Dividing Residual Control over AI Learning Assets: Complementary Investment Across Firm Boundaries

Shaohui Wang* Siwei Wang†

August 2026

Abstract

Enterprise AI often becomes useful only after a firm organizes its proprietary knowledge and an external provider adapts the technology to that knowledge. Their collaboration creates a persistent, customer-specific *AI learning asset* that neither party brings to the relationship in finished form. We develop a bilateral incomplete-contracting model to examine how two forms of control over this asset affect the investments that create it. Cross-client isolation limits the provider's reuse of customer-induced learning. Post-exit continuity determines how much of the learned state the firm can keep operating after separation. Because neither party captures the full return to its contribution, both invest below the jointly efficient level. Greater continuity encourages firm investment, but it may reduce provider adaptation enough to lower the performance of the deployed AI system. When isolation is incomplete and the parties' contributions are complementary, more valuable cross-client reuse can increase both investments, AI performance, and joint value, even as the firm's investment shortfall widens. Assigning isolation and continuity separately can therefore achieve outcomes that a bundled rule excludes. The paper extends property-rights theory from control over existing assets to the governance of learning created through joint use.

Keywords: artificial intelligence; AI learning assets; property rights; hold-up; complementary investment; firm boundaries

JEL classification: D23; L22; M10; O32

*J. Mack Robinson College of Business, Georgia State University, Atlanta, GA, USA. Email: swang83@gsu.edu

†School of Law, Wake Forest University, Winston-Salem, NC, USA. Email: wangs25@wfu.edu

1. Introduction

Consider the customer-support deployment scenario studied by Brynjolfsson et al. (2025). An AI software firm supplied a GPT-based conversational assistant to a Fortune 500 business-process software company and its support agents. The provider fine-tuned the system on prior customer–agent conversations labeled by resolution and handling outcomes. During live chats, the assistant suggested replies and linked agents to the client's internal documentation, while agents retained final authority over what they sent.[1] The arrangement makes the division of productive inputs concrete. The client contributes conversation histories, internal documents, and the practices of experienced agents; the provider contributes the general model, fine-tuning, and operating software. The collaboration can create an asset that neither party brings to the relationship in finished form and whose productive value depends on complementary contributions from both sides.[2] Because the asset is formed only through use, its eventual technical components and economically consequential uses may be difficult to anticipate, specify, and verify when the parties contract. This is the central setting of incomplete-contract theory: when not every future use can be allocated ex ante, residual rights of control determine what each party can preserve after disagreement and thereby shape investment incentives before the asset is created (Grossman & Hart, 1986; Hart & Moore, 1990).[3] Depending on the contract and technical architecture, the provider may be able to apply customer-specific learning to other clients even when the firm cannot keep the adapted configuration operating after exit. For the firm, repository integration, evaluation design, and workflow redesign are relationship-specific investments whose value may be difficult to recover outside the focal deployment. Once made, much of this expenditure is sunk; if the firm expects to retain little control over post-exit operation, it captures less of the return to its contribution and will invest less (Grossman & Hart, 1986; Grout, 1984; Hart & Moore, 1990; Klein et al., 1978; Williamson, 1979). The provider faces a parallel incentive problem. A broad restriction on cross-client reuse may eliminate an outside return that would otherwise justify customer-specific engineering and adaptation, reducing the provider's willingness to invest (Che & Hausch, 1999; Kang et al., 2009; Lo et al., 2022). Different allocations of residual control therefore change not only who can use the completed asset after separation, but also how much each party is willing to contribute before that asset exists. This tension leads to

[1] Brynjolfsson et al. (2025) study the staggered rollout of this system to 5,172 customer-support agents. Their study estimates productivity and worker-learning effects; it does not report how rights over customer-specific learning were allocated. We use the setting to motivate the governance problem.

[2] We later define this persistent, customer-specific state as an *AI learning asset*.

[3] Following Grossman and Hart (1986) and Hart and Moore (1990), *residual rights of control* refer to authority over uses of an asset that are not fully allocated by the initial contract. In the present setting, allocating residual control does not mean physically dividing an informational asset. It means determining which party can exercise particular uses of the customer-specific learning state when continued cooperation is unavailable. Evidence from IT outsourcing likewise shows that task complexity limits how fully future contingencies can be specified and that even detailed contracts may remain incomplete in their coverage or enforcement (Y. Chen & Bharadwaj, 2009).

the central question of this paper: *How does dividing control over cross-client reuse and post-exit operation shape what each party is willing to contribute and, ultimately, how well the AI system they jointly create performs?*

Drawing on organizational learning, we call the persistent, customer-specific state retained in evaluation histories, retrieval structures, adapted configurations, and operating rules an *AI learning asset* (Argote & Ingram, 2000). The asset is one component of situated AI, the broader operating arrangement in which a general model is grounded in firm knowledge and integrated with surrounding work (Kemp, 2024). Performance also depends on current provider adaptation and other organizational complements (Grant, 1996).[4] The asset can therefore persist even when the firm lacks the interfaces, documentation, support, or routines needed to operate the deployed system independently.

We study two residual rights over this asset. *Cross-client isolation* limits provider reuse of customer-induced learning and the firm's exposure to diffusion. *Post-exit continuity* determines how much of the learned state the firm can preserve and operate after separation. The rights are separable. A no-improvement covenant can restrict reuse without delivering a runnable configuration, whereas learned-state delivery, compatible interfaces, documentation, or transition support can improve continuity without preventing reuse. Raw-data export alone is insufficient because it transfers an input rather than the configuration that uses it (Vijairaghavan et al., 2025).[5]

This allocation creates what we call the *reverse information paradox*[6]. Arrow's classic paradox places the seller at risk of revealing an idea before payment (Arrow, 1962). Here the buyer must reveal and codify its own context to make the purchased technology useful, while the resulting learning may remain reusable by the provider and unavailable for independent operation by the buyer. This outcome requires neither misconduct nor contractual breach. Because both contributions are sunk before the asset is complete, disagreement uses determine their marginal returns. The setting combines hold-up (Grossman & Hart, 1986; Grout, 1984; Hart & Moore, 1990) with cooperative investment: strengthening one party's claim can weaken the other's, while an outside use can both expose the firm and sustain provider effort (Che & Hausch, 1999; Kang et al., 2009; Lo et al., 2022).

[4]An AI learning asset is created or materially improved through the relationship, persists across uses, reflects the focal customer's context, and affects later operation or bargaining. It excludes unchanged source data, the provider's general model, one-off outputs, and current performance.

[5]We use "right" in an effective economic sense: contract terms, technical architecture, access credentials, and operating arrangements jointly determine what a party can actually do (Aghion & Tirole, 1997; Boudreau, 2010). Appendix C maps observable provisions to isolation and continuity. The evidence establishes feasible separability, not prevalence, enforceability, treatment strength, or causal effects.

[6]The phrase *reverse information paradox* is taken from Nadella, CEO of Microsoft, who argues that an AI buyer may disclose proprietary knowledge simply to make purchased intelligence useful. https://snscratchpad.com/posts/reverse-information-paradox We use the label more narrowly for the incomplete-contracting problem created by costly proprietary-knowledge investment, complementary provider adaptation, and divided residual control over cross-client reuse and post-exit operation.

We formalize this mechanism in a bilateral incomplete-contracting model. Before investing, the parties assign isolation and continuity. The firm chooses *proprietary-knowledge investment*, costly effort that makes its knowledge usable by the AI system; the provider chooses *provider adaptation*, complementary customer-specific technical and implementation effort. After these inputs form the AI learning asset, the parties bargain over focal operation. The model separates situated-AI performance from bilateral value net of investment costs and compares split control with a bundled rule that ties the two rights together.

Under the maintained stability conditions, both parties invest below the bilateral benchmark. Stronger continuity narrows the firm's hold-up exposure and raises proprietary-knowledge investment, but it shifts part of the return from continued focal operation away from the provider and can reduce adaptation. Situated-AI performance falls when the lost adaptation outweighs the added firm knowledge. Bilateral value can turn at a different point.

Productive reuse works through a different channel. Under incomplete isolation and positive complementarity, a more valuable cross-client use raises provider adaptation, which increases the return to firm investment. Both inputs, situated-AI performance, and bilateral value can rise even as the firm's shortfall from the bilateral first best widens. The joint opportunity expands faster than the firm's private incentive.

Separating the rights expands the feasible governance set. Isolation changes reuse, diffusion, and total value per unit of learning; continuity reallocates the claim on focal operation. A bundled rule can therefore exclude profitable combinations of the two. Split control weakly raises the best attainable bilateral value and yields a strict gain when the bundled optimum blocks an improving adjustment in one dimension.

The paper develops a property-rights theory of AI learning assets formed through joint use. Governance shapes the asset before allocating its returns: rights over reuse and continuity determine what the parties are willing to build across the firm boundary.

2. Joint Formation and Divided Residual Control

Enterprise AI develops through complementary contributions from the firm and the provider, which together create the learning asset studied here.

2.1. Complementary conversion situates a general technology

Possessing specialized knowledge does not make it available for productive use. Knowledge-based theories of the firm instead emphasize organizing arrangements that integrate dispersed expertise and make it usable in action (Grant, 1996; Kogut & Zander, 1992); which arrangement is appropriate depends on the knowledge problem being solved (Nickerson & Zenger, 2004). The distinction is visible in project work. From longitudinal records of a global software-services firm, Ethiraj et al.

(2005) distinguish knowledge of a particular client accumulated through repeated interaction from the more general project-management systems built through deliberate investment. In semiconductor manufacturing, Hatch and Dyer (2004) connect persistent learning differences to firm-specific investment in human capital. Both studies show why common access to a technology or production method need not remove the need for local systems that make knowledge repeatedly usable.

AI makes that conversion unusually visible. Employees must identify relevant repositories, resolve conflicts among sources, encode exceptions, construct evaluations, and redesign workflows. We call the costly firm-side work *proprietary-knowledge investment* and denote its date-specific flow by d. It is not the quantity of data disclosed. A large archive may remain unusable, whereas a small set of carefully adjudicated cases can change how a system behaves. The chess study by Krakowski et al. (2023) isolates how the arrangement for working with AI matters apart from raw human or engine strength. Yokoi et al. (2026) observe the representational work through which perfume experts translate tacit judgment for generative AI. We use d to aggregate these heterogeneous forms of proprietary-knowledge investment.

The provider confronts a different bottleneck. A general model does not arrive with the customer's connectors, retrieval design, routing rules, reliability tests, tool integrations, or implementation support. We denote this dated technical and implementation effort by e, *provider adaptation*. The effort flow is distinct from any adapter, index, or configuration retained as an output of that effort. In the model, the direct term in e represents current provider support, while the complementarity term captures how adaptation helps encode firm knowledge durably. Firm investment organizes what the system must learn about the firm; provider adaptation makes the deployment responsive to that organized knowledge. Their complementarity follows the relational view of jointly productive partner resources (Dyer & Singh, 1998). Research on algorithmic decision authority similarly shows that returns to a common technology depend on its surrounding allocation of authority (Kim et al., 2024).

2.2. Investment flows create an AI learning asset

Conversion and adaptation can affect more than the current transaction. Experience is retained in people, tools, tasks, routines, and the connections among them (Argote & Ingram, 2000; Levitt & March, 1988). Organizational learning consequently changes later performance through creation, retention, and transfer, even after the episode that produced the learning has ended (Argote & Miron-Spektor, 2011). The distinction between flow and stock is essential. Dierickx and Cool (1989) show that strategic asset stocks are built through investment flows and may be subject to time compression, erosion, and other accumulation frictions. In industrial settings, Argote et al. (1990) document both learning and depreciation; Darr et al. (1995) find that productivity knowledge transfers within, but not across, franchise organizations; and Benkard (2000) estimates economically consequential forgetting in aircraft production. Persistence, transferability, and independent operation must therefore be established separately.

The AI learning asset may be distributed across retrieval indexes, adapter weights, evaluation histories, memory, and workflow logic. Its identity follows from the history and future consequence of those components, not from a particular technical label.[7]

Situated AI is the operating arrangement in which the AI learning asset combines with current provider adaptation and the firm's other organizational complements (Kemp, 2024). We use Q for its operating performance. Persistence of the asset does not guarantee independent operation of the deployment because some complements may disappear after exit. The ByteDance study illustrates how cumulative AI learning can extend across tasks under unified authority (Wan et al., 2026); when an external provider participates, authority over the asset is divided across the firm boundary.

Property-rights reasoning enters at precisely this boundary. Control of an informational resource can alter incentives even when the resource is not a conventional physical asset (Brynjolfsson, 1994). Rajan and Zingales (1998) show that access to a critical resource can allocate power and incentives when conventional ownership is unavailable. Aghion and Tirole (1997) distinguish formal from real authority, while Foss et al. (2021) argue that returns to ownership depend on competence in exercising it. For enterprise AI, these distinctions imply that title to source data or to a general model cannot be presumed to allocate effective uses of the later AI learning asset. The relevant control may instead lie in contract terms, system architecture, access credentials, documentation, and the practical ability to run or reuse what the relationship has produced.

2.3. Disagreement uses and investment incentives

Much of the conversion work is costly to redeploy. Curated repositories, adjudicated evaluations, and redesigned routines are built around a particular deployment; provider engineering is similarly tailored to the customer's environment. Once these investments are sunk, the return to each contribution depends on what the investor can preserve without the other's cooperation. That difference between relationship value and next-best use creates a quasi-rent and exposes the investment to hold-up (Grout, 1984; Klein et al., 1978; Williamson, 1979).[8]

Empirical strategy research supports the importance of specificity while also warning against treating it as a sufficient statistic for governance. In 277 coal contracts, Joskow (1987) finds that greater relationship-specific investment is associated with longer commitments. Comparing

[7] Customer specificity is relational. A routing rule may look portable in isolation yet derive its value from the firm's taxonomy, approval process, and accumulated exceptions. A component bearing a customer name, by contrast, may contain no learning from use.

[8] Hold-up does not require contractual breach or opportunistic conduct. When an investment is sunk and its marginal quality cannot be fully contracted ex ante, later bargaining may leave the investor with less than the value created by its contribution. Anticipating that return, the investor may reduce its initial contribution.

Table 1. Construct boundaries and illustrative manifestations

Construct	Definition and boundary	Illustrative manifestations
Proprietary organizational knowledge	Knowledge held before the focal relationship creates customer-specific learning. It is an input to the learning asset, not the asset itself.	Internal documents and data; employee judgment; process knowledge; business rules; exceptions.
Proprietary-knowledge investment (d)	Costly firm work that makes proprietary knowledge usable by the AI system. It is not equivalent to data volume or a one-time disclosure.	Curating sources; connecting repositories; coding exceptions; constructing evaluations; adjudicating corrections; redesigning workflows.
Provider adaptation (e)	Customer-specific technical and implementation effort by the provider, distinct from the quality of its general model.	Engineering connectors and retrieval; configuring tools, routes, and adapters; testing reliability; supporting implementation.
AI learning asset ($\mathcal{L}$)	Persistent, customer-specific state created or materially improved through use and consequential for future performance or bargaining. It excludes unchanged source data, the generic model, one-off outputs that leave no state, and current performance itself.	Retained retrieval indexes (Lewis et al., 2020); evaluation histories; persistent memory (Park et al., 2023); trained adapter weights (Hu et al., 2022); routing and workflow configurations (Sculley et al., 2015).
Situated AI (performance Q)	A general AI technology grounded in firm knowledge, bounded to the focal domain, and integrated with the surrounding work. Q records the repeatable task performance produced when the learning asset operates with current adaptation and other required resources.	Reliable diagnosis, exception handling, tool use, escalation, and repeatable workflow execution.
Cross-client isolation (x)	Effective restriction on provider reuse of customer-induced learning and on associated diffusion beyond the focal relationship.	No-training/no-improvement restrictions; tenant segregation; retention limits; monitoring; technical separation.
Post-exit continuity (y)	Fraction of focal operating value supported by the AI learning asset that the firm can preserve and operate after disagreement. It is more demanding than raw-data export.	Learned-state delivery; documentation; transition support; self-hosting; escrow; independent operation.

Notes: The manifestations are illustrative rather than exhaustive. The model uses effective economic rights: a formal clause changes x only insofar as it changes feasible reuse or diffusion exposure, and changes y only insofar as it changes post-exit operation.

