

Strategic Outsourcing to Shared AI Platforms^{*}

Byung-Cheol Kim[†] Keeyoung Rhee[‡]

September 22, 2026

Abstract

Firms increasingly outsource innovation to AI platforms serving their competitors, allowing one client's problem solving to improve what the platform delivers to another. We study a duopoly in which firms choose between in-house innovation and a shared AI platform. Cross-client learning occurs when both firms use the platform. Stronger learning can initially support joint use, yet eventually destabilize it as each firm's problem solving increasingly benefits its rival. Too many or too few firms may adopt the platform relative to a coordinated benchmark that maximizes their combined payoff subject to symmetric effort. Blocking cross-client learning can reduce competitive harm yet leave the firms collectively worse off.

JEL Classification: D43, L22, L24, O32

Keywords: shared AI platforms, outsourcing, coopetition, cross-client learning, technology adoption

^{*}We are indebted to Pietro Peretto for extensive discussions on this research agenda, and to Nina Baranchuk, Jay Pil Choi, and participants in the SKKU Department of Economics brownbag seminar for valuable comments. We gratefully acknowledge financial support from the NET Institute, www.netinst.org, through its 2026 summer research grant.

[†]Department of Economics, Finance and Legal Studies, Culverhouse College of Business, University of Alabama, Tuscaloosa, AL 35487. Email: bkim34@ua.edu.

[‡]Sungkyunkwan University (SKKU), Department of Economics, 25-2 Sungkyunkwan-ro, Jongno-gu, Seoul, 03063, Korea. Email: ky.rhee829@gmail.com.

1 Introduction

AI services increasingly support a wide range of business activities. A recent survey of nearly 6,000 senior executives finds that about 70 percent of companies in the United States, the United Kingdom, Germany, and Australia actively use some form of AI, with adoption concentrated among the largest and most productive firms (Yotzov et al., 2026). Field and experimental evidence also documents substantial productivity gains on knowledge-intensive work, with especially large gains for lower-performing workers in some settings (Noy and Zhang, 2023; Brynjolfsson, Li and Raymond, 2025; Dell’Acqua et al., 2023). The technology is also pushing the frontier of innovation itself, from materials discovery to drug development (Merchant et al., 2023; Xu et al., 2025). Many of these capabilities are supplied by third-party platforms that serve many firms at once, including firms that compete with each other.

A shared platform can make a firm more productive in two ways. The first is the platform’s own capability. The second arises when the platform learns across its clients: methods worked out for one client can be reused for another, feedback can be aggregated, and experience can improve the tools used for other clients. When clients compete, this cross-client learning cuts both ways. A firm benefits from knowledge generated by its rival’s effort on the platform, but its own effort also strengthens that rival. Such learning need not disclose any client’s confidential information, so protecting confidentiality does not by itself remove this tension.

Firms already treat what they supply to external AI services as strategically sensitive. Samsung, for example, restricted employee use of generative AI after proprietary source code was uploaded, citing the risk that material held on external servers could be disclosed to other users.¹ Business-facing AI services also offer contractual and technical safeguards, including commitments not to train general models on business data.² Such commitments restrict one channel of cross-client learning; whether others remain active depends on the platform’s design.

¹“Samsung Bans Staff’s AI Use After Spotting ChatGPT Data Leak,” *Bloomberg*, May 2, 2023. A related concern arose in the September 2026 dispute over AI-assisted work on the Navier–Stokes problem, when a researcher questioned whether unpublished material supplied to an AI provider had influenced the provider’s own work; the provider denied such influence. See “How an AI Math Breakthrough Ignited a Controversy,” *Science*, September 8, 2026, and OpenAI, <https://openai.com/index/navier-stokes-solution/>. The episode concerns possible reuse of client material, not evidence of cross-client learning among competing firms.

²OpenAI, *Enterprise privacy*, <https://openai.com/enterprise-privacy/>; Anthropic, *Commercial Terms of Service* §B, <https://www.anthropic.com/legal/commercial-terms>; Microsoft, *Data, privacy, and security for Foundry Models sold by Azure*, <https://learn.microsoft.com/en-us/azure/foundry/responsible-ai/openai/data-privacy>. Accessed September 2026.

This paper therefore asks how the prospect of cross-client learning in shared AI services shapes firms' decisions to use shared AI for innovation. When does a firm outsource a competitively important task to a shared platform, and when does it keep the task in-house? Can two rivals sustain the use of the same platform, knowing that each one's effort may sharpen the other's capability? And when they do, is the effort they put into innovation too much or too little relative to the coordinated benchmark?

To address these questions, we treat the extent of cross-client learning as a feature of the AI platform's learning architecture rather than as a disclosure choice by the client: firms choose whether to use the shared platform, and the learning channel is active only when both do. Two symmetric firms choose, for a hard innovation task, between keeping it in-house and outsourcing to the shared platform at a fixed fee. Each firm allocates scarce effort between a routine task with a safe return and the hard innovation task, whose prize is eroded when both firms succeed, and chooses its organizational form and its effort simultaneously.

The model's central feature is that participation changes the technology of competition. In a conventional adoption problem, joining an AI platform changes payoffs, and complementarity typically takes the form of an exogenous network effect; here, joint participation activates cross-client learning, after which each firm's effort enters its rival's innovation technology.

The analysis delivers three sets of results. First, outsourcing to a shared AI platform is not a standard make-or-buy decision. When both firms keep the task in-house, a regime we call autarky, each allocates excessive effort to the hard innovation task because it ignores the rent destruction effect its rival suffers. Joint outsourcing introduces a second externality in the opposite direction: each firm's effort improves the AI platform's problem-solving capability available to its rival, a benefit for which the firm is not compensated. These forces generate four symmetric-equilibrium regions: unique autarky, multiple symmetric equilibria, unique joint outsourcing, and a residual region with no symmetric pure-strategy equilibrium.

The second result concerns the sustainability of joint outsourcing. Cross-client learning does not affect a firm's incentive to outsource while its rival remains in-house, but it becomes central to the incentive to stay once both firms use the platform. Stronger learning can initially make staying more attractive. Beyond a certain point, however, the competitive spillover discourages effort enough to undermine the value of remaining on the platform. Joint outsourcing therefore cannot be sustained beyond a finite learning intensity, although it may be sustainable even without cross-client learning.

Greater stand-alone platform productivity can also make joint outsourcing harder to sustain. When cross-client learning is weak and firms retain enough rents from joint success,

higher productivity strengthens the incentive to remain on the platform. Under intense competition, however, the resulting increase in joint success can dissipate enough innovation rent to weaken that incentive. The platform’s ability to retain both clients therefore depends on how its productivity gains interact with competition.

Third, decentralized organizational choice can produce too little or too much outsourcing relative to the coordinated benchmark.³ Firms may remain in-house when the benchmark prefers outsourcing, or jointly outsource when it prefers a single platform client. A distinct conflict arises even within private payoffs: joint outsourcing can be the unique symmetric equilibrium although both firms would earn more under mutual in-house production, yielding a Prisoner’s Dilemma ordering of the regime payoffs. Conditional on joint outsourcing, effort is excessive relative to the benchmark when rent dissipation dominates and insufficient when the benefit conferred through cross-client learning dominates.

These findings make learning architecture central to policy toward shared AI platforms. Lower access costs encourage both the first and the second client to outsource; stronger cross-client learning changes the value of joint use without changing the incentive to outsource alone. Restricting learning can preserve joint use by limiting competitive spillovers, but can also undermine it by removing the productivity benefits that keep clients on the platform. Sustaining joint use is itself an incomplete policy objective: stronger learning can reduce equilibrium success as firms cut effort, and the firms’ combined payoff excludes consumer benefits. Policy must therefore assess how access and learning rules jointly shape participation, innovation, and the gains reaching consumers.

Related literature. The paper relates first to the theory of outsourcing under information leakage. The closest predecessor is [Baccara \(2007\)](#), who studies outsourcing with cross-client leakage in a monopolistically competitive setting. In a leakage model, information transferred to rivals is primarily a loss to the originating firm. In our setting the cross-client flow is also a productive input: the same cross-client learning that exposes a firm to strategic spillovers raises the effectiveness of the shared problem-solving technology. Beyond this, we study a strategic duopoly in which organizational choices interact directly, and we separate the platform’s private productivity gain from the cross-client learning rate, which we treat as a feature of the platform’s architecture rather than a technological constant.

The model also connects to cooperative R&D and research joint ventures, which study the

³Throughout, *decentralized* describes choices and equilibrium outcomes when firms act separately, while *private* describes firm-level payoffs and incentives. The coordinated benchmark maximizes the firms’ combined payoff, imposing equal effort on firms in the same organizational role. It does not include consumer surplus.

interaction between knowledge sharing and product-market competition (d’Aspremont and Jacquemin, 1988; Kamien, Muller and Zang, 1992). Kamien, Muller and Zang (1992) distinguish sharing research results from coordinating research investment, and joint outsourcing here has the same character: firms benefit from one another’s effort while continuing to choose that effort independently. The difference lies in where the sharing resides. In a research joint venture it is an agreement among the firms themselves; here it runs through a third party whose learning architecture the firms take as given, and joint use must survive unilateral withdrawal. That location also separates two margins: the platform’s stand-alone productivity gain is available to a sole client, whereas cross-client learning requires both rivals to participate.

The paper relates second to work on platforms that learn from their customer base. Hagiu and Wright (2023) distinguish learning across users from learning within a user and show how data-enabled learning creates competitive advantage. Their providers compete for customers whose use improves them; our competing firms share a single provider. Schaefer and Sapi (2023) provide empirical evidence from web search, where richer within-user information improves the efficiency of learning across users; their users are individuals rather than competing firms.

This learning can create resistance to sharing data. Jones and Tonetti (2020) show that firms may hoard nonrival data for fear of creative destruction, while Gans (2026) examines how supplying data to AI intermediaries can strengthen rivals and discourage contributions. On the buying side, Liu, Ma and Veldkamp (2025) show how a seller’s inability to restrict future sales to competitors reduces clients’ willingness to pay for data. We study an organizational margin: productive spillovers arise only when both rivals use the same platform. Joint participation exposes each firm to stronger competition but also allows it to benefit from the rival’s effort. This combination determines whether joint outsourcing is sustainable and how much effort firms supply.

The shared-provider structure also connects our paper to Brown and MacKay (2023), who study competition among firms adopting the same third-party pricing algorithm. Our focus is innovation effort, where a common intermediary links competitive incentives to productive spillovers. More broadly, the model studies an AI-mediated form of knowledge organization. Whereas knowledge hierarchies route difficult problems to internal experts (Garicano, 2000; Garicano and Rossi-Hansberg, 2006), our firms delegate problem solving to a shared external provider whose capability partly derives from rival use.

The paper relates third to the IO literature on excess inertia and excess momentum in

technology adoption. Classic models of standardization and network externalities show that decentralized adoption decisions can diverge from the benchmark in either direction: adopters may remain with an incumbent technology when coordinated switching is socially desirable, or may move too quickly to a new standard when adoption strands an installed base (Farrell and Saloner, 1985, 1986; Katz and Shapiro, 1986, 1992). Choi (1994a) extends this framework to irreversible technology choice under uncertainty, and Choi (1994b) treats compatibility as a strategic choice that shapes network value; related work shows that both directions of inefficiency can arise under strategic timing and consumer heterogeneity (de Bijl and Goyal, 1995; Choi and Thum, 1998). Our excessive in-house production and excessive joint outsourcing are parallels to inertia and momentum, but the underlying mechanism differs. The classical literature emphasizes demand-side compatibility, installed-base value, and coordination among adopters; here the externality is supply-side and architectural. A shared AI platform converts one client’s effort into knowledge that may improve a rival’s capability, while product-market competition determines whether that transfer raises or lowers the firms’ combined payoff.

Finally, the paper relates to recent work on the economic externalities of AI. He (2026) shows that adverse selection can lead to excessive AI adoption. Acemoglu, Kong and Ozdaglar (2026) study dynamic externalities in knowledge production for AI systems, while Chen (2026) examines how AI displacement of entry-level tasks contracts the future pool of skilled workers. Our model likewise features undercompensated contributions to shared productive capacity, alongside an opposing externality through product-market competition. These forces can generate either insufficient or excessive effort.

Closer to our competitive setting, Choi, Jeon and Menicucci (2026) show that first-mover advantages and learning from rivals’ experimentation can induce premature or insufficient AI deployment relative to joint-profit maximization. Their spillover informs a waiting firm about technological risk; ours raises productive capability and requires both rivals to use the same intermediary. This bilateral participation requirement links learning to organizational stability: stronger cross-client learning can raise the coordinated value of joint outsourcing while making it harder to sustain.

Roadmap. Section 2 presents the model. Section 3 characterizes the symmetric equilibrium regions using two deviation gaps. Sections 4 and 5 analyze autarky and joint outsourcing, in each case asking whether the regime is sustained as a decentralized equilibrium and whether effort is distorted within it. Section 6 compares the three organizational forms and asks when decentralized choices diverge from the coordinated benchmark. Section 7 draws the implications and Section 8 concludes. The main appendix contains the proofs; the Supplementary

Appendix contains additional derivations and numerical analysis.⁴

2 Model

Consider two ex ante symmetric firms, indexed by $i \in \{1, 2\}$. Each firm has one unit of organizational effort. Firm i allocates $r_i \in [0, 1]$ of its effort to a “hard” innovation task and the remaining $1 - r_i$ to a “routine” task. The routine task generates a safe payoff $\phi(1 - r_i)$, where $\phi > 0$ is the marginal payoff of effort allocated to the routine task.⁵ The hard task is a winner-takes-most innovation contest. If only one firm succeeds, it obtains monopoly value $V > 0$ in a new market. If both firms succeed, downstream competition erodes the prize, and each firm earns βV , where $\beta \in (0, 1)$ denotes the fraction of the monopoly value that survives competition. Thus, a smaller β corresponds to more intense product-market competition in the new market.

Unlike the routine task, the hard task may fail. Let $q_i \in [0, 1]$ denote firm i ’s probability of solving the hard task. We assume that success events are independent across firms. Given rival firm j ’s success probability q_j , firm i earns V if it succeeds alone and βV if both firms succeed. Thus, firm i ’s expected hard-task payoff is

$$Vq_i[1 - (1 - \beta)q_j],$$

where $(1 - \beta)q_j$ captures the rent-dissipation force: the more likely the rival is to succeed, the more firm i ’s own prize is eroded by competition.

Firm i ’s success probability q_i depends on its effort allocation r_i and on whether it uses a shared artificial intelligence (AI) platform to enhance its innovation capability. Simultaneously with its effort choice r_i , firm i chooses $a_i \in \{0, 1\}$, where $a_i = 1$ means that firm i uses the external AI platform to perform or support the hard task at a fixed fee $c > 0$, and $a_i = 0$ means that it relies on its internal technology alone. The platform augments the firm’s own effort rather than replacing it: the firm continues to choose r_i under either organizational form, and what changes is how productive that effort is. We nonetheless refer to $a_i = 1$ as outsourcing throughout. Given rival firm j ’s strategy (a_j, r_j) , firm i ’s effective input into the

⁴The Supplementary Appendix is available at https://sites.ua.edu/bkim/files/2026/09/AI_Outsourcing_OA.pdf.

⁵Reading hard-task effort as carrying an opportunity cost extends to any increasing, weakly concave routine-task payoff $g(1 - r_i)$. The linear specification gives a constant marginal opportunity cost and keeps the equilibrium characterization transparent.

hard task, denoted by x_i , is

$$x_i := (1 + \eta a_i) r_i + \lambda a_i a_j r_j. \quad (1)$$

Thus, AI outsourcing affects firms' ability to solve the hard task through two channels:

1. A direct productivity gain $\eta > 0$, which scales up the marginal productivity of firm i 's own effort. The private benefit accrues to firm i whenever it outsources the hard task to the AI platform ($a_i = 1$), regardless of the rival's outsourcing decision.
2. A cross-client learning effect $\lambda \geq 0$. Cross-client learning operates only when both firms outsource to the same AI platform ($a_i = a_j = 1$): the platform aggregates knowledge across clients, so rival firm j 's effort contributes to firm i 's effective capability through the term λr_j . We interpret λ as the intensity of cross-client learning that the platform's architecture permits.⁶ At fixed effort, cross-client learning weakly raises each firm's effective input; the model does not consider negative transfer that degrades technical performance.

The case $\lambda = 0$ corresponds to a fully segregated AI platform architecture, in which outsourcing provides only a private productivity gain and generates no cross-client learning.

Firm i 's success probability is $q_i = q(x_i)$, where $q : [0, \infty) \rightarrow [0, 1]$ is twice continuously differentiable, strictly increasing, strictly concave, and satisfies $q(0) = 0$. The distinction between the private productivity gain and cross-client learning does not depend on the functional form of q . To obtain closed-form expressions and the comparative statics stated with each result, we specialize throughout to the exponential specification

$$q(x) = 1 - e^{-\gamma x}. \quad (2)$$

The parameter $\gamma > 0$ in (2) captures the tractability of the hard task. A higher γ means that any given effective input is converted into successful innovation with a higher probability. Tractability scales success under every organizational regime, so it is an inherent property of the hard task rather than of the AI platform. Two normalized composites recur throughout: the normalized opportunity cost $\mu \equiv \phi/(V\gamma)$ and the fee $\kappa \equiv c/V$. We use the level fee c in payoff and policy expressions and κ in closed forms and numerical calibrations.⁷

⁶Cross-client learning is a reduced-form feature of the technology over the relevant production horizon; it need not occur automatically or immediately after each interaction.

