

Research Article

Vertical Integration in the Academic Production Chain: A Hold-Up Theory of Inappropriate Advisor–Student Relationships

Zuzheng Wang¹

1. Independent researcher

We model the advisor–student relationship as a bilateral production relationship with relationship-specific investments, non-contractible division of academic credit, and asymmetric control rights. We introduce a continuous “vertical integration” parameter θ capturing the degree to which the advisor absorbs the student’s personal autonomy—romantic involvement, personal dependence, social isolation—and show that integration operates economically by degrading the student’s outside option in ex-post bargaining. Three results follow. First, even absent any direct taste for integration, the advisor’s privately optimal credit share strictly exceeds the surplus-maximizing share, so integration is privately rational and socially inefficient; the loss is transmitted through the student’s rational underinvestment. Second, comparative statics show that equilibrium integration is decreasing in the student’s outside option and in expected sanctions, generating the testable prediction that exploitation concentrates on students with the weakest exit options. Third, in a two-period extension without commitment, the advisor’s extraction motive is time-inconsistent under essentially arbitrary production technologies—the result uses only revealed preference, with no concavity, differentiability, or functional form: once the student’s specific investment is sunk, the advisor escalates integration until the student’s interim participation constraint binds—a ratchet. The student rationally remains in arrangements she would have rejected ex ante, because her exit option has been endogenously destroyed; anticipating this, she underinvests to a level pinned down solely by the salvage value of her human capital—the production function drops out entirely—so the advisor himself can gain from institutional commitment devices. The model rationalizes portable funding, mandatory co-advising, credit-verification regimes, and enforcement as instruments that operate on threat points rather than on preferences.

1. Introduction

The advisor–graduate–student relationship combines three features that contract theory identifies as generative of opportunism: both parties make investments that are far more valuable inside the match than outside it; the division of the joint output—authorship order, credit for ideas, patent shares—cannot be contracted upon *ex ante* and is settled by *ex-post* bargaining; and the control rights over the relationship’s essential assets (funding, laboratory access, recommendation letters, approval of graduation) rest almost entirely with one party. We take these features seriously and ask what they imply about a phenomenon usually discussed in purely ethical terms: inappropriate personal relationships between advisors and students, including romantic involvement and broader forms of personal dependence.

Our central modeling device is to treat such entanglement as *vertical integration* of the student into the advisor’s domain of control. The defining economic effect of integration, we argue, is not on production but on *threat points*: a student romantically involved with or personally dependent on her advisor faces far higher exit costs, and a degraded exit option translates directly into a smaller share of the joint surplus in *ex-post* bargaining. Integration converts the student’s outside option into a hostage.

Three results organize the paper. (i) In the static model, the advisor’s privately optimal degree of integration is strictly positive even when he derives no direct utility from it, while total surplus is maximized at balanced credit shares; the wedge is a standard hold-up externality (Section 2). (ii) Comparative statics deliver sharp, testable predictions: integration falls with the student’s outside option and with expected sanctions (Section 3). (iii) In a two-period version without commitment, extraction is time-inconsistent: *ex ante* the advisor moderates because he must elicit the student’s investment, but once that investment is sunk his payoff is globally increasing in the credit share, so he escalates until the student’s interim participation constraint binds (Section 4). The escalation result formalizes the “ratchet” or entrapment dynamic and shows that “she could have left” is analytically wrong: remaining is sequentially rational precisely because the exit option has been endogenously destroyed.

1.1. Related literature

The framework belongs to the property-rights theory of the firm ^{[1][2]} and the hold-up literature ^{[3][4]}. The interim-expropriation logic echoes the ratchet effect in dynamic incentive problems ^{[5][6]} and the

literature on specific human capital and employer monopsony power ^{[7][8]}. Reference-point and grievance effects ^[9] offer a complementary channel we abstract from. On the application side, the paper relates to the economics of science and credit allocation ^{[10][11]} and to empirical work on harassment and power in academia (e.g., the NASEM 2018 consensus report).