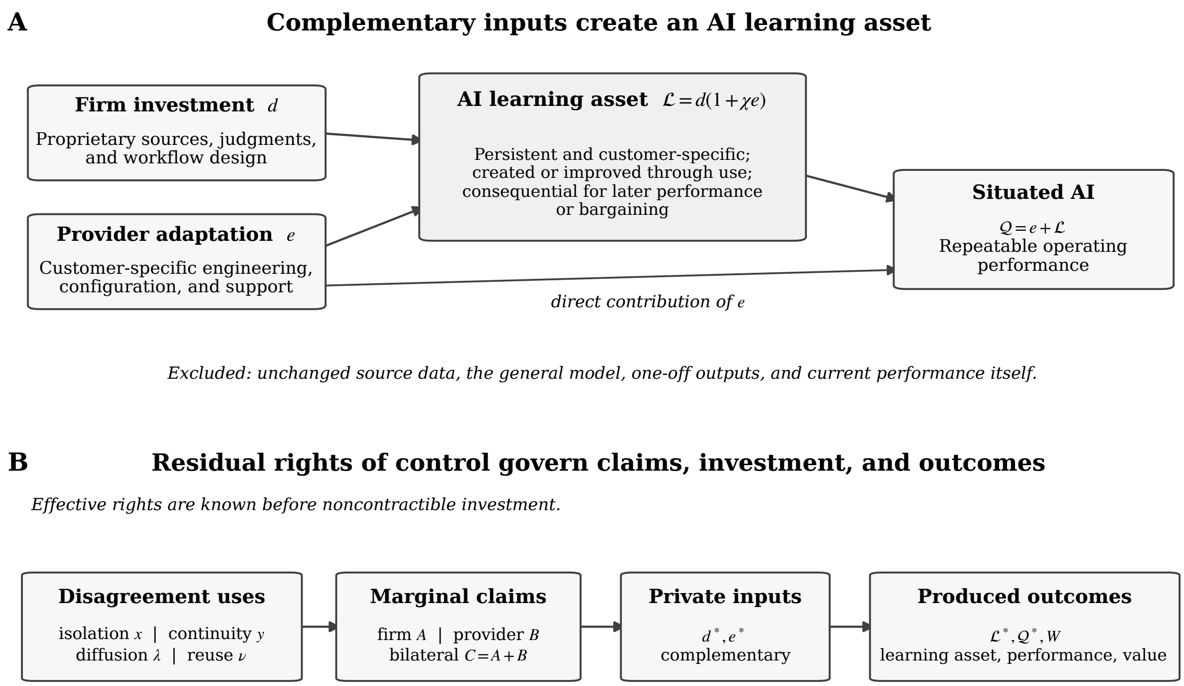

Figure 1. The AI learning asset and divided residual control. Panel A separates proprietary-knowledge investment and provider adaptation from the AI learning asset they jointly create and from situated AI performance. Provider adaptation enters performance both through the asset and directly. Panel B traces the model's sequence from residual rights of control to marginal claims, complementary investment, and produced outcomes.

automotive systems, Dyer (1996) shows how specialized supplier networks can combine relationship-specific investment with competitive advantage. Yet the systematic review by David and Han (2004) finds that empirical support for transaction-cost predictions varies with the prediction and its operationalization. Together, this evidence makes the uses available after disagreement central to any investment prediction.

Incomplete-contracting theory supplies that link. Grossman and Hart (1986) and Hart and Moore (1990) show that residual control changes investment by changing disagreement payoffs and therefore the marginal return expected before bargaining. A fixed transfer can divide expected surplus but cannot restore the return to an unverifiable marginal improvement. Nor does the classic ownership prescription yield an unconditional ranking: alternative bargaining assumptions can reverse comparative statics (de Meza & Lockwood, 1998), option arrangements can implement efficient investment in particular environments (Nöldeke & Schmidt, 1995), and vertical integration predictions require restrictions on how investments and assets interact (Whinston, 2003). The relevant question is thus which disagreement uses a governance arrangement changes.

Enterprise AI is also a cooperative-investment problem. Proprietary-knowledge investment makes provider adaptation more productive, and adaptation helps encode firm knowledge into the AI learning asset. Che and Hausch (1999) show why contracting can be weak when one party's investment directly benefits the other and surplus is later renegotiated. Strategy research adds that specificity need not uniformly discourage investment. Reusable knowledge and reputation can sustain commitments beyond the focal exchange, while competition, contract form, and relational commitment alter the return to adaptation and appropriation (Hann et al., 2025; Jia, 2013; Kang et al., 2009; Lo et al., 2022). Here those forces operate on an asset formed after contracting whose focal and cross-client uses can be controlled separately.

Two measures keep the mechanism precise. The firm's *marginal hold-up wedge* is the difference between the bilateral value created by another unit of learning and the marginal value the firm expects to capture from that unit. An *equilibrium investment shortfall* compares the chosen firm contribution with the bilateral benchmark after both parties' responses are allowed to adjust. A positive marginal wedge implies underinvestment conditional on provider effort under standard concavity; an equilibrium ordering additionally requires the response maps to be monotone and stable. This separation matters in a complementary system because a change that enlarges the wedge can simultaneously induce more provider effort and raise firm investment in levels.

2.4. Isolation and continuity as distinct dimensions of control

The choice of x and y begins with property-rights theory, not with a statistical reduction of contract terms. Grossman and Hart (1986) and Hart and Moore (1990) direct attention to decisions that an incomplete contract leaves open and to what each party can do if cooperation later breaks down.

Applied to an AI learning asset, this inquiry isolates at least two economically different disagreement uses. One is the provider's ability to carry customer-induced learning beyond the focal relationship; the other is the firm's ability to preserve and operate the customer-specific learned state after exit. The first is consequential because data are nonrival and learning can occur across users (Hagiu & Wright, 2023; Jones & Tonetti, 2020). The second is consequential because raw-data portability need not preserve derived inferences or other relational features that gave the data productive value (Vijairaghavan et al., 2025). These dimensions serve the mechanism developed here; they are not an exhaustive taxonomy of enterprise-AI governance.

The contract review in Appendix C.2 gives these theoretical dimensions observable content. Restrictions on training, service improvement, retention, cross-client use, and diffusion map to cross-client isolation. Learned-state delivery, compatible interfaces, self-hosting or escrow, documentation, transition assistance, and independent operation map to post-exit continuity. The reviewed documents combine these provision families in different ways, which supports treating them as separable dimensions of effective control. This clause-level mapping follows research showing that IT-outsourcing contracts combine distinct families of governance provisions (Y. Chen & Bharadwaj, 2009). Related evidence treats termination rights and vendor rights to reuse know-how as separately delineated decisions (Susarla, 2012). [9]

Accordingly, *cross-client isolation*, x, measures the effective restriction on provider reuse of customer-induced learning and on the focal firm's associated diffusion exposure. The baseline uses one isolation index because the same contractual and technical safeguards often affect both consequences; Appendix B.4 separates them. *Post-exit continuity*, y, records the share of focal operating value supported by the AI learning asset that the firm can preserve and run after disagreement. Continuity is more demanding than possession of raw data because the state may require compatible interfaces, documentation, transition support, and current complements before it is operable. Both x and y are effective economic rights: they summarize bundles of legal, technical, and organizational provisions without assigning a natural cardinal value to any single clause. In this sense, formal and effective authority can diverge, and granting access need not devolve control (Aghion & Tirole, 1997; Boudreau, 2010). Standardization can also redistribute uncertainty across ecosystem participants rather than simply eliminate it (Toh & Pyun, 2024).

Continuity improves the firm's fallback but reduces the continuation surplus over which the provider bargains. With complementary inputs, its net production and value effects are therefore

[9] The contract corpus grounds the constructs; it does not produce or measure the continuous variables x and y. It is purposive and single-coded, and its 0/1/2 scores are ordinal descriptions of located provisions. The codebook also records adaptation allocation and assurance. Adaptation allocation concerns rights in particular outputs, models, configurations, or other artifacts, but possession of an artifact need not determine either cross-client reuse or post-exit operation. Assurance concerns segregation, monitoring, auditability, and enforcement, and therefore determines whether a nominal restriction becomes effective rather than adding a third residual-control dimension. The evidence documents observable counterparts and feasible separability, not market prevalence, enforceability, or causal effects.

theoretically ambiguous and must be derived jointly rather than inferred from the firm's appropriation alone.

Isolation presents a different trade-off. Preventing diffusion protects the focal firm, but preventing reuse can remove a productive outside application. Pacheco-de-Almeida and Zemsky (2012) show that knowledge disclosure can alter competitive dynamics rather than act as a simple loss, while Kang et al. (2009) show how uses beyond a focal exchange can support specific investment. The tension here is therefore between diffusion protection and an outside opportunity that may strengthen provider adaptation.

Because continuity reallocates focal operation while isolation changes reuse and diffusion, the rights cannot generally be placed on one provider-to-firm scale. Differentiated safeguards already appear in interfirm governance (Hoetker & Mellewigt, 2009), and digital resources make the configuration of authority and redeployment especially consequential (Giustiziero et al., 2023; Sakhartov & Helfat, 2026). The model asks which investment configurations become implementable when these rights can be chosen separately.

3. A Model of Situated AI Governance

Consider a firm that knows its business but cannot, by itself, turn that knowledge into a working AI deployment. An external provider has the general model, infrastructure, and engineering expertise, but lacks the firm's repositories, operating judgments, and task-specific feedback. Each party therefore contributes different knowledge. This division follows Grant (1996), who treats the firm as an institution for integrating specialized knowledge, and Nickerson and Zenger (2004), who connect governance to the organization of knowledge search. In software services, Ethiraj et al. (2005) show that client-specific capabilities develop through repeated work, while Weigelt (2013) shows that the value of supplier capabilities depends partly on their interaction with client operational capabilities. We represent these contributions as noncontractible inputs made before bargaining over continued operation. Their interaction creates the AI learning asset defined in the preceding section. The model asks how effective rights over its later uses change the inputs used to create it.

3.1. Parties, rights, and timing

A focal firm F contracts with an AI provider P. The firm initially controls the proprietary knowledge required for the task; the provider controls the general model, technical infrastructure, and adaptation expertise. Both have quasilinear payoffs, and an unrestricted ex ante transfer separates participation from marginal investment incentives (Che & Hausch, 1999; Grossman & Hart, 1986; Hart & Moore, 1990). Because source curation and customer-specific engineering cannot be specified and verified exhaustively, control matters through the value each party can preserve if cooperation later breaks down (Grossman & Hart, 1986; Hart & Moore, 1990).[10]

At date 0, the parties agree on an up-front transfer T_0 and at least two effective rights. *Cross-client isolation* $x \in [0, 1]$ measures how fully the provider is prevented from reusing customer-induced learning elsewhere. *Post-exit continuity* $y \in [0, 1]$ measures the share of operating value supported by that learning which the firm can preserve and operate after disagreement. Thus $x = 1$ denotes full isolation and $y = 1$ full operational continuity. Treating access and control as separable follows Boudreau (2010), who shows that granting access to a digital platform and devolving control activate distinct mechanisms. It is also consistent with Y. Chen and Bharadwaj (2009), whose analysis of IT-outsourcing contracts identifies several distinct classes of governance provisions.

The unit interval is a normalization of effective control, not a cardinal scale for individual clauses. Legal terms, technical architecture, and operating arrangements can all contribute to either right.[11] Both parties know the effective (x, y) before they invest; this follows the standard timing in which control is allocated before noncontractible investment (de Meza & Lockwood, 1998; Grossman & Hart, 1986; Hart & Moore, 1990). The baseline therefore contains no rights uncertainty; Appendix B separately studies written isolation, hidden reuse, and assurance.

At date 1, the firm chooses proprietary-knowledge investment $d \geq 0$ and the provider simultaneously chooses customer-specific adaptation $e \geq 0$. At date 2, the inputs form the AI learning asset. At date 3, after both investments are sunk, the parties bargain over focal operation. If bargaining fails, each exercises the uses permitted by (x, y). The rights fixed at date 0 therefore affect date-1 investment through date-3 disagreement values. This contract–investment–bargaining sequence is standard in property-rights and cooperative-investment models (Che & Hausch, 1999; Grossman & Hart, 1986; Grout, 1984; Hart & Moore, 1990). The simultaneous choice abstracts from strategic timing between the two investments. Figure 2 depicts the sequence.

[10]The maintained contractibility restriction does not require the contributions to be hidden from the parties. They may observe enough about the realized system to bargain ex post while a court or other enforcer cannot verify the marginal quality of each contribution, or all of its later uses, well enough to condition an ex ante payment schedule on it. This distinction between observability and verifiability is central to incomplete contracting. In the closely related setting of IT outsourcing, Susarla et al. (2010) document how nonprice contract provisions respond to relationship-specific investment and hold-up risk.

[11]Following the distinction between formal and real authority in Aghion and Tirole (1997), x and y measure uses that a party can actually exercise, not legal title or clause counts. The continuous indexes permit comparative statics; the endpoints retain substantive interpretations. In particular, continuity is broader than statutory portability. Data portability may allow a user to download and transfer raw data (Vijairaghavan et al., 2025), whereas y also requires the interfaces, documentation, permissions, and complementary operating resources needed to keep the AI learning asset working.

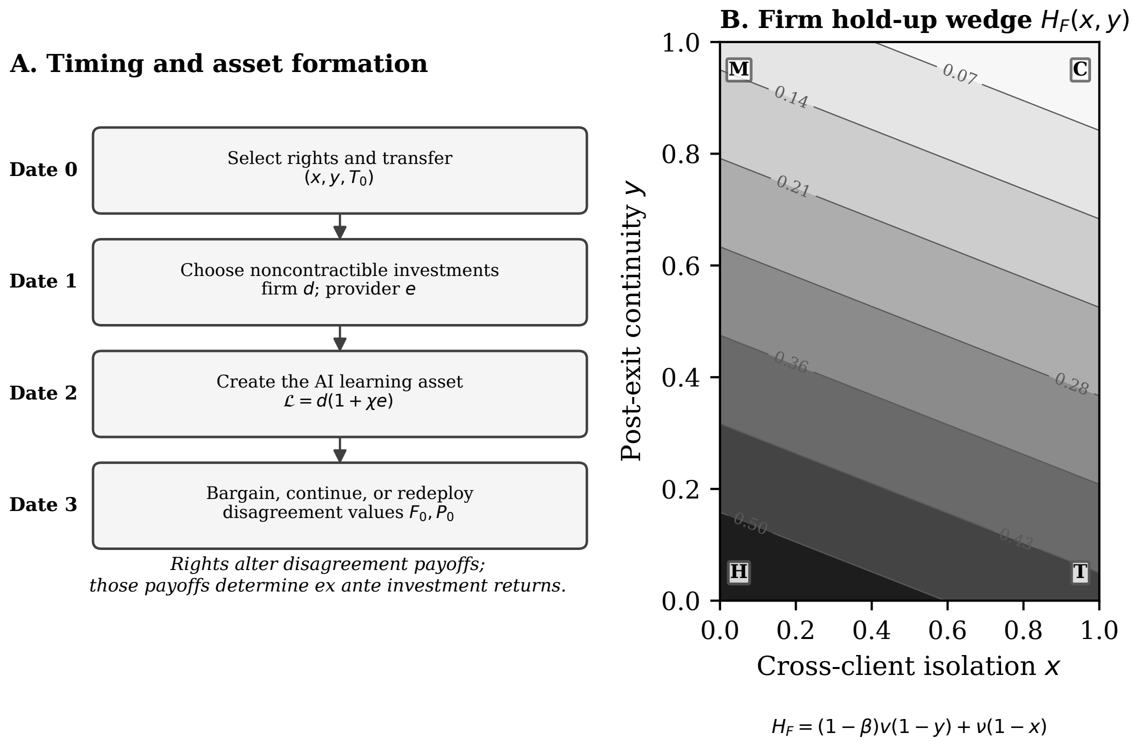

Figure 2. Timing and the two channels of proprietary-knowledge hold-up. Panel A shows the contracting, investment, asset-formation, and bargaining sequence. Panel B plots the closed-form wedge $H_F(x, y) = (1 - \beta)\nu(1 - y) + \nu(1 - x)$ over isolation and continuity.

3.2. The AI learning asset and situated AI performance

Firm investment organizes proprietary sources and supplies task-specific judgments. Provider adaptation determines how effectively the system retrieves and acts on that knowledge. These roles correspond to the grounding and representational work described in situated-AI research (Kemp, 2024; Yokoi et al., 2026). They also follow evidence that technical investment becomes productive through complementary organizational and user-side changes (Bartel et al., 2007; Bresnahan et al., 2002; Weigelt, 2013). We denote the AI learning asset by $\mathcal{L}$ and use the seeded bilinear technology below.[12]

$$\mathcal{L}(d, e) = d(1 + \chi e), \qquad \chi \geq 0, \tag{1}$$

where χ measures complementarity. Firm input seeds the asset, since $\mathcal{L}(0, e) = 0$. Adaptation raises the return to firm input whenever $\chi > 0$, because $\mathcal{L}_{de} = \chi$.

[12] The substantive restriction is complementarity: $\mathcal{L}_{de} > 0$ means that provider adaptation raises the marginal return to proprietary-knowledge investment (Milgrom & Roberts, 1990; Milgrom & Shannon, 1994). The specification gives transparent closed forms. Appendix A treats a general twice-differentiable technology, and Appendix B reports Cobb–Douglas and CES alternatives; see also Chatain and Plaksenkova (2026).