⁷The exponential success function is standard in the R&D contest literature; see [Loury \(1979\)](#); [Lee and Wilde \(1980\)](#); [Dasgupta and Stiglitz \(1980\)](#); [Reinganum \(1989\)](#); [Aghion and Howitt \(1992\)](#). It provides tractability, not mechanism. The vertical outsourcing boundary holds under any success technology. Un-

We consider a normal-form game where each firm $i \in \{1, 2\}$ simultaneously chooses a strategy $(a_i, r_i) \in \{0, 1\} \times [0, 1]$ to maximize its total expected payoff, taking rival firm j 's strategy (a_j, r_j) as given:

$$\pi_i(a_i, r_i; a_j, r_j) = \underbrace{\phi(1 - r_i)}_{\text{routine-task payoff}} + \underbrace{Vq(x_i)[1 - (1 - \beta)q(x_j)]}_{\text{hard-task expected payoff}} - \underbrace{c a_i}_{\text{AI platform fee}}. \quad (3)$$

The two firms' outsourcing decisions determine the following organizational regimes:

- Autarky $(a_i, a_j) = (0, 0)$ yields

$$q_i = q(r_i), \quad q_j = q(r_j);$$

- Unilateral outsourcing $(a_i, a_j) = (1, 0)$ yields

$$q_i = q((1 + \eta)r_i), \quad q_j = q(r_j);$$

- Joint outsourcing $(a_i, a_j) = (1, 1)$ yields

$$q_i = q((1 + \eta)r_i + \lambda r_j), \quad q_j = q((1 + \eta)r_j + \lambda r_i).$$

We analyze Nash equilibria of this normal-form game.⁸ In Sections 4 and 5, we focus primarily on symmetric pure-strategy equilibria, in which the two firms make the same outsourcing decision, $a_i = a_j \in \{0, 1\}$. These equilibria correspond to the two central organizational regimes of interest. The first is *autarky*, $(a_i, a_j) = (0, 0)$, in which both firms keep the hard task in-house. The second is *joint outsourcing*, $(a_i, a_j) = (1, 1)$, in which both firms outsource the hard task to the shared AI platform. Throughout, a two-digit superscript ab records the pair $(a_i, a_j) = (a, b)$: u^{00} and u^{11} are on-path payoffs under the two symmetric regimes, while

der any strictly concave success function and weakly concave routine-task payoff, the productivity gain and cross-client spillover both persist; only the explicit closed-form expressions and comparative-static formulas require additional regularity.

⁸A natural alternative to our simultaneous-move formulation is sequential timing: firms first decide whether to outsource and then choose effort after observing each other's outsourcing decisions. Conditional on a given symmetric organizational regime, the effort conditions coincide with the corresponding continuation-game best responses, so the effort characterizations do not turn on simultaneity. The organizational-deviation conditions need not coincide, however, because an outsourcing deviation would let the rival adjust its effort in the continuation game, whereas here the rival's effort is held at its on-path level. We therefore treat simultaneous choice as the maintained timing assumption.

d^{10} is firm i 's payoff when it outsources and the rival stays in-house, and d^{01} its payoff when it withdraws and the rival remains on the platform.

Remark 1 (Alternative Formulation). *A simpler model could make AI use the default and let firms choose only how much activity is exposed to cross-client learning. We do not take this route because it would eliminate organizational choice: there would be no distinction among autarky, unilateral outsourcing, and joint outsourcing, and hence no Prisoner's Dilemma in which joint outsourcing is the unique symmetric equilibrium even though both firms prefer autarky. It would also erase the clean separation between the fee, which governs outsourcing, and cross-client learning, which governs effort.*

3 Equilibrium Characterization

In this section, we characterize symmetric pure-strategy equilibria using two organizational deviation gaps, Δ^0 and Δ^1 , together with a within-regime effort condition specific to joint outsourcing. The deviation gaps determine whether a firm wishes to change its organizational choice, taking the rival's organizational choice and on-path effort as given. Under joint outsourcing, cross-client learning introduces a further requirement: the symmetric effort candidate must itself survive all effort deviations while both firms remain on the platform.

We first consider the symmetric autarky equilibrium profile, $(a_i, a_j) = (0, 0)$. In this regime, firm i 's effort choice satisfies

$$-\phi + \left. \frac{dq_i}{dr_i} \right|_{r_i=r^I} V[1 - (1 - \beta)q_j] = 0.$$

Effort is chosen on $[0, 1]$ and each firm's payoff is strictly concave in its own effort, so the constrained optimum is unique; denote it by r^I . Since $q_i = q_j = q^I \equiv q(r^I)$ in the associated symmetric pure-strategy equilibrium, r^I satisfies, whenever it is interior,

$$\underbrace{\phi}_{\text{opportunity cost}} = \underbrace{q'(r^I)}_{\text{marginal success prob.}} \cdot \underbrace{V[1 - (1 - \beta)q(r^I)]}_{\text{expected prize}}, \quad (\text{I})$$

and is otherwise pinned at the boundary of $[0, 1]$ at which the corresponding Kuhn–Tucker inequality holds. The payoff of each firm at that effort, denoted by u^{00} , is

$$u^{00} := \phi(1 - r^I) + Vq^I[1 - (1 - \beta)q^I]. \quad (4)$$

Next, consider the symmetric joint-outsourcing candidate, $(a_i, a_j) = (1, 1)$. Whether it

is an equilibrium is settled below by two conditions; until then we call r^J the candidate effort and u^{11} the candidate payoff. In this regime, firm i 's effort affects not only its own effective input $x_i = (1 + \eta)r_i + \lambda r_j$ but also the rival's effective input $x_j = (1 + \eta)r_j + \lambda r_i$. Thus, raising r_i has two marginal effects on firm i 's expected payoff: it increases firm i 's own success probability through the direct productivity channel, but it also raises the rival's success probability through the cross-client learning channel. The first-order condition is therefore

$$-\phi + \left. \frac{dq_i}{dr_i} \right|_{r_i=r^J} V[1 - (1 - \beta)q_j] - q_i V(1 - \beta) \left. \frac{dq_j}{dr_i} \right|_{r_i=r^J} = 0.$$

At the symmetric on-path effort r^J , the two effective inputs coincide:

$$x_i = x_j = x^J \equiv (1 + \eta + \lambda) r^J,$$

and hence

$$q(x_i) = q(x_j) = q^J \equiv q(x^J),$$

where the superscript J refers to joint AI outsourcing. Thus $1 + \eta + \lambda$ is the effective-input scale for a symmetric increase in common effort: raising both firms' effort by one unit raises each firm's effective input by $1 + \eta$ from its own effort and by λ from the rival's. By contrast, a unilateral increase in firm i 's effort alone raises x_i by $1 + \eta$ and x_j by λ ; the combined scale $1 + \eta + \lambda$ applies only when both firms move together. Under unilateral outsourcing, where only one firm is a platform client, the scale is $1 + \eta$, since cross-client learning is then inactive. The first-order condition can then be written as

$$\phi = \underbrace{q'(x^J)}_{\text{marginal success prob.}} \cdot V \left[\underbrace{1 + \eta}_{\text{own channel}} - \underbrace{(1 - \beta)(1 + \eta + \lambda)}_{\text{competition wedge}} q^J \right]. \quad (\text{J})$$

Because effort is bounded by the unit endowment, (J) characterizes r^J only when its solution is interior; when the endowment binds the symmetric candidate is $r^J = 1$, and under the maintained condition $(1 + \eta) > \mu$, zero effort is not a symmetric best response.⁹ Comparing the first-order conditions (I) and (J) reveals that the marginal incentive under joint outsourcing differs from that under autarky in two ways. First, the direct productivity channel is amplified from 1 to $(1 + \eta)$, reflecting the direct productivity gain from AI outsourcing. Second, the competition wedge is also amplified: each unit of the firms' own effort raises the probability of

⁹At $r_i = r_j = 0$ the marginal payoff from raising own effort is $V\gamma(1 + \eta - \mu) > 0$; see Lemma A.1.

joint success not only through the direct productivity channel η , but also through the cross-client learning channel λ . Thus, relative to autarky, joint outsourcing strengthens both the private productivity benefit and the rent-dissipation force. Finally, the payoff of each firm at the symmetric joint-outsourcing candidate, denoted by u^{11} , is

$$u^{11} := \phi(1 - r^J) - c + Vq^J[1 - (1 - \beta)q^J]. \quad (5)$$

Turning to deviations, under autarky a firm may prefer to unilaterally outsource to the AI platform ($a'_i = 1$) while the rival stays in-house ($a_j = 0$). By doing so, the deviating firm's effective input increases to $x'_i = (1 + \eta)r_i$ while the rival's input remains $x_j = r^I$. Thus, the deviation payoff from autarky, denoted by d^{10} , is

$$d^{10} := \max_{r \in [0,1]} \phi(1 - r) - c + Vq((1 + \eta)r)[1 - (1 - \beta)q^I]. \quad (6)$$

Next, consider a firm's deviation from joint outsourcing. If firm i brings the hard task back in-house ($a'_i = 0$), its effective input becomes $x'_i = r_i$. Since its rival firm j continues to outsource and allocates effort $r_j = r^J$, the rival's effective input becomes $(1 + \eta)r^J$: firm i 's withdrawal removes its contribution to the rival's success probability through cross-client learning. The deviation payoff from joint outsourcing is therefore

$$d^{01} := \max_{r \in [0,1]} \phi(1 - r) + Vq(r)[1 - (1 - \beta)q((1 + \eta)r^J)]. \quad (7)$$

Write r^{01} for the maximizer in (7), the effort the withdrawing firm supplies in-house, and let

$$q^{01} \equiv q(r^{01}), \quad \tilde{q}^J \equiv q((1 + \eta)r^J)$$

denote its own in-house success probability and that of the rival who stays on the AI platform. The rival's effort is unchanged at r^J , but it has lost the pooled contribution λr^J , so $\tilde{q}^J < q^J$ whenever $\lambda > 0$: a withdrawal weakens the rival it leaves behind, and the more so the stronger is cross-client learning. Both probabilities are used in the comparative statics of Section 5. The closed-form derivations of r^I , r^J , q^I , q^J , d^{10} , and d^{01} under the exponential specification are presented in Supplementary Appendix S1.

Combining the payoff objects in (4)–(7), we define two deviation gaps that determine the

stability of autarky and joint outsourcing:

$$\begin{aligned}\Delta^0 &:= u^{00} - d^{10}, \\ \Delta^1 &:= u^{11} - d^{01}.\end{aligned}\tag{8}$$

The first gap, Δ^0 , measures the payoff gain from remaining in autarky rather than deviating to unilateral outsourcing when the rival stays in-house. The second gap, Δ^1 , measures the payoff gain from staying in joint outsourcing rather than deviating back to in-house production when the rival continues to outsource to the AI platform. Under the exponential specification in (2), both gaps admit closed-form expressions, which can be written in terms of β , η , λ , the normalized opportunity cost $\mu \equiv \phi/(V\gamma)$, and the normalized fee $\kappa \equiv c/V$. We use these composites throughout the paper and in all numerical examples.¹⁰

One further condition is specific to joint outsourcing. Because cross-client learning makes each firm's effort affect the rival's success probability, a firm's payoff need not be globally concave in its own effort. Consequently the symmetric candidate r^J , determined by (J) when interior and by the corresponding Kuhn–Tucker condition at the upper boundary, need not be a global best response. Even if no small effort deviation is profitable, a firm may prefer a substantially different effort level elsewhere in $[0, 1]$.

When both firms outsource and the rival exerts r^J , one can write firm i 's payoff as

$$\pi^{11}(r_i; r^J) := \phi(1 - r_i) - c + Vq((1 + \eta)r_i + \lambda r^J) \left[1 - (1 - \beta)q((1 + \eta)r^J + \lambda r_i) \right]. \tag{9}$$

Define the maximum gain from changing effort while remaining on the platform, and the set of cross-client learning intensities at which that gain vanishes, by

$$G^J(\beta, \lambda, \eta) := \max_{r \in [0, 1]} \{ \pi^{11}(r; r^J) - \pi^{11}(r^J; r^J) \} \geq 0, \quad \mathcal{G}^J(\beta, \eta) := \{ \lambda \geq 0 : G^J(\beta, \lambda, \eta) = 0 \}, \tag{10}$$

where the inequality holds because $r = r^J$ is feasible. Thus r^J is a global best response if and only if $\lambda \in \mathcal{G}^J(\beta, \eta)$. The platform fee cancels from the payoff difference, so $\mathcal{G}^J$ is independent of c . Section 5 characterizes this set and establishes uniqueness of the symmetric candidate

¹⁰Writing effort in transformed units $z = \gamma r$ removes any separate γ term from the interior first-order conditions and the normalized payoff differences, so γ enters only through the constraint $z \leq \gamma$. Holding μ fixed, tractability therefore has no independent effect when all relevant efforts are interior. Holding ϕ and V fixed, a change in γ also changes μ and can affect the deviation gaps even on an interior branch: $\Delta_\gamma^1 = -(\mu/\gamma)\Delta_\mu^1$, so because Δ_μ^1 has no primitive sign, neither does Δ_γ^1 .

(Lemma A.1). No analogous condition is required under autarky or unilateral outsourcing, where own effort does not affect the rival's effective input and the firm's payoff is strictly concave in its own effort.

We can now state the equilibrium classification. Autarky, $(a_1, a_2) = (0, 0)$, is a symmetric pure-strategy equilibrium if and only if $\Delta^0 \geq 0$. Joint outsourcing, $(a_1, a_2) = (1, 1)$, is a symmetric pure-strategy equilibrium if and only if $\lambda \in \mathcal{G}^J(\beta, \eta)$ and $\Delta^1 \geq 0$. The first condition rules out effort deviations that leave the organizational choice unchanged; the second rules out withdrawal from the platform accompanied by an optimal adjustment of in-house effort. Henceforth (μ, κ, γ) are held fixed and suppressed from the arguments of all effort and organizational-support objects. Writing the second condition as a set alongside the first,

$$\mathcal{D}^1(\beta, \eta) := \{\lambda \geq 0 : \Delta^1(\beta, \lambda, \eta) \geq 0\}, \quad \Lambda^1(\beta, \eta) := \mathcal{G}^J(\beta, \eta) \cap \mathcal{D}^1(\beta, \eta), \quad (11)$$

joint outsourcing is a symmetric pure-strategy equilibrium exactly when $\lambda \in \Lambda^1(\beta, \eta)$. Conditional on $\lambda \in \mathcal{G}^J(\beta, \eta)$, the symmetric pure-strategy equilibria are classified as follows:

- (i) If $\Delta^0 \geq 0$ and $\Delta^1 < 0$, autarky is the unique symmetric equilibrium.
- (ii) If $\Delta^0 \geq 0$ and $\Delta^1 \geq 0$, both autarky and joint outsourcing are symmetric equilibria.
- (iii) If $\Delta^0 < 0$ and $\Delta^1 \geq 0$, joint outsourcing is the unique symmetric equilibrium.
- (iv) If $\Delta^0 < 0$ and $\Delta^1 < 0$, no symmetric pure-strategy equilibrium exists.

Outside $\mathcal{G}^J(\beta, \eta)$, joint outsourcing cannot be a symmetric equilibrium regardless of the sign of Δ^1 , which is then an organizational gap evaluated at an effort candidate that is not itself an equilibrium. Autarky is the unique symmetric pure-strategy equilibrium when $\Delta^0 \geq 0$; when $\Delta^0 < 0$, none exists.

Cases (i)–(iii) admit at least one symmetric pure-strategy equilibrium. Case (iv) means that no *symmetric pure-strategy* equilibrium exists; it is not an equilibrium nonexistence result for the full game. Within $\mathcal{G}^J$, the sign pattern $\Delta^0 < 0$ and $\Delta^1 < 0$ describes an anti-coordination region: a firm prefers outsourcing when its rival remains in-house, but prefers in-house production when its rival outsources. Supplementary Appendix S6 examines equilibria beyond symmetric pure strategies.¹¹ The benchmark and policy analysis below therefore

¹¹ Asymmetric pure-strategy equilibria may exist, subject to deviation conditions evaluated at the asymmetric on-path effort pair. A symmetric mixed-strategy equilibrium also exists throughout the strict residual region when the conditional effort fixed point is unique and continuous in the organizational mixing probability. When asymmetric and mixed equilibria coexist we impose no selection criterion.

focuses on the symmetric pure-strategy regions.

Figure 1 plots the two organizational boundaries, $\Delta^0 = 0$ and $\Delta^1 = 0$, over a domain contained in $\mathcal{G}^J$; globally the frontier of $\mathcal{G}^J$ is a third equilibrium boundary. The two organizational boundaries are drawn in (β, λ) -space for the numerical example with $\mu = 0.3$, $\eta = 0.5$, $\kappa = 0.1$, and $\gamma = 1$. Two observations are worth highlighting. First, the region $\Delta^0 \geq 0$, in which the symmetric autarky equilibrium is sustained, is independent of the cross-client learning intensity λ . Second, the region $\Delta^1 \geq 0$, in which the symmetric joint-outsourcing equilibrium is sustained, need not expand monotonically with λ : for sufficiently high values of λ , stronger cross-client learning can make joint outsourcing less stable. The next sections formalize these comparative statics and then compare the equilibrium organizational outcomes with the coordinated benchmark allocation.




Figure 1 – Symmetric equilibrium regime map in (β, λ) -space, $(\mu, \eta, \kappa, \gamma) = (0.3, 0.5, 0.1, 1)$. Solid: $\Delta^0 = 0$. Dashed: $\Delta^1 = 0$.