2. The static model

2.1. Primitives

Assumption 1 (Agents). *Two risk-neutral agents, an advisor A and a student P , maximize expected payoffs. There is no discounting within the static model.*

Assumption 2 (Production). *Joint academic output has publication value $R = R(a, p)$, where $a \geq 0$ is the advisor's investment (guidance, ideas, reputation, funding access) and $p \geq 0$ is the student's investment (research effort, match-specific human capital). R is strictly increasing, strictly concave, and supermodular ($R_{ap} > 0$), with $R(0, p) = R(a, 0) = 0$.*

Assumption 3 (Specificity). *Both investments are relationship-specific: their value outside the match is a fraction $\delta \in [0, 1]$ of their inside value (δ plays a role only in the dynamic extension; in the static model exit simply yields the exogenous outside options below).*

Assumption 4 (Contractual incompleteness). *The division of R (authorship order, credit, patent shares) is not contractible ex ante and is determined by ex-post bargaining. Investments a, p are mutually observable but not verifiable by third parties.*

Assumption 5 (Control rights). *The advisor holds residual control rights over the relationship's assets: funding, laboratory access, recommendation letters, and approval of graduation.*

Assumption 6 (Costs). *Investment costs are $c_A(a) = a^2/2$ and $c_P(p) = p^2/2$.*

2.2. Vertical integration

Let $\theta \in [0, \bar{\theta}]$ denote the degree of vertical integration of the student into the advisor's personal domain: $\theta = 0$ is an arm's-length professional relationship; larger θ encompasses escalating personal dependence, social isolation within the group, and romantic or marital entanglement.

Integration affects the environment in three ways.

(i) *Threat points.* The student's disagreement payoff in ex-post bargaining is $d_P(\theta) = (1 - \theta)\bar{d}_P$, with $\bar{d}_P > 0$ her baseline outside option; the advisor's disagreement payoff $\bar{d}_A$ is unaffected (he can re-staff the project). With generalized Nash bargaining (advisor weight β), the advisor's realized share of R is

$$\lambda = \frac{\bar{d}_A + \beta [R - \bar{d}_A - (1 - \theta)\bar{d}_P]}{R}, \frac{\partial \lambda}{\partial \theta} = \frac{\beta \bar{d}_P}{R} > 0. \quad (1)$$

For tractability we work with the reduced form

$$\lambda(\theta) = \lambda_0 + \gamma\theta, \lambda_0 \in (0, 1), \gamma > 0, \lambda(\bar{\theta}) < 1, \quad (2)$$

where λ_0 is the share implied by professional norms and γ is the extraction rate of integration.

(ii) *Private benefit.* The advisor derives a direct non-pecuniary benefit $B(\theta)$, increasing and concave with $B(0) = 0$. This captures the consumption value of the personal relationship.

(iii) *Student cost.* The student bears $\kappa(\theta)$, increasing and convex with $\kappa(0) = 0$: psychological cost, career risk, stigma.

Expected sanctions are $\pi(\theta)F$, with π increasing convex, $\pi(0) = 0$: detection probability times the penalty (dismissal, legal and reputational costs), borne by the advisor.

2.3. Timing and payoffs

1. **Offer.** The advisor offers an (implicit) arrangement θ ; the student accepts or takes her ex-ante outside option $\bar{u}_P$.
2. **Investment.** Both simultaneously choose a, p .
3. **Bargaining.** $R(a, p)$ is realized; shares follow (2).

Payoffs are

$$U_A = \lambda(\theta)R(a, p) - \frac{a^2}{2} + B(\theta) - \pi(\theta)F, \quad (3)$$

$$U_P = (1 - \lambda(\theta))R(a, p) - \frac{p^2}{2} - \kappa(\theta). \quad (4)$$

2.4. Equilibrium with Cobb–Douglas production

Let $R = 2a^{1/2}p^{1/2}$.

Proposition 1 (Investment equilibrium). *Given $\lambda \in (0, 1)$, the unique Nash equilibrium of the investment subgame is*

$$a^*(\lambda) = \lambda^{3/4}(1-\lambda)^{1/4}, p^*(\lambda) = (1-\lambda)^{3/4}\lambda^{1/4},$$

yielding realized output and continuation payoffs

$$R^* = 2\sqrt{\lambda(1-\lambda)}, S^*(\lambda) \equiv R^* - \frac{(a^*)^2 + (p^*)^2}{2} = \frac{3}{2}\sqrt{\lambda(1-\lambda)},$$

$$u_A(\lambda) = \frac{3}{2}\lambda^{3/2}(1-\lambda)^{1/2}, u_P(\lambda) = \frac{3}{2}(1-\lambda)^{3/2}\lambda^{1/2}.$$