The operating performance of situated AI is

$$Q(d,e) = e + \mathcal{L}(d,e). \tag{2}$$

The direct component e is useful while the provider supports the system; $\mathcal{L}$ is the AI learning asset. The additive form keeps a current input and an accumulated asset visible in the same performance units, following the flow–stock distinction in research on asset accumulation and organizational learning (Argote & Miron-Spektor, 2011; Benkard, 2000; Dierickx & Cool, 1989). A unit of either component produces focal operating value $v > 0$. The distinction allows a right to induce more firm knowledge yet lower performance if it discourages enough current provider adaptation.

3.3. Disagreement uses and marginal claims

When isolation is incomplete, the AI learning asset has two cross-client consequences. The provider can reuse it without depleting the focal asset, consistent with the nonrivalry of data and with models that distinguish cross-user learning from repeated learning within a user (Hagiu & Wright, 2023; Jones & Tonetti, 2020). Diffusion can also erode the focal firm's appropriability, a risk documented in interfirm knowledge exchange and software outsourcing (Kale et al., 2000; Tiwana, 2008). Let $\nu(1-x)$ be the provider's productive reuse value per unit of learning and $\lambda(1-x)$ the focal firm's diffusion loss, with $\nu, \lambda \geq 0$. The common factor $1-x$ captures safeguards that reduce both consequences over the range being compared; their values remain distinct because disclosure can create productive opportunities as well as appropriability hazards (Pacheco-de-Almeida & Zemsky, 2012). One unit of the AI learning asset then creates bilateral value

$$C(x) = v + (\nu - \lambda)(1-x) > 0. \tag{3}$$

Isolation protects focal value but gives up reuse. It raises bilateral value per unit of learning exactly when $\lambda > \nu$. Bilateral value includes only the payoffs of F and P; rival and consumer effects lie outside the model. This follows the strategy distinction between value created and captured inside a focal relationship and broader welfare (Chatain & Zemsky, 2011).

If bargaining fails, the firm preserves the operating share y and bears the permitted diffusion loss, while the provider retains the permitted reuse value. Their marginal disagreement values for the AI learning asset are

$$F_0(x,y) = vy - \lambda(1-x), \qquad P_0(x) = \nu(1-x). \tag{4}$$

These are marginal disagreement values per unit of learning, not the parties' complete outside payoffs.[13] Their form follows property-rights models in which control changes investment through disagreement payoffs rather than through title alone (de Meza & Lockwood, 1998; Grossman & Hart, 1986; Hart & Moore, 1990). The role of y also follows switching-cost research: dependence is shaped by what remains usable after exit, not merely by ownership of the original input (P.-Y. Chen & Hitt, 2002; Vijairaghavan et al., 2025). The incremental surplus from maintaining focal operation is $C - F_0 - P_0 = v(1-y)$. Informational nonrivalry permits reuse to remain available after focal operation ends (Jones & Tonetti, 2020); when reuse instead requires continued access, P_0 also depends on relationship survival.

Direct adaptation may also retain some value after disagreement. Specific investment can generate reusable knowledge or reputation (Kang et al., 2009), and the value of redeployment depends on how capabilities and control are configured (Ethiraj et al., 2005; Lo et al., 2022; Sakhartov & Helfat, 2026). Let $\omega \in [0, v)$ be the provider's salvage value from one unit of direct adaptation when focal cooperation ends. Because that redeployment substitutes for focal use, $\omega < v$ preserves gains from cooperation. Appendix B separately considers nonrival use of direct adaptation. Setting $\omega = 0$ recovers the no-salvage baseline; v continues to measure nonrival reuse of the AI learning asset.

Let $\beta \in (0,1)$ be the firm's generalized Nash bargaining weight. We use generalized Nash bargaining as a reduced-form efficient division of ex post surplus with transferable utility (Binmore et al., 1986; Nash, 1950; Rubinstein, 1982); Grout (1984) applies the same bargaining logic to sunk investment and hold-up.[14] Efficient bargaining over the learning margin gives the firm and provider claims

$$A(x,y) = F_0 + \beta(C - F_0 - P_0) = v[\beta + (1-\beta)y] - \lambda(1-x), \tag{5}$$

$$B(x,y) = P_0 + (1-\beta)(C - F_0 - P_0) = (1-\beta)v(1-y) + v(1-x). \tag{6}$$

They exhaust bilateral marginal value:

$$A(x,y) + B(x,y) = C(x). \tag{7}$$

For direct adaptation, the firm's claim is $g_F \equiv \beta(v-\omega)$ and the provider's claim is

$$g_P \equiv \omega + (1-\beta)(v-\omega) = (1-\beta)v + \beta\omega, \qquad g_F + g_P = v. \tag{8}$$

[13]Omitted constants may keep total disagreement payoffs positive even if $F_0 < 0$, which can occur when another unit of learning causes more diffusion harm than preserved operating value. This convention lets the model focus on the marginal return that governs investment.

[14]We hold β fixed to isolate the effects of x and y on disagreement values; alternative bargaining protocols can change property-rights predictions (de Meza & Lockwood, 1998). At fixed sunk (d, e), agreement creates $ve + C\mathcal{L}$ and disagreement creates $\omega e + (F_0 + P_0)\mathcal{L}$, so the aggregate gain is $(v-\omega)e + v(1-y)\mathcal{L} \geq 0$. Assigning Nash shares to this gain yields g_F, g_P, A, and B.

The claim identity reveals the first hold-up object.

Lemma 1 (Marginal hold-up decomposition)**.** *Suppose the bargaining surplus satisfies* $C - F_0 - P_0 \geq 0$*. For arbitrary marginal disagreement values* F_0 *and* P_0*, the identity*

$$H_F \equiv C - A = (1-\beta)(C - F_0 - P_0) + P_0 = (1-\beta)(C - F_0) + \beta P_0 \tag{9}$$

implies $H_F \geq 0$ *whenever, in addition,* $P_0 \geq 0$*. In the baseline,*

$$H_F = C - A = B = (1-\beta)v(1-y) + v(1-x). \tag{10}$$

Consequently $H_{F,y} = -(1-\beta)v < 0$*,* $H_{F,v} = 1 - x$*,* $H_{F,x} = -v$*, and* $H_{F,\lambda} = 0$*.*

The last derivative is informative. Diffusion harm lowers the firm's claim and bilateral value by the same amount in this baseline, so it does not enter the gap $C - A$. Continuity narrows the wedge by returning focal operating value to the firm's fallback. Productive reuse widens it because the provider preserves an outside value that the firm does not capture. The transfer T_0 can divide expected rents or satisfy participation constraints, but it is sunk when investment is chosen and cannot change any marginal claim (Grossman & Hart, 1986; Hart & Moore, 1990). This result applies to a fixed, noncontingent T_0; contracts tied to verifiable outcomes can alter marginal incentives (Baker et al., 2002; Edlin & Reichelstein, 1996; Nöldeke & Schmidt, 1995).

3.4. Private investment and the bilateral benchmark

Investment costs are increasing and convex, as in canonical models of noncontractible and cooperative investment (Che & Hausch, 1999; Grossman & Hart, 1986; Hart & Moore, 1990). We use the tractable quadratic specialization $k_F d^2/2$ and $k_P e^2/2$, with $k_F, k_P > 0$ (d'Aspremont & Jacquemin, 1988); Appendix A gives the corresponding general convex-cost conditions. After the bargain is substituted, private payoffs are

$$U_F(d, e; x, y) = g_F e + A(x, y)\mathcal{L}(d, e) - \frac{k_F}{2}d^2 - T_0, \tag{11}$$

$$U_P(d, e; x, y) = g_P e + B(x, y)\mathcal{L}(d, e) - \frac{k_P}{2}e^2 + T_0. \tag{12}$$

An interior Nash equilibrium satisfies

$$k_F d = A(1 + \chi e), \qquad k_P e = g_P + \chi B d. \tag{13}$$

Proprietary-knowledge investment is governed by A; provider adaptation is governed by its direct return g_P and its claim B on the AI learning asset. Generically, each leaves part of the bilateral marginal return uninternalized, the two-sided investment problem emphasized in property-rights and cooperative-investment models (Chatain & Plaksenkova, 2026; Che & Hausch, 1999; Hart & Moore, 1990).

Write

$$K \equiv k_F k_P, \qquad \Delta(x, y) \equiv K - \chi^2 A(x, y) B(x, y), \qquad M_P \equiv k_P + \chi g_P. \tag{14}$$

We maintain $A > 0$, $B \geq 0$, and $K > \chi^2 C(x)^2$ over the rights being compared. This strong-stability condition implies $\Delta > 0$, gives the private and bilateral response maps spectral radii below one (and hence contraction representations in suitable weighted norms), and guarantees strict concavity of the bilateral benchmark. Curvature restrictions of this kind are standard in the analysis of concave games (Rosen, 1965); the displayed inequality is sufficient for the comparisons below.[15] The private investment levels are

$$d^*(x, y) = \frac{A M_P}{\Delta}, \tag{15}$$

$$e^*(x, y) = \frac{k_F g_P + \chi A B}{\Delta}. \tag{16}$$

Because each payoff is strictly concave in the party's own action, the first-order conditions are sufficient. Moreover, $A > 0$, $g_P > 0$, $M_P > 0$, and $\Delta > 0$ make both displayed choices strictly positive; the solution is therefore the unique interior Nash equilibrium in the maintained region. Equilibrium learning and situated AI performance are

$$\mathcal{L}^* = d^*(1 + \chi e^*), \qquad Q^* = e^* + \mathcal{L}^*. \tag{17}$$

The bilateral first best provides the diagnostic benchmark used in property-rights and cooperative-investment models to identify the marginal returns that decentralized investors leave uninternalized (Chatain & Plaksenkova, 2026; Che & Hausch, 1999; Grossman & Hart, 1986; Hart & Moore, 1990).[16] Joint surplus before governance cost is

$$S(d, e; x) = ve + C(x)\mathcal{L}(d, e) - \frac{k_F}{2} d^2 - \frac{k_P}{2} e^2, \tag{18}$$

[15] The product of the private best-response slopes is $\chi^2 AB/K$, while the bilateral response product is $\chi^2 C^2/K$; the maintained condition keeps both below one. The bilateral surplus Hessian has diagonal entries $-k_F$ and $-k_P$ and off-diagonal entry χC, so the same inequality makes it negative definite. Boundary solutions require the corresponding one-sided Kuhn–Tucker conditions.

[16] The benchmark maximizes the combined payoff of F and P for a given rights configuration. Consumer, rival, worker, and other third-party effects lie outside this bilateral comparison.

with first-order conditions

$$k_F d^{FB} = C(1 + \chi e^{FB}), \qquad k_P e^{FB} = \nu + \chi C d^{FB}. \tag{19}$$

Comparing these conditions with Equation (13) separates the firm's marginal hold-up from the full equilibrium investment shortfall.

3.5. Split and bundled governance sets

Implementing rights consumes real resources. Isolation can require restricted retention, tenant separation, monitoring, and enforceable safeguards; continuity can require learned-state export, compatible tooling, documentation, transition support, and internal operating capacity. Research on transaction costs and alliances treats governance mechanisms as costly responses whose effectiveness depends on the assets being governed (Argyres & Mayer, 2007; Hoetker & Mellewigt, 2009; Williamson, 1991). Work on data portability likewise models investment that lowers implementation or compliance cost (Vijairaghavan et al., 2025). Let $G(x, y)$ be a continuous governance-cost function; the quadratic specification used in the figures is

$$G(x, y) = f_x x + f_y y + \frac{c_x}{2}x^2 + \frac{c_y}{2}y^2 + c_{xy}xy, \tag{20}$$

with a positive-semidefinite cost Hessian.[17] Define equilibrium bilateral value and its net counterpart as

$$W(x, y) = S(d^*(x, y), e^*(x, y); x), \qquad \Phi(x, y) = W(x, y) - G(x, y). \tag{21}$$

Organizational architecture need not move every boundary choice together. Jacobides and Billinger (2006) describe vertical architecture as a configuration of multiple transactional choices, while Boudreau (2010) distinguishes digital access from control and Hoetker and Mellewigt (2009) show that governance mechanisms should be matched to the assets involved. Under split control the feasible set is $\Gamma_S = [0, 1]^2$.[18] A bundled or unitary arrangement instead ties the two effective rights to a pre-specified regular rule

$$\Gamma_U = \{\gamma(z) : z \in [0, 1]\} \subseteq \Gamma_S, \tag{22}$$

[17] G contains real implementation resources. The quadratic displays rising marginal difficulty, and c_{xy} allows shared infrastructure or interference between the two rights (Poppo & Zenger, 2002). Appendix B permits a general twice-differentiable cost. Positive semidefiniteness keeps cost convex; monotonicity additionally requires nonnegative marginal cost over the rights region being compared.

[18] For the global governance comparison, the admissible-region restrictions are imposed at every point of Γ_S. A primitive sufficient condition is $\lambda < \beta\nu$ and $K > \chi^2[\nu + \max\{\nu - \lambda, 0\}]^2$: the first gives $\min_{\Gamma_S} A = \beta\nu - \lambda > 0$, while $B \geq 0$ throughout the box, and the second evaluates strong stability at $\max_{\Gamma_S} C = \nu + \max\{\nu - \lambda, 0\}$. Under these restrictions Φ is continuous on the compact rights box.

where γ is Lipschitz and piecewise continuously differentiable.[19] The diagonal $\gamma(z) = (z, z)$ is the transparent benchmark in which isolation and continuity must move together; more general regular paths allow unequal rates of co-movement.

The four corners provide descriptive reference bundles:

$$\mathrm{H} = (0, 0), \qquad \mathrm{T} = (1, 0), \qquad \mathrm{M} = (0, 1), \qquad \mathrm{C} = (1, 1).$$

They support comparison of discrete and hybrid organizational arrangements in the comparative-governance tradition (Parmigiani, 2007; Williamson, 1991). The governance comparison is between the best net bilateral value implementable over Γ_S and that implementable over the constrained path Γ_U.

[19]The rule is fixed before optimization, so it cannot be redrawn around the desired solution. Lipschitz, piecewise-C^1 regularity also excludes a space-filling encoding of two independent choices as one scalar.

Table 2. Core notation

Symbol	Name	Meaning
F, P	Parties	Focal firm and external AI provider.
d, e	Investments	Firm proprietary-knowledge investment and provider customer-specific adaptation.
χ	Complementarity	Strength with which provider adaptation raises the productivity of firm knowledge investment.
$\mathcal{L}(d, e)$	AI learning asset	Persistent customer-specific state created by investment; baseline $\mathcal{L} = d(1 + \chi e)$.
$Q(d, e)$	Situated AI performance	Current operating performance; baseline $Q = e + \mathcal{L}$.
x, y	Governance rights	Effective cross-client isolation and post-exit continuity, each normalized to $[0, 1]$.
v	Focal value	Operating value per unit of current situated AI performance.
λ, ν	Cross-client effects	Firm diffusion loss and provider productive reuse value per unit of nonisolated learning.
β	Bargaining weight	Firm's generalized Nash bargaining weight; provider weight is $1 - \beta$.
ω	Adaptation salvage value	Provider value from a unit of direct adaptation if focal cooperation ends; $0 \leq \omega < v$.
F_0, P_0	Disagreement values	Firm and provider values per unit of learning after disagreement.
A, B, C	Marginal claims	Firm claim, provider claim, and combined bilateral value per unit of learning; $A + B = C$.
g_F, g_P	Direct claims	Firm and provider claims on direct adaptation; $g_F + g_P = v$.
H_F	Marginal hold-up wedge	Learning-state value not captured by the firm's next unit of conversion, $H_F = C - A = B$.
$d^{FB} - d^*$	Equilibrium investment shortfall	Difference between bilateral first-best and private-equilibrium firm investment; distinct from H_F.
k_F, k_P	Cost curvature	Marginal-cost curvature of d and e in the quadratic specialization.
Δ, M_P	Solution terms	$\Delta = k_F k_P - \chi^2 AB$ and $M_P = k_P + \chi g_P$.
W, Φ	Governance value	Equilibrium bilateral value before governance cost and net value $\Phi = W - G$.
Γ_S, Γ_U	Feasible governance sets	Split rights box $[0, 1]^2$ and a pre-specified regular bundled-rights path contained in it.

Notes. Appendix D provides notation for extensions and evidence variables. The table distinguishes the marginal hold-up mechanism H_F from its equilibrium investment consequence $d^{FB} - d^*$.

4. Hold-Up, Investment Reallocation, and Split Control

Following Grout (1984) and Hart and Moore (1990), sunk, noncontractible investment responds to the claim an investor can realize ex post rather than to total surplus. With cooperative investment, one party's claim distortion also changes the return to the other's complementary contribution (Chatain & Plaksenkova, 2026; Che & Hausch, 1999). Even with $A + B = C$, governance changes the marginal claims that guide each investment. We first establish the underinvestment benchmark, then compare continuity and productive reuse, and finally ask what split control can implement that a bundled rule cannot.

