^M/v = 0.5$ ) can reduce training to one-eighth of the social optimum for  $a = 0.5$ .*

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

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

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 (see Supplemental Appendix A for derivation).

### 3.4 Mechanism Design: The Training Levy

A *training levy mechanism* imposes a per-senior levy and reimburses firms for each junior that matures and is retained, offsetting the poaching externality.

**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.*

However, the levy requires firms to accept immediate cost for long-run benefit:

**Proposition 5** (Participation Failure Under Myopia). *If offered voluntarily, a firm with discount factor  $\beta^M$  participates if and only if  $\beta^M \geq \bar{\beta}$ . For  $\beta^M < \bar{\beta}$ , the firm’s myopic valuation of future senior services is insufficient to justify the levy cost.*

This reveals a fundamental asymmetry: the agents whose myopia necessitates intervention are the same agents who refuse to participate. Proofs are in the Supplemental Appendix A.

### 3.5 The University as Mechanism Designer

The university has a structural advantage: its planning horizon far exceeds the typical firm’s. We now show that this asymmetry allows the university to design a mechanism that succeeds where the Pigouvian levy fails.

The key insight is that AI destroyed the *substrate* for learning-by-doing. The university provides a *synthetic substitute*: structured professional experiences in which juniors work on real problems under senior supervision, generating tangible deliverables while developing judgment as a byproduct. This is the logic of the medical residency and clinical legal education, 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  $v^R > 0$ ; 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 condition  $v^R \geq f$  requires that supervised students produce deliverables whose value covers the fee—satisfied in settings like teaching hospitals, law clinics, and consulting practica, but an empirical condition.

**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  $\beta^M \in (0, 1)$ , any beliefs about other firms’ actions, and any realization of the poaching shock, each firm weakly prefers participation.*

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

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

Both terms are non-negative when  $v^R \geq f$ , so  $\pi_i^R \geq 0$ . This involves no discount factor, no beliefs about other firms, and no realization of  $\lambda$ . Participation is a dominant strategy.  $\square$

*Remark.* Proposition 6 is the paper’s central result. 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 residents not because they are altruistic about medical education, but because residents produce clinical services exceeding supervision costs.

### 3.6 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.*

See proof in the Supplemental Appendix A. Panel A of Table 1 summarizes the model’s inner-working.

Table 1: Equilibrium regimes and welfare ranking.

<b>Panel A: Model Insights</b>					
Regime	Training	Senior Stock	Welfare	Key Friction	
Social planner	$t^*$	$S^*$	$W^*$	None	
Nash (patient firms)	$t^{NE} < t^*$	$S^{NE} < S^*$	$W^{NE}$	Poaching ( $\lambda$ )	
Micro-residency	$t^R$ (embedded)	$S^R \in (0, S^*)$	$W^R$	Poaching (mitigated)	
Myopic Nash	$t^{MN} \ll t^*$	$S^{MN} \rightarrow 0$	$W^{MN}$	Poaching $\times$ myopia	
<b>Panel B: Numerical Illustration</b>					
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%

Panel B provides a numerical illustration using  $\lambda = 0.30$ ,  $a = 0.5$ ,  $\delta = 0.10$ ,  $\beta = 0.95$ ,  $\beta^M = 0.60$ ,  $\sigma = 1$ ,  $\alpha = 0.20$ ,  $q_0 = 1$ ,  $\tau = -0.10$ , and  $J = 1$  (derivations in the Supplemental Appendix B). The poaching externality alone destroys ten percent of first-best welfare. Combined with myopia, training falls to five percent of the optimum, leaving the senior stock at one-sixth of its efficient level. 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  $\beta^M$  is low.

## 4. Implications for University Education

### 4.1 What Universities Should Do

Proposition 6 shows that the micro-residency succeeds where market forces and Pigouvian intervention fail, because it 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 judgment development with tangible deliverables: embedded consulting projects under faculty supervision, AI-audit residencies, and cross-functional simulations replicating 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 post-AI: its pre-AI function (screening and foundational human capital) persists, but the comparative static  $\partial t^*/\partial q_0 > 0$  implies an additional role—higher  $q_0$  raises the marginal return to whatever firm-side training does occur. Investments in university quality are thus more socially valuable post-AI than pre-AI.

### 4.2 The Role of Government Funding

As of 2023, U.S. university revenues derive primarily from government funding, with state, federal, and local sources accounting for more than 50% of all university funding ([Visual Capitalist](#)). When firms are patient, the training levy implements the first-best. When firms are myopic, the levy must be mandatory. The optimal levy rate  $r^* = \lambda \cdot v \cdot m^{\text{post}}(t^*)$  provides a concrete calibration target estimable from observable quantities (senior wages, turnover rates, maturation durations). Government can also mandate 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.

### 4.3 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.

#### 4.4 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.<sup>13</sup>

**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 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 rela-

---

<sup>13</sup>See [Dreyfus and Dreyfus \(1986\)](#) on the role of consequential decision-making in expert development.

tively 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.

## 5. Conclusion

Generative AI is not only displacing jobs during the transition period but also substantially disrupting the developmental infrastructure through which novices become experts. This paper identifies the mechanism: AI makes junior labor uncompetitive relative to the technology in routine cognitive tasks, prompting firms to reassign that work and thereby destroying the Arrow learning-by-doing channel that sustained professional career ladders at zero marginal training cost. As a result, the economy is forced from a regime of costless, bundled human capital formation into one of costly, explicit Ben-Porath investment, under labor-market frictions that activate the classical Beckerian poaching externality.