Proposition 2 (Private vs. social optimum). (i) Joint surplus $S^*(\lambda)$ is single-peaked with maximum at $\lambda = 1/2$: balanced credit balances the two hold-up problems. (ii) The advisor's continuation payoff $u_A(\lambda)$ is single-peaked with maximum at $\lambda = 3/4$. (iii) Hence for any $\lambda_0 < 3/4$, and even with $B \equiv 0$, $\pi \equiv 0$, the advisor strictly prefers $\theta > 0$; every unit of extraction beyond $\lambda = 1/2$ destroys surplus, the loss being transmitted through the student's underinvestment, $\partial p^* / \partial \theta = \gamma p^{*'}(\lambda) < 0$ for $\lambda > 1/4$.

The advisor's problem is

$$\max_{\theta \in [0, \bar{\theta}]} u_A(\lambda(\theta)) + B(\theta) - \pi(\theta)F \quad \text{s.t.} \quad u_P(\lambda(\theta)) - \kappa(\theta) \geq \bar{u}_P. \quad (\text{PC})$$

3. Comparative statics

Two regimes arise, indexed by whether (PC) binds.

Proposition 3 (Unconstrained regime). Suppose (PC) is slack at the optimum and the objective in (PC) is strictly concave in θ . The interior optimum θ^* solves $\gamma u'_A(\lambda(\theta^*)) + B'(\theta^*) = \pi'(\theta^*)F$, and

$$\frac{\partial \theta^*}{\partial F} < 0, \frac{\partial \theta^*}{\partial B'} > 0.$$

With $B \equiv 0$ and $\pi \equiv 0$, the chosen share is $\lambda(\theta^*) = \min \{3/4, \lambda_0 + \gamma \bar{\theta}\}$.

Proposition 4 (Binding-PC regime). Suppose (PC) binds, so θ^* solves $u_P(\lambda(\theta)) - \kappa(\theta) = \bar{u}_P$ with $\lambda(\theta^*) > 1/4$. Then

$$\frac{\partial \theta^*}{\partial \bar{u}_P} < 0, \frac{\partial \theta^*}{\partial \lambda_0} < 0, \frac{\partial \theta^*}{\partial \gamma} < 0, \frac{\partial \theta^*}{\partial k} < 0 \text{ for } \kappa(\theta) = k\theta^2.$$

Corollary 1 (Selection into exploitation). Equilibrium integration is decreasing in the student's outside option. Exploitation therefore concentrates on students with the weakest exit options—those with non-portable funding, visa constraints, thin external networks, or in fields dominated by a single laboratory.

Remark 1 (Welfare). For $\lambda(\theta) > 1/2$, $dS^*/d\theta = \gamma S^{*'}(\lambda) < 0$: the deadweight loss $S^*(1/2) - S^*(\lambda(\theta))$ is strictly increasing in θ . Extraction past balanced shares destroys strictly more surplus than it transfers.

4. The two-period ratchet

4.1. Setup

There are two dates. At $t = 1$ the advisor offers θ_1 and, upon acceptance, the student sinks p at cost $c_P(p)$. At $t = 1.5$ the advisor proposes an escalation $\theta_2 \geq \theta_1$ on a take-it-or-leave-it basis; if the student rejects, she exits with interim outside payoff

$$w(\theta_1, p) = \omega + \delta\psi(p) - x\theta_1, \quad (5)$$

where ω is her raw interim option, $\delta \in [0, 1]$ is the salvage rate on match-specific capital (ψ increasing, $\psi(0) = 0$), and $x \geq 0$ measures the direct exit cost of unwinding personal entanglement already accumulated at level θ_1 . At $t = 2$ the advisor invests a_2 , output $R(a_2, p)$ is realized, and shares follow $\lambda(\theta_2)$.

In this section we drop the regularity imposed in Assumption 2 and require only:

Assumption 7 (General technology). $R : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ is continuous and nondecreasing in each argument, with $R(a, 0) = 0$; c_A and c_P are continuous, strictly increasing, with $c_A(0) = c_P(0) = 0$, and c_A is coercive relative to R : for each p , $R(a, p) - c_A(a) \rightarrow -\infty$ as $a \rightarrow \infty$. No concavity, differentiability, supermodularity, or homotheticity of R is assumed.