4.1. Marginal hold-up and equilibrium underinvestment

Holding provider adaptation fixed, the firm's private and bilateral conditional choices solve

$$d^F(e) = \frac{A(1 + \chi e)}{k_F}, \qquad d^{FB}(e) = \frac{C(1 + \chi e)}{k_F}. \tag{23}$$

Their difference is the marginal wedge translated into investment units:

$$d^{FB}(e) - d^F(e) = \frac{H_F(1 + \chi e)}{k_F}. \tag{24}$$

Holding proprietary-knowledge investment fixed, the corresponding provider choices are

$$e^P(d) = \frac{g_P + \chi B d}{k_P}, \qquad e^{FB}(d) = \frac{v + \chi C d}{k_P}, \tag{25}$$

so that

$$e^{FB}(d) - e^P(d) = \frac{\beta(v - \omega) + \chi A d}{k_P} > 0. \tag{26}$$

Adaptation produces direct focal value as well as learning, so the provider's conditional shortfall contains both the uncaptured direct return $\beta(v - \omega)$ and the firm's claim A on the learning margin.

Proposition 1 (Marginal hold-up and equilibrium underinvestment)**.** *Suppose $A > 0$, $B \geq 0$, $0 \leq \omega < v$, and $K > \chi^2 C^2$. Then the private and bilateral response maps have unique fixed points. The conditional gaps are given by Equations* (24) *and* (26)*. At equilibrium,*

$$e^* < e^{FB}, \qquad d^* \leq d^{FB}. \tag{27}$$

The firm inequality is strict whenever $B > 0$ or $\chi > 0$. Equality $d^* = d^{FB}$ occurs only in the separable knife edge $B = 0$ and $\chi = 0$.

Proof sketch. The strong-stability condition makes both increasing response maps contractions. Equations (24) and (26) show that the bilateral map weakly dominates the private map in the firm component and strictly dominates it in the provider component. Monotone fixed-point comparison yields Equation (27). If $B > 0$, the firm's conditional gap is strict; if $B = 0$ but $\chi > 0$, the strictly larger first-best provider input raises the first-best firm response. Appendix A gives the fixed-point proof and the corresponding result for general convex costs and a general increasing learning technology.

The proposition distinguishes a marginal claim wedge from an equilibrium investment shortfall. In cooperative-investment models, private claims need not coincide with joint marginal returns (Che & Hausch, 1999); with complementary inputs, one party's distortion also travels through the other's response (Chatain & Plaksenkova, 2026). Thus $H_F = C - A$ depresses proprietary-knowledge investment conditional on e, whereas $d^{FB} - d^*$ compares two fixed points after both inputs adjust. Even when $H_F = 0$, provider underinvestment can lower firm investment through $\chi > 0$; a parameter that widens H_F may also draw in provider adaptation. The wedge identifies a mechanism, while the equilibrium shortfall records both responses.

In the baseline, $H_F = (1-\beta)v(1-y) + v(1-x)$. The continuity component vanishes as the firm obtains the full bargaining weight, while productive reuse remains:

$$\lim_{\beta \uparrow 1} H_F = v(1-x). \tag{28}$$

de Meza and Lockwood (1998) show that property-rights predictions depend on how outside options enter bargaining. Consistent with that logic, an independently realizable provider use cannot be eliminated by relabeling the firm's bargaining weight: actual disagreement uses, rather than β alone, drive the limit in Equation (28).

4.2. Continuity relieves firm hold-up but can reduce performance

Let $a_y \equiv (1-\beta)v > 0$. Continuity moves marginal claim from the provider to the firm while leaving bilateral marginal value fixed:

$$(A_y, B_y, C_y) = (a_y, -a_y, 0), \qquad H_{F,y} = -a_y < 0. \tag{29}$$

Continuity directly narrows the firm's hold-up wedge. Its production consequence depends on how the complementary provider input responds. Continuity operates through the outside-option channel emphasized by de Meza and Lockwood (1998): it reallocates value preserved after disagreement and thereby changes ex ante investment incentives.[20]

Proposition 2 (Continuity-induced reallocation)**.** *Under the maintained interior and strong-stability*

[20]The baseline holds $C_y = 0$ and defines continuity as focal operating value preserved after disagreement, a narrower margin than the supply assurance studied by Bolton and Whinston (1993). Appendix B allows continuity to affect transition and breakdown value.

conditions, stronger post-exit continuity changes private investment according to

$$d_y^* = \frac{a_y M_P (K - \chi^2 A^2)}{\Delta^2} > 0, \tag{30}$$

$$e_y^* = \frac{a_y \chi k_F M_P (B - A)}{\Delta^2}. \tag{31}$$

For $\chi > 0$, the sign of e_y^* is the sign of $B - A$; for $\chi = 0$, provider adaptation is unchanged. Situated AI performance satisfies

$$Q_y^* = (1 + \chi e^*) d_y^* + (1 + \chi d^*) e_y^*, \tag{32}$$

and falls if and only if

$$\chi k_F (A - B)(1 + \chi d^*) > (K - \chi^2 A^2)(1 + \chi e^*). \tag{33}$$

Proof sketch. Differentiate the closed-form investments along the claim path in Equation (29). The common positive factor is $a_y M_P / \Delta^2$; the firm response leaves $K - \chi^2 A^2 > 0$, while the provider response leaves $\chi k_F (B - A)$. Substitution into the total derivative of $Q = e + d(1 + \chi e)$ yields the necessary and sufficient condition in Equation (33). Appendix A derives the claim derivatives and shows why general complementarity alone is insufficient: outside the seeded multiplicative technology, additional curvature conditions are required for these signs.

Continuity thus narrows the firm's hold-up wedge and raises proprietary-knowledge investment in the maintained region.[21] A performance reversal requires more. The provider must begin with the smaller claim on the AI learning asset, $B < A$; inputs must be complementary; and the lost direct and complementary contribution of adaptation must exceed the performance created by added firm knowledge. At $\chi = 0$, continuity raises performance; at $A = B$, it leaves provider adaptation unchanged and also raises performance. This dependence on the interaction of claims and complementarity follows the increasing-differences logic of Milgrom and Shannon (1994), while the exact threshold remains specific to the model's technology (Chatain & Plaksenkova, 2026).

Operating performance and bilateral value have distinct reversal thresholds. At the private equilibrium define the bilateral marginal returns not internalized by the investors:

$$S_d \equiv B(1 + \chi e^*), \qquad S_e^{\omega} \equiv \beta(v - \omega) + \chi A d^*. \tag{34}$$

Corollary 1 (A distinct bilateral-value reversal)**.** *Under Proposition 2, equilibrium bilateral value*

[21] The derivatives are local, hold the other primitives fixed, and apply along admissible directions; comparisons at $y \in \{0, 1\}$ are one-sided.

falls with continuity if and only if

$$\chi k_F(A-B)S_e^{\omega} > B(1+\chi e^*)(K-\chi^2 A^2). \tag{35}$$

Neither this condition nor Equation (33) *generally implies the other.*

The distinction is substantive. Q records operating output, whereas W values the relationship's bilateral return net of the investments used to produce it. This parallels the strategy distinction between value creation and value capture (Chatain & Zemsky, 2011). The two reversal thresholds need not be nested.

Figure 3 displays one admissible parameter slice. To the right of the computed performance peak y_Q, firm input continues to rise while provider adaptation and situated AI performance fall. Bilateral value reaches a distinct peak y_W in this illustration. The existence, uniqueness, and ordering of the two peaks are properties of the chosen slice, not general theorems.

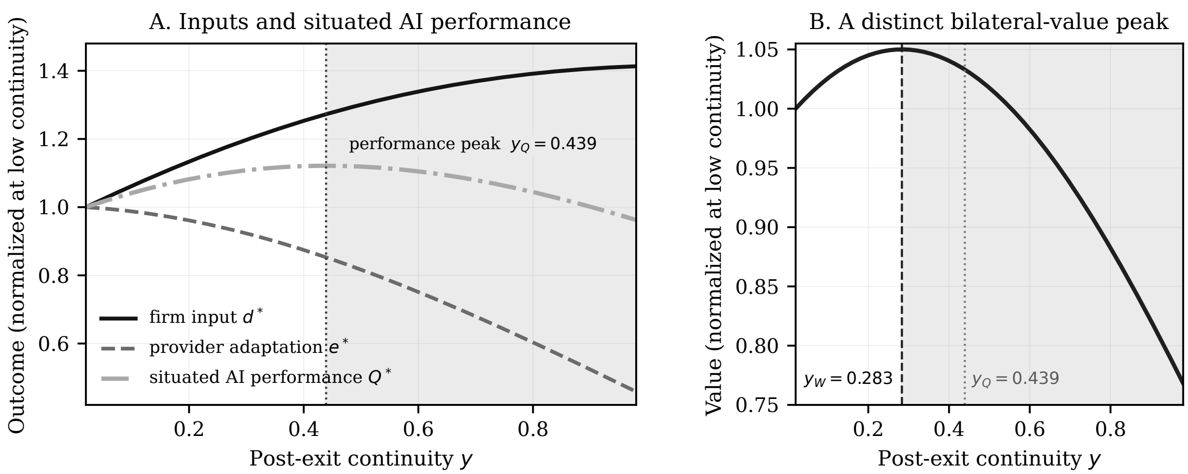

Figure 3. Continuity-induced decoupling in one admissible parameter slice with $\omega = 0$. Panel A shows firm input, provider adaptation, and situated AI performance; shading begins at the slice-specific performance peak y_Q. Panel B shows bilateral value and marks its distinct peak y_W. The two thresholds are computed from the exact derivatives in Proposition 2; their ordering is specific to the plotted parameter slice. Outcomes are normalized at the lowest plotted continuity.

4.3. Productive reuse can widen the gap while raising every level

Unlike a transfer of a fixed claim, productive reuse creates an additional bilateral opportunity. This distinction follows the nonrival-use logic of Jones and Tonetti (2020) and the separation of cross-user from within-user learning in Hagiu and Wright (2023). With $r \equiv 1 - x$, a marginal increase in v follows

$$(A_v, B_v, C_v) = (0, r, r), \qquad H_{F,v} = r. \tag{36}$$

It raises the provider's private return and the bilateral opportunity by the same amount while leaving the firm's marginal claim unchanged.[22]

Proposition 3 (Productive-reuse duality)**.** *Under the maintained interior and strong-stability conditions, if $x < 1$ and $\chi > 0$,*

$$d_v^* > 0, \qquad e_v^* > 0, \qquad Q_v^* > 0. \tag{37}$$

For every $x < 1$, including $\chi = 0$, along a local change in v that remains in the admissible region,

$$W_v > 0, \qquad \frac{\partial(d^{FB} - d^*)}{\partial v} > 0. \tag{38}$$

The exact gap derivative is

$$(1-x)\left[\frac{(k_P + \chi v)(K + \chi^2 C^2)}{(K - \chi^2 C^2)^2} - \frac{A^2 \chi^2 M_P}{\Delta^2}\right] > 0. \tag{39}$$

All reuse derivatives vanish at full isolation.

Proof sketch. Along Equation (36), A is fixed and B rises. The closed forms give $d_B = A^2\chi^2 M_P/\Delta^2$ and $e_B = A\chi k_F M_P/\Delta^2$, which are positive with complementarity. The performance sign follows immediately. The envelope derivative of W contains the direct gain $r\mathcal{L}^*$ and nonnegative investment-response terms. The first-best response to C exceeds the private response to B, yielding Equation (39). Appendix A supplies the factorization establishing strict positivity.

Kang et al. (2009) show that reusable knowledge and reputation can sustain unilateral specific investment. Proposition 3 extends that logic to a state formed by complementary contributors: an outside use draws in provider adaptation, which raises the productivity of proprietary-knowledge investment. The resulting configuration and its bilateral value rise even as the bilateral benchmark expands faster than decentralized investment. The widening gap records how much faster the bilateral opportunity grows than realized investment.[23]

A pure transfer toward the provider behaves differently. Let $(A_\tau, B_\tau, C_\tau) = (-1, 1, 0)$. Then

$$d_\tau^* = -\frac{M_P(K - \chi^2 A^2)}{\Delta^2} < 0, \tag{40}$$

while e_τ^* and Q_τ^* have no unconditional sign. The positive reuse result comes from the new productive outside use. A transfer toward the provider within fixed surplus has different effects.

[22] This derivative holds isolation x fixed. Changing x also moves diffusion harm through Equation (41) and may have a different sign.

[23] The absolute gap $d^{FB} - d^*$ is measured against the bilateral benchmark; Proposition 3 separately signs realized investment and bilateral value.

The marginal-claim transfer τ differs from the fixed payment T_0. A fixed, noncontingent transfer drops out of both investment first-order conditions; contingent options or breach remedies can instead change marginal incentives (Edlin & Reichelstein, 1996; Nöldeke & Schmidt, 1995).

4.4. Split control expands the implementable governance set

Isolation follows the claim path

$$(A_x, B_x, C_x) = (\lambda, -\nu, \lambda - \nu), \tag{41}$$

whereas continuity follows Equation (29). The two movements are generically independent.

Lemma 2 (Local irreducibility of rights and investment)**.** *Let $a_y = (1-\beta)\nu$. The claim-map Jacobian is*

$$J_{AB} = \begin{bmatrix} \lambda & a_y \\ -\nu & -a_y \end{bmatrix}, \qquad \det J_{AB} = a_y(\nu - \lambda). \tag{42}$$

At an interior rights bundle, if $\lambda \neq \nu$, the split claim image contains a local open set. If also $\chi > 0$, the investment response to claims satisfies

$$\det \frac{\partial(d^*, e^*)}{\partial(A, B)} = \frac{A\chi k_F M_P^2}{\Delta^3} > 0, \tag{43}$$

so the split investment image is locally two-dimensional as well.

The rank result establishes local noncompressibility. Proposition 4 next compares net bilateral value over the two feasible sets introduced in Equation (22).

Proposition 4 (Split-control implementation)**.** *Let $\Gamma_S = [0,1]^2$ and let the bundled regime be a pre-specified Lipschitz, piecewise-C^1 path $\Gamma_U = \gamma([0,1]) \subseteq \Gamma_S$. Maintain the model's admissibility restrictions throughout Γ_S. Then Φ is continuous, Γ_S and Γ_U are compact, and the maxima below exist.*

(a) Feasible-set value.

$$\max_{(x,y)\in\Gamma_S} \Phi(x,y) \geq \max_{(x,y)\in\Gamma_U} \Phi(x,y). \tag{44}$$

(b) Off-path improvement. *Suppose $z_U \in \arg\max_{z\in[0,1]} \Phi(\gamma(z))$ is interior, $\gamma(z_U)$ lies in the interior of the rights box, γ is differentiable with $\gamma'(z_U) \neq 0$, and Φ is continuously differentiable in a neighborhood of $\gamma(z_U)$. Then*

$$\nabla\Phi(\gamma(z_U)) \cdot \gamma'(z_U) = 0. \tag{45}$$

If $\nabla\Phi(\gamma(z_U)) \neq 0$, split control admits a feasible off-path direction that strictly raises Φ.

(c) Diagonal benchmark. *Suppose Φ is differentiable at the relevant rights bundle. For*

$\gamma(z) = (z, z)$*, an interior bundled optimum obeys*

$$\Phi_x(z_U, z_U) + \Phi_y(z_U, z_U) = 0. \tag{46}$$

It can also be an interior split optimum only if $\Phi_x(z_U, z_U) = \Phi_y(z_U, z_U) = 0$.

(d) Unique split optimum. *If* Φ *is strictly concave and its unique split optimum is not on* Γ_U*, the inequality in Equation* (44) *is strict.*

Proof sketch. Part (a) follows from $\Gamma_U \subseteq \Gamma_S$. At a regular interior path optimum, the directional derivative along the path is zero. If the full gradient is nonzero, it is normal to the path and provides an improving direction in the two-dimensional feasible set, proving part (b). Part (c) specializes the path derivative. Part (d) follows because a unique global split maximizer off the bundled path cannot be attained in Γ_U. Appendix A gives the tangent-cone statement for rights-box and path boundaries.

Proposition 4 compares the best configuration available under split control with the best configuration permitted by the bundled rule. It formalizes the architectural premise that governance configures separable choices rather than a single make-or-buy scalar (Boudreau, 2010; Jacobides & Billinger, 2006). Appendix B.4 examines the local-rank condition when reuse restriction and diffusion protection are separated.

Figure 4 illustrates the diagonal benchmark. The bundled regime can move only along $x = y$; split control can move away from that path. At the illustrated bundled optimum the path derivative is zero but the two partial derivatives are not, so the off-path split direction improves net bilateral value.