4 Symmetric Autarky Equilibrium

We begin with symmetric autarky, in which both firms keep the hard task in-house. Each retains its own innovation technology but forgoes the shared platform's productivity gain and cross-client learning. This regime provides the starting point for analyzing whether a firm outsources when its rival remains in-house. We first characterize when neither firm wishes to outsource unilaterally, and how the autarky-supporting region, $\Delta^0 \geq 0$, varies with platform

productivity (η), product-market competition (β), and the platform fee (c). We then compare, conditional on autarky, decentralized hard-task effort with the coordinated benchmark. That comparison identifies the within-regime effort distortion, which arises because each firm ignores the rent its effort destroys for its rival. Whether autarky is the organizational form the benchmark would choose is an across-regime question, taken up in Section 6.

4.1 The Equilibrium-Supporting Region

The autarky-deviation gap Δ^0 determines whether autarky is self-sustaining: $(a_i, a_j) = (0, 0)$ is an equilibrium if and only if $\Delta^0 \geq 0$.

Proposition 1 (Comparative statics for the autarky-supporting region). *Suppose the relevant optimization problems admit interior solutions. Then:*

- (i) $\frac{\partial \Delta^0}{\partial \lambda} = 0$, that is, the autarky-supporting region is independent of λ .
- (ii) $\frac{\partial \Delta^0}{\partial \eta} < 0$, that is, the autarky-supporting region shrinks as η increases.
- (iii) $\frac{\partial \Delta^0}{\partial \beta} < 0$, that is, the autarky-supporting region shrinks as β increases.
- (iv) $\frac{\partial \Delta^0}{\partial c} > 0$, that is, the autarky-supporting region expands as c increases.

The incentive to deviate from autarky to unilateral AI outsourcing does not depend on the cross-client learning effect (Proposition 1-(i)): neither the firm's payoff under autarky, u^{00} , nor its payoff from a unilateral deviation to AI outsourcing, d^{10} , depends on λ . Thus, $\Delta^0 = u^{00} - d^{10}$ is independent of the cross-client learning effect.

Parts (ii) and (iii) of Proposition 1 show that each firm has a stronger incentive to deviate from autarky to AI outsourcing when either the direct AI productivity gain is larger or post-innovation competition is less severe. First, a deviating firm can raise its effective input from r_i to $(1+\eta)r_i$ by using the AI platform, thereby increasing its expected payoff from the hard task. Second, a higher β raises both the autarky payoff and the payoff from unilateral outsourcing, since more of the prize survives when both firms succeed. However, the effect is stronger on the deviation path, because AI assistance raises the deviating firm's success probability above its autarky level and so leaves that path more exposed to the event in which both firms succeed. Hence d^{10} rises by more than u^{00} , which shrinks the region in which autarky is sustained as an equilibrium. Lastly, as the AI service fee increases, each firm becomes less willing to outsource, expanding the parameter region that supports autarky as an equilibrium (Proposition 1-(iv)).

The findings in Proposition 1 are illustrated in Figure 1 for the numerical example. Since Δ^0 is strictly decreasing in β , the autarky-supporting set, when nonempty, is a lower interval within the range considered. In the running calibration it has an interior endpoint $\bar{\beta}_0$, determined by $\Delta^0(\bar{\beta}_0, \eta, c) = 0$, so autarky is sustained if and only if $\beta \leq \bar{\beta}_0$; monotonicity alone does not guarantee that such a crossing exists for every calibration.

4.2 Decentralized versus Coordinated Effort

How does decentralized effort compare with benchmark effort under autarky? Each firm's equilibrium effort r^I solves the first-order condition (I):

$$\phi = Vq'(r^I)[1 - (1 - \beta)q(r^I)].$$

Comparison with a coordinated benchmark requires the firms' combined payoff. For an organizational profile $(a_1, a_2) \in \{0, 1\}^2$ and effort pair $(r_1, r_2) \in [0, 1]^2$, we define the combined payoff function as

$$W(a_1, r_1; a_2, r_2) \equiv \phi(1 - r_1) + \phi(1 - r_2) - c(a_1 + a_2) + Vq_1[1 - (1 - \beta)q_2] + Vq_2[1 - (1 - \beta)q_1], \quad (12)$$

Here $q_i = q(x_i)$, with x_i given by (1). Thus W sums routine-task and expected innovation payoffs, net of total fees paid for the AI platform service.¹² The combined payoff under autarky is then equal to $W(0, r_1; 0, r_2)$.

The problem is symmetric in the two firms, but symmetry of the primitives does not by itself make a global maximizer symmetric. At an interior solution, the first-order conditions are

$$\phi = Vq'(r_i)[1 - 2(1 - \beta)q(r_j)], \quad i \neq j, \quad i, j \in \{1, 2\}.$$

Since $q(r) = 1 - e^{-\gamma r}$ and $q'(r) = \gamma e^{-\gamma r}$, the two conditions read

$$\frac{\phi}{V\gamma} = e^{-\gamma r_1}[(2\beta - 1) + 2(1 - \beta)e^{-\gamma r_2}] = e^{-\gamma r_2}[(2\beta - 1) + 2(1 - \beta)e^{-\gamma r_1}],$$

implying $(2\beta - 1)(e^{-\gamma r_1} - e^{-\gamma r_2}) = 0$. Hence, for $\beta \neq 1/2$, any interior stationary point satisfies $r_1 = r_2$; at $\beta = 1/2$ a symmetric stationary point exists but need not be unique. This

¹²The fee c enters the firms' combined payoff whether it reflects a real service cost or a transfer to the platform, since the client firms pay it either way. A social-welfare calculation would instead count the platform's real resource cost and include consumer surplus.

rules out asymmetric *interior* stationary points, but not asymmetric boundary optima, and for $\beta < 1/2$ the symmetric point need not be globally optimal, because concentrating effort on one firm can reduce rent dissipation.¹³ Consistent with the paper's focus on symmetric decentralized equilibria, we therefore restrict the coordinated autarky problem to symmetric effort and write

$$W(0, r; 0, r) = 2\phi(1 - r) + 2Vq(r)[1 - (1 - \beta)q(r)]. \quad (13)$$

Let r_B^I denote its maximizer on $[0, 1]$. When it is interior it satisfies

$$\phi = Vq'(r_B^I)[1 - 2(1 - \beta)q(r_B^I)], \quad (\text{B-I})$$

and otherwise it is characterized by the corresponding Kuhn–Tucker condition at the relevant boundary of $[0, 1]$. Because the platform fee is fixed conditional on the organizational regime, it drops out of the first-order condition for effort, though it remains part of W itself. The following result covers both symmetric regimes at once, autarky at $a = 0$ and joint outsourcing at $a = 1$.

Lemma 1 (Single crossing of the coordinated objective). *Fix $\beta \in (0, 1)$ and $a \in \{0, 1\}$. The derivative of $r \mapsto W(a, r; a, r)$ crosses zero at most once on $[0, 1]$, and any crossing is from positive to negative. Hence this map has a unique maximizer on $[0, 1]$: r_B^I when $a = 0$ and r_B^J when $a = 1$.*

Finally, we compare the decentralized equilibrium effort under autarky with the symmetric benchmark effort.

Proposition 2 (Excessive effort under autarky). *Under autarky, decentralized hard-task effort is weakly greater than the coordinated benchmark effort:*

$$r^I \geq r_B^I.$$

The inequality is strict whenever $r^I \in (0, 1)$; at a boundary, equality may occur.

Comparing (I) and (B-I) shows that the wedge between r^I and r_B^I arises because the benchmark internalizes rent dissipation whereas individual firms do not. A marginal increase

¹³At $(\mu, \eta, \kappa, \gamma) = (0.3, 0.5, 0.1, 1)$ and $\beta = 0.42$, the unrestricted maximizer is the asymmetric corner $(r_1, r_2) = (0, 1)$, whose value exceeds the symmetric benchmark value; the gap widens as β falls and vanishes for $\beta \geq 1/2$.

in r_i raises firm i 's probability of success, but it also reduces firm j 's expected prize when both firms succeed. This rent loss imposed on the rival is ignored by each firm in the decentralized problem, while the coordinated benchmark internalizes it. As a result, each firm allocates excessive effort to the hard task relative to the coordinated benchmark level. This excess is the intensive-margin counterpart of the business-stealing effect in [Mankiw and Whinston \(1986\)](#).¹⁴

5 Symmetric Joint-Outsourcing Equilibrium

We now turn to joint outsourcing, $(a_i, a_j) = (1, 1)$, where both firms use the AI platform and cross-client learning becomes active on the equilibrium path. This changes the technology of competition: a firm's effort improves its own innovation prospects but also contributes to the rival's effective input. Joint outsourcing therefore creates a knowledge spillover from shared learning together with rent dissipation from strengthening the rival.

We characterize the learning intensities that sustain this regime, how competition and platform productivity move its withdrawal boundaries, and whether effort conditional on joint outsourcing lies above or below the coordinated benchmark. Whether joint outsourcing is also the organizational regime preferred by the benchmark is a separate across-regime question, taken up in [Section 6](#).

5.1 Cross-Client Learning and Joint Outsourcing

Joint outsourcing is sustained only if two conditions hold. The symmetric effort r^J must be a global best response while both firms remain on the platform, and neither firm may prefer to withdraw and re-optimize its effort in-house. The sets $\mathcal{G}^J(\beta, \eta)$ and $\mathcal{D}^1(\beta, \eta)$ collect the cross-client learning intensities satisfying the first and second conditions, respectively, so the equilibrium-supporting set is their intersection,

$$\Lambda^1(\beta, \eta) = \mathcal{G}^J(\beta, \eta) \cap \mathcal{D}^1(\beta, \eta).$$

We characterize $\mathcal{D}^1$ first and impose the effort condition afterwards.

Proposition 3 (When cross-client learning is necessary). *Suppose $(1 + \eta) > \mu$ and that at $\lambda = 0$ both the symmetric platform effort and the effort chosen after withdrawal are interior*

¹⁴In their setting, equilibrium entry is excessive because entrants ignore the profit destruction they impose on incumbents; here, equilibrium effort is excessive because each firm ignores the rent destruction its effort imposes on the rival.

for every β in the range considered (Appendix A states the interiority conditions). Then $\Delta^1(\beta, 0, \eta)$ is weakly increasing in β , strictly when $\eta > 0$. Hence the set of β for which the withdrawal condition holds at $\lambda = 0$ is an upper interval. If this interval has a lower endpoint $\hat{\beta}(\eta)$ interior to the range considered, then:

- (i) for $\beta \geq \hat{\beta}(\eta)$, joint outsourcing is sustained without cross-client learning, provided $0 \in \mathcal{G}^J(\beta, \eta)$;
- (ii) for $\beta < \hat{\beta}(\eta)$, if $\mathcal{D}^1(\beta, \eta)$ is nonempty, sustaining joint outsourcing requires strictly positive cross-client learning.

At $\lambda = 0$, no knowledge passes between clients, so staying on the platform amounts to trading its stand-alone productivity advantage against the fee. Competition determines how valuable that advantage is. When β is high, joint success destroys less of the innovation rent, so improving the hard task is worth more and the productivity advantage is more likely to justify the fee. When β is low, joint success dissipates more of the rent, and the same advantage becomes less valuable. Cross-client learning is therefore needed, if at all, under sufficiently intense competition: not because a shared platform inherently requires learning across clients, but because learning must compensate for the weaker private value of staying. In the running example, the crossing is interior, at $\hat{\beta} \approx 0.538$.

The organizational condition also fails once cross-client learning becomes strong enough.

Proposition 4 (When cross-client learning becomes self-defeating). *Suppose $(1 + \eta) > \mu$ and $\beta \in (0, 1)$. Along the symmetric joint-outsourcing branch, as $\lambda \rightarrow \infty$, the candidate effort and the associated success probability vanish, $r^J \rightarrow 0$ and $q^J \rightarrow 0$, and $\Delta^1(\beta, \lambda, \eta)$ converges to a strictly negative limit that depends on neither β nor η . Hence, for every such β and η , $\mathcal{D}^1(\beta, \eta)$ and $\Lambda^1(\beta, \eta)$ are bounded above: sufficiently strong cross-client learning defeats joint outsourcing at every level of competition and platform productivity.*

The mechanism is the endogenous effort response. As λ grows, firms adjust how much effort they contribute to the increasingly strong learning link. Holding the firms' effort fixed, stronger cross-client learning makes each firm's problem solving more valuable to its rival, but it also makes contributing effort strategically more costly: a unit of effort increasingly raises the rival's probability of success and thereby erodes the innovator's own rent. If the success probability remained bounded away from zero as λ grew, this competitive harm would grow without bound while the direct productivity effect of own effort stayed bounded. Along the symmetric branch, effort therefore falls sharply: as is shown analytically in Appendix A, q^J

vanishes at rate $1/\lambda$, whereas r^J vanishes at rate $1/\lambda^2$. Hence the transferred contribution λr^J also vanishes.¹⁵ An increasingly powerful learning link is thus met by an even stronger withdrawal of effort, until there is almost nothing left for the technology to amplify. Staying eventually means paying the fee for a hard task that has effectively disappeared, while leaving avoids the fee and restores the option to innovate in-house.

The two propositions characterize the organizational condition at the two ends of the learning range but do not determine its global shape. The next lemma gives a sufficient condition under which $\Delta^1(\beta, \lambda, \eta)$ is weakly single-peaked in λ : it may decline from $\lambda = 0$ onward, or first rise and then decline, but it does not reverse direction more than once.

Lemma 2 (Single-peaked withdrawal gap). *Under the exponential specification, suppose $1 + \eta > \mu$, $\beta \in [\frac{1}{2}, 1)$, and $c > 0$. Then $\lambda \mapsto \Delta^1(\beta, \lambda, \eta)$ is weakly single-peaked on $[0, \infty)$, and the withdrawal-support set*

$$\mathcal{D}^1(\beta, \eta) = \{\lambda \geq 0 : \Delta^1(\beta, \lambda, \eta) \geq 0\}$$

is empty, a singleton, or a compact interval.

The condition on β is sufficient rather than necessary; it limits rent dissipation enough to rule out repeated reversals in the withdrawal incentive. Its economic content is that total innovation rent under joint success, $2\beta V$, is at least the monopoly prize V , so raising the common success probability never lowers the firms' gross innovation payoff. The maintained condition $1 + \eta > \mu$ makes the symmetric joint-outsourcing candidate strictly positive; it does not by itself imply $r^J < 1$, which requires the separate feasibility condition in Lemma A.1, and the lemma covers both the interior branch and the upper corner. For $\beta < \frac{1}{2}$ the conclusion can fail, and we verify single-peakedness numerically at every parameterization reported below that violates the condition.¹⁶ By Lemma 2, $\mathcal{D}^1(\beta, \eta)$ is either empty or a compact interval; when it is nonempty, we write its endpoints as

$$\mathcal{D}^1(\beta, \eta) = [\underline{\lambda}_D(\beta, \eta), \bar{\lambda}_D(\beta, \eta)],$$

as distinct from the effort cutoff introduced below. Every parameterization reported in the paper falls in the regular case.

¹⁵Appendix A establishes both rates, together with a strictly positive limit for λq^J ; for large λ the constrained candidate coincides with the interior branch, so no separate interiority assumption is needed.

¹⁶Across a grid spanning $\mu \in [0.1, 0.8]$, $\eta \in [0.2, 2]$, $\gamma \in \{0.5, 1, 2\}$, and $\kappa \in \{0.05, 0.1, 0.2\}$, we found no violation of single-peakedness at any $\beta \geq \frac{1}{2}$ and isolated violations at $\beta < \frac{1}{2}$.

It remains to impose the equilibrium effort condition. Whenever $\mathcal{G}^J(\beta, \eta)$ is an initial interval $[0, \lambda^G(\beta, \eta)]$, the equilibrium-supporting set, if nonempty, is

$$\Lambda^1(\beta, \eta) = \left[\underline{\lambda}_D(\beta, \eta), \min\{\bar{\lambda}_D(\beta, \eta), \lambda^G(\beta, \eta)\} \right]. \quad (14)$$

Our running example has this structure; the computation of λ^G is described in Supplementary Appendix S3.

The two cutoffs represent different ways in which joint outsourcing can fail. At $\lambda = \bar{\lambda}_D$, a firm would rather leave the platform and return to in-house innovation. At $\lambda = \lambda^G$, the firm may still prefer to remain on the platform, but the symmetric joint-outsourcing effort is no longer a global best response. Since joint outsourcing requires both conditions, its upper boundary is whichever cutoff is reached first. Figure 2 displays the resulting set Λ^1 in the running example.




Figure 2 – The equilibrium-supporting set Λ^1 (shaded) under $(\mu, \eta, \kappa, \gamma) = (0.3, 0.5, 0.1, 1)$ and $V = 1$. Panel (a): Δ^1 against λ at selected β ; dots mark $\Delta^1 = 0$. Panel (b): the axis tick marks $\hat{\beta} \approx 0.538$; the dotted vertical line and open circle mark $\beta^* \approx 0.752$; in the hatched region the withdrawal condition holds but the effort condition fails.

In the running example, the binding constraint switches at $\beta^* \approx 0.752$. Panel (b) shows that when $\beta < \beta^*$, competition is relatively intense and withdrawal becomes attractive first, so $\bar{\lambda}_D$ determines the upper boundary. As β rises, joint success destroys less of a firm's innovation rent, making withdrawal less attractive and pushing $\bar{\lambda}_D$ outward. Beyond β^* , the

global-effort condition is reached first, so λ^G determines the upper boundary. Figure 2 thus shows a change in the nature of the deviation that limits joint outsourcing: from leaving the platform to deviating in effort while remaining on it.