Given (θ_2, p) , define the advisor's continuation value and (a selection of) his optimal investment,

$$v_A(\lambda; p) = \max_{a \geq 0} [\lambda R(a, p) - c_A(a)], a^*(\lambda, p) \in \arg \max_{a \geq 0} [\lambda R(a, p) - c_A(a)], \quad (6)$$

and the student's gross continuation value $v_P(\lambda; p) = (1 - \lambda)R(a^*(\lambda, p), p)$.

4.2. Time inconsistency under (almost) no assumptions

Lemma 1 (Time inconsistency, technology-free). Under Assumption 7, for every sunk p the value $v_A(\lambda; p)$ is nondecreasing in λ on $(0, 1]$, and strictly increasing at every λ at which optimal output is positive, $R(a^*(\lambda, p), p) > 0$. The result uses only revealed preference: it holds for non-concave, non-differentiable, non-supermodular technologies, for any tie-breaking among maximizers, and does not require first-order conditions to characterize a^* .

The contrast with Proposition 2 is the economics of the section: ex ante, the advisor moderates because he must elicit the student's investment; ex post, with p sunk, his payoff rises with his share for every

production function, so any moderation must come from the student's participation constraint or from institutions—never from the technology.

4.3. Dissipation and the general ratchet

Escalation stops where the student's interim participation constraint binds. For that stopping point to be well-behaved we need the student's continuation value, net of integration costs, to be eventually decreasing:

Assumption 8 (Dissipation). $\Lambda(\theta) \equiv v_P(\lambda(\theta); p) - \kappa(\theta)$ is strictly decreasing on the relevant range $[\theta_1, \theta_{\max}]$, where $\theta_{\max} = \sup \{\theta : \lambda(\theta) < 1\}$.

Assumption 8 is not innocuous—raising λ also improves the advisor's investment incentive, which benefits the student—so we give a sharp sufficient condition.

Lemma 2 (Sufficient conditions for dissipation). Let R and c_A be twice differentiable with an interior a^* . Then

$$\frac{\partial v_P}{\partial \lambda} = \frac{(1 - \lambda)R_a(a^*, p)^2}{c_A''(a^*) - \lambda R_{aa}(a^*, p)} - R(a^*, p), \quad (7)$$

where the denominator is positive by the second-order condition. In particular, for $R = Aa^\alpha p^\beta$ and $c_A(a) = a^{1+\eta}/(1 + \eta)$ with $\eta \geq 1$,

$$\frac{\partial v_P}{\partial \lambda} < 0 \Leftrightarrow \lambda > \frac{\alpha}{1 + \eta}.$$

Since $\alpha < 1$ and $\eta \geq 1$, any $\lambda_0 \geq 1/2$ places the entire range $[\theta_1, \theta_{\max}]$ in the dissipation region; and $\kappa' > 0$ relaxes the condition further.

Assumption 8 rules out exactly one economic channel, which is worth naming. Raising the advisor's share does two things to the student: it shrinks her slice of the pie one-for-one (the direct effect, the $-R$ term in (7)), and it may enlarge the pie by strengthening the advisor's incentive to invest (the indirect effect, the positive first term). Call the indirect effect *trickle-down engagement*: the idea that ceding more of the credit “activates” the advisor and thereby benefits the student. Dissipation says the direct effect dominates. The following examples show when—and, more importantly, where in the parameter space—this is the empirically relevant case.

Example 1 (The attention-saturated advisor). Suppose the advisor's usable input is capacity-constrained: $R(a, p) = F(\min\{a, \bar{a}\}, p)$ for some capacity $\bar{a} > 0$ and any technology F satisfying Assumption 7. The

capacity $\bar{a}$ is the advisor's finite attention, divided across many students, teaching, grant-writing, and administration. For λ large enough that $a^*(\lambda, p) = \bar{a}$ —the advisor is already fully engaged with this project—further increases in λ leave a^* unchanged, and (7) degenerates to

$$\frac{\partial v_P}{\partial \lambda} = -F(\bar{a}, p) < 0 \quad \text{exactly.}$$

Dissipation holds globally past saturation, for *any* F : once the binding margin is the advisor's attention rather than his willingness, extraction cannot buy engagement, so it is a pure transfer plus incentive damage to the student. In the benchmark calibration ($F = 2a^{1/2}p^{1/2}$, $\bar{a} = 0.4$, quadratic costs), saturation obtains for all $\lambda \gtrsim 0.5$ —that is, from the professional-norm baseline λ_0 onward, exactly the range the escalation path traverses. We regard this as the empirically typical case: the practical test of Assumption 8 is whether ceding additional credit to one's advisor visibly buys additional engagement, and for a busy advisor at capacity the answer is no.