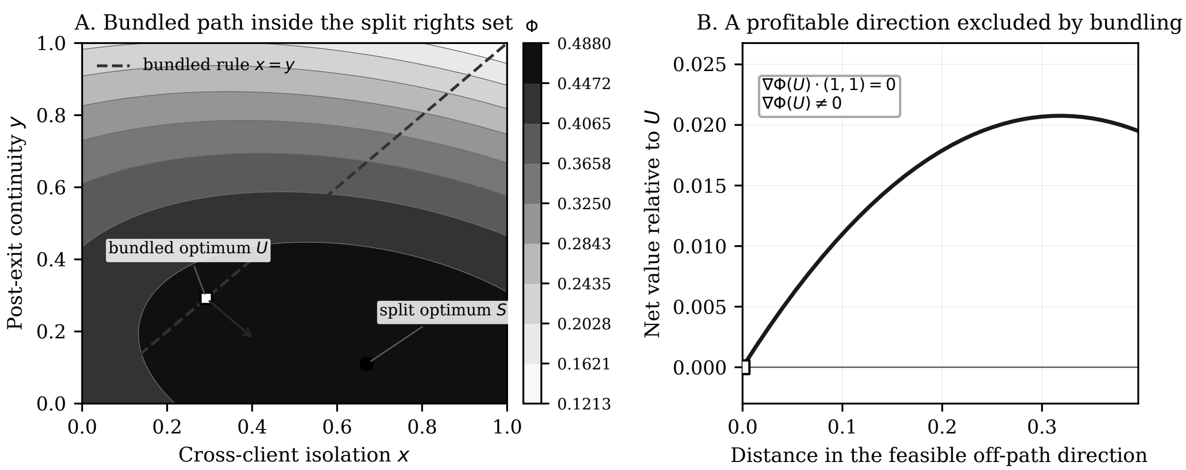

Figure 4. Split control versus a bundled diagonal rule in a numerical illustration. Panel A plots net equilibrium bilateral value $\Phi(x, y)$ over the split rights box. The dashed diagonal is the bundled rule $x = y$; its optimum U differs from the split optimum S. Panel B follows the feasible gradient direction from U. The directional derivative along the bundled path is zero, while the full gradient is nonzero and the off-path direction initially raises net value.

The results replace a scalar prescription for “more control” with a comparison among the distinct uses that governance reallocates.

5. Discussion

The strategic content of an AI governance right lies in how it changes the marginal claims attached to learning created through the relationship. This perspective extends property-rights reasoning to an asset formed after contracting, clarifies the cross-side investment effects of different rights, and recasts the firm boundary as a configuration of economically distinct uses.

5.1. Residual control and endogenous asset formation

Canonical property-rights theory studies how control over an asset changes disagreement payoffs and, through them, noncontractible investment (Grossman & Hart, 1986; Hart & Moore, 1990). In enterprise AI, the economically consequential asset is not fully present when control is assigned. The firm enters with proprietary knowledge and the provider with a general technology, while the customer-specific learning state emerges only after both parties have invested. Ownership of these antecedent inputs does not settle who can exercise the later uses of the resulting state. Residual control therefore reaches backward into the formation of the asset itself.

This production structure also changes how underinvestment should be interpreted. The marginal hold-up wedge measures the portion of the bilateral return from another unit of firm investment that the firm does not expect to capture. The equilibrium investment shortfall includes an additional effect: the response of the provider’s complementary input. Provider underinvestment can depress the firm’s contribution even when the firm’s own marginal wedge is small. Conversely, an outside opportunity that strengthens provider adaptation can raise realized firm investment while increasing its distance from the bilateral first best. A local claim distortion and an equilibrium investment shortfall are therefore distinct objects in a cooperative- investment setting (Che & Hausch, 1999).

The analysis gives situated AI an incentive foundation. Existing research shows that a general technology becomes productive only when it is grounded in local knowledge, integrated with work, and supported by an appropriate allocation of authority (Grant, 1996; Kemp, 2024; Kim et al., 2024; Krakowski et al., 2023; Yokoi et al., 2026). Our model explains why the inputs required for that grounding may vary across otherwise similar deployments. Rights over future reuse and post-exit operation affect how much proprietary knowledge the firm makes usable and how much customer-specific adaptation the provider supplies. Firms with access to similar models and similar source data can consequently develop different situated AI capabilities because their governance arrangements induce different paths of joint investment.

5.2. Claim paths and cross-side investment incentives

The results distinguish two economic functions that are often compressed into a general notion of control. Post-exit continuity reallocates the claim on focal operation. In the baseline, it raises the firm's disagreement position and lowers the provider's claim while leaving bilateral value per unit of learning unchanged. The firm responds by increasing proprietary-knowledge investment. The provider may reduce adaptation because a smaller share of the return from continued focal operation remains attached to its contribution. When that response is sufficiently large, stronger continuity lowers situated AI performance even as the firm invests more.

This result identifies a cross-side incentive condition; it does not imply that continuity is generally harmful. Continuity can preserve value after breakdown, reduce dependence, and encourage firm investment. Its net effect depends on the complementary contribution it displaces, the claims initially held by the parties, and the cost of implementing independent operation. A governance safeguard should therefore be evaluated through the responses of both contributors rather than through the protected party's appropriation alone. This logic complements evidence that adaptation and appropriation jointly shape the investment consequences of contractual governance (Lo et al., 2022).

Productive reuse follows a different claim path. A more valuable cross-client application increases both the provider's private return and the bilateral opportunity created by another unit of learning. Provider adaptation rises, and complementarity makes the firm's knowledge investment more productive. Realized investment, performance, and bilateral value can therefore increase together. The firm's shortfall from first best may nevertheless widen because the bilateral opportunity grows faster than its privately captured return.

The positive reuse result depends on reuse creating an additional productive use. Accidental leakage, a transfer of bargaining rents, and a general model upgrade follow different economic paths. The mechanism instead extends the logic that reusable knowledge and capabilities can sustain relationship-specific investment (Dyer, 1996; Jia, 2013; Kang et al., 2009). Productive spillovers and appropriation hazards can coexist: the same outside use that weakens the firm's relative claim can increase the amount of learning that the relationship actually produces.

These comparisons yield a broader implication. The strategic content of a right is the vector of claims and productive opportunities that it changes. An ordinal measure of "more firm control" or "more provider control" discards this information. Continuity transfers part of an existing focal claim, whereas productive reuse can expand the opportunity that the parties are investing to create. Their effects cannot be inferred from the identity of the party whose position appears stronger.

5.3. Rights architecture and the firm boundary

Cross-client isolation and post-exit continuity generate different claim movements. A firm can be protected from cross-client diffusion while remaining dependent on the provider for operation. It can also obtain a runnable learned state while leaving some provider reuse permitted. Treating these arrangements as points on a single ownership scale would collapse economically relevant variation.

Split control makes the two uses separate design variables. It allows the parties to adjust continuity without mechanically changing reuse and to restrict diffusion without necessarily transferring the full operating configuration. The resulting gain is configurational: the additional dimension allows one investment incentive to be strengthened without automatically changing the other in the same direction.

The value result in Proposition 4 has a precise interpretation. Split control weakly improves the attainable objective because every bundled arrangement remains feasible within the larger rights set. The substantive content lies in the condition for a strict gain. A strict improvement arises when the best bundled arrangement leaves a profitable off-path adjustment unavailable and the benefit of that adjustment exceeds its implementation cost. Split control has no intrinsic advantage when the preferred configuration already lies on the bundled path or when legal, technical, and coordination costs make separate implementation uneconomic. The result therefore supports conditional unbundling of rights rather than universal fragmentation of control.

This configurational view connects AI governance to research on vertical architecture and ownership competence (Boudreau, 2010; Foss et al., 2021; Jacobides & Billinger, 2006). Formal title matters only insofar as the holder can exercise the relevant uses. Contract terms, system architecture, credentials, documentation, and operating capability jointly determine effective control. Standardization may reduce coordination difficulty while redistributing exposure among participants (Toh & Pyun, 2024); an architecture that simplifies focal operation can simultaneously concentrate control over reuse or transition.

Persistence increases the strategic importance of these choices. Current underinvestment leaves a smaller learning asset for subsequent periods. After an asset has accumulated, continuity can also preserve realized value when breakdown or transition becomes possible. The dynamic extension in Appendix B.11 shows that a sufficiently large inherited state can justify the fixed cost of changing architecture. Rights therefore shape both current investment and the stock on which future boundary choices are made. This channel links capability accumulation and boundary design (Argyres & Zenger, 2012; Jacobides & Winter, 2005) and suggests that governance may change as parties learn to contract and the value of the retained state grows (Mayer & Argyres, 2004).

5.4. Empirical implications

The model yields three observable patterns. A continuity-induced reversal requires post-exit continuity and firm proprietary-knowledge investment to rise while provider adaptation falls; situated AI performance declines only when the loss of adaptation dominates. More valuable productive reuse under incomplete isolation instead raises provider adaptation and, through complementarity, firm investment and performance, even as the firm's shortfall from the bilateral first best may widen. The split-control result applies when isolation and continuity can vary independently and a bundled arrangement constrains that choice. Distinguishing these mechanisms therefore requires separate observation of the effective rights, both investments, and operating performance.

Table 3. Three empirical distinctions implied by the theory

Governance change	Model implication	What would distinguish it empirically
Greater post-exit continuity	Proprietary-knowledge investment rises, but provider adaptation—and therefore situated AI performance—can fall under Equation (33).	Observe the two investments and repeatable task performance separately, rather than treating portability or switching alone as performance.
More valuable productive reuse under incomplete isolation	Both investments and bilateral value rise, while the firm's distance from its bilateral first-best investment can widen.	Distinguish productive cross-client learning from a transfer of bargaining rents, accidental leakage, or a general model upgrade.
Separate choice of isolation and continuity	The best split-control arrangement can outperform every arrangement available on a bundled rights path.	Recover the constraint linking the effective rights and their implementation costs, and establish that the clauses create distinct effective rights.

The contract analysis in Appendix C shows that cross-client isolation and post-exit continuity correspond to different contractual provisions and appear in different combinations. Its purpose is to ground the two rights in observable practice. Appendix C also describes the sample, coding procedure, and limitations of this evidence.

6. Conclusion

Enterprise AI relationships create an asset that neither party brings in finished form. The firm contributes proprietary knowledge, while the provider contributes a general model and customer-specific adaptation. Their joint work produces an AI learning asset. The *reverse information paradox* arises when the provider can reuse that learning elsewhere while the firm cannot keep the resulting AI system operating after exit.

Our model shows that two distinct rights govern this problem. Cross-client isolation limits provider reuse, while post-exit continuity determines what the firm can continue to operate. Greater continuity encourages firm investment, but it can also reduce provider adaptation and lower AI performance. Productive reuse follows a different channel. It can raise both parties' investments, AI performance, and bilateral value, even as the gap between the firm's actual and efficient investment grows. Because the two rights affect different incentives, assigning them separately can achieve outcomes that a bundled rule cannot.

The paper extends property-rights theory to learning created through use. Governance divides the returns to that learning and shapes how much of it the parties are willing to create. In enterprise AI, the firm boundary is partly defined by where customer-induced learning can travel and who can keep it operating after the relationship ends.

References

Aghion, P., & Tirole, J. (1997). Formal and real authority in organizations. *Journal of Political Economy*, *105*(1), 1–29. https://doi.org/10.1086/262063

Argote, L., Beckman, S. L., & Epple, D. (1990). The persistence and transfer of learning in industrial settings. *Management Science*, *36*(2), 140–154. https://doi.org/10.1287/mnsc.36.2.140

Argote, L., & Ingram, P. (2000). Knowledge transfer: A basis for competitive advantage in firms. *Organizational Behavior and Human Decision Processes*, *82*(1), 150–169. https://doi.org/10.1006/obhd.2000.2893

Argote, L., & Miron-Spektor, E. (2011). Organizational learning: From experience to knowledge. *Organization Science*, *22*(5), 1123–1137. https://doi.org/10.1287/orsc.1100.0621

Argyres, N. S., & Mayer, K. J. (2007). Contract design as a firm capability: An integration of learning and transaction cost perspectives. *Academy of Management Review*, *32*(4), 1060–1077. https://doi.org/10.5465/amr.2007.26585739

Argyres, N. S., & Zenger, T. R. (2012). Capabilities, transaction costs, and firm boundaries. *Organization Science*, *23*(6), 1643–1657. https://doi.org/10.1287/orsc.1110.0736

Arrow, K. J. (1962). Economic welfare and the allocation of resources for invention. In R. R. Nelson (Ed.), *The rate and direction of inventive activity: Economic and social factors* (pp. 609–626). Princeton University Press. https://doi.org/10.1515/9781400879762-024

Baker, G., Gibbons, R., & Murphy, K. J. (2002). Relational contracts and the theory of the firm. *Quarterly Journal of Economics*, *117*(1), 39–84. https://doi.org/10.1162/003355302753399445

Bartel, A., Ichniowski, C., & Shaw, K. (2007). How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills. *Quarterly Journal of Economics*, *122*(4), 1721–1758. https://doi.org/10.1162/qjec.2007.122.4.1721

Benkard, C. L. (2000). Learning and forgetting: The dynamics of aircraft production. *American Economic Review*, *90*(4), 1034–1054. https://doi.org/10.1257/aer.90.4.1034

Binmore, K., Rubinstein, A., & Wolinsky, A. (1986). The Nash bargaining solution in economic modelling. *RAND Journal of Economics*, *17*(2), 176–188. https://doi.org/10.2307/2555382

Bolton, P., & Whinston, M. D. (1993). Incomplete contracts, vertical integration, and supply assurance. *Review of Economic Studies*, *60*(1), 121–148. https://doi.org/10.2307/2297815

Boudreau, K. J. (2010). Open platform strategies and innovation: Granting access vs. devolving control. *Management Science*, *56*(10), 1849–1872. https://doi.org/10.1287/mnsc.1100.1215

Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics*, *117*(1), 339–376. https://doi.org/10.1162/003355302753399526

Brynjolfsson, E. (1994). Information assets, technology and organization. *Management Science*, *40*(12), 1645–1662. https://doi.org/10.1287/mnsc.40.12.1645

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *Quarterly Journal of Economics*, *140*(2), 889–942. https://doi.org/10.1093/qje/qjae044

Chatain, O., & Plaksenkova, E. (2026). Incentive alignment in ecosystems: The role of complementarity types and multi-sided bargaining. *Strategic Management Journal*, *47*(1), 228–256. https://doi.org/10.1002/smj.70012

Chatain, O., & Zemsky, P. (2011). Value creation and value capture with frictions. *Strategic Management Journal*, *32*(11), 1206–1231. https://doi.org/10.1002/smj.939

Che, Y.-K., & Hausch, D. B. (1999). Cooperative investments and the value of contracting. *American Economic Review*, *89*(1), 125–147. https://doi.org/10.1257/aer.89.1.125

Chen, P.-Y., & Hitt, L. M. (2002). Measuring switching costs and the determinants of customer retention in internet-enabled businesses: A study of the online brokerage industry. *Information Systems Research*, *13*(3), 255–274. https://doi.org/10.1287/isre.13.3.255.78

Chen, Y., & Bharadwaj, A. (2009). An empirical analysis of contract structures in IT outsourcing. *Information Systems Research*, *20*(4), 484–506. https://doi.org/10.1287/isre.1070.0166

Darr, E. D., Argote, L., & Epple, D. (1995). The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises. *Management Science*, *41*(11), 1750–1762. https://doi.org/10.1287/mnsc.41.11.1750

d'Aspremont, C., & Jacquemin, A. (1988). Cooperative and noncooperative R&D in duopoly with spillovers. *American Economic Review*, *78*(5), 1133–1137. https://doi.org/10.2307/1807173

David, R. J., & Han, S.-K. (2004). A systematic assessment of the empirical support for transaction cost economics. *Strategic Management Journal*, *25*(1), 39–58. https://doi.org/10.1002/smj.359

de Meza, D., & Lockwood, B. (1998). Does asset ownership always motivate managers? outside options and the property rights theory of the firm. *Quarterly Journal of Economics*, *113*(2), 361–386. https://doi.org/10.1162/003355398555621

Dierickx, I., & Cool, K. (1989). Asset stock accumulation and sustainability of competitive advantage. *Management Science*, *35*(12), 1504–1511. https://doi.org/10.1287/mnsc.35.12.1504

Dyer, J. H. (1996). Specialized supplier networks as a source of competitive advantage: Evidence from the auto industry. *Strategic Management Journal*, *17*(4), 271–291. https://doi.org/10.1002/(SICI)1097-0266(199604)17:4<271::AID-SMJ807>3.0.CO;2-Y

Dyer, J. H., & Singh, H. (1998). The relational view: Cooperative strategy and sources of interorganizational competitive advantage. *Academy of Management Review*, *23*(4), 660–679. https://doi.org/10.5465/amr.1998.1255632

Edlin, A. S., & Reichelstein, S. (1996). Holdups, standard breach remedies, and optimal investment. *American Economic Review*, *86*(3), 478–501.