The comparative statics below study the withdrawal condition, so they apply only to boundaries created by $\Delta^1 = 0$. In the running example, this means the positive lower boundary and the upper boundary for $\beta < \beta^*$. They do not describe a lower boundary pinned at zero or an upper boundary determined instead by the global-effort condition.

5.2 Competition and Platform Productivity

We now ask how competition and platform productivity move the boundaries at which firms become willing to withdraw from the platform. The same boundary logic applies to both. At a boundary of $\mathcal{D}^1$ where Δ^1 crosses zero with nonzero slope, the implicit function theorem gives, for any primitive z ,

$$\frac{\partial \underline{\lambda}_D}{\partial z} = -\frac{\Delta_z^1}{\Delta_\lambda^1} \Big|_{\lambda=\underline{\lambda}_D}, \quad \frac{\partial \bar{\lambda}_D}{\partial z} = -\frac{\Delta_z^1}{\Delta_\lambda^1} \Big|_{\lambda=\bar{\lambda}_D}. \quad (15)$$

Since Δ^1 crosses zero from below at $\underline{\lambda}_D$ and from above at $\bar{\lambda}_D$, a primitive that raises the payoff from staying relative to withdrawing at both boundaries lowers the lower cutoff and raises the upper one, that is, $\mathcal{D}^1$ expands. If it lowers the relative payoff at both boundaries, $\mathcal{D}^1$ contracts. If the effects have opposite signs, the two cutoffs move in the same direction. When λ^G rather than $\Delta^1 = 0$ determines a boundary of Λ^1 , that boundary instead follows the global effort condition.¹⁷

Competition. Section 5.1 showed that weaker competition raises the value of staying when there is no cross-client learning. At $\lambda = 0$, withdrawal leaves the rival's capability unchanged, and the withdrawing firm, having given up the productivity boost, succeeds less often than a staying firm, so it is less exposed to joint success and benefits less when joint success becomes less costly.

By contrast, positive cross-client learning ($\lambda > 0$) creates a new margin: withdrawal now also weakens the rival, who loses the transferred input. The direct gain from a higher β is proportional to $(q^J)^2$ for a firm that stays and to $q^{01}\tilde{q}^J$ for one that withdraws, and a higher β

¹⁷Equation (15) applies at a regular zero, where $\Delta_\lambda^1 \neq 0$, and so does not cover $\underline{\lambda}_D = 0$. If $\Delta^1(\beta, 0, \eta; z) > 0$, continuity of Δ^1 in z makes that strict inequality persist under small perturbations and the lower boundary stays at zero. That $\lambda = 0$ continues to belong to $\mathcal{G}^J(\beta, \eta)$ must be checked separately: membership is the equality $G^J = 0$, which continuity does not preserve.

also moves the rival's equilibrium effort. With both channels active, the net effect of weaker competition is no longer straightforward.

Proposition 5 (When weaker competition stabilizes joint outsourcing). *Fix the remaining primitives and suppose r^J is interior. If*

$$\beta \geq \frac{1 + \lambda}{1 + \eta + \lambda}, \quad (16)$$

then $d\Delta^1/d\beta \geq 0$, with strict inequality whenever $\lambda > 0$ and the withdrawing firm's optimal effort is positive.

Condition (16) can be rewritten as $\eta \geq (1 - \beta)(1 + \eta + \lambda)$. It compares the platform's stand-alone productivity advantage with a competitive-loss term, which is larger when joint success destroys more rent and when a firm's effort has a stronger total effect on innovation. The condition therefore becomes easier to satisfy as η rises and harder as λ rises, because stronger cross-client learning makes the firm's effort more consequential for its rival. Because the threshold increases in λ , it suffices to check the condition at $\bar{\lambda}_D$. If it holds there and r^J is interior at both regular boundaries, it also holds at $\underline{\lambda}_D$, so $d\Delta^1/d\beta \geq 0$ at both. By (15), a marginal increase in β then lowers $\underline{\lambda}_D$ and raises $\bar{\lambda}_D$, expanding $\mathcal{D}^1$ from both ends.

Note that condition (16) is sufficient, not necessary: it requires $\beta \geq 1/(1 + \eta)$ at $\lambda = 0$, yet Proposition 3 delivers monotonicity there without this sufficient condition. Moreover, in the running example, $\mathcal{D}^1$ expands as competition weakens even where (16) fails at both boundaries, as depicted by the left panel of Figure 2.¹⁸

Platform productivity. Recall that a higher η makes the AI platform more effective at the hard task without learning, thereby raising the value of staying. By contrast, a higher η also raises the rival's success probability under joint outsourcing, and therefore, whenever $\beta < 1$, intensifies dissipation of rents from innovation through a higher likelihood that both firms succeed. Both channels influence the firms' equilibrium efforts, and neither unconditionally dominates.

Without cross-client learning ($\lambda = 0$), the contest between the two countervailing forces is transparent. To see this, let $y_0 \in (0, 1)$ denote the symmetric failure probability, characterized by

$$(1 + \eta) y_0 [\beta + (1 - \beta)y_0] = \mu,$$

¹⁸A parameterization in which the condition holds at both boundaries, with the boundaries moving as predicted, is reported in Supplementary Appendix S4.

which balances the marginal return to hard-task effort against its normalized opportunity cost $\mu = \phi/(V\gamma)$.¹⁹ Differentiating $\Delta^1(\beta, 0, \eta)$ with respect to η gives

$$\frac{1}{V} \frac{d\Delta^1}{d\eta} \Big|_{\lambda=0} = \frac{\mu}{(1+\eta)^2} \left[\underbrace{\ln \frac{1}{y_0}}_{\text{platform-value effect}} - \underbrace{\frac{\eta(1-\beta)y_0}{\beta + 2(1-\beta)y_0}}_{\text{competitive effect}} \right]. \quad (17)$$

The two terms in brackets capture the opposing effects of higher stand-alone AI productivity on firms' incentives to sustain joint outsourcing. The first term is positive, since $y_0 \in (0, 1)$: a more productive AI platform makes staying more valuable relative to withdrawing and giving up the AI platform's advantage. The second term is negative: greater AI platform productivity also strengthens the rival, making joint success more costly when competition dissipates innovation rents. This effect is proportional to $1 - \beta$ and vanishes as $\beta \rightarrow 1$. Which force dominates therefore depends on the intensity of product-market competition.

Proposition 6 (When platform productivity stabilizes joint outsourcing). *Suppose $(1 + \eta) > \mu$. For β sufficiently close to one and λ sufficiently close to zero, $d\Delta^1/d\eta > 0$.*

Proposition 6 identifies a region in which greater AI platform productivity unambiguously strengthens the incentive to remain jointly outsourced. When β is close to one, little rent is lost when both firms succeed, so making the rival more productive is not costly. The dominant effect of a higher η is then the AI platform advantage itself: a firm that stays becomes more productive, while a firm that withdraws gives up that advantage and continues to compete against a rival that retains it. When cross-client learning is also weak, this force dominates, and Δ^1 rises with η .

The implication for the withdrawal-support interval follows from (15), which converts the sign of $d\Delta^1/d\eta$ at each boundary into the direction that boundary moves. In the running example, numerical calculation confirms $\Delta_\eta^1 > 0$ at both regular boundaries of $\mathcal{D}^1$, so greater AI platform productivity lowers $\underline{\lambda}_D$ and raises $\bar{\lambda}_D$, and $\mathcal{D}^1$ expands from both ends. Whether Λ^1 expands also depends on the global-effort condition $\mathcal{G}^J$.

Nevertheless, the stabilizing effect of a higher η on joint outsourcing is not global. When competition is sufficiently intense and the AI platform is already sufficiently productive, a

¹⁹The marginal return is the product of the platform's productivity, $1 + \eta$; the marginal effect of effective input on success, y_0 , normalized by γ (since $q'(x) = \gamma y_0$); and the expected prize share retained upon success, $\beta + (1 - \beta)y_0$: the whole prize when the rival fails and only the fraction β when it also succeeds.

further increase in stand-alone AI productivity can strengthen the rival enough that the resulting rent-dissipation effect outweighs the firm's own productivity gain, yielding $d\Delta^1/d\eta < 0$ near $\lambda = 0$. More generally, $d\Delta^1/d\eta$ need not have the same sign at the two boundaries of the region supporting joint outsourcing. Supplementary Appendix S4 examines these two possibilities separately: it derives primitive sufficient conditions under which $d\Delta^1/d\eta < 0$ for sufficiently small λ , and then provides a parameterization in which $\Delta_\eta^1 < 0$ at the lower boundary but $\Delta_\eta^1 > 0$ at the upper boundary. In the latter case, an increase in η shifts both cutoffs toward stronger cross-client learning, rather than expanding the joint-outsourcing region in both directions as in the running example.

5.3 Decentralized versus Coordinated Effort

Section 4.2 compared the decentralized and coordinated effort conditions under autarky, as expressed by (I) and (B-I), where the only wedge arises from the rent-dissipation externality: each firm ignores the reduction in its rival's expected innovation rent caused by a higher probability of joint success. Joint outsourcing adds a second channel through cross-client learning: a firm's hard-task effort also raises the rival's effective input and success probability through the shared AI platform. The knowledge benefit adds to the rent-dissipation externality already present under autarky. The two external effects have opposite signs, so decentralized effort can lie above or below the coordinated benchmark.

Recall from (J) that, when interior, the decentralized joint-outsourcing effort satisfies

$$\phi = q'(x^J) V[(1 + \eta) - (1 - \beta)(1 + \eta + \lambda)q^J].$$

Each firm holds the rival's effort fixed, while recognizing that its own effort raises the rival's success probability through cross-client learning. Let r_B^J denote the maximizer of $r \mapsto W(1, r; 1, r)$ on $[0, 1]$. When interior, it satisfies

$$\phi = q'((1 + \eta + \lambda)r_B^J) V[(1 + \eta + \lambda) - 2(1 - \beta)(1 + \eta + \lambda)q((1 + \eta + \lambda)r_B^J)], \quad (\text{B-J})$$

and otherwise it is characterized by the corresponding Kuhn–Tucker condition at the relevant boundary of $[0, 1]$.

Comparing (J) with (B-J) isolates the two external effects. The coordinator values the contribution that firm i 's effort makes to firm j 's success through λ , which raises its marginal valuation of effort relative to the firm's. It also internalizes the rent loss that firm i 's success imposes on firm j , which lowers that valuation. Each firm registers the first effect only

through the reduction in its own innovation rent, and ignores the second entirely. Evaluating the coordinated marginal payoff at the decentralized effort gives the wedge

$$\begin{aligned}\Omega &\equiv \lambda - (1 - \beta)(1 + \eta + \lambda) q^J \\ &= \underbrace{\lambda[1 - (1 - \beta)q^J]}_{\text{knowledge benefit conferred on the rival}} - \underbrace{(1 - \beta)(1 + \eta) q^J}_{\text{rent loss imposed on the rival}},\end{aligned}\tag{18}$$

so that

$$\Omega > 0 \iff \frac{\lambda}{1 + \eta + \lambda} > (1 - \beta) q^J.\tag{19}$$

The left side of (19) is the share of a firm's effective input contributed by the rival's effort; the right side is the expected share of the rival's prize dissipated by the firm's own success.

Proposition 7 (Rent dissipation versus knowledge spillover). *Suppose $\lambda \in \mathcal{G}^J(\beta, \eta)$ and $\Delta^1 \geq 0$ (equivalently, $\lambda \in \Lambda^1$), and suppose the decentralized symmetric joint-outsourcing effort is interior. Evaluated at that effort, the coordinated marginal payoff is*

$$\left. \frac{\partial}{\partial r_i} W(1, r_i; 1, r_j) \right|_{r_i=r_j=r^J} = V q'(x^J) \Omega.$$

This wedge ranks decentralized and coordinated efforts under joint outsourcing: the decentralized effort falls short of the coordinated benchmark when $\Omega > 0$, exceeds it when $\Omega < 0$, and equals it when $\Omega = 0$.

At $\lambda = 0$, on the interior stationary branch, the wedge equals $-(1 - \beta)(1 + \eta)q^J < 0$. Whenever joint outsourcing is sustained at $\lambda = 0$, decentralized effort is therefore excessive: without cross-client learning only rent dissipation remains, exactly as under autarky.

Corollary 1 (Monotonicity of the effort wedge). *Suppose $(1 + \eta) > \mu$. Along the interior symmetric stationary branch, Ω is strictly increasing in λ and therefore crosses zero at most once. At supported values $\lambda \in \Lambda^1$ on this branch, decentralized effort exceeds coordinated effort below any such crossing and falls short above it.*

Along the stationary branch, both movements make the condition for insufficient effort in (19) easier to satisfy: the rival's contribution to the effective input rises with λ , while the associated success probability falls. This decline is an equilibrium response rather than a direct technological effect of stronger cross-client learning. Holding effort fixed, an increase in λ mechanically raises each firm's effective input through stronger pooling. In equilibrium,

however, firms respond by reducing their hard-task effort. When this endogenous effort reduction outweighs the direct pooling gain, both effective input and the probability of success decline as λ increases.

A fixed per-client fee enters neither marginal condition, and since it is paid in the candidate and in every on-platform deviation alike, it cancels out of the global comparison defining $\mathcal{G}^J(\beta, \eta)$. It therefore leaves r^J , q^J , Ω , and $\mathcal{G}^J(\beta, \eta)$ unchanged, affecting organizational payoffs but not effort conditional on joint outsourcing. Two qualifications matter. The crossing in Corollary 1 need not lie in Λ^1 , so the reversal is observed only if Λ^1 contains values of λ on both sides of it. And the sign of Ω ranks decentralized and coordinated effort at a given architecture without signing the combined-payoff effect of changing λ , which also moves q^J , the coordinated benchmark, and possibly the organizational regime.

6 Decentralized and Coordinated Organizational Choice

Sections 4 and 5 asked whether each symmetric regime is sustainable in equilibrium and how decentralized effort compares with benchmark effort within it. We now compare the organizational regimes under two criteria. In Section 6.1, we analyze how autarky, unilateral outsourcing, and joint outsourcing are ranked at the coordinated benchmark once effort is optimized within each regime. We further study how that ranking changes with the AI service fee, platform productivity, cross-client learning, and competition. In Section 6.2, we rank the same regimes as sustained by decentralized choice and identify where equilibrium delivers too little or too much outsourcing relative to the coordinated benchmark.

6.1 The Coordinated Benchmark

In our model, organizational choices are decentralized: each firm independently decides whether to use the AI platform, taking the rival's decision as given. By contrast, the coordinated benchmark maximizes the firms' combined payoff, comparing autarky, unilateral outsourcing, and joint outsourcing while optimally choosing the firms' effort allocations within each regime. The decentralized and coordinated outcomes may therefore differ because an individual firm does not internalize how its outsourcing decision affects the rival through product-market competition or cross-client learning, the latter operating only when both firms use the shared AI platform.

Recall the combined payoff $W(a_1, r_1; a_2, r_2)$ from (12). For each regime $m \in \{00, 10, 11\}$,

define its coordinated benchmark value by

$$\begin{aligned}
W_B(00) &\equiv \max_{r \in [0,1]} W(0, r; 0, r), \\
W_B(10) &\equiv \max_{(r_O, r_N) \in [0,1]^2} W(1, r_O; 0, r_N), \\
W_B(11) &\equiv \max_{r \in [0,1]} W(1, r; 1, r).
\end{aligned} \tag{20}$$

Since regimes 10 and 01 are symmetric, we focus throughout on regime 10 without loss of generality. We use r_O and r_N to denote the effort levels of the outsourcing and non-outsourcing firms, respectively. Under the coordinated benchmark, the optimal effort choices are r_B^I under autarky, (r_B^O, r_B^N) under unilateral outsourcing, and r_B^J under joint outsourcing. Each benchmark value depends on the full set of primitives; we display the arguments (β, λ) only when their dependence on these parameters is relevant to the analysis.

The comparison in (20) operates on two margins. Within each organizational regime, the coordinated benchmark chooses effort to balance the expected innovation payoff against the opportunity cost of effort. Across regimes, it compares the resulting optimized values net of the platform fees: zero under autarky, c under unilateral outsourcing, and $2c$ under joint outsourcing. Because the platform fee is fixed and does not depend on effort, changes in c do not alter the within-regime optimal effort choices. Instead, a higher c shifts the coordinated ranking toward regimes with fewer platform clients, lowering the combined payoff by c for each additional client.

The coordinated benchmark in (20) imposes symmetric effort in autarky and in joint outsourcing, the two regimes in which the firms occupy the same organizational role. It is therefore a restricted benchmark rather than the unrestricted coordinated optimum. Unilateral outsourcing is the only regime in which the two firms occupy different organizational roles—one firm outsources to the AI platform while the other keeps the hard task in-house—so the benchmark may assign them different effort levels. This restriction keeps the coordinated benchmark directly comparable to the symmetric equilibria analyzed above. An unrestricted coordinator, however, may prefer an asymmetric effort allocation even under a symmetric organizational regime. In particular, when competition is sufficiently intense, concentrating hard-task effort on one firm can reduce the probability of joint success and thereby mitigate rent dissipation.

Autarky and joint outsourcing differ in the effort required to generate a given effective input. Under autarky, effective input equals effort, $x = r$. Under symmetric joint outsourcing, each firm's effective input is $x = (1 + \eta + \lambda)r$. Because joint outsourcing lowers the effort

required to attain any given x , the coordinated benchmark chooses weakly more effective input under joint outsourcing, which implies a weakly higher success probability. The same comparison ranks the two values before fees: any effective input feasible under autarky is feasible under joint outsourcing at lower effort cost, so $W_B(11) + 2c \geq W_B(00)$, and autarky can be the benchmark regime only when the two fees outweigh the effort saved.