Example 2 (Field calibration of the threshold). For unconstrained power technologies the boundary of the dissipation region is $\lambda = \alpha/(1 + \eta)$: the advisor's output elasticity relative to (one plus) the convexity of his effort costs. In student-effort-intensive experimental fields (bench science, data-heavy empirical work), the advisor's elasticity is modest, say $\alpha \in [0.3, 0.4]$, and his marginal time is steeply convex, $\eta \geq 1$: the threshold is at most 0.2, so every share above 0.2 lies in the dissipation region. In idea-driven theoretical fields with a dominant senior scholar, $\alpha \in [0.7, 0.8]$ and $\eta = 1$ put the threshold at 0.35–0.40. In both cases the threshold sits strictly below any professional-norm baseline $\lambda_0 \approx 1/2$, and indeed $\alpha/(1 + \eta) \leq \alpha/2 < 1/2$ for all $\alpha < 1$, $\eta \geq 1$: under quadratic-or-steeper costs, *no* calibration of a power technology places professional-norm shares outside the dissipation region.

Remark 2 (Trickle-down engagement does not rescue integration). The only region where Assumption 8 can fail is $\lambda < \alpha/(1 + \eta)$: a star advisor currently holding a *small* share, where ceding credit genuinely activates his effort and benefits the student. Two observations confine this regime to irrelevance for the ratchet. First, by Example 2 it lies strictly below λ_0 , so the escalation path $\lambda(\theta) \in [\lambda(\theta_1), 1)$ never enters it: escalation starts where trickle-down is already exhausted. Second, and more substantively, a student voluntarily conceding credit to activate a star coauthor is an arm's-length transaction operating through λ directly; it requires no degradation of her exit option and is not integration in the sense of this paper. What Assumption 8 rules out is therefore precisely the defense one hears of such arrangements—that deepening personal entanglement benefits the student “because the advisor invests more in her career.” Past the threshold, and always past professional-norm shares, that defense is quantitatively false: the

engagement it invokes is already saturated (Example 1) or already paid for (Example 2), and what remains is pure dissipation.

Proposition 5 (Ratchet, general form). *Under Assumptions 7 and 8, in the subgame-perfect equilibrium without commitment the advisor escalates to $\theta_2^* = \hat{\theta}(\theta_1, p)$, where $\hat{\theta}$ solves the binding interim participation constraint*

$$v_P(\lambda(\hat{\theta}); p) - \kappa(\hat{\theta}) = \omega + \delta\psi(p) - x\theta_1 \quad (8)$$

(or $\hat{\theta} = \theta_{\max}$ if no solution exists in the range). Escalation is strict, $\theta_2^* > \theta_1$, whenever the interim constraint is slack at θ_1 , i.e. whenever quasi-rents exist ($\delta < 1$ and/or $x > 0$). Moreover

$$\frac{\partial \hat{\theta}}{\partial \delta} < 0, \frac{\partial \hat{\theta}}{\partial \omega} < 0, \frac{\partial \hat{\theta}}{\partial x} > 0, \frac{\partial \hat{\theta}}{\partial \theta_1} > 0.$$

The last inequality is the ratchet complementarity: integration already in place lowers the exit option and thereby enables further integration. None of these signs depends on the production function.