Ethiraj, S. K., Kale, P., Krishnan, M. S., & Singh, J. V. (2005). Where do capabilities come from and how do they matter? A study in the software services industry. *Strategic Management Journal*, *26*(1), 25–45. https://doi.org/10.1002/smj.433

Foss, N. J., Klein, P. G., Lien, L. B., Zellweger, T., & Zenger, T. (2021). Ownership competence. *Strategic Management Journal*, *42*(2), 302–328. https://doi.org/10.1002/smj.3222

Giustiziero, G., Kretschmer, T., Somaya, D., & Wu, B. (2023). Hyperspecialization and hyperscaling: A resource-based theory of the digital firm. *Strategic Management Journal*, *44*(6), 1391–1424. https://doi.org/10.1002/smj.3365

Grant, R. M. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal*, *17*(S2), 109–122. https://doi.org/10.1002/smj.4250171110

Grossman, S. J., & Hart, O. D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, *94*(4), 691–719. https://doi.org/10.1086/261404

Grout, P. A. (1984). Investment and wages in the absence of binding contracts: A Nash bargaining approach. *Econometrica*, *52*(2), 449–460. https://doi.org/10.2307/1911498

Hagiu, A., & Wright, J. (2023). Data-enabled learning, network effects, and competitive advantage. *RAND Journal of Economics*, *54*(4), 638–667. https://doi.org/10.1111/1756-2171.12453

Hann, R. N., Subasi, M., & Zheng, Y. (2025). Building credible commitments via board ties: Evidence from the supply chain. *Strategic Management Journal*, *46*(11), 2839–2873. https://doi.org/10.1002/smj.3740

Hart, O., & Moore, J. (1990). Property rights and the nature of the firm. *Journal of Political Economy*, *98*(6), 1119–1158. https://doi.org/10.1086/261729

Hatch, N. W., & Dyer, J. H. (2004). Human capital and learning as a source of sustainable competitive advantage. *Strategic Management Journal*, *25*(12), 1155–1178. https://doi.org/10.1002/smj.421

Hoetker, G., & Mellewigt, T. (2009). Choice and performance of governance mechanisms: Matching alliance governance to asset type. *Strategic Management Journal*, *30*(10), 1025–1044. https://doi.org/10.1002/smj.775

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., & Chen, W. (2022). LoRA: Low-rank adaptation of large language models. *International Conference on Learning Representations*.

Jacobides, M. G., & Billinger, S. (2006). Designing the boundaries of the firm: From make, buy, or ally to the dynamic benefits of vertical architecture. *Organization Science*, *17*(2), 249–261. https://doi.org/10.1287/orsc.1050.0167

Jacobides, M. G., & Winter, S. G. (2005). The co-evolution of capabilities and transaction costs: Explaining the institutional structure of production. *Strategic Management Journal*, *26*(5), 395–413. https://doi.org/10.1002/smj.460

Jia, N. (2013). Competition, governance, and relationship-specific investments: Theory and implications for strategy. *Strategic Management Journal*, *34*(13), 1551–1567. https://doi.org/10.1002/smj.2077

Jones, C. I., & Tonetti, C. (2020). Nonrivalry and the economics of data. *American Economic Review*, *110*(9), 2819–2858. https://doi.org/10.1257/aer.20191330

Joskow, P. L. (1987). Contract duration and relationship-specific investments: Empirical evidence from coal markets. *American Economic Review*, *77*(1), 168–185.

Kale, P., Singh, H., & Perlmutter, H. (2000). Learning and protection of proprietary assets in strategic alliances: Building relational capital. *Strategic Management Journal*, *21*(3), 217–237. https://doi.org/10.1002/(SICI)1097-0266(200003)21:3<217::AID-SMJ95>3.0.CO;2-Y

Kang, M.-P., Mahoney, J. T., & Tan, D. (2009). Why firms make unilateral investments specific to other firms: The case of OEM suppliers. *Strategic Management Journal*, *30*(2), 117–135. https://doi.org/10.1002/smj.730

Kemp, A. (2024). Competitive advantage through artificial intelligence: Toward a theory of situated AI. *Academy of Management Review*, *49*(3), 618–635. https://doi.org/10.5465/amr.2020.0205

Kim, H., Glaeser, E. L., Hillis, A., Kominers, S. D., & Luca, M. (2024). Decision authority and the returns to algorithms. *Strategic Management Journal*, *45*(4), 619–648. https://doi.org/10.1002/smj.3569

Klein, B., Crawford, R. G., & Alchian, A. A. (1978). Vertical integration, appropriable rents, and the competitive contracting process. *Journal of Law and Economics*, *21*(2), 297–326. https://doi.org/10.1086/466922

Kogut, B., & Zander, U. (1992). Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science*, *3*(3), 383–397. https://doi.org/10.1287/orsc.3.3.383

Krakowski, S., Luger, J., & Raisch, S. (2023). Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal*, *44*(6), 1425–1452. https://doi.org/10.1002/smj.3387

Levitt, B., & March, J. G. (1988). Organizational learning. *Annual Review of Sociology*, *14*, 319–340. https://doi.org/10.1146/annurev.so.14.080188.001535

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, *33*, 9459–9474.

Lo, D. H.-F., Zanarone, G., & Ghosh, M. (2022). Contracting to (dis)incentivize? an integrative transaction-cost approach on how contracts govern specific investments. *Strategic Management Journal*, *43*(8), 1528–1555. https://doi.org/10.1002/smj.3376

Mayer, K. J., & Argyres, N. S. (2004). Learning to contract: Evidence from the personal computer industry. *Organization Science*, *15*(4), 394–410. https://doi.org/10.1287/orsc.1040.0074

Milgrom, P., & Roberts, J. (1990). The economics of modern manufacturing: Technology, strategy, and organization. *American Economic Review*, *80*(3), 511–528. https://doi.org/10.2307/2006681

Milgrom, P., & Shannon, C. (1994). Monotone comparative statics. *Econometrica*, *62*(1), 157–180. https://doi.org/10.2307/2951479

Nash, J. F. (1950). The bargaining problem. *Econometrica*, *18*(2), 155–162. https://doi.org/10.2307/1907266

Nickerson, J. A., & Zenger, T. R. (2004). A knowledge-based theory of the firm—the problem-solving perspective. *Organization Science*, *15*(6), 617–632. https://doi.org/10.1287/orsc.1040.0093

Nöldeke, G., & Schmidt, K. M. (1995). Option contracts and renegotiation: A solution to the hold-up problem. *The RAND Journal of Economics*, *26*(2), 163–179. https://doi.org/10.2307/2555911

Pacheco-de-Almeida, G., & Zemsky, P. B. (2012). Some like it free: Innovators' strategic use of disclosure to slow down competition. *Strategic Management Journal*, *33*(7), 773–793. https://doi.org/10.1002/smj.971

Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, 1–22. https://doi.org/10.1145/3586183.3606763

Parmigiani, A. (2007). Why do firms both make and buy? an investigation of concurrent sourcing. *Strategic Management Journal*, *28*(3), 285–311. https://doi.org/10.1002/smj.580

Poppo, L., & Zenger, T. (2002). Do formal contracts and relational governance function as substitutes or complements? *Strategic Management Journal*, *23*(8), 707–725. https://doi.org/10.1002/smj.249

Rajan, R. G., & Zingales, L. (1998). Power in a theory of the firm. *Quarterly Journal of Economics*, *113*(2), 387–432. https://doi.org/10.1162/003355398555630

Rosen, J. B. (1965). Existence and uniqueness of equilibrium points for concave N-person games. *Econometrica*, *33*(3), 520–534. https://doi.org/10.2307/1911749

Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica*, *50*(1), 97–109. https://doi.org/10.2307/1912531

Sakhartov, A. V., & Helfat, C. E. (2026). Resource redeployment in multi-business firms: Centralized or decentralized? *Strategic Management Journal*, *47*(7), 1878–1906. https://doi.org/10.1002/smj.70077

Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, *28*, 2503–2511.

Susarla, A. (2012). Contractual flexibility, rent seeking, and renegotiation design: An empirical analysis of information technology outsourcing contracts. *Management Science*, *58*(7), 1388–1407. https://doi.org/10.1287/mnsc.1110.1493

Susarla, A., Subramanyam, R., & Karhade, P. (2010). Contractual provisions to mitigate holdup: Evidence from information technology outsourcing. *Information Systems Research*, *21*(1), 37–55. https://doi.org/10.1287/isre.1080.0204

Tiwana, A. (2008). Does interfirm modularity complement ignorance? a field study of software outsourcing alliances. *Strategic Management Journal*, *29*(11), 1241–1252. https://doi.org/10.1002/smj.694

Toh, P. K., & Pyun, E. (2024). Risky business: How standardization as coordination tool in ecosystems impacts firm-level uncertainty. *Strategic Management Journal*, *45*(4), 649–679. https://doi.org/10.1002/smj.3562

Vijairaghavan, V., Hidaji, H., & Nault, B. R. (2025). Consequences of resorting to fines and investments to regulate data portability [Published online August 26, 2025]. *Management Science*, *72*(5), 3761–3779. https://doi.org/10.1287/mnsc.2024.08404

Wan, F., Yang, T., Shi, X., Rong, K., & Ansari, S. (2026). Scaling high and wide: How firms leverage AI and organizational design to overcome the scale–scope trade-off [Early View]. *Strategic Management Journal*. https://doi.org/10.1002/smj.70090

Weigelt, C. (2013). Leveraging supplier capabilities: The role of locus of capability deployment. *Strategic Management Journal*, *34*(1), 1–21. https://doi.org/10.1002/smj.1998

Whinston, M. D. (2003). On the transaction cost determinants of vertical integration. *The Journal of Law, Economics, and Organization*, *19*(1), 1–23. https://doi.org/10.1093/jleo/19.1.1

Williamson, O. E. (1979). Transaction-cost economics: The governance of contractual relations. *Journal of Law and Economics*, *22*(2), 233–261. https://doi.org/10.1086/466942

Williamson, O. E. (1991). Comparative economic organization: The analysis of discrete structural alternatives. *Administrative Science Quarterly*, *36*(2), 269–296. https://doi.org/10.2307/2393356

Yokoi, T., Laureiro-Martinez, D., Magni, F., & Brusoni, S. (2026). Organizing across cognitive asymmetry in human–AI collaboration: A study of perfume creation [Early View]. *Strategic Management Journal*. https://doi.org/10.1002/smj.70089

A. Mathematical Foundations and Proofs

This appendix proves the four results in the order in which they appear in the main text. It first separates a marginal hold-up wedge from the resulting equilibrium investment shortfall. It then derives the continuity and productive-reuse comparative statics. The final section proves the local-rank lemma and the split-control implementation result. Throughout, a star denotes the private equilibrium and an FB superscript denotes the bilateral first best.

A.1. Common notation and admissibility

Write

$$r \equiv 1 - x, \qquad K \equiv k_F k_P, \qquad g_F \equiv \beta(v - \omega), \qquad g_P \equiv (1 - \beta)v + \beta\omega,$$

and

$$M_P \equiv k_P + \chi g_P, \qquad A = v[\beta + (1 - \beta)y] - \lambda r,$$

$$B = (1 - \beta)v(1 - y) + vr, \qquad C = A + B = v + (v - \lambda)r.$$

The direct claims exhaust focal operating value, $g_F + g_P = v$, just as the AI-learning-asset claims exhaust bilateral marginal value, $A + B = C$.

Assumption 1 (Admissible parameter region)**.** *For every rights bundle and local parameter path under consideration,*

$$A > 0, \qquad B \geq 0, \qquad C = A + B > 0, \qquad K > \chi^2 C^2. \tag{47}$$

For Proposition 4, "under consideration" means every point of the full rights box Γ_S.

Because $A, B \geq 0$, the final inequality implies

$$\Delta \equiv K - \chi^2 AB > 0, \qquad \Delta^{FB} \equiv K - \chi^2 C^2 > 0,$$

as well as $K - \chi^2 A^2 > 0$. The condition is stronger than existence of the private equilibrium alone. It places the private and bilateral linear best-response systems in a common stable region and makes the bilateral objective strictly concave.

All comparative statics below are local and are evaluated along parameter paths that remain in the admissible region. At $x, y \in \{0, 1\}$ and at a primitive lower bound such as $v = 0$, a derivative means the corresponding feasible one-sided derivative. Inverse-function statements apply only at interior rights bundles.

A.1.1. Closed forms and a derivative basis

The private first-order conditions can be written

$$\begin{bmatrix} k_F & -A\chi \\ -B\chi & k_P \end{bmatrix} \begin{bmatrix} d \\ e \end{bmatrix} = \begin{bmatrix} A \\ g_P \end{bmatrix}. \tag{48}$$

Cramer's rule gives

$$d^* = \frac{AM_P}{\Delta}, \tag{49}$$

$$e^* = \frac{k_F g_P + \chi AB}{\Delta}. \tag{50}$$

The bilateral benchmark instead satisfies

$$\begin{bmatrix} k_F & -C\chi \\ -C\chi & k_P \end{bmatrix} \begin{bmatrix} d^{FB} \\ e^{FB} \end{bmatrix} = \begin{bmatrix} C \\ v \end{bmatrix}, \tag{51}$$

and hence

$$d^{FB} = \frac{C(k_P + \chi v)}{\Delta^{FB}}, \tag{52}$$

$$e^{FB} = \frac{k_F v + \chi C^2}{\Delta^{FB}}. \tag{53}$$

Holding g_P fixed, differentiation of Equations (49)–(50) with respect to the two AI-learning-asset claims gives

$$d_A = \frac{KM_P}{\Delta^2}, \qquad d_B = \frac{A^2\chi^2 M_P}{\Delta^2}, \tag{54}$$

$$e_A = \frac{B\chi k_F M_P}{\Delta^2}, \qquad e_B = \frac{A\chi k_F M_P}{\Delta^2}. \tag{55}$$

The bilateral marginal returns left uninternalized at the private equilibrium are

$$S_d \equiv \left.\frac{\partial S}{\partial d}\right|_{(d^*,e^*)} = B(1 + \chi e^*), \tag{56}$$

$$S_e^{\omega} \equiv \left.\frac{\partial S}{\partial e}\right|_{(d^*,e^*)} = g_F + \chi A d^*. \tag{57}$$

The first identity follows from $k_F d^* = A(1 + \chi e^*)$ and $C - A = B$. The second follows from $k_P e^* = g_P + \chi B d^*$, $v - g_P = g_F$, and $C - B = A$. Therefore, for a primitive θ that changes (A, B, C)

while leaving the direct claims fixed,

$$W_\theta = C_\theta \mathcal{L}^* + S_d d_\theta^* + S_e^\omega e_\theta^*. \tag{58}$$

A.2. A general comparative-statics result

The bilinear technology makes the main results transparent, but it is useful to state exactly which parts survive outside that specialization. Let the AI learning asset be a twice continuously differentiable function $f(d, e)$. The private first-order conditions are

$$c_F'(d) = A f_d(d, e), \qquad c_P'(e) = h_e(e; \theta) + B f_e(d, e), \tag{59}$$

where h is the provider's direct-return function. Holding the technology and cost functions fixed as θ changes, define

$$D_F \equiv c_F'' - A f_{dd}, \qquad D_P \equiv c_P'' - h_{ee} - B f_{ee}, \qquad m \equiv f_{de},$$

and suppose

$$D_F > 0, \qquad D_P > 0, \qquad \Omega \equiv D_F D_P - ABm^2 > 0. \tag{60}$$

Implicit differentiation of Equation (59) yields

$$d_\theta = \frac{D_P A_\theta f_d + Am(h_{e\theta} + B_\theta f_e)}{\Omega}, \tag{61}$$

$$e_\theta = \frac{BmA_\theta f_d + D_F(h_{e\theta} + B_\theta f_e)}{\Omega}. \tag{62}$$

Along a continuity-like pure claim reallocation, $(A_\theta, B_\theta, h_{e\theta}) = (a, -a, 0)$ with $a > 0$, these expressions become

$$d_\theta = \frac{a(D_P f_d - Am f_e)}{\Omega}, \tag{63}$$

$$e_\theta = \frac{a(Bm f_d - D_F f_e)}{\Omega}. \tag{64}$$

Thus complementarity $m \geq 0$ by itself does not sign either response. For a seeded multiplicative technology $f(d, e) = d\psi(e)$ with $\psi(e) > 0$, an interior firm choice, and quadratic firm cost, the firm's first-order condition $k_F d = A\psi(e)$ implies

$$\operatorname{sgn}(e_\theta) = \operatorname{sgn}(B - A) \quad \text{when } \psi'(e) > 0, \tag{65}$$

and

$$d_\theta > 0 \quad \Longleftrightarrow \quad k_F D_P > A^2[\psi'(e)]^2. \tag{66}$$

For $\psi(e) = 1 + \chi e$, $h(e) = g_P e$, and quadratic provider cost, $D_P = k_P$; Assumption 1 supplies the last inequality.

Assumption 2 (General learning technology)**.** *$f_d > 0$ in the relevant region, f is weakly increasing in e, and f is weakly concave in each input. Investment costs are twice continuously differentiable and strictly convex, and the relevant conditional optima are unique.*

Theorem 1 (General conditional underinvestment)**.** *Suppose the firm's private return to the AI learning asset is $Af(d, e)$ and bilateral return is $Cf(d, e)$, where $0 \le A \le C$. Under Assumption 2, for every fixed e, $d^F(e) \le d^{FB}(e)$, with strict inequality whenever $C > A$ and the bilateral choice is positive. If the full private and bilateral best-response maps are isotone contractions on a common invariant order interval and the latter weakly dominates the former componentwise, their fixed points inherit the componentwise order.*

The ordered-map argument draws on the monotone-comparative-statics and stable concave-game traditions (Milgrom & Shannon, 1994; Rosen, 1965). Those methods motivate the sufficient ordering and stability conditions; the particular Ω restriction and fixed-point comparison above are specific to this model.