The coordination failure 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.

## References

- Acemoglu, D. (2024). The simple macroeconomics of AI. NBER Working Paper No. 32487.
- 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., D.H. Autor, and S. Johnson (2026). Building pro-worker artificial intelligence. NBER Working Paper No. 34854.
- Acemoglu, D., D. Kong, and A. Ozdaglar (2026). AI, human cognition and knowledge collapse. NBER Working Paper No. 34910.
- 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.
- Aghion, P., B.F. Jones, and C.I. Jones (2018). Artificial intelligence and economic growth. In A.K. Agrawal, J. Gans, and A. Goldfarb (eds.), *The Economics of Artificial Intelligence: An Agenda*, pp. 237–282. University of Chicago Press.
- 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.
- 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.
- Brynjolfsson, E., D. Li, and L. Raymond (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2):889–942.
- 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. Kraymer, 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.
- Gans, J. S. (2025). *The Microeconomics of Artificial Intelligence*. MIT Press.
- Gans, J. S. (2025). Growth in AI knowledge. NBER Working Paper No. 33907.
- 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.
- Jones, C. I. (2026). A.I. and our economic future. NBER Working Paper No. 34779.
- 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.

# Supplemental Appendix

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

Miao Ben Zhang

### A. Proofs and Technical Details

#### A.1 Proof of Proposition 1 (Cooperative Optimum)

The planner's objective 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.  $\square$

#### A.2 Proof of Proposition 2 (Nash Equilibrium)

From (4), firm  $i$ 's FOC with respect to  $t_i$  is:

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

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

#### 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  $(t^{\text{NE}})$  yields (5). The multiplicative structure follows from the power-function form of the solution.  $\square$

#### A.4 Proof of Corollary 2 (Pipeline Shutdown)

From (4), the FOC for  $J_i$  is  $v(1 - \lambda)m^{\text{post}}(t_i) + (\tau - t_i)/(1 - \beta) = 0$ . The first term is positive; the second is negative for  $\tau < t_i$ . At  $\bar{\tau} = t^{\text{MN}} - v^M(1 - \beta^M)(1 - \lambda)m^{\text{post}}(t^{\text{MN}})$ , the two terms cancel and  $J_i = 0$  is optimal.  $\square$

### A.5 Proof of Proposition 4 (Training Levy)

Under the levy, firm  $i$ 's modified payoff includes a reimbursement  $r(1 - \lambda)m^{\text{post}}(t_i)J_i$  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  $\sigma/\delta$ , replicating the planner's FOC.  $\square$

### 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  $v^M(1 - \lambda)m^{\text{post}}(t^*)J/\delta \geq \tau^{L^*}S^*/(\sigma - \tau^{L^*})$ . Solving for the critical  $\beta^M$  yields  $\bar{\beta}$ ; for  $\beta^M < \bar{\beta}$ , autarky dominates.  $\square$

### A.7 Supermodular Game Structure

Define the reduced game  $\hat{\Gamma}$  in which each firm chooses  $t_i \in \mathbb{R}_+$ . In symmetric equilibrium, the game reduces to  $N$  identical single-agent problems. We verify supermodularity:  $\partial^2 \Pi_i / \partial t_i \partial J_i \propto \alpha a t_i^{a-1} q_0^b > 0$ , confirming that training and hiring are complements. Existence and uniqueness then follow from Tarski's fixed-point theorem (Topkis, 1998) and the contraction property of the best-response mapping.  $\square$

### A.8 Proof of Proposition 7 (Welfare Ordering)

The welfare function under the planner's objective (treating poaching as zero-sum in aggregate) is:

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

Since  $W(t)$  is strictly concave in  $t$  and  $t^*$  is the unique maximizer,  $W(t^*) > W(t)$  for any  $t \neq t^*$ . The ranking  $t^* > t^{\text{NE}} > t^R > t^{\text{MN}}$  (from Propositions 2–3 and the definition of  $t^R$ ) and the concavity of  $W$  yield  $W^* > W^{\text{NE}} > W^R > W^{\text{MN}}$ .  $\square$

## B. Numerical Illustration: Derivations

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

**Discount-adjusted senior values.**

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

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.**

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

**Nash equilibrium (patient firms).**

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

**Myopic Nash equilibrium.**

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

**Micro-residency.** The university sets  $t^R = 0.25$ , a level that myopic firms would never choose voluntarily but participate in because  $v^R \geq f$ .

**Maturation rates.**  $m(t) = 0.20 \cdot t^{0.5}$ :

$$m(t^*) = 0.200, \quad m(t^{NE}) = 0.140, \quad m(t^R) = 0.100, \quad m(t^{MN}) \approx 0.044.$$

**Steady-state senior stocks.** From  $\delta S = (1 - \lambda) m(t) J$ :

$$S^* = 2.00, \quad S^{NE} = 0.98, \quad S^R = 0.70, \quad S^{MN} = 0.31.$$

**Welfare.**  $W(t) = 2t^{0.5} - 0.10 - t$ :

$$W^* = 0.90, \quad W^{NE} = 0.81, \quad W^R = 0.65, \quad W^{MN} \approx 0.30.$$

The ordering  $W^* > W^{NE} > W^R > W^{MN}$  is confirmed.