Proposition 6 (Entrapment). *Let θ^{acc} denote the largest integration level the student would accept ex ante at a committed offer. For $x > 0$ or δ sufficiently small, there exist parameters such that $\theta_2^* = \hat{\theta}(\theta_1, p) > \theta^{\text{acc}}$ while the student's interim continuation payoff still (weakly) exceeds $w(\theta_1, p)$. Remaining in an arrangement she would have rejected ex ante is sequentially rational, because her exit option has been endogenously degraded by her own sunk investment and by θ_1 .*

4.4. Technology-free underinvestment and the demand for commitment

Rational expectations close the model: at $t = 1$ the student foresees (8), hence knows her period-2 payoff equals her exit value (5). Her investment problem is therefore $\max_p \omega + \delta\psi(p) - x\theta_1 - c_P(p)$, with solution

$$\delta\psi'(p^{\text{nc}}) = c'_P(p^{\text{nc}}). \quad (9)$$

Remark 3 (The production function drops out). Equation (9) contains no R , no $\lambda(\cdot)$, no κ , no γ : absent commitment, the student's investment is pinned down entirely by the salvage market for her human capital. However productive the collaboration could be, the technology is irrelevant to her incentive at the margin, because every quasi-rent it generates will be expropriated. This is the sense in which the cautionary conclusion is maximally general.

Under a binding commitment to $\theta_2 = \theta_1$, her investment p^c instead solves

$$(1 - \lambda_1) \frac{d}{dp} R(a^*(\lambda_1, p), p) = c'_P(p), \quad (10)$$

which does depend on the technology, and exceeds R_p term-by-term when $R_{ap} \geq 0$ (her effort crowds in his).

Proposition 7 (Underinvestment and the value of commitment). (i) $p^{\text{nc}} < p^c(\lambda_1)$ whenever $\delta\psi'(p) < (1 - \lambda_1) \frac{d}{dp} R(a^*(\lambda_1, p), p)$ on the relevant range; in particular for δ small, since $p^{\text{nc}} \rightarrow 0$ as $\delta \rightarrow 0$ while p^c is bounded away from zero whenever the match has positive committed value. (ii) The student accepts at $t = 1$ only if $\delta\psi(p^{\text{nc}}) - c_P(p^{\text{nc}}) \geq (\bar{u}_P - \omega) + x\theta_1$; as $\delta \rightarrow 0$ the match unravels entirely. (iii) There exist binding caps $\bar{\theta}$ (institutional commitment devices) that strictly raise both parties' ex-ante payoffs relative to the no-commitment equilibrium.

Remark 4 (Benchmark example). With $R = 2a^{1/2}p^{1/2}$, $c_A(a) = a^2/2$, $c_P(p) = p^2/2$, $\psi(p) = p$: $a^* = \lambda^{2/3}p^{1/3}$, $v_P = 2(1 - \lambda)\lambda^{1/3}p^{2/3}$ (dissipative for $\lambda > 1/4$, consistent with Lemma 2 at $\alpha = 1/2$, $\eta = 1$), $p^{\text{nc}} = \delta$, and $p^c(\lambda_1) = [\frac{4}{3}(1 - \lambda_1)\lambda_1^{1/3}]^{3/4} \approx 0.62$ at $\lambda_1 = 1/2$.

5. Policy implications

The comparative statics map one-to-one into institutional design; the instruments operate on threat points, not preferences. Portable fellowships raise $\bar{u}_P$, ω , and δ (Propositions 4, 5); mandatory co-advising and empowered thesis committees split the control rights of Assumption 5, lowering γ and capping λ ; structured contribution statements (e.g., the CRediT taxonomy) shrink the non-contractible component of credit, attenuating Assumption 4; reporting channels and enforcement raise $\pi(\theta)F$ (Proposition 3); and information institutions—climate surveys, exit interviews, “whisper networks”—convert the one-shot game into a reputational repeated game.

6. Discussion and limitations

The model abstracts from several forces. Asymmetric information about the advisor's type at the matching stage would generate a screening equilibrium in which extractive types pool with professional ones early, rationalizing gradual revelation and the informational role of whisper networks. Risk aversion would add an insurance motive for dependence. Multi-student groups introduce tournament effects in the allocation of recommendation letters. Finally, we model a single match; general equilibrium in the market for advisors, with reputations, is a natural next step.