Proposition 8 (Competition and the benchmark ranking). *At the coordinated benchmark, weaker competition makes joint outsourcing weakly more attractive relative to autarky:*

$$\frac{\partial}{\partial \beta} [W_B(00) - W_B(11)] = 2V [q(r_B^I)^2 - q((1 + \eta + \lambda)r_B^J)^2] \leq 0.$$

A unit increase in β raises each firm's payoff by V in the state where both succeed and leaves every other state unchanged, so at fixed effort the combined payoff rises by $2V$ times the probability of joint success; the effort response has no first-order effect on the optimized value. Since joint outsourcing succeeds more often, it gains more from softer competition and loses more from harder competition. Intense competition therefore works in autarky's favor even though autarky is the less productive regime, because it dissipates more of the rent in the regime that succeeds more often.

No analogous ranking is available between autarky and unilateral outsourcing. Under unilateral outsourcing the two firms have different success probabilities, and the probability that both succeed, $q_O q_N$, may be either higher or lower than under autarky, so neither the effective-input ordering nor the sign of the derivative carries over.

The proposition holds η and λ fixed, and their incidence differs across regimes. Autarky depends on neither parameter. Unilateral outsourcing is weakly improved by η and unaffected by λ , since cross-client learning is inactive when only one firm is a client. Joint outsourcing is weakly improved by both, which enter only through their sum.²⁰

6.2 Decentralized Choice versus the Coordinated Benchmark

Ranking the three values in (20) gives the regime preferred by the coordinated benchmark,

$$m_B(\beta, \lambda) \in \arg \max_{m \in \{00, 10, 11\}} W_B(m; \beta, \lambda),$$

²⁰Under autarky $x_i = r_i$, so $W_B(00)$ involves neither parameter. Under unilateral outsourcing, raising η to η' and replacing any feasible r_O by $(1 + \eta)r_O / (1 + \eta') \leq r_O$ leaves the outsourcing firm's effective input unchanged and lowers its effort, so $W_B(10)$ weakly rises. The claims involving λ are Proposition 9; and since $x_i = (1 + \eta + \lambda)r$ under joint outsourcing, part (ii) of that proposition also gives monotonicity in η there.

and the three pairwise indifference loci, $W_B(11) = W_B(10)$, $W_B(11) = W_B(00)$, and $W_B(10) = W_B(00)$, contain its regime boundaries; a locus is active where the two tied regimes weakly dominate the third.²¹ Among the symmetric decentralized equilibria, let $m_E \in \{00, 11\}$ denote the selected regime. To display a single equilibrium outcome, we use a status-quo selection convention: whenever autarky and joint outsourcing are both sustained, which requires $\Delta^0 \geq 0$, $\Delta^1 \geq 0$, and $\lambda \in \mathcal{G}^J(\beta, \eta)$, we select autarky. This is consistent with best-response adjustment starting from the in-house organization, since no firm then outsources unilaterally. It is an equilibrium-selection convention, not a welfare ranking, and it affects the classification of the multiplicity region. For $m \in \{00, 11\}$, let $W_E(m)$ denote the firms' combined payoff evaluated at the decentralized equilibrium effort in regime m .²²

The shortfall from the coordinated benchmark decomposes as

$$W_B(m_B) - W_E(m_E) = \underbrace{[W_B(m_B) - W_B(m_E)]}_{\text{regime-selection gap}} + \underbrace{[W_B(m_E) - W_E(m_E)]}_{\text{within-regime effort gap}}.$$

Both terms are nonnegative, since m_B maximizes W_B and $W_B(m)$ maximizes over effort within regime m . The first measures the loss from selecting an organizational regime different from the benchmark regime. The second measures the loss from decentralized effort within the selected regime, and is characterized by Proposition 2 under autarky and Proposition 7 under joint outsourcing. The regime-selection gap is what remains. Cross-client learning is not its only source, but it creates a distinctive asymmetry: it enters the coordinated comparison of regimes while leaving the private incentive to outsource first untouched.

Proposition 9 (Cross-client learning and first outsourcing). *Hold the remaining primitives fixed.*

- (i) *The private incentive to become the first platform client, Δ^0 , and the benchmark values $W_B(00)$ and $W_B(10)$ do not depend on λ .*

²¹Details of the benchmark computation under the exponential specification are in Supplementary Appendix S5.

²²The two objects range over different sets by construction: the coordinated benchmark compares all three organizational forms, whereas the decentralized comparison is restricted to the symmetric equilibria studied in Sections 4 and 5. Where $\Delta^0 < 0$ and either $\Delta^1 < 0$ or $\lambda \notin \mathcal{G}^J(\beta, \eta)$, neither symmetric regime is an equilibrium, m_E is undefined, and the decomposition below does not apply. Supplementary Appendix S6 studies that region: it gives the conditions under which unilateral outsourcing is sustained as an asymmetric pure-strategy equilibrium, and shows that a symmetric mixed-strategy equilibrium exists under the stated uniqueness and continuity condition in which each firm outsources with probability $\rho^* \in (0, 1)$ and is indifferent between the two organizational modes.

(ii) $W_B(11)$ is weakly increasing in λ , strictly wherever $r_B^J > 0$.

For each $m \in \{00, 10\}$ the coordinated advantage $W_B(11) - W_B(m)$ is therefore weakly increasing in λ , while the decentralized first-outsourcing boundary $\Delta^0 = 0$ does not move with λ .

The learning channel is inactive until both firms outsource, so it offers nothing to a firm considering becoming the first client. Raising λ strengthens the coordinated case for joint outsourcing without helping the market reach it.

Decentralized and coordinated organizational choices can diverge in either direction. Whenever a symmetric equilibrium is selected, there is too little platform use if firms remain in autarky, $m_E = 00$, while the benchmark prefers at least one client, $m_B \in \{10, 11\}$, and too much platform use if both firms outsource, $m_E = 11$, while the benchmark prefers at most one, $m_B \in \{00, 10\}$.

The running calibration exhibits both disagreements. The benchmark never ranks autarky first, since $W_B(10) > W_B(00)$ for every β , so its only active boundary is $W_B(11) = W_B(10)$: at low cross-client learning the benchmark prefers unilateral outsourcing, at high learning joint outsourcing. For $\beta \leq \bar{\beta}_0$, where $\Delta^0 \geq 0$ and neither firm gains from becoming the first platform client, firms stay in autarky though the benchmark prefers platform use; stronger learning does not change that, since Δ^0 does not depend on λ . For β just above $\bar{\beta}_0$ and low λ , firms jointly outsource though the benchmark prefers a single client: the second client adds another fee and raises the probability of joint success, which dissipates rents, while contributing little cross-client learning. The two regions need not meet, and here they are separated by a band in which joint outsourcing is not sustained. Figure 3 shows the configuration.²³

A different organizational conflict arises without any comparison to the benchmark. Suppose joint outsourcing is sustained, $\lambda \in \mathcal{G}^J(\beta, \eta)$, and $\Delta^0 < 0 < \Delta^1$. At the relevant symmetric effort profiles, outsourcing is then a strict best response to either symmetric action of the rival, so joint outsourcing is the unique symmetric equilibrium. If $u^{00} > u^{11}$, however, both firms would prefer mutual autarky. The regime-payoff ordering then has a Prisoner's Dilemma form: outsourcing is preferred against either symmetric rival action, but mutual outsourcing leaves both firms worse off.

²³The figure is restricted to the region where symmetric pure-strategy equilibria exist and applies the status-quo selection convention. Autarky and joint outsourcing are both sustained only on a thin subset of the plotted range; an alternative selection would reclassify that subset alone, leaving the rest of the map unchanged.




Figure 3 – Benchmark regimes and decentralized-coordinated conflict, $(\mu, \eta, \kappa, \gamma) = (0.3, 0.5, 0.1, 1)$. Panel (a) shows the benchmark-preferred regime; the dashed curve is the coordinated boundary $W_B(11) = W_B(10)$. Panel (b) overlays this boundary and the vertical line $\beta = \bar{\beta}_0$, where $\Delta^0 = 0$, on the selected symmetric-equilibrium map. Both forms of decentralized-coordinated disagreement appear, separated by a region in which joint outsourcing is not sustained.

Example 1 (A Prisoner’s Dilemma in organizational choice). *In the running example, set $\beta = 0.50$ and $\lambda = 0.15$. The joint effort $r^J = 0.691$ is a global best response, so $\lambda \in \mathcal{G}^J(\beta, \eta)$. The private payoffs are*

$$u^{00} = 0.452, d^{10} = 0.457, u^{11} = 0.442, d^{01} = 0.433,$$

and hence

$$d^{10} > u^{00} > u^{11} > d^{01}.$$

The outer inequalities give $\Delta^0 = -0.005 < 0 < \Delta^1 = 0.009$, so outsourcing is preferred against each of the two symmetric on-path profiles. The middle inequality shows that both firms would nevertheless prefer mutual autarky, an outcome neither can reach through a profitable unilateral deviation.

The dilemma is not confined to this point. Across the computed grid, the ordering $d^{10} > u^{00} > u^{11} > d^{01}$ holds for β roughly in $[0.43, 0.75]$ and for λ up to about 0.3. It holds even at $\lambda = 0$, so cross-client learning is not necessary for it: stand-alone platform productivity makes unilateral outsourcing attractive, while the fees and rent dissipation of mutual outsourcing

leave both firms worse off. This is a comparison of private payoffs; it neither implies that the coordinated benchmark prefers autarky nor concerns the within-regime effort distortion.

Taken together, the analysis separates three margins: outsourcing by the first firm, outsourcing by the second, and effort after joint outsourcing. A uniform change in the per-client cost moves both outsourcing incentives in the same direction and therefore cannot target the first and second clients separately. Changing λ leaves the first-outsourcing incentive unchanged, while the post-outsourcing effort distortion is governed by Ω . Aligning one margin therefore need not align the others.

7 Implications for Shared AI Platforms

Policy may affect two features of shared AI platforms: access conditions shape the effective per-client cost c , while cross-client-learning rules shape λ .²⁴ An intervention may affect both features and may also change platform productivity η . The discussion below traces consequences for firms' outsourcing and effort.

Consumer benefits and the coordinated benchmark. The coordinated benchmark maximizes the two client firms' combined payoff, subject to symmetric effort when they occupy the same organizational role. The labels excessive and insufficient therefore refer to that objective. The model specifies no demand and therefore does not determine consumer surplus. Consumers may nevertheless benefit both when innovation creates a new product and when a second successful firm competes in supplying it.²⁵ Under that assumption, adding these consumer benefits to the firms' combined payoff weakly raises the optimal common effort within a fixed symmetric regime: any effort below the unique coordinated maximizer yields a lower combined payoff and no greater expected consumer surplus, so it cannot become optimal once those benefits are included. Effort exceeding the coordinated benchmark therefore need not exceed the level selected by this consumer-inclusive criterion.

The distinction is especially transparent when $\beta < \frac{1}{2}$: industry innovation profit falls from V under single success to $2\beta V$ under joint success. If competition transfers at least this lost profit to consumers, so that $CS_D - CS_M \geq (1 - 2\beta)V$, joint success generates weakly greater

²⁴The separation associated with c relies on its being a fixed charge. A usage-based charge would enter the effort first-order conditions.

²⁵For illustration, normalize consumer surplus to zero when neither firm succeeds, and let $CS_M \geq 0$ and $CS_D \geq CS_M$ denote consumer surplus when one and two firms succeed, respectively. Expected consumer surplus is $\mathbb{E}[CS] = CS_M[q_1(1 - q_2) + q_2(1 - q_1)] + CS_D q_1 q_2$, whose derivative with respect to q_i is $(1 - q_j)CS_M + q_j(CS_D - CS_M) \geq 0$. These state-contingent benefits are additional primitives; neither their levels nor their dependence on β is specified by the model.

total surplus despite yielding lower industry profit, holding innovation expenditures and platform costs fixed. This interpretation illustrates a potential divergence between the coordinated benchmark and social welfare; the required mapping from competition to consumer benefits is additional to the model.

These observations do not determine the welfare effect of competition or cross-client learning. The model does not specify how β affects consumer benefits. Moreover, firms adjust effort: by Lemma A.2, stronger cross-client learning lowers the common success probability along the supported interior joint-outsourcing branch. Under the fixed consumer-benefit assumptions above, expected consumer surplus therefore weakly falls. Sustaining joint outsourcing is consequently not sufficient for a welfare improvement. A full evaluation would also include the platform’s resource costs and treat payments between clients and the platform as transfers.

Access and switching costs. A firm may hesitate to rely on an outside AI platform if moving to another provider later would be costly. Policy addresses that cost in three ways: by capping or eliminating the switching charges a client pays on leaving, by requiring portability so that its data and models can be retrieved and transferred in usable form, and by requiring interoperability to reduce technical barriers to using those assets with another provider. The EU Data Act does all three. OMB guidance on federal AI procurement likewise directs agencies to address vendor lock-in through contract terms governing knowledge transfer, data and model portability, licensing, and pricing.²⁶

The model does not represent these switching frictions separately. To the extent that the preceding measures reduce the effective cost of using an outside platform, we capture their effect through a lower c . A lower c makes platform use more attractive whether or not the rival has already outsourced, and may therefore sustain joint outsourcing even where the coordinated benchmark prefers only one platform client.

These access provisions do not by themselves govern cross-client learning. The OMB guidance also restricts one channel through which such learning may occur: absent agency consent, contracts must prohibit vendors from using nonpublic agency data or outputs to train publicly or commercially available AI algorithms. The portability provisions correspond most closely to the access-cost margin, while the training restriction constrains one potential

²⁶Regulation (EU) 2023/2854, Arts. 25 and 29–30, applicable from 12 September 2025. Article 29 limits switching charges to costs directly linked to switching before 12 January 2027 and prohibits them from that date. Office of Management and Budget, Memorandum M-25-22, *Driving Efficient Acquisition of Artificial Intelligence in Government* (Apr. 3, 2025), §§ 3(e)(iv) and 4(d)(iii)(C).

channel of cross-client learning.

Cross-client learning and data-use restrictions. Competition enforcement against common intermediaries offers the closest analogue. RealPage supplies landlords with revenue-management software that recommends rents. The Department of Justice alleged that competing landlords provided RealPage with nonpublic data on rents, occupancy, lease terms, and concessions, which RealPage pooled and used to generate rent recommendations for its clients. A recommendation to one landlord could therefore reflect information obtained from its competitors, even though their underlying data were not directly disclosed to it.

The resulting Final Judgment limits the use of competitors’ data both in generating rent recommendations and in training the underlying models. For model training, such data generally must be at least twelve months old and cannot relate to active leases. The judgment also limits the geographic detail that models trained on such data may use, requires changes to product features, and establishes independent monitoring. The European Commission’s horizontal-cooperation guidelines identify similar safeguards for data sharing among competitors. The safeguards include using older or aggregated data, limiting access, and assigning data processing to a “clean team” separated from commercial decision-making or to an independent trustee.²⁷

The analogy has an important limit. These measures address the exchange of competitively sensitive market information and the resulting risk of coordinated conduct. In RealPage, competitors’ information informed pricing recommendations. The mechanism studied here does not require conveying information about a rival’s circumstances: solving one client’s problem can improve the platform’s ability to solve another’s, without disclosing the original information or coordinating market conduct.

These examples show that restrictions on cross-client data use need not be all or nothing. In the model, joint outsourcing is sustained only over a bounded range of cross-client learning. Too much learning can destabilize joint use, while complete isolation can do the same when some learning is needed to sustain it. A tighter restriction can therefore either stabilize or destabilize joint outsourcing, depending on the initial level of cross-client learning. Restrictions may also reduce stand-alone productivity η . Their evaluation must account for these effects on participation and effort; preserving joint platform use alone is not a sufficient policy criterion.

²⁷ *United States v. RealPage, Inc.*, No. 1:24-cv-00710 (M.D.N.C. May 19, 2026), Final Judgment, Doc. 194, especially §§ IV, V, and VII; European Commission, Guidelines on horizontal co-operation agreements, 2023/C 259/01, paras. 393–394, 406–408, and 418.

8 Conclusion

Shared AI platforms do more than improve productivity. When the same platform serves competing firms, it may learn from one client’s problem solving and use that knowledge to assist another. A firm’s outsourcing decision then depends not only on what the platform can do for it, but also on what its own use enables the platform to do for its rival.

Joint outsourcing is sustained only over a bounded range of cross-client learning: its productivity gains can support joint use, but once each firm’s problem solving benefits its rival enough, remaining on the platform becomes unattractive. Effort under joint outsourcing may be excessive or insufficient relative to the coordinated benchmark. Each firm ignores the rent its own success destroys for its rival, pushing effort up; it also goes uncompensated for the capability its effort confers on that rival, pushing effort down. Decentralized organizational choice can also depart from the coordinated ranking in both directions: firms may stay in-house when their joint payoff would be higher with platform use, or both may outsource when it would be higher with a single client.

These results do not by themselves rank policies. The coordinated benchmark maximizes the firms’ combined payoff subject to symmetric effort whenever they occupy the same organizational role; it is neither the unrestricted coordinated optimum nor a measure of social welfare, and the model specifies neither consumer demand nor endogenous platform behavior. The analysis identifies how access conditions and learning rules affect firms’ outsourcing and effort, not which regulation is socially optimal.