Proof. For fixed e, private and bilateral marginal objectives are respectively $Af_d - c'_F$ and $Cf_d - c'_F$. The second weakly exceeds the first at every d. Concavity of f in d and strict convexity of cost give the conditional ordering. If the bilateral choice is positive, its first-order condition is interior; when $C > A$, the private marginal objective evaluated there is $(A - C)f_d < 0$, so the unique private optimum is strictly smaller (and may be the zero corner). For the equilibrium statement, iterate both maps from a common lower point in the invariant interval. Componentwise dominance preserves the order at every iteration; contraction gives convergence to the unique ordered fixed points. The additional map conditions are essential: one-dimensional conditional underinvestment alone does not order a two-input equilibrium. □

A.3. Proof of the hold-up and underinvestment result

For arbitrary disagreement claims F_0 and P_0, the Nash bargaining solution assigns (Binmore et al., 1986; Grout, 1984; Nash, 1950)

$$A = F_0 + \beta(C - F_0 - P_0), \qquad B = P_0 + (1 - \beta)(C - F_0 - P_0).$$

Consequently

$$C - A = (1 - \beta)(C - F_0 - P_0) + P_0 = (1 - \beta)(C - F_0) + \beta P_0, \tag{67}$$

so nonnegative bargaining surplus and $P_0 \geq 0$ are primitive sufficient conditions for $C - A \geq 0$. This proves Lemma 1. Substitution of $F_0 = \nu y - \lambda r$ and $P_0 = \nu r$ gives

$$H_F = C - A = B = (1-\beta)\nu(1-y) + \nu r.$$

The derivatives and bargaining-weight limit stated in the main text follow immediately.

Holding e fixed and subtracting the firm conditions gives

$$d^{FB}(e) - d^F(e) = \frac{(C-A)(1+\chi e)}{k_F} = \frac{H_F(1+\chi e)}{k_F}.$$

Holding d fixed and subtracting the provider conditions gives

$$e^{FB}(d) - e^P(d) = \frac{\nu - g_P + \chi(C-B)d}{k_P} = \frac{\beta(\nu-\omega) + \chi A d}{k_P} > 0.$$

To order the fixed points, define

$$T^P(d,e) = \left(\frac{A(1+\chi e)}{k_F}, \frac{g_P + \chi B d}{k_P}\right),$$
$$T^{FB}(d,e) = \left(\frac{C(1+\chi e)}{k_F}, \frac{\nu + \chi C d}{k_P}\right).$$

Both maps are increasing. Their nonnegative slope matrices have spectral radii $\chi\sqrt{AB/K}$ and $\chi C/\sqrt{K}$, respectively, both below one under Assumption 1; equivalently, each is a contraction in a suitable weighted norm. Moreover, $T^{FB} \geq T^P$ componentwise and the provider component is strictly larger everywhere. Monotone fixed-point comparison first gives $(d^{FB}, e^{FB}) \geq (d^*, e^*)$. Evaluating the provider maps at these ordered fixed points then yields

$$e^{FB} = T_2^{FB}(d^{FB}, e^{FB}) \geq T_2^{FB}(d^*, e^*) > T_2^P(d^*, e^*) = e^*,$$

where the strict difference is at least $g_F/k_P > 0$. If $B > 0$, then $C > A$ and the firm component is directly strict. If $B = 0$ but $\chi > 0$, the strict provider ordering feeds into the firm response at the next iteration. Only when $B = 0$ and $\chi = 0$ does the firm choose $A/k_F = C/k_F$ in both systems. This proves Proposition 1.

A fixed transfer T_0 does not enter any of these derivatives. It shifts the levels of the parties' payoffs with derivatives $T_{0,d} = T_{0,e} = 0$; it therefore cannot close either marginal investment gap.[24]

[24]This statement is limited to a transfer fixed before investment and independent of nonverifiable actions or outcomes, as in the narrow contract space of Grossman and Hart (1986) and Hart and Moore (1990). Verifiable milestones, options, breach remedies, or relational enforcement can alter marginal incentives (Baker et al., 2002; Edlin & Reichelstein, 1996; Nöldeke & Schmidt, 1995).

A.4. Proof of continuity-induced reallocation

Let $a_y = (1-\beta)v$. Continuity follows $(A_y, B_y, C_y) = (a_y, -a_y, 0)$. Applying Equations (54)–(55) gives

$$d_y^* = a_y(d_A - d_B) = \frac{a_y M_P (K - \chi^2 A^2)}{\Delta^2} > 0, \tag{68}$$

$$e_y^* = a_y(e_A - e_B) = \frac{a_y \chi k_F M_P (B - A)}{\Delta^2}. \tag{69}$$

Because $Q = e + d(1 + \chi e)$,

$$Q_y^* = \frac{a_y M_P}{\Delta^2} \left[(1 + \chi e^*)(K - \chi^2 A^2) + \chi k_F (1 + \chi d^*)(B - A) \right]. \tag{70}$$

The prefactor is positive. Moving the second term to the other side gives Equation (33), proving Proposition 2.

Since $C_y = 0$, the envelope in Equation (58) gives

$$W_y = \frac{a_y M_P}{\Delta^2} \left[B(1 + \chi e^*)(K - \chi^2 A^2) + \chi k_F S_e^\omega (B - A) \right]. \tag{71}$$

Again the prefactor is positive. Rearranging proves Equation (35) and Corollary 1. The two thresholds are not nested in general. The performance comparison weights the adaptation response by $1 + \chi d^*$, whereas the bilateral-value comparison weights it by the uninternalized return S_e^ω and weights the firm response by $B(1 + \chi e^*)$.

Two admissible parameter points establish both directions of nonimplication. Set $v = 1$ and list the remaining primitives as $(\beta, \omega, \lambda, \nu, x, y, \chi, k_F, k_P)$. At

$$(0.3382, 0.9386, 0.1001, 0.5329, 0.3726, 0.9033, 1.1368, 17.5496, 0.1329),$$

$A = 0.8732 > B = 0.3983$, $\Delta^{FB} = 0.2429 > 0$, $Q_y = -0.0123$, and $W_y = 0.1397$. At

$$(0.2399, 0.3456, 0.0905, 0.5300, 0.4357, 0.9322, 0.9052, 2.9287, 0.5664),$$

$A = 0.8974 > B = 0.3506$, $\Delta^{FB} = 0.3826 > 0$, $Q_y = 0.1078$, and $W_y = -0.1323$. Thus a performance reversal can occur without a bilateral-value reversal, and conversely.

At $\chi = 0$, $d_y^* = a_y / k_F > 0$ and $e_y^* = 0$, so both performance and bilateral value weakly rise. At $A = B$, adaptation is locally unchanged and performance rises. At $y = 1$, the formulas are left derivatives; a further increase in continuity is not feasible.

A.5. Proof of productive-reuse duality

Let $r = 1 - x$. Productive reuse follows $(A_v, B_v, C_v) = (0, r, r)$. Hence

$$d_v^* = r d_B = \frac{r A^2 \chi^2 M_P}{\Delta^2}, \tag{72}$$

$$e_v^* = r e_B = \frac{r A \chi k_F M_P}{\Delta^2}. \tag{73}$$

Both derivatives are strict for $r > 0$ and $\chi > 0$; the total derivative of $Q = e + d(1 + \chi e)$ is then strict as well.

Equation (58) yields

$$W_v = r\left[\mathcal{L}^* + S_d d_B + S_e^{\omega} e_B\right] > 0 \quad \text{for } r > 0. \tag{74}$$

The direct $C_v \mathcal{L}^*$ term makes the inequality strict even at $\chi = 0$. Differentiating the first-best firm input with respect to C gives

$$d_C^{FB} = \frac{(k_P + \chi v)(K + \chi^2 C^2)}{(K - \chi^2 C^2)^2}.$$

Therefore

$$\frac{\partial(d^{FB} - d^*)}{\partial v} = r\left[\frac{(k_P + \chi v)(K + \chi^2 C^2)}{(K - \chi^2 C^2)^2} - \frac{A^2 \chi^2 M_P}{\Delta^2}\right]. \tag{75}$$

The bracket is strictly positive. First, $k_P + \chi v > M_P$ because $v - g_P = g_F > 0$ whenever $\chi > 0$ (and the remaining comparison is still strict when $\chi = 0$). Second, $K + \chi^2 C^2 > A^2 \chi^2$. Finally,

$$0 < K - \chi^2 C^2 \leq \Delta,$$

because $C^2 - AB = A^2 + AB + B^2 > 0$. Thus the first fraction has a larger numerator and a weakly smaller positive denominator. Multiplication by $r > 0$ proves the strict gap result.

For comparison, a pure claim transfer toward the provider follows $(A_\tau, B_\tau, C_\tau) = (-1, 1, 0)$. Equations (54) and (55) give

$$d_\tau^* = -d_A + d_B = -\frac{M_P(K - \chi^2 A^2)}{\Delta^2} < 0,$$

$$e_\tau^* = -e_A + e_B = \frac{\chi k_F M_P (A - B)}{\Delta^2}.$$

The provider and performance effects are therefore not signed. This completes the proof of Proposition 3 and shows why productive reuse cannot be represented as a redistribution of a fixed surplus.

A.6. Proof of local irreducibility and split implementation

Isolation and continuity generate the claim Jacobian

$$J_{AB} = \frac{\partial(A, B)}{\partial(x, y)} = \begin{bmatrix} \lambda & (1-\beta)\nu \\ -\nu & -(1-\beta)\nu \end{bmatrix}.$$

Its determinant is $(1-\beta)\nu(\nu-\lambda)$. At an interior rights bundle, if this determinant is nonzero, the inverse-function theorem implies that the image of an interior neighborhood in the rights box contains an open neighborhood in claim space. No differentiable scalar index can locally represent both claims there, because such a representation would make their gradients collinear.

Using Equations (54)–(55), the determinant of the investment response to claims is

$$\det \frac{\partial(d^*, e^*)}{\partial(A, B)} = \frac{A\chi k_F M_P^2 (K - \chi^2 AB)}{\Delta^4} = \frac{A\chi k_F M_P^2}{\Delta^3} > 0$$

when $\chi > 0$. This proves Lemma 2. The condition $\lambda \neq \nu$ is sufficient for this particular two-right parameterization; it is not necessary for all decompositions of isolation. Appendix B separates reuse restriction from diffusion protection. By the chain rule, the composite local determinant at an interior rights bundle is

$$\det \frac{\partial(d^*, e^*)}{\partial(x, y)} = \frac{A\chi k_F M_P^2}{\Delta^3}(1-\beta)\nu(\nu-\lambda),$$

which is nonzero under the two conditions stated in the lemma.

We now prove Proposition 4. Assumption 1 holds throughout Γ_S, so the closed-form equilibrium and Φ are continuous there. The rights box is compact, and the Lipschitz image $\Gamma_U = \gamma([0, 1])$ is compact as well; both maxima therefore exist. Because $\Gamma_U \subseteq \Gamma_S$, every bundled choice is feasible under split control. Maximizing the same objective over the two sets gives

$$\max_{q \in \Gamma_S} \Phi(q) \geq \max_{q \in \Gamma_U} \Phi(q),$$

which proves part (a).

For part (b), set $q_U = \gamma(z_U)$. At an interior regular optimum of the bundled path, the chain rule gives

$$\left. \frac{d\Phi(\gamma(z))}{dz} \right|_{z=z_U} = \nabla\Phi(q_U) \cdot \gamma'(z_U) = 0.$$

If q_U is interior to the rights box and $\nabla\Phi(q_U) \neq 0$, the vector $h = \nabla\Phi(q_U)$ is a feasible local direction. Differentiability gives

$$\Phi(q_U + th) - \Phi(q_U) = t\|\nabla\Phi(q_U)\|^2 + o(t) > 0$$

for sufficiently small $t > 0$. The direction is normal to the bundled tangent and cannot be obtained by a first-order movement along that path.

For the diagonal path $\gamma(z) = (z, z)$, the path condition reduces to $\Phi_x + \Phi_y = 0$. An interior split optimum must instead satisfy $\Phi_x = \Phi_y = 0$, proving part (c). Finally, if Φ is strictly concave, its split maximizer is unique. When that maximizer is absent from Γ_U, no point on the bundled path can attain the same objective value; the inequality is strict. This proves part (d).

For completeness, the boundary version uses the tangent cone

$$T_S(q) = \mathrm{cl}\{t(z - q) : t \geq 0, z \in \Gamma_S\}.$$

If a bundled optimum q_U lies on the boundary of either the path or the rights box, its first-order condition applies only to feasible path tangents. Split control admits a strict local improvement whenever there exists $h \in T_S(q_U)$ with $\nabla\Phi(q_U) \cdot h > 0$. If no such direction exists, the first-order comparison is inconclusive; the weak feasible-set inequality still holds. This tangent-cone statement prevents an interior two-sided derivative from being applied mechanically at $x, y \in \{0, 1\}$.

A.7. Boundary and interpretation notes

The closed forms concern the positive-investment region. If $A \leq 0$, the firm's best response is the corner $d = 0$ and the interior derivatives do not apply. At $B = 0$, $A = C$ but provider adaptation remains positive because $g_P > 0$. At full isolation, $r = 0$ and all derivatives along the productive reuse path vanish. Bargaining feasibility holds because $C - F_0 - P_0 = v(1 - y) \geq 0$.

The split-control theorem therefore compares the best attainable bilateral objectives over the two governance sets.

B. Extensions, Alternative Specifications, and Boundary Conditions

The baseline is deliberately spare. This appendix considers what changes when the AI learning asset has several components, written rights are imperfectly enforced, isolation separates reuse from diffusion, a contractor also invests, or learning carries into a later period. The bilateral investment problem in the main text remains the point of comparison.

B.1. A multi-component learning asset

The main text uses a scalar learning asset $\mathcal{L}$. Let the underlying state instead be

$$\mathbf{z} = (z_I, z_A, z_M, z_O), \tag{76}$$

where z_I is informational state, z_A adaptive technical state, z_M interaction and memory state, and z_O organizational coupling.[25] Let current productive learning be

$$\mathcal{L} = \Phi(\mathbf{z}), \tag{77}$$

Assume that Φ is weakly increasing in every component. At a smooth point, this condition is $\Phi_j \geq 0$; strict inequality is not required for a component that is currently slack. The scalar baseline can be understood as an aggregator of these components.

If the components are strict complements, a deliberately strict bottleneck representation is

$$\Phi(\mathbf{z}) = \min_j \{a_j z_j\}, \qquad z_j \geq 0, \quad a_j > 0. \tag{78}$$

A bottleneck can then make delivery of one component ineffective without delivery of the others. If components are imperfect substitutes, a CES aggregator is

$$\Phi(\mathbf{z}) = \left(\sum_j \omega_j z_j^{\rho_z} \right)^{1/\rho_z}, \qquad z_j > 0, \quad \omega_j > 0, \quad \sum_j \omega_j = 1, \quad \rho_z \leq 1, \quad \rho_z \neq 0. \tag{79}$$

At $\rho_z = 0$, Equation (79) is understood by its Cobb–Douglas limit $\prod_j z_j^{\omega_j}$.[26] The main hold-up identity applies to the marginal value of Φ at smooth points. For the weakest-link form, let $I(\mathbf{z}) = \{j : a_j z_j = \Phi(\mathbf{z})\}$. Its one-sided directional derivative along $\mathbf{h}$ is $D\Phi(\mathbf{z}; \mathbf{h}) = \min_{j \in I(\mathbf{z})} a_j h_j$; a slack component therefore has zero local effect, while a tie requires a supergradient or this directional-derivative characterization. The representation clarifies that transferring one artifact can leave the binding state elsewhere.

B.2. Continuity as a weakest-link right

Let

$$y_D = \text{data and output delivery}, \qquad y_S = \text{learned-state delivery}, \qquad y_O = \text{independent operation},$$

[25]The four components classify economically distinct carriers of retained learning; they are not an exhaustive technical taxonomy. The distinction follows work that locates organizational learning in heterogeneous repositories and treats productive stocks as accumulated rather than purchased all at once (Argote & Miron-Spektor, 2011; Dierickx & Cool, 1989). Technical-debt research similarly shows that model code alone does not exhaust a deployed machine- learning system's consequential state (Sculley et al., 2015).

[26]The O-ring logic in Kremer (1993) motivates strong task complementarity but does not supply the exact minimum operator used here. The CES form follows Arrow et al. (1961); ρ_z governs substitutability and the $\rho_z \to 0$ limit is Cobb–Douglas. Both forms are alternative specifications rather than estimates of an enterprise-AI production function.

with each component in $[0, 1]$. When all three are necessary to preserve a unit of learning after exit, effective continuity is

$$y = \min\{y_D, y_S, y_O\}. \tag{80}$$

Lemma 3 (Learning-asset continuity bottleneck)**.** *Increasing a nonbinding component of Equation* (80) *does not change the firm's disagreement payoff. Complete raw-data export does not eliminate continuity hold-up when learned-state delivery or independent operation remains incomplete.*

Proof. If $y_j > \min\{y_D, y_S, y_O\}$, any feasible increase in y_j leaves the minimum unchanged. In the main model, the firm's disagreement value per unit of learning is $F_0 = \nu y - \lambda(1 - x)$, so the learning asset contributes $F_0\mathcal{L}$ to its disagreement payoff. The continuity-dependent part is unchanged when the minimum is unchanged. Setting $y_D = 1$ is insufficient whenever $\min\{y_S, y_O\} < 1$. □

A smooth CES continuity aggregator smooths the kink at a binding tie. The strategic implication is unchanged: data portability can support switching (Vijairaghavan et al., 2025), but possession of source files is not equivalent to preserving the complementary architecture and productive state that use them (Boudreau, 2010; Sculley et al., 2015).

B.3. Written isolation, hidden reuse, and assurance

The main text interprets x as effective isolation. Let $\bar{x} \in [0, 1]$ be promised isolation and let $h \in [0, \bar{x}]$ denote hidden circumvention, so

$$x = \bar{x} - h. \tag{81}$$

Assurance is fixed at date 0. After the AI learning asset is formed but before the disagreement uses are realized, the provider chooses h. For $\mathcal{L} > 0$, its total incremental payoff from circumvention is $\mathcal{L}$ times the per-unit problem

$$\max_{0 \le h \le \bar{x}} \left\{ \nu h - a\left(qh + \frac{\psi}{2}h^2\right)\right\}, \qquad a, q, \psi > 0, \tag{82}$$

where a scales the provider-borne expected cost of circumvention and q and ψ determine its level and curvature.[27] The solution is

$$h^*(\bar{x}, a) = \min\left\{\bar{x}, \max\left[0, \frac{\nu/a - q}{\psi}\right]\right\}. \tag{83}$$

[27] The assurance index is a reduced-form composite, not an observed enforcement probability. Research on IT outsourcing, governance safeguards, and relational enforcement motivates its ingredients (Baker et al., 2002; Oxley, 1997; Poppo & Zenger, 2002; Susarla et al., 2010); the linear–quadratic cost is chosen for tractability. The accounting below treats the cost as a real provider loss or a payment that leaves the focal dyad. If a sanction were transferred to the firm, its incidence would also enter F_0 and the claims would have to be recomputed. Fixed assurance implementation resources instead belong in $G(x, y)$.

If $a \leq \nu/(q + \psi\bar{x})$, promised isolation is fully circumvented and effective isolation is zero. If $\nu/(q + \psi\bar{x}) < a < \nu/q$, circumvention is partial. If $a \geq \nu/q$, hidden reuse is fully deterred. The choice is unique for $\mathcal{L} > 0$ and is independent of its level because $\mathcal{L}$ multiplies the entire objective; the date-1 investors can therefore anticipate h^*. At $\mathcal{L} = 0$, the choice of h is payoff irrelevant.

The cost that induces compliance must also enter the claim accounting. Define

$$c_a(h) \equiv a\left(qh + \frac{\psi}{2}h^2\right).$$

When this is a provider-borne loss, the disagreement and bilateral values per unit of learning are

$$\begin{aligned} F_0^h &= \nu y - \lambda(1 - x), \\ P_0^h &= \nu(1 - x) - c_a(h^*), \\ C^h &= \nu + (\nu - \lambda)(1 - x) - c_a(h^*). \end{aligned} \tag{84}$$

They leave agreement surplus $C^h - F_0^h - P_0^h = \nu(1 - y)$ and imply

$$A^h = \nu[\beta + (1 - \beta)y] - \lambda(1 - x), \qquad B^h = (1 - \beta)\nu(1 - y) + \nu(1 - x) - c_a(h^*),$$

with $A^h + B^h = C^h$. Thus assurance can affect investment through both effective isolation and the net circumvention cost; the interior baseline formulas apply only when $A^h > 0$, $B^h \geq 0$, and the corresponding stability condition holds. This is why the contract evidence examines legal and technical provisions together rather than relying on clause labels alone.

B.4. Separating reuse restriction from diffusion protection

The baseline uses one effective isolation measure because the same technical or contractual arrangement often affects both provider reuse and focal diffusion exposure. Yet nonrival cross-user learning and competitive diffusion are economically distinct (Choe et al., 2023; Hagiu & Wright, 2023; Jones & Tonetti, 2020; Pacheco-de-Almeida & Zemsky, 2012). Let $x_R \in [0, 1]$ restrict productive cross-client reuse and let $x_D \in [0, 1]$ protect the focal firm from diffusion. The marginal bilateral and disagreement values become

$$C(x_R, x_D) = \nu + \nu(1 - x_R) - \lambda(1 - x_D), \tag{85}$$

$$F_0(x_D, y) = \nu y - \lambda(1 - x_D), \tag{86}$$

$$P_0(x_R) = \nu(1 - x_R). \tag{87}$$

The negotiated claims are therefore

$$A(x_D, y) = v[\beta + (1-\beta)y] - \lambda(1 - x_D), \tag{88}$$

$$B(x_R, y) = (1-\beta)v(1-y) + v(1 - x_R). \tag{89}$$

They still satisfy $A + B = C$.

This decomposition removes any concern that two-dimensional control in the baseline is an artifact of $\lambda \neq v$. With $a_y = (1-\beta)v$, the claim Jacobian for reuse restriction and continuity is

$$\det \frac{\partial(A, B)}{\partial(x_R, y)} = a_y v > 0 \qquad \text{when } v > 0, \tag{90}$$

whereas the corresponding Jacobian for diffusion protection and continuity is

$$\det \frac{\partial(A, B)}{\partial(x_D, y)} = -a_y \lambda < 0 \qquad \text{when } \lambda > 0. \tag{91}$$

Thus, at an interior rights bundle that remains in the admissible region, continuity and at least one isolation margin generate locally independent claims whenever either outside-use consequence is economically active. The baseline determinant $(1-\beta)v(v-\lambda)$ arises only because a single x moves the two isolation margins together in opposite-value directions.

B.5. Alternative learning technologies

The conditional hold-up result in Theorem 1 requires proprietary investment to have positive marginal product and the firm to capture less than its full marginal value, together with the theorem's convex-cost, uniqueness, and relevant interiority conditions. Here we record two alternatives to Equation (1).

A Cobb–Douglas technology is

$$\mathcal{L} = d^{\alpha}(e + \bar{e})^{1-\alpha}, \qquad 0 < \alpha < 1, \quad \bar{e} > 0. \tag{92}$$

The seed $\bar{e}$ avoids a zero-investment coordination equilibrium.[28] The firm condition is

$$A\alpha d^{\alpha-1}(e + \bar{e})^{1-\alpha} = c_F'(d),$$

while the planner replaces A with C. For $A > 0$ and finite marginal cost, the infinite marginal

[28] The seeds $\bar{d}$ and $\bar{e}$ below are analytic devices that remove the zero-input coordination corner; they are not measured endowments. Cobb–Douglas and CES forms provide standard alternatives for testing the role of input complementarity (Arrow et al., 1961; Milgrom & Roberts, 1990). Once the substitution structure changes, the continuity response must be rederived; the baseline closed-form signs cannot simply be carried over (Milgrom & Shannon, 1994).

product at $d \downarrow 0$ puts the displayed choice in the positive region; at $A = 0$ the appropriate solution is the zero corner. Conditional underinvestment then follows from the general theorem.

A CES technology is

$$\mathcal{L} = \left[\alpha(d+\bar{d})^{\rho} + (1-\alpha)(e+\bar{e})^{\rho}\right]^{1/\rho} - K_0, \qquad 0<\alpha<1, \quad \bar{d},\bar{e}>0, \quad \rho \le 1, \quad \rho \neq 0, \tag{93}$$

where

$$K_0 = \left[\alpha\bar{d}^{\rho} + (1-\alpha)\bar{e}^{\rho}\right]^{1/\rho}$$

normalizes zero investments. At $\rho = 0$, the corresponding normalized Cobb–Douglas limit is $(d+\bar{d})^{\alpha}(e+\bar{e})^{1-\alpha} - \bar{d}^{\alpha}\bar{e}^{1-\alpha}$. As ρ approaches one, the inputs become closer substitutes. The investment conditions still assign different marginal claims to the two inputs, but the sign of the continuity-induced situated AI performance response must be rederived from the CES marginal products; it does not follow from the baseline closed form.

B.6. Provider fallback value of direct adaptation

The baseline already incorporates $\omega \in [0, v)$ as substitutive salvage: the provider realizes it if focal operation ends, rather than alongside focal use. It therefore enters the provider's direct claim $g_P = (1-\beta)v + \beta\omega$ but not the AI-learning-asset wedge. This is an outside-option distinction in the property-rights tradition (Hart & Moore, 1990); reusable knowledge or capability is different because it can create a productive use beyond the focal relationship (Kang et al., 2009; Sakhartov & Helfat, 2026).

A genuinely nonrival outside use of direct adaptation is different (Jones & Tonetti, 2020). If a unit of e creates an additional outside return $\rho_e > 0$ that coexists with focal operation, survives disagreement, and causes no additional focal diffusion loss, bilateral direct value becomes $v + \rho_e$, the provider's private direct claim becomes $g_P + \rho_e$, and the first-best return also rises by ρ_e. Adding ρ_e only to the provider payoff would therefore misclassify a productive outside use as a redistribution. If the extra return existed only under continued cooperation, the provider would instead capture the Nash share $(1-\beta)\rho_e$. This is the same accounting distinction that separates productive reuse v from a pure transfer of a claim on the AI learning asset.

B.7. A triadic implementation extension

Some deployments involve an external systems integrator that contributes implementation work distinct from both provider adaptation and firm knowledge (Ethiraj et al., 2005; Parmigiani & Mitchell, 2009; Weigelt, 2013). Let the contractor choose $u \ge 0$ and generalize the learning technology to

$$\mathcal{L}(d,e,u) = d(1+\chi e + \eta u), \qquad \eta \ge 0. \tag{94}$$

Let A_F, A_P, and A_C be the parties' marginal claims on the learning asset, with

$$A_F + A_P + A_C = C.$$

Let g_P and g_C be direct returns to provider adaptation and contractor implementation. With quadratic costs and $k_F, k_P, k_M > 0$, an interior equilibrium satisfies

$$k_F d = A_F(1 + \chi e + \eta u), \tag{95}$$

$$k_P e = g_P + \chi A_P d, \tag{96}$$

$$k_M u = g_C + \eta A_C d. \tag{97}$$

These first-order conditions characterize a unique positive equilibrium under, for example,

$$A_F > 0, \quad A_P, A_C \geq 0, \quad g_P, g_C > 0, \qquad k_F > A_F\left(\frac{\chi^2 A_P}{k_P} + \frac{\eta^2 A_C}{k_M}\right). \tag{98}$$

The last inequality makes the spectral radius of the nonnegative response matrix smaller than one; equivalently, substituting the e and u responses leaves a strictly positive coefficient on d. Without it, the displayed equations are only stationary conditions, and corners require the corresponding Kuhn–Tucker inequalities.

The incidence of a larger contractor claim depends on the rights change that creates it. Lower continuity reduces A_F; weaker isolation reduces A_F through diffusion exposure; and more valuable productive reuse raises the provider and bilateral claims while leaving A_F unchanged. A triadic rights rule must likewise specify whether an increase in A_C comes from A_F, from A_P, or from an additional outside use. Holding d fixed, a larger A_C strengthens the contractor's incentive to choose u. Its equilibrium effects on u and proprietary-knowledge investment depend on that claim path.

The baseline d–e system is recovered when the contractor is absent or technologically separable and the remaining claims satisfy $A_F = A$, $A_P = B$, and the provider's direct return is g_P. Internalizing implementation need not reproduce those conditions unless it also removes the contractor's separate claim and payoff interaction.

A complete triadic model would endogenize access, improvement, exclusivity, and continuity rights among all three parties and would ask whether the implementation layer should be internalized.[29] That broader boundary question is distinct from the present paper's buyer-side proprietary-investment mechanism.

[29] The extension isolates the additional claim-allocation problem created by a third complementary investor; see Chatain and Plaksenkova (2026) for the cooperative-investment motivation.

B.8. Downstream competition and the diffusion-loss parameter

The main text represents competitive harm by $\lambda(1-x)\mathcal{L}$. One microfoundation is a differentiated-product market in which provider reuse lowers a rival's marginal cost. Let focal and rival profits after learning be

$$\pi_F = \pi_F(q_F, q_R), \qquad \pi_R = \pi_R(q_F, q_R),$$

with $\pi_{F,q_R} < 0$. Suppose provider reuse raises rival situated AI performance by

$$q_R = q_R^0 + \zeta(1-x)\mathcal{L}, \qquad \zeta > 0.$$

Then, holding q_F fixed and evaluating the derivative at the reference market outcome, the focal firm's first-order competitive loss coefficient is

$$\lambda = -\pi_{F,q_R}\zeta > 0. \tag{99}$$

The reduced-form λ is therefore larger when customer tasks are strategically similar, provider learning generalizes strongly, and rival situated AI performance has a large negative effect on focal profit. Formal models of disclosure and customer-data sharing show that diffusion can reshape downstream competition rather than impose a mechanical loss on the disclosing firm (Choe et al., 2023; Hagiu & Wright, 2023; Pacheco-de-Almeida & Zemsky, 2012). This extension therefore interprets λ as the local net competitive effect relevant to the focal relationship.

B.9. Relationship-finance interpretation

The relationship-banking analogy can also be represented within the general wedge. Relationship-lending research shows how proprietary information accumulated by an incumbent intermediary can improve service while strengthening that intermediary's informational position (Boot, 2000; Rajan, 1992; Sharpe, 1990). Let a customer-specific information stock I increase the incumbent intermediary's monitoring advantage. The borrower's outside value falls by $\ell(I)$, while the informed lender's outside or bargaining value rises by $b(I)$, with $\ell', b' \geq 0$. Holding bilateral marginal value fixed, Equation (67) gives the local wedge response $(1-\beta)\ell'(I) + \beta b'(I) \geq 0$. If the borrower must make effort to generate verifiable records or relationship-specific information, its marginal wedge therefore contains both lost outside value and lender fallback. The AI model differs in two respects. The learning asset directly enters production, and the provider can nonrivally redeploy parts of it across customers. Thus, the finance analogy supplies intuition for informational monopoly but not the full situated-AI production mechanism.

B.10. Governance-cost variants

The main model uses a linear–quadratic $G(x, y)$ so that implementation costs and their interaction are explicit. Governance consumes real design, monitoring, and coordination resources (Argyres & Mayer, 2007; Hoetker & Mellewigt, 2009; Williamson, 1991); the exact linear–quadratic form is a tractability choice. The same problem can be written for any twice continuously differentiable cost:

$$\max_{(x,y)\in[0,1]^2} \{W(x, y) - G(x, y)\} . \tag{100}$$

At an interior solution, $W_x = G_x$ and $W_y = G_y$. At a boundary, the corresponding one-sided conditions apply. Global concavity of $W - G$ is a sufficient condition under which the interior conditions, together with the boundary Kuhn–Tucker conditions, characterize global optima. Without such a global shape restriction, they are only necessary conditions and candidate values must be compared over the feasible set. A more general cost function therefore changes the gradients and curvature to be checked; it does not alter the distinction between isolation's value-changing margin and continuity's claim-reallocation margin.

B.11. A two-period retained-state extension

The one-period model makes the investment distortion visible but does not yet show what repeated use leaves behind. Consider two operating periods. Let $\mathcal{L}_t = d_t(1 + \chi e_t)$ be new customer-specific learning created in period t, let $\sigma \in [0, 1]$ be the retained share of the prior state, and let $\delta \in (0, 1]$ be the discount factor. This flow-to-stock representation follows research on resource accumulation, organizational learning, and experience depreciation (Argote & Miron-Spektor, 2011; Benkard, 2000; Besanko et al., 2010; Dierickx & Cool, 1989).[30] The productive state evolves as

$$S_1 = \mathcal{L}_1, \qquad S_2 = \sigma S_1 + \mathcal{L}_2. \tag{101}$$

Write A_t, B_t, and C_t for the period-t marginal claims on one unit of productive state. When cooperation is certain to continue into period 2, one unit of learning created in period 1 carries the effective claims

$$\widetilde{A}_1 = A_1 + \delta\sigma A_2, \qquad \widetilde{B}_1 = B_1 + \delta\sigma B_2, \qquad \widetilde{C}_1 = C_1 + \delta\sigma C_2. \tag{102}$$

Because only the learning stock is retained in this extension, the period-1 direct returns remain $g_{P,1}$ privately and v_1 bilaterally. The period-1 investment problem otherwise has the baseline form with

[30] The single retention rate