Appendix A. Proofs

Proof of Proposition 1

Given λ , the advisor solves $\max_a \lambda 2a^{1/2}p^{1/2} - a^2/2$, with FOC $\lambda a^{-1/2}p^{1/2} = a$, i.e. $a^{3/2} = \lambda p^{1/2}$; symmetrically $p^{3/2} = (1 - \lambda)a^{1/2}$. Solving, $a^* = \lambda^{3/4}(1 - \lambda)^{1/4}$ and $p^* = (1 - \lambda)^{3/4}\lambda^{1/4}$. Substitution gives $R^* = 2\sqrt{\lambda(1 - \lambda)}$, $(a^*)^2 + (p^*)^2 = \sqrt{\lambda(1 - \lambda)}$, hence $S^* = \frac{3}{2}\sqrt{\lambda(1 - \lambda)}$, and the stated u_A, u_P . Uniqueness follows from strict concavity of each best response and the contraction property of the FOC system on $(0, \infty)^2$. $\square$

Proof of Proposition 2

$S^{*'}(\lambda) = \frac{3}{4}(1 - 2\lambda)/\sqrt{\lambda(1 - \lambda)}$ vanishes at $\lambda = 1/2$. Differentiating u_A : $u'_A(\lambda) = \frac{3}{4}\lambda^{1/2}(1 - \lambda)^{-1/2}(3 - 4\lambda)$, which vanishes at $\lambda = 3/4$ and is positive below. For (iii), $u'_A(\lambda_0) > 0$ when $\lambda_0 < 3/4$, so a marginal increase in θ raises u_A ; $p^{*'}(\lambda) = \frac{1}{4}\lambda^{-3/4}(1 - \lambda)^{-1/4}(1 - 4\lambda) < 0$ for $\lambda > 1/4$. $\square$

Proof of Proposition 3

Implicit differentiation of the FOC; the second-order condition gives the sign of $\partial\theta^*/\partial F = \pi'(\theta^*)/\text{SOC} < 0$. The last claim follows from $u'_A(\lambda) \geq 0$ as $\lambda \leq 3/4$. $\square$

Proof of Proposition 4

Let $\Phi(\theta) \equiv u_P(\lambda_0 + \gamma\theta) - \kappa(\theta)$. For $\lambda > 1/4$, $u'_P(\lambda) < 0$, so $\Phi'(\theta) = \gamma u'_P - \kappa' < 0$; Φ is strictly decreasing. The implicit function theorem on $\Phi(\theta^*) = \bar{u}_P$ gives $\partial\theta^*/\partial \bar{u}_P = 1/\Phi' < 0$, $\partial\theta^*/\partial \lambda_0 = -u'_P/\Phi' < 0$, $\partial\theta^*/\partial \gamma = -\theta^*u'_P/\Phi' < 0$, and $\partial\theta^*/\partial k = \theta^{*2}/\Phi' < 0$. $\square$

Proof of Lemma 1

Existence of a^* follows from continuity and coercivity (Weierstrass on a compact sublevel set). Take $\lambda' > \lambda$ and any $a^* \in \arg \max_a [\lambda R(a, p) - c_A(a)]$. By revealed preference,

$$v_A(\lambda'; p) \geq \lambda' R(a^*, p) - c_A(a^*) \geq \lambda R(a^*, p) - c_A(a^*) = v_A(\lambda; p),$$

with the middle inequality strict whenever $R(a^*, p) > 0$. No differentiability, concavity, or uniqueness of a^* is used. $\square$

Proof of Lemma 2

The envelope theorem is not available for v_P (the student does not choose a), so differentiate directly: $\partial v_P / \partial \lambda = -R(a^*, p) + (1 - \lambda)R_a(a^*, p)\partial a^* / \partial \lambda$. Implicit differentiation of the interior FOC $\lambda R_a(a^*, p) = c'_A(a^*)$ gives $\partial a^* / \partial \lambda = R_a / (c''_A - \lambda R_{aa}) > 0$, yielding (7). For the power family, the FOC gives $\lambda \alpha R / a^* = (a^*)^\eta$; substituting $R_a = \alpha R / a$, $R_{aa} = \alpha(\alpha - 1)R / a^2$, and $c''_A = \eta \lambda \alpha R / a^2$ into (7) yields $\partial v_P / \partial \lambda = R \left[\frac{(1-\lambda)\alpha}{\lambda(1+\eta-\alpha)} - 1 \right]$, which is negative iff $\alpha < \lambda(1 + \eta)$. $\square$