As shared AI systems become embedded in R&D, product design, legal strategy, and pricing, competing firms will increasingly depend on the same learning systems. For those firms, choosing a shared AI platform is not merely a make-or-buy decision. It is also a choice about how much of their own problem solving may become a rival’s capability.

References

- Acemoglu, Daron, Dingwen Kong, and Asuman Ozdaglar.** 2026. “AI, Human Cognition and Knowledge Collapse.” National Bureau of Economic Research Working Paper 34910.
- Aghion, Philippe, and Peter Howitt.** 1992. “A Model of Growth Through Creative Destruction.” *Econometrica*, 60(2): 323–351.
- Baccara, Mariagiovanna.** 2007. “Outsourcing, Information Leakage, and Consulting Firms.” *RAND Journal of Economics*, 38(1): 269–289.
- Brown, Zach Y., and Alexander MacKay.** 2023. “Competition in Pricing Algorithms.” *American Economic Journal: Microeconomics*, 15(2): 109–156.
- Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond.** 2025. “Generative AI at Work.” *Quarterly Journal of Economics*, 140(2): 889–942.
- Chen, Alvin.** 2026. “The Organizational Cost of Artificial Intelligence.” Preliminary draft, Stockholm School of Economics and Swedish House of Finance.
- Choi, Jay Pil.** 1994a. “Irreversible Choice of Uncertain Technologies with Network Externalities.” *RAND Journal of Economics*, 25(3): 382–401.
- Choi, Jay Pil.** 1994b. “Network Externality, Compatibility Choice, and Planned Obsolescence.” *Journal of Industrial Economics*, 42(2): 167–182.
- Choi, Jay Pil, and Marcel Thum.** 1998. “Market Structure and the Timing of Technology Adoption with Network Externalities.” *European Economic Review*, 42(2): 225–244.
- Choi, Jay Pil, Doh-Shin Jeon, and Domenico Menicucci.** 2026. “AI Safety and Competition.” Toulouse School of Economics TSE Working Paper 26-1745.
- Dasgupta, Partha, and Joseph Stiglitz.** 1980. “Industrial Structure and the Nature of Innovative Activity.” *Economic Journal*, 90(358): 266–293.
- d’Aspremont, Claude, and Alexis Jacquemin.** 1988. “Cooperative and Noncooperative R&D in Duopoly with Spillovers.” *American Economic Review*, 78(5): 1133–1137.
- de Bijl, Paul W. J., and Sanjeev Goyal.** 1995. “Technological Change in Markets with Network Externalities.” *International Journal of Industrial Organization*, 13(3): 307–325.

- Dell’Acqua, Fabrizio, Edward McFowland, III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Kraymer, François Candelon, and Karim 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 24-013.
- Farrell, Joseph, and Garth Saloner.** 1985. “Standardization, Compatibility, and Innovation.” *RAND Journal of Economics*, 16(1): 70–83.
- Farrell, Joseph, and Garth Saloner.** 1986. “Installed Base and Compatibility: Innovation, Product Preannouncements, and Predation.” *American Economic Review*, 76(5): 940–955.
- Gans, Joshua S.** 2026. “Market Power in Artificial Intelligence.” *Annual Review of Economics*.
- Garicano, Luis.** 2000. “Hierarchies and the Organization of Knowledge in Production.” *Journal of Political Economy*, 108(5): 874–904.
- Garicano, Luis, and Esteban Rossi-Hansberg.** 2006. “Organization and Inequality in a Knowledge Economy.” *Quarterly Journal of Economics*, 121(4): 1383–1435.
- Hagi, Andrei, and Julian Wright.** 2023. “Data-enabled learning, network effects, and competitive advantage.” *RAND Journal of Economics*, 54(4): 638–667.
- He, Guojun.** 2026. “The AI Lemons Problem.” SSRN Working Paper.
- Jones, Charles I., and Christopher Tonetti.** 2020. “Nonrivalry and the Economics of Data.” *American Economic Review*, 110(9): 2819–2858.
- Kamien, Morton I., Eitan Muller, and Israel Zang.** 1992. “Research Joint Ventures and R&D Cartels.” *American Economic Review*, 82(5): 1293–1306.
- Katz, Michael L., and Carl Shapiro.** 1986. “Technology Adoption in the Presence of Network Externalities.” *Journal of Political Economy*, 94(4): 822–841.
- Katz, Michael L., and Carl Shapiro.** 1992. “Product Introduction with Network Externalities.” *Journal of Industrial Economics*, 40(1): 55–83.
- Lee, Tom, and Louis L. Wilde.** 1980. “Market Structure and Innovation: A Reformulation.” *Quarterly Journal of Economics*, 94(2): 429–436.

- Liu, Ernest, Song Ma, and Laura Veldkamp.** 2025. “Data Sales and Data Dilution.” *Journal of Financial Economics*, 169: 104053.
- Loury, Glenn C.** 1979. “Market Structure and Innovation.” *Quarterly Journal of Economics*, 93(3): 395–410.
- Mankiw, N. Gregory, and Michael D. Whinston.** 1986. “Free Entry and Social Inefficiency.” *RAND Journal of Economics*, 17(1): 48–58.
- Merchant, Amil, Simon Batzner, Samuel S. Schoenholz, Muratahan Aykol, Gowoon Cheon, and Ekin Dogus Cubuk.** 2023. “Scaling Deep Learning for Materials Discovery.” *Nature*, 624(7990): 80–85.
- Milgrom, Paul, and Ilya Segal.** 2002. “Envelope Theorems for Arbitrary Choice Sets.” *Econometrica*, 70(2): 583–601.
- Noy, Shakked, and Whitney Zhang.** 2023. “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.” *Science*, 381(6654): 187–192.
- Reinganum, Jennifer F.** 1989. “The Timing of Innovation: Research, Development, and Diffusion.” In *Handbook of Industrial Organization*. Vol. 1, , ed. Richard Schmalensee and Robert D. Willig, Chapter 14, 849–908. Amsterdam:North-Holland.
- Schaefer, Maximilian, and Geza Sapi.** 2023. “Complementarities in Learning from Data: Insights from General Search.” *Information Economics and Policy*, 65: 101063.
- Xu, Zuojun, Feng Ren, Peng Wang, et al.** 2025. “A Generative AI-Discovered TNIK Inhibitor for Idiopathic Pulmonary Fibrosis: A Randomized Phase 2a Trial.” *Nature Medicine*, 31: 2602–2610.
- Yotzov, Ivan, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven J. Davis, Kevin M. Foster, Aaron Jalca, Brent H. Meyer, Paul Mizen, Michael A. Navarrete, Pawel Smietanka, Gregory Thwaites, and Ben Zhe Wang.** 2026. “Firm Data on AI.” National Bureau of Economic Research Working Paper 34836.

A Proofs

This appendix states two auxiliary lemmas used in the body and then proves the main-text results in order.

Auxiliary Lemmas for the Joint-Outsourcing Candidate

Lemma A.1 (Symmetric hard-task effort: candidate and existence). *Suppose $(1 + \eta) > \mu$. Write $y \equiv e^{-\gamma(1+\eta+\lambda)r}$ for the symmetric failure probability at common effort r , so that the symmetric first-order condition (J) is $F^J(y) = 0$ with*

$$F^J(y) \equiv (1 - \beta)(1 + \eta + \lambda)y^2 + [(1 + \eta)\beta - (1 - \beta)\lambda]y - \mu,$$

and let $\hat{y}^J$ denote its unique positive root and $\hat{r}^J \equiv -\ln \hat{y}^J / (\gamma(1 + \eta + \lambda))$ the corresponding unconstrained stationary effort.

- (i) *The profile $r^J = \min\{\hat{r}^J, 1\}$ is the unique symmetric candidate for a pure-strategy hard-task effort equilibrium. It satisfies $r^J = \hat{r}^J \in (0, 1)$ when $F^J(e^{-\gamma(1+\eta+\lambda)}) < 0$ and $r^J = 1$ when $F^J(e^{-\gamma(1+\eta+\lambda)}) \geq 0$; at equality $\hat{r}^J = 1$, so the two descriptions agree there.*
- (ii) *At $r = 0$, each firm has a profitable upward deviation, so the symmetric zero-effort profile is never an equilibrium.*
- (iii) *A symmetric pure-strategy hard-task effort equilibrium exists if and only if $\lambda \in \mathcal{G}^J(\beta, \eta)$, and it is then unique and equal to r^J .*

Proof. Direct substitution into (3) shows that a firm's marginal payoff with respect to its own effort, evaluated at the symmetric profile $r_i = r_j = r$, is

$$\left. \frac{\partial \pi_i(1, r_i; 1, r_j)}{\partial r_i} \right|_{r_i=r_j=r} = V\gamma F^J(y), \quad y = e^{-\gamma(1+\eta+\lambda)r}, \quad (\text{A.1})$$

so every statement about the sign of F^J is a statement about the sign of the marginal payoff.

For $\beta < 1$, the quadratic function F^J has a positive leading coefficient $(1 - \beta)(1 + \eta + \lambda) > 0$, so it is strictly convex. Furthermore, $F^J(0) = -\mu < 0$ and $F^J(1) = (1 + \eta) - \mu > 0$. Combining these together, F^J crosses zero only once at $\hat{y}^J < 1$, hence $\hat{r}^J > 0$ is the unique symmetric stationary effort. For completeness, at $\beta = 1$, $F^J(y) = (1 + \eta)y - \mu$ has the unique root $\mu/(1 + \eta) \in (0, 1)$.

Next, define $\underline{y} \equiv e^{-\gamma(1+\eta+\lambda)}$. Since $y(r) = e^{-\gamma(1+\eta+\lambda)r}$ is strictly decreasing in r , $\hat{r}^J \leq 1$ if and only if $\hat{y}^J \geq \underline{y}$. From the derivation of $\hat{y}^J$ above, $\hat{y}^J \geq \underline{y}$ is equivalent to $F^J(\underline{y}) \leq 0$,

while we assign $\hat{y}^J = \underline{y}$ and thus $\hat{r}^J = 1$ under equality $F^J(\underline{y}) = 0$. If $F^J(\underline{y}) > 0$, then we have $F^J(y) > 0$ for all $y \in [\underline{y}, 1]$, implying that the optimal y should be a corner solution: $r^J = \min\{\hat{r}^J, 1\}$. This proves part (i).

Furthermore, we have $y = 1$ at $r = 0$, so (A.1) gives a marginal payoff of $V\gamma[(1+\eta)-\mu] > 0$: raising effort from zero is profitable, which proves part (ii).

Lastly, any symmetric pure-strategy equilibrium is therefore either an interior outcome where each firm's stationary effort is $\hat{r}^J$ or an upper corner outcome which is a candidate only if $\hat{r}^J \geq 1$. In both cases, the equilibrium effort is r^J , and it is an equilibrium exactly when no deviation in $[0, 1]$ gives a strictly higher payoff, which is the case when $G^J = 0$, or equivalently, $\lambda \in \mathcal{G}^J(\beta, \eta)$. This proves part (iii). Q.E.D.

Lemma A.2 (Learning-intensity inversion). *Fix $\beta \in (0, 1)$, suppose $(1 + \eta) > \mu$, and restrict attention to the interior stationary branch $r^J = \hat{r}^J < 1$. Let y denote the failure probability as defined in (A.1). Then we have the following.*

(i) *The joint-outsourcing first-order condition is equivalent to*

$$\lambda(y; \beta, \eta) = \frac{(1 + \eta)y[\beta + (1 - \beta)y] - \mu}{(1 - \beta)y(1 - y)}. \quad (\text{A.2})$$

(ii) *$\lambda(\cdot; \beta, \eta)$ is strictly increasing on $(0, 1)$ and $\lim_{y \rightarrow 1} \lambda(y; \beta, \eta) = \infty$. Therefore, the interior stationary condition determines a unique failure probability y^J , which is strictly increasing in λ and converges to one as $\lambda \rightarrow \infty$. The corresponding hard-task effort r^J is strictly decreasing in λ and converges to zero as $\lambda \rightarrow \infty$.*

Proof. Part (i) is immediate from rearranging the symmetric first-order condition $\mu = y\{(1 + \eta)[\beta + (1 - \beta)y] - \lambda(1 - \beta)(1 - y)\}$ for λ .

For part (ii), differentiating (A.2) and simplifying gives

$$\lambda'(y) = \frac{(1 + \eta)y^2 - 2\mu y + \mu}{(1 - \beta)y^2(1 - y)^2}.$$

The numerator is strictly positive on $(0, 1)$: completing the square gives

$$(1 + \eta)y^2 - 2\mu y + \mu = (1 + \eta) \left(y - \frac{\mu}{1 + \eta} \right)^2 + \mu \left(1 - \frac{\mu}{1 + \eta} \right) > 0,$$

since $1 + \eta > \mu > 0$. The denominator is always positive on $(0, 1)$. Hence, $\lambda'(y) > 0$ for every $y \in (0, 1)$. Furthermore, the numerator of (A.2) remains positive as $y \rightarrow 1$, while the

denominator of (A.2) converges to zero, implying that $\lim_{y \rightarrow 1} \lambda(y; \beta, \eta) = \infty$.

Since $\lambda(y; \beta, \eta)$ is strictly increasing, it admits a well-defined inverse under the interior stationary condition. Therefore, we have

$$\frac{dy^J}{d\lambda} = \left(\frac{\partial \lambda(y; \beta, \eta)}{\partial y} \Big|_{y=y^J} \right)^{-1} > 0.$$

Moreover, since $\lambda(y; \beta, \eta) \rightarrow \infty$ as $y \rightarrow 1$, we have $\lim_{\lambda \rightarrow \infty} y^J = 1$.

The monotonicity of effort requires a separate argument because y^J depends on both r^J and the effective-input coefficient $1 + \eta + \lambda$. Recall that

$$r^J = -\frac{\ln y^J}{\gamma(1 + \eta + \lambda)},$$

which yields

$$\frac{dr^J}{d\lambda} = \frac{\ln y^J - (1 + \eta + \lambda) \frac{1}{y^J} \frac{dy^J}{d\lambda}}{\gamma(1 + \eta + \lambda)^2}.$$

Since $y^J \in (0, 1)$ along the interior stationary branch, we have $\ln y^J < 0$. Moreover, the preceding argument establishes $dy^J/d\lambda > 0$. Taken together, both terms in the numerator are negative, implying $dr^J/d\lambda < 0$. Thus, stronger cross-client learning strictly reduces the equilibrium hard-task effort along the interior stationary branch. *Q.E.D.*

Proof of Proposition 1

To prove part (i), recall that cross-client learning remains inactive under autarky whether any firm outsources or not. Hence, from (4) and (6), neither d^{10} nor u^{00} has λ in its objective function, implying $\partial u^{00}/\partial \lambda = \partial d^{10}/\partial \lambda = 0$ and thus $\partial \Delta^0/\partial \lambda = \partial u^{00}/\partial \lambda - \partial d^{10}/\partial \lambda = 0$.

For the remaining parts, define $y^I \equiv e^{-\gamma r^I}$ and $q^I = 1 - y^I$. With these notations, we can express the share of the prize each firm expects to keep conditional on its own success under autarky as $1 - (1 - \beta)q^I = \beta + (1 - \beta)y^I$. Then the first-order condition (I) under autarky is expressed for y^I as follows:

$$y^I [\beta + (1 - \beta)y^I] = \mu. \tag{A.3}$$

Suppose a firm deviates to outsourcing while the other firm stays in-house. The deviating firm's optimal payoff will be then d^{10} under the effort r^{10} derived from the interior first-order condition of (6):

$$\phi = V\gamma(1 + \eta)e^{-\gamma(1 + \eta)r^{10}} [\beta + (1 - \beta)y^I].$$

Comparing it with (A.3), we have

$$e^{-\gamma(1+\eta)r^{10}} = \frac{y^I}{1+\eta}, \quad r^{10} = -\frac{1}{\gamma(1+\eta)} \ln\left(\frac{y^I}{1+\eta}\right), \quad q^{10} = 1 - \frac{y^I}{1+\eta}. \quad (\text{A.4})$$

Substituting (A.3) and (A.4) into $u^{00} - d^{10}$ and collecting terms, we have

$$\frac{\Delta^0}{V} = \frac{\mu}{1+\eta} \left[\eta(\ln y^I - 1) + \ln(1+\eta) \right] + \kappa. \quad (\text{A.5})$$

Proof of part (ii). y^I is determined by the first-order condition (A.3) without η . Hence, differentiating (A.5) with respect to η gives

$$\frac{1}{V} \frac{\partial \Delta^0}{\partial \eta} = \frac{\mu}{(1+\eta)^2} \ln\left(\frac{y^I}{1+\eta}\right) < 0,$$

where the last strict inequality follows from $y^I < 1 \leq 1+\eta$.

Proof of part (iii). (A.3) can be rewritten as

$$(1-\beta)(y^I)^2 + \beta y^I - \mu = 0.$$

Then the implicit function theorem implies

$$\frac{\partial y^I}{\partial \beta} = -\frac{y^I(1-y^I)}{\beta + 2(1-\beta)y^I} < 0.$$

Since β implicitly affects (A.5) through y^I , we have

$$\frac{1}{V} \frac{\partial \Delta^0}{\partial \beta} = \frac{\mu\eta}{1+\eta} \frac{1}{y^I} \frac{\partial y^I}{\partial \beta} \leq 0,$$

where the strict negativity follows from $\eta > 0$.

Proof of part (iv). Since $\partial d^{10}/\partial c = -1$ and $\partial u^{00}/\partial c = 0$, we have $\partial \Delta^0/\partial c = \partial u^{00}/\partial c - \partial d^{10}/\partial c = 1 > 0$. Q.E.D.