Proof of Proposition 5

By Lemma 1 (and B nondecreasing), the advisor's period-2 objective $v_A(\lambda(\theta_2); p) + B(\theta_2)$ is increasing in θ_2 ; his feasible set is $\{\theta_2 \in [\theta_1, \theta_{\max}] : \Lambda(\theta_2) \geq w(\theta_1, p)\}$, which under Assumption 8 is the interval $[\theta_1, \hat{\theta}]$ with $\hat{\theta}$ the unique solution of (8) (or $\theta_{\max}$). The optimum is the right endpoint. Strict escalation follows because at $\theta_2 = \theta_1$ the constraint is slack when $\Lambda(\theta_1) > \omega + \delta\psi(p) - x\theta_1$. Comparative statics follow from the implicit function theorem applied to $\Lambda(\hat{\theta}) = w$: with $\Lambda' < 0$, $\partial \hat{\theta} / \partial w = 1 / \Lambda' < 0$, and w is increasing in δ and ω and decreasing in x and θ_1 . $\square$

Proof of Proposition 6

As $\delta \rightarrow 0$ or $x\theta_1$ grows, w falls, so $\hat{\theta}$ rises (Proposition 5), approaching $\theta_{\max}$, while θ^{acc} is fixed (it is defined from ex-ante objects only). Hence $\hat{\theta} > \theta^{\text{acc}}$ for $(\delta, 1/x)$ small. Sequential rationality of remaining is by construction of (8): at $\theta_2 = \hat{\theta}$ the student is exactly indifferent to exit and strictly prefers staying at any $\theta_2 < \hat{\theta}$. $\square$

Proof of Proposition 7

(i) Compare (9) and (10): both first-order conditions equate $c'_P(p)$ to a marginal benefit; when $\delta\psi'(p) < (1 - \lambda_1) \frac{d}{dp} R(a^*(\lambda_1, p), p)$ pointwise on the relevant range, convexity of c_P orders the solutions. As $\delta \rightarrow 0$, (9) forces $c'_P(p^{\text{nc}}) \rightarrow 0$, hence $p^{\text{nc}} \rightarrow 0$, while p^c solves (10), which does not involve δ . (ii) Substitute p^{nc} into the ex-ante participation condition $\omega + \delta\psi(p) - x\theta_1 - c_P(p) \geq \bar{u}_P$; as $\delta \rightarrow 0$ the left side tends to $\omega - x\theta_1 < \bar{u}_P$ under the maintained assumption that the raw interim option does not dominate the ex-ante one. (iii) A verifiable cap $\bar{\theta} = \theta_1$ implements the commitment outcome, in which the student earns rents above her exit value and invests $p^c > p^{\text{nc}}$. The advisor's committed payoff $\lambda_1 R(a^*(\lambda_1, p^c), p^c) - c_A(a^*)$ is bounded away from zero whenever the match has positive committed

value, while his no-commitment payoff $v_A(\lambda(\hat{\theta}); p^{nc})$ tends to $v_A(\cdot; 0) = 0$ as $\delta \rightarrow 0$ (using $R(a, 0) = 0$ and $c_A \geq 0$). Hence for δ small the cap strictly benefits both parties. $\square$

References

1. [△]Grossman SJ, Hart OD (1986). "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *J Polit Econ.* **94**(4):691–719.
2. [△]Hart O, Moore J (1990). "Property Rights and the Nature of the Firm." *J Polit Econ.* **98**(6):1119–1158.
3. [△]Klein B, Crawford RG, Alchian AA (1978). "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process." *J Law Econ.* **21**(2):297–326.
4. [△]Williamson OE (1985). *The Economic Institutions of Capitalism*. Free Press.
5. [△]Freixas X, Guesnerie R, Tirole J (1985). "Planning Under Incomplete Information and the Ratchet Effect." *Rev Econ Stud.* **52**(2):173–191.
6. [△]Laffont JJ, Tirole J (1988). "The Dynamics of Incentive Contracts." *Econometrica.* **56**(5):1153–1175.
7. [△]Acemoglu D, Pischke JS (1999). "The Structure of Wages and Investment in General Training." *J Polit Econ.* **107**(3):539–572.
8. [△]Manning A (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press.
9. [△]Hart O, Moore J (2008). "Contracts as Reference Points." *Q J Econ.* **123**(1):1–48.
10. [△]Merton RK (1968). "The Matthew Effect in Science." *Science.* **159**(3810):56–63.
11. [△]Stephan P (2012). *How Economics Shapes Science*. Harvard University Press.

Supplementary data: available at <https://doi.org/10.32388/Y07U36>

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.