Proof of Lemma 1

At $a_i = a_j = a$ and $r_i = r_j = r$, the effective inputs of two firms in (1) also coincide at $x_i = x_j = x := (1 + (\eta + \lambda)a)r$, which gives

$$W(a, r; a, r) = 2\phi(1-r) - 2ca + 2Vq(x)[1 - (1-\beta)q(x)].$$

Define $y := e^{-\gamma x}$, which is strictly decreasing in r . Differentiating with respect to r gives

$$\frac{1}{2} \partial_r W(a, r; a, r) = -\phi + V\gamma(1 + (\eta + \lambda)a)y[(2\beta - 1) + 2(1 - \beta)y].$$

Since $\phi > 0$, $\partial_r W \geq 0$ only if $y[(2\beta - 1) + 2(1 - \beta)y] > 0$, which is equivalent to $(2\beta - 1) + 2(1 - \beta)y > 0$ since $y > 0$. For such y , we have

$$\frac{d}{dy} \left\{ y[(2\beta - 1) + 2(1 - \beta)y] \right\} = (2\beta - 1) + 4(1 - \beta)y = [(2\beta - 1) + 2(1 - \beta)y] + 2(1 - \beta)y > 0.$$

Furthermore, since $dy/dr = -\gamma(1 + (\eta + \lambda)a)y < 0$ and $(2\beta - 1) + 4(1 - \beta)y > (2\beta - 1) + 2(1 - \beta)y > 0$, we have

$$\partial_r^2 W(a, r; a, r) = 2V\gamma(1 + (\eta + \lambda)a)[(2\beta - 1) + 4(1 - \beta)y] \frac{dy}{dr} < 0,$$

Hence, the coordinated marginal payoff $\partial_r W(a, r; a, r)$ is strictly decreasing whenever it is nonnegative; in particular, $\partial_r^2 W(a, r; a, r) < 0$ whenever $\partial_r W(a, r; a, r) = 0$.

Together with the preceding argument, this implies that $\partial_r W(a, r; a, r)$ can cross zero at most once, and any such crossing is from positive to negative. Therefore, the coordinated payoff $W(a, r; a, r)$ is either monotone in r or single-peaked, and consequently has a unique global maximizer. Specifically, the maximizer is $r = 0$ if $\partial_r W(a, 0; a, 0) \leq 0$, $r = 1$ if $\partial_r W(a, 1; a, 1) \geq 0$, and otherwise the unique interior stationary point $r \in (0, 1)$ satisfying $\partial_r W(a, r; a, r) = 0$. *Q.E.D.*

Proof of Proposition 2

Under symmetric autarky $a_i = a_j = 0$, each firm's payoff is $\pi_i(0, r_i; 0, r_j)$ from (3). Let

$$M^I(r) \equiv \left. \frac{\partial \pi_i(0, r_i; 0, r_j)}{\partial r_i} \right|_{r_i=r_j=r}$$

denote the firm's marginal payoff of its own effort, evaluated at the symmetric profile, the autarky counterpart of the joint-outsourcing marginal payoff in (A.1).

Comparing M^I with the marginal coordinated payoff $\partial_r W$ gives, at every $r \in [0, 1]$,

$$\frac{1}{2} \partial_r W(0, r; 0, r) = M^I(r) - V(1 - \beta)q'(r)q(r) \leq M^I(r), \quad (\text{A.6})$$

where the weak inequality follows from $q' \geq 0$, $q \geq 0$, and $\beta < 1$. The wedge $M^I(r) -$

$\frac{1}{2} \partial_r W(0, r; 0, r)$ is the competitive externality: a firm does not internalize that raising its own success probability lowers the rival's expected innovation rent.

We next show that both $\partial_r W$ and $M^I(r)$ cross zero at most once, and only from above. For $\partial_r W$, this follows directly from Lemma 1 evaluated at $a = 0$. For the firm's marginal payoff, using $y = e^{-\gamma r}$, we can write

$$M^I(r) = -\phi + V\gamma y[\beta + (1 - \beta)y].$$

The term $y[\beta + (1 - \beta)y]$ is strictly increasing in y , while y is strictly decreasing in r . Hence $M^I(r)$ is strictly decreasing in r and therefore crosses zero at most once, necessarily from positive to negative. The pointwise ordering in (A.6) then implies single-crossing: $\partial_r W(0, r; 0, r) \leq 0$ if $M^I(r) = 0$, that is, the coordinated marginal payoff reaches zero at a weakly lower r than the private marginal payoff $M^I(r)$. Since each objective is either monotone or single-peaked, their respective maximizers satisfy $r_B^I \leq r^I$.

Note that the strict inequality $r_B^I < r^I$ holds for interior cases $r^I \in (0, 1)$, so that $M^I(r^I) = 0$. In such cases, (A.6) gives

$$\frac{1}{2} \partial_r W(0, r; 0, r) \Big|_{r=r^I} = -V(1 - \beta)q'(r^I)q(r^I) < 0,$$

where the strict negativity follows from $q(r^I) > 0$, $q'(r^I) > 0$, and $\beta < 1$. Since $\partial_r W(0, r; 0, r)$ crosses zero at most once and from above, it is therefore negative at and beyond r^I , implying that $r_B^I < r^I$.

Lastly, equality in $r_B^I \leq r^I$ can hold in boundary cases. If $r^I = 0$, then $q(r^I) = 0$, the wedge in (A.6) vanishes, and $r_B^I = 0 = r^I$. If $r^I = 1$, the ordering yields only $r_B^I \leq 1 = r^I$ which allows equality. Q.E.D.

Proof of Proposition 3

First, consider $\lambda = 0$, in which case the symmetric joint-outsourcing first-order condition with respect to the corresponding failure probability y^J is

$$(1 + \eta) y^J [\beta + (1 - \beta)y^J] = \mu, \tag{A.7}$$

implying $y^J = y_0$. Differentiating (A.7) in β at $\lambda = 0$, we have

$$\frac{\partial y_0}{\partial \beta} = -\frac{y_0(1 - y_0)}{\beta + 2(1 - \beta)y_0} < 0. \tag{A.8}$$

At $\lambda = 0$, even if one firm deviates and takes the hard task in-house, the rival's failure probability remains y_0 . Hence the deviating firm will retain the share $\beta + (1 - \beta)y_0$ of the prize of the success in the hard task. The deviating firm's first-order condition will be therefore

$$e^{-\gamma r^{01}} [\beta + (1 - \beta)y_0] = \mu. \quad (\text{A.9})$$

Combining with (A.7), we have $e^{-\gamma r^{01}} = (1 + \eta)y_0$. The respective efforts r^J and r^{01} are interior if and only if

$$y_0 > e^{-\gamma(1+\eta)} \text{ and } e^{-\gamma} < (1 + \eta)y_0 < 1. \quad (\text{A.10})$$

Since $\frac{\partial y_0}{\partial \beta} < 0$ by (A.8), the set of β satisfying (A.10) is an interval, so the hypothesis that (A.10) holds over the range of β considered can be verified at the endpoints of that range.

By definition $y^J = \exp[-\gamma(1 + \eta + \lambda)r^J]$ we have

$$\gamma r^J = -\frac{\ln y_0}{1 + \eta}$$

at $\lambda = 0$. Moreover, combining (A.7) and (A.9) gives $e^{-\gamma r^{01}} = (1 + \eta)y_0$, and hence

$$\gamma r^{01} = -\ln((1 + \eta)y_0).$$

Since the rival's success probability remains $1 - y_0$ under both joint outsourcing and deviation at $\lambda = 0$, the deviating firm faces the same expected prize share $1 - (1 - \beta)(1 - y_0) = \beta + (1 - \beta)y_0$. The deviating firm's success probability is $1 - y_0$ under joint outsourcing, whereas its success probability from the in-house deviation is $1 - (1 + \eta)y_0$. Therefore, Δ^1 is rewritten as

$$\frac{\Delta^1(\beta, 0, \eta)}{V} = \mu\gamma(r^{01} - r^J) + [(1 - y_0) - (1 - (1 + \eta)y_0)] [\beta + (1 - \beta)y_0] - \kappa.$$

Using $\gamma r^J = -\frac{\ln y_0}{1 + \eta}$ and $\gamma r^{01} = -\ln((1 + \eta)y_0)$, the effort term becomes

$$\mu\gamma(r^{01} - r^J) = \mu \left[-\ln((1 + \eta)y_0) + \frac{\ln y_0}{1 + \eta} \right] = -\mu \ln(1 + \eta) - \frac{\mu\eta}{1 + \eta} \ln y_0.$$

Furthermore, the difference in success probabilities is

$$(1 - y_0) - (1 - (1 + \eta)y_0) = \eta y_0.$$

Moreover, $(1 + \eta)y_0[\beta + (1 - \beta)y_0] = \mu$ follows from (A.7), implying

$$\eta y_0[\beta + (1 - \beta)y_0] = \frac{\mu\eta}{1 + \eta}.$$

Combining these expressions yields

$$\frac{\Delta^1(\beta, 0, \eta)}{V} = \frac{\mu\eta}{1 + \eta}(1 - \ln y_0) - \mu \ln(1 + \eta) - \kappa. \quad (\text{A.11})$$

Therefore, we have

$$\frac{1}{V} \frac{\partial \Delta^1(\beta, 0, \eta)}{\partial \beta} = -\frac{\mu\eta}{1 + \eta} \frac{1}{y_0} \frac{\partial y_0}{\partial \beta} \geq 0,$$

where the last inequality follows from $\frac{\partial y_0}{\partial \beta} < 0$ from (A.8) and holds strictly whenever $\eta > 0$. Over the range of primitives on which (A.10) holds, the set of β with $\Delta^1(\beta, 0, \eta) \geq 0$ is an upper truncated interval; when its lower endpoint is interior to that range, it is the cutoff $\hat{\beta}(\eta)$.

For case (i), take $\beta \geq \hat{\beta}(\eta)$, so $\Delta^1(\beta, 0, \eta) \geq 0$. Under the additional hypothesis $0 \in \mathcal{G}^J(\beta, \eta)$, both conditions hold at $\lambda = 0$, so joint outsourcing is sustained without cross-client learning. For case (ii), take $\beta < \hat{\beta}(\eta)$, so $\Delta^1(\beta, 0, \eta) < 0$. If $\mathcal{D}^1(\beta, \eta)$ is nonempty, continuity of Δ^1 in λ implies that $\mathcal{D}^1(\beta, \eta)$ is bounded away from zero, so sustaining joint outsourcing requires strictly positive cross-client learning. *Q.E.D.*

Proof of Lemma 2

See Supplementary Appendix S2 for the complete proof.

Proof of Proposition 4

Suppose the equilibrium effort under joint outsourcing satisfies $0 < r^J < 1$. Then, in terms of the corresponding failure probability y^J , the symmetric first-order condition can be written as

$$\lambda(y) = \frac{(1 + \eta)y[\beta + (1 - \beta)y] - \mu}{(1 - \beta)y(1 - y)}.$$

Since $\lambda(y)$ is strictly increasing and $\lim_{y \rightarrow 1} \lambda(y) = +\infty$, the failure probability satisfying the first-order condition must satisfy $y^J \rightarrow 1$ as $\lambda \rightarrow \infty$. Then it follows immediately that $q^J = 1 - y^J \rightarrow 0$. Moreover, it follows from the definition of y^J that $r^J = -\frac{\ln y^J}{\gamma(1 + \eta + \lambda)}$, so $y^J \rightarrow 1$ implies that $r^J \rightarrow 0$ as $\lambda \rightarrow \infty$. Hence, for all sufficiently large λ , the effort level satisfying the first-order condition lies strictly below the upper bound $r = 1$. The equilibrium effort

therefore coincides with this solution, so the above limits apply to the constrained problem as well.

Next, we characterize the rates at which q^J and r^J converge to zero. The first-order condition can be rearranged as

$$(1 - \beta)(1 + \eta + \lambda)q^J = (1 + \eta) - \frac{\mu}{1 - q^J}.$$

Since $q^J \rightarrow 0$ as $\lambda \rightarrow \infty$, the right-hand side converges to $(1 + \eta) - \mu$. Because the first-order condition holds for every λ , the left-hand side must converge to the same limit. Hence,

$$\lim_{\lambda \rightarrow \infty} (1 + \eta + \lambda)q^J = \frac{(1 + \eta) - \mu}{1 - \beta}.$$

Furthermore, to obtain the corresponding rate for effort, recall that $1 - q^J = e^{-\gamma(1+\eta+\lambda)r^J}$, which is equivalent to

$$\gamma(1 + \eta + \lambda)r^J = -\ln(1 - q^J) \iff r^J = \frac{-\ln(1 - q^J)}{\gamma(1 + \eta + \lambda)}.$$

Multiplying both sides of the last identity by $(1 + \eta + \lambda)^2$ yields

$$(1 + \eta + \lambda)^2 r^J = \frac{1}{\gamma}(1 + \eta + \lambda)[- \ln(1 - q^J)] = \frac{1}{\gamma}[(1 + \eta + \lambda)q^J] \frac{[- \ln(1 - q^J)]}{q^J}.$$

Because $\frac{-\ln(1-q^J)}{q^J} \rightarrow 1$ as $q^J \rightarrow 0$, we obtain

$$\lim_{\lambda \rightarrow \infty} (1 + \eta + \lambda)^2 r^J = \frac{(1 + \eta) - \mu}{\gamma(1 - \beta)}.$$

Finally, since $r^J \rightarrow 0$ and $q^J \rightarrow 0$, we have

$$u^{11} = \phi(1 - r^J) - c + Vq^J[1 - (1 - \beta)q^J] \longrightarrow \phi - c.$$

If one firm deviates by withdrawing from the joint-outsourcing regime, the rival's success probability will become $\tilde{q}^J = 1 - e^{-\gamma(1+\eta)r^J} \rightarrow 0$. Then the withdrawing firm's objective function is

$$\phi(1 - r) + V(1 - e^{-\gamma r})[1 - (1 - \beta)\tilde{q}^J],$$

which differs from $\phi(1 - r) + V(1 - e^{-\gamma r})$ by at most $V(1 - \beta)\tilde{q}^J$ uniformly in $r \in [0, 1]$. Since

the maxima of two functions differ by at most the supremum of their absolute difference,

$$d^{01} \longrightarrow \phi + V \max_{r \in [0,1]} \{1 - e^{-\gamma r} - \mu \gamma r\},$$

and that maximum is nonnegative because $r = 0$ is feasible and yields zero. Subtracting the two limits,

$$\lim_{\lambda \rightarrow \infty} \Delta^1(\beta, \lambda, \eta) = -c - V \max_{r \in [0,1]} \{1 - e^{-\gamma r} - \mu \gamma r\},$$

which depends on neither β nor η and is strictly negative because $c > 0$. Hence $\Delta^1 < 0$ for all sufficiently large λ , so $\mathcal{D}^1(\beta, \eta)$, and hence its subset $\Lambda^1(\beta, \eta)$, is bounded above. *Q.E.D.*

Proof of Proposition 5

Recall the following notations:

- $y^J = 1 - q^J$: the failure probability under the symmetric joint-outsourcing equilibrium;
- $y^{01} = 1 - q^{01}$: the failure probability of the deviating firm;
- $\tilde{y}^J = 1 - \tilde{q}^J$: the failure probability of the rival firm left alone in the AI platform when the deviating firm takes the hard task in-house;

Since r^J is assumed to be interior, the joint first-order condition is

$$\mu = y^J [(1 + \eta) - (1 - \beta)(1 + \eta + \lambda)(1 - y^J)]. \quad (\text{A.12})$$

Suppose one firm deviates from joint outsourcing to in-house production while the rival remains on the AI platform and chooses effort r^J . The deviating firm's payoff from effort r is

$$\pi^{01}(r; r^J) = \phi(1 - r) + V(1 - e^{-\gamma r}) \left[1 - (1 - \beta)(1 - e^{-\gamma(1+\eta)r^J}) \right].$$

Letting $y = e^{-\gamma r}$ and $\tilde{y}^J = e^{-\gamma(1+\eta)r^J}$, the payoff function is rewritten as

$$\pi^{01}(r; r^J) = \phi(1 - r) + V(1 - y) [\beta + (1 - \beta)\tilde{y}^J].$$

Hence, the deviating firm's marginal payoff is

$$\frac{\partial \pi^{01}(r; r^J)}{\partial r} = V\gamma [y\{\beta + (1 - \beta)\tilde{y}^J\} - \mu].$$

Therefore, if the optimal deviation effort r^{01} is interior, its first-order condition is

$$y^{01}[\beta + (1 - \beta)\tilde{y}^J] = \mu.$$

Then the associated Kuhn–Tucker conditions yield

$$(i) \ r^{01} = 0: y^{01} = 1 \text{ and } \beta + (1 - \beta)\tilde{y}^J \leq \mu;$$

$$(ii) \ r^{01} \in (0, 1): y^{01}[\beta + (1 - \beta)\tilde{y}^J] = \mu;$$

$$(iii) \ r^{01} = 1: y^{01} = e^{-\gamma} \text{ and } e^{-\gamma}[\beta + (1 - \beta)\tilde{y}^J] \geq \mu.$$

Step 1: $y^{01} \geq y^J$. We can write $(1 + \eta) - (1 - \beta)(1 + \eta + \lambda) = \beta(1 + \eta + \lambda) - \lambda$. Since $\beta(1 + \eta + \lambda) \geq 1 + \lambda$ from (16), we have $(1 + \eta) - (1 - \beta)(1 + \eta + \lambda) \geq 1$. Since $y^J \leq 1$, the bracket on the right-hand side of (A.12) satisfies

$$(1 + \eta) - (1 - \beta)(1 + \eta + \lambda)(1 - y^J) \geq 1 \geq \beta + (1 - \beta)\tilde{y}^J,$$

where the last inequality follows from $\tilde{y}^J \leq 1$.

We prove $y^{01} \geq y^J$ for the three cases of the Kuhn-Tucker condition separately. First, if $r^{01} = 0$, then $y^{01} = 1 \geq y^J$ trivially. Furthermore, if $r^{01} \in (0, 1]$, conditions (ii) and (iii) give

$$y^{01}[\beta + (1 - \beta)\tilde{y}^J] \geq \mu \geq y^J[\beta + (1 - \beta)\tilde{y}^J],$$

with equality in the first comparison when r^{01} is interior, which implies $y^{01} \geq y^J$ since $\beta + (1 - \beta)\tilde{y}^J > 0$.

Step 2: $(q^J)^2 \geq q^{01}\tilde{q}^J$. Since $1 + \eta \leq 1 + \eta + \lambda$, we have $\tilde{y}^J \geq y^J$ which is strict whenever $\lambda > 0$. By Step 1, $\tilde{y}^J \geq y^J$ gives $q^{01} \leq q^J$ and $\tilde{q}^J \leq q^J$, implying $(q^J)^2 - q^{01}\tilde{q}^J \geq 0$.

Step 3: $\frac{d\Delta^1}{d\beta} \geq 0$. Define $\mathcal{D} := (1 - \beta)(1 + \eta + \lambda)y^J + \frac{\mu}{y^J}$, which is strictly positive. Implicit differentiation of the first-order condition under symmetric joint outsourcing (A.12) gives

$$\frac{dr^J}{d\beta} = \frac{q^J}{\gamma\mathcal{D}} > 0,$$

To differentiate $\Delta^1 = u^{11} - d^{01}$ with respect to β , recall that each firm i under joint outsourcing chooses r_i to solve

$$\max_{r_i \in [0, 1]} \left\{ \phi(1 - r_i) - c + Vq((1 + \eta)r_i + \lambda r_j) \left[1 - (1 - \beta)q((1 + \eta)r_j + \lambda r_i) \right] \right\},$$

where $r_i = r_j = r^J$ in the symmetric equilibrium. If firm i deviates to in-house production, it instead chooses r to solve

$$\max_{r \in [0,1]} \left\{ \phi(1-r) + Vq(r) \left[1 - (1-\beta)q((1+\eta)r^J) \right] \right\},$$

with solution r^{01} .

The envelope theorem eliminates the effect of β through each firm's own optimal effort, but not the effect through the rival's equilibrium effort r^J . Hence,

$$\begin{aligned} \frac{1}{V} \frac{d\Delta^1}{d\beta} &= (q^J)^2 - q^{01}\tilde{q}^J \\ &\quad + q'((1+\eta+\lambda)r^J) \left[\lambda - (1-\beta)(1+\eta+\lambda)q^J \right] \frac{dr^J}{d\beta} \\ &\quad + (1-\beta)(1+\eta)q^{01}q'((1+\eta)r^J) \frac{dr^J}{d\beta}. \end{aligned}$$

Using $q(x) = 1 - \exp(-\gamma x)$ for effective input x , together with $\frac{dr^J}{d\beta} = \frac{q^J}{\gamma \mathcal{D}}$, the above expression simplifies to

$$\frac{1}{V} \frac{d\Delta^1}{d\beta} = [(q^J)^2 - q^{01}\tilde{q}^J] + \frac{q^J \mathcal{E}_\beta}{\mathcal{D}}, \quad (\text{A.13})$$

where the second term capturing the induced-effort channel has a numerator

$$\mathcal{E}_\beta := \lambda y^J [\beta + (1-\beta)y^J] - (1+\eta)(1-\beta) [y^J q^J - \tilde{y}^J q^{01}].$$

If $\mathcal{E}_\beta \geq 0$, both terms of (A.13) are nonnegative by Step 2. Suppose instead $\mathcal{E}_\beta < 0$. Since $\lambda y^J [\beta + (1-\beta)y^J] \geq 0$, we have

$$-\mathcal{E}_\beta \leq (1+\eta)(1-\beta) [y^J q^J - \tilde{y}^J q^{01}].$$

In particular, the term $y^J q^J - \tilde{y}^J q^{01}$ inside the brackets is strictly positive. Since $\mathcal{D} \geq (1-\beta)(1+\eta+\lambda)y^J \geq (1-\beta)(1+\eta)y^J$, we have

$$\frac{q^J(-\mathcal{E}_\beta)}{\mathcal{D}} \leq q^J \left[q^J - \frac{\tilde{y}^J}{y^J} q^{01} \right].$$

Combining this condition with (A.13), we have

$$\frac{1}{V} \frac{d\Delta^1}{d\beta} \geq (q^J)^2 - q^{01}\tilde{q}^J - (q^J)^2 + q^J \frac{\tilde{y}^J}{y^J} q^{01} = q^{01} \frac{\tilde{y}^J - y^J}{y^J} \geq 0,$$

where the last inequality follows from $\tilde{y}^J \geq y^J$. When $\lambda > 0$ and $r^{01} > 0$, we have $\tilde{y}^J > y^J$ and $q^{01} > 0$, implying that the inequality is strict. If $r^{01} = 0$, the bound degenerates and only the weak inequality survives. *Q.E.D.*

Proof of Proposition 6

We first evaluate the derivative at the boundary point $(\beta, \lambda) = (1, 0)$. At $\beta = 1$, a firm's payoff does not depend on whether its rival succeeds or not, so each firm under joint outsourcing solves the following maximization problem:

$$\max_{r \in [0,1]} \{ \phi(1-r) - c + Vq((1+\eta)r) \}.$$

Since the objective function under joint outsourcing is strictly concave in r and its derivative at $r = 0$ is

$$-\phi + V\gamma(1+\eta) = V\gamma[(1+\eta) - \mu] > 0,$$

the unique maximizer, denoted by $r_{1,0}^J$ to indicate the joint-outsourcing effort at $(\beta, \lambda) = (1, 0)$, is strictly positive. At $(\beta, \lambda) = (1, 0)$, the payoff under joint outsourcing is therefore

$$u^{11} = \max_{r \in [0,1]} \{ \phi(1-r) - c + Vq((1+\eta)r) \}.$$

By contrast, a firm that deviates to in-house production solves

$$d^{01} = \max_{r \in [0,1]} \{ \phi(1-r) + Vq(r) \},$$

which is independent of η . Hence,

$$\Delta^1(1, 0, \eta) = \phi(1 - r_{1,0}^J) - c + Vq((1+\eta)r_{1,0}^J) - d^{01}.$$

Differentiating with respect to η , the direct effect of η on the joint-outsourcing payoff is $Vq'((1+\eta)r_{1,0}^J)r_{1,0}^J$. Although the optimal effort $r_{1,0}^J$ also varies with η , the envelope theorem implies that its induced effect on the optimized payoff can be ignored. Moreover, d^{01} does not

depend on η . Therefore,

$$\left. \frac{d\Delta^1}{d\eta} \right|_{(\beta, \lambda) = (1, 0)} = V\gamma r_{1,0}^J e^{-\gamma(1+\eta)r_{1,0}^J} > 0$$

because $r_{1,0}^J > 0$. Both payoff problems depend continuously on (β, λ) , and strict concavity at $(\beta, \lambda) = (1, 0)$ implies that their optimal effort choices remain locally unique and vary continuously with (β, λ) . It follows that $d\Delta^1/d\eta$ is continuous in a one-sided neighborhood of $(1, 0)$. Since $\frac{d\Delta^1}{d\eta} > 0$ at $(\beta, \lambda) = (1, 0)$, the same strict inequality continues to hold for β sufficiently close to one and λ sufficiently close to zero. *Q.E.D.*

Proof of Proposition 7

Let $x^J = (1 + \eta + \lambda)r^J$ and $q^J = q(x^J)$, where r^J is the symmetric equilibrium effort under joint outsourcing. A representative firm's payoff as a function of its own effort, with the rival's effort fixed at r^J , is $\pi^{11}(\cdot; r^J)$ as defined in (9); the rival's effective input depends on that effort through cross-client learning. Differentiating $\pi^{11}(r; r^J)$ with respect to r , we have

$$\begin{aligned} \partial_r \pi^{11}(r; r^J) &= -\phi + V(1 + \eta) q'((1 + \eta)r + \lambda r^J) \left[1 - (1 - \beta)q((1 + \eta)r^J + \lambda r) \right] \\ &\quad - V(1 - \beta)\lambda q((1 + \eta)r + \lambda r^J) q'((1 + \eta)r^J + \lambda r). \end{aligned}$$

At $r = r^J$, both firms' effective inputs equal x^J . Since r^J is interior, the symmetric decentralized first-order condition at $r = r^J$ is

$$\phi = Vq'(x^J) [(1 + \eta) - (1 - \beta)(1 + \eta + \lambda)q^J]. \quad (\text{A.14})$$

The coordinated benchmark chooses a common effort to maximize

$$W(1, r; 1, r) = 2\phi(1 - r) - 2c + 2Vq((1 + \eta + \lambda)r) [1 - (1 - \beta)q((1 + \eta + \lambda)r)],$$

whose marginal payoff can be written as

$$\frac{1}{2} \partial_r W(1, r; 1, r) = -\phi + V(1 + \eta + \lambda)q'(x) [1 - 2(1 - \beta)q(x)], \text{ where } x = (1 + \eta + \lambda)r.$$

Since W is symmetric in its two firm arguments, $\frac{1}{2} \partial_r W(1, r; 1, r)$ equals the unilateral coordinated marginal payoff $\partial_{r_i} W(1, r_i; 1, r_j)$ evaluated at $r_i = r_j = r$. Evaluating at $r = r^J$ and

substituting (A.14) for ϕ , we have

$$\begin{aligned}\frac{1}{2} \partial_r W(1, r; 1, r) \Big|_{r=r^J} &= Vq'(x^J) \left[(1 + \eta + \lambda)(1 - 2(1 - \beta)q^J) \right. \\ &\quad \left. - ((1 + \eta) - (1 - \beta)(1 + \eta + \lambda)q^J) \right] \\ &= Vq'(x^J) [\lambda - (1 - \beta)(1 + \eta + \lambda)q^J] = Vq'(x^J) \Omega.\end{aligned}$$

Since $V > 0$ and $q'(x^J) > 0$, the coordinated marginal payoff at the decentralized effort has the sign of Ω . By Lemma 1 at $a = 1$ the coordinated objective has a unique maximizer r_B^J on $[0, 1]$ and its derivative crosses zero at most once, from positive to negative, so the sign of Ω locates r_B^J relative to r^J even when the maximizer is at a boundary. *Q.E.D.*

Proof of Corollary 1

By Lemma A.2-(ii), whenever the equilibrium effort under joint outsourcing satisfies $0 < r^J < 1$ and is characterized by the symmetric first-order condition, the corresponding failure probability y^J is strictly increasing in λ . Since $q^J = 1 - y^J$, we have $\frac{dq^J}{d\lambda} < 0$. Furthermore, differentiating (18) with respect to λ gives

$$\frac{d\Omega}{d\lambda} = 1 - (1 - \beta) \left[q^J + (1 + \eta + \lambda) \frac{dq^J}{d\lambda} \right] > 1 - (1 - \beta)q^J > 0,$$

where the first strict inequality follows from $\frac{dq^J}{d\lambda} < 0$ and the last inequality follows from $\beta \in (0, 1)$ and $q^J < 1$. Thus, Ω is strictly increasing in λ and can cross zero at most once. For values of λ at which joint outsourcing is supported as equilibrium and $0 < r^J < 1$, Proposition 7 then translates the sign of Ω into the stated ranking between the decentralized and coordinated effort levels. *Q.E.D.*

Proof of Proposition 8

Step 1: effective-input ordering. We first prove $(1 + \eta + \lambda)r_B^J \geq r_B^I$, the ordering asserted in the text preceding the proposition. Write each benchmark problem in terms of the effective input X : under autarky, the common effort equals X ; under joint outsourcing, an effective input X requires effort $X/(1 + \eta + \lambda)$. Substituting $r_1 = r_2$ into (12) and expressing the objective in terms of effective input X , both regimes reduce to maximizing

$$2Vq(X)[1 - (1 - \beta)q(X)] - tX,$$

up to the common additive constant 2ϕ and, under joint outsourcing, the additional fixed cost $-2c$. Under autarky, $X = r$ and hence $t = 2\phi$ with $X \in [0, 1]$. Under joint outsourcing, $X = (1 + \eta + \lambda)r$, so the corresponding unit cost of effective input is $t = \frac{2\phi}{1+\eta+\lambda}$ with $X \in [0, 1 + \eta + \lambda]$.

Let $x_B^I = r_B^I$ and $x_B^J = (1 + \eta + \lambda)r_B^J$ denote the respective effective inputs in two benchmarks; each is the unique maximizer of its problem by Lemma 1. Suppose $x_B^J < x_B^I$. Then we have $x_B^J < x_B^I \leq 1 \leq 1 + \eta + \lambda$, implying that x_B^J is feasible under autarky and x_B^I is feasible under joint outsourcing. Define $G(X) \equiv 2Vq(X)[1 - (1 - \beta)q(X)]$ as the gross innovation payoff, optimality in each problem gives

$$G(x_B^I) - 2\phi x_B^I \geq G(x_B^J) - 2\phi x_B^J \text{ and } G(x_B^J) - \frac{2\phi}{1 + \eta + \lambda} x_B^J \geq G(x_B^I) - \frac{2\phi}{1 + \eta + \lambda} x_B^I,$$

which implies

$$2\phi \left(1 - \frac{1}{1 + \eta + \lambda} \right) (x_B^J - x_B^I) \geq 0.$$

Since $\phi > 0$, the leading factor is strictly positive when $\eta + \lambda > 0$, thereby $x_B^J \geq x_B^I$. This contradicts the supposition. If $\eta + \lambda = 0$, the two benchmark problems coincide and thus $x_B^J = x_B^I$. Hence $(1 + \eta + \lambda)r_B^J \geq r_B^I$, and since q is increasing, $q((1 + \eta + \lambda)r_B^J) \geq q(r_B^I)$.

Step 2: derivative with respect to competition. The aggregate payoff function $W(a_1, r_1; a_2, r_2)$ in (12) in any coordinated benchmark can be rewritten as

$$W(a_1, r_1; a_2, r_2) = \phi(2 - (r_1 + r_2)) - c(a_1 + a_2) + V(q_1 + q_2) - 2V(1 - \beta)q_1q_2.$$

Hence, β enters only through the last term $-2V(1 - \beta)q_1q_2$. At any fixed effort pair, we have $\frac{\partial W}{\partial \beta} = 2Vq_1q_2$.

For $a \in \{0, 1\}$, the map $(r, \beta) \mapsto W(a, r; a, r)$ is continuously differentiable. Furthermore, the feasible set of the firms' effort choice $[0, 1]$ is compact and does not depend on β . Lastly, the optimal effort is unique by Lemma 1. By the envelope theorem (Milgrom and Segal, 2002, Theorem 3), $W_B(00)$ and $W_B(11)$ are differentiable in β ; both derivatives are equal to the partial derivative evaluated at the maximizer:

$$\frac{\partial W_B(00)}{\partial \beta} = 2Vq(r_B^I)^2, \quad \frac{\partial W_B(11)}{\partial \beta} = 2Vq((1 + \eta + \lambda)r_B^J)^2.$$

Subtracting these gives the displayed expression in the proposition. Since $q \geq 0$, Step 1 implies $q(r_B^I)^2 \leq q((1 + \eta + \lambda)r_B^J)^2$, so the difference is non-positive. Q.E.D.

Proof of Proposition 9

Proof of part (i). By (1), $x_i = (1 + \eta a_i)r_i + \lambda a_i a_j r_j$, so λ enters only through the product $a_i a_j$. That product is zero under autarky and under unilateral outsourcing, so λ appears in neither $W(0, r_1; 0, r_2)$ nor $W(1, r_O; 0, r_N)$, and hence in neither $W_B(00)$ nor $W_B(10)$. It is absent from u^{00} and d^{10} for the same reason, so $\partial \Delta^0 / \partial \lambda = 0$, as in Proposition 1-(i).

Proof of part (ii). Fix β and η , let $\lambda' > \lambda$, and let r maximize $s \mapsto W(1, s; 1, s)$ at λ . Set

$$r' = \frac{1 + \eta + \lambda}{1 + \eta + \lambda'} r \leq r,$$

which is feasible since $0 \leq r' \leq r \leq 1$. The common effective input is unchanged, $(1 + \eta + \lambda')r' = (1 + \eta + \lambda)r$, so both success probabilities, and with them the entire innovation component of W , take the same values at (λ', r') as at (λ, r) ; the fee $2c$ is also unchanged. Only the routine payoff differs, rising by $2\phi(r - r')$. Hence

$$W_B(11; \beta, \lambda') \geq W(1, r'; 1, r')|_{\lambda'} = W_B(11; \beta, \lambda) + 2\phi(r - r'),$$

so $W_B(11)$ is weakly increasing in λ , and strictly increasing whenever $r > 0$, since then $r' < r$.

For the final claim, $W_B(m)$ with $m \in \{00, 10\}$ does not depend on λ by part (i), so $W_B(11) - W_B(m)$ inherits the monotonicity of $W_B(11)$; and Δ^0 does not depend on λ , so the locus $\Delta^0 = 0$ is unchanged as λ varies. *Q.E.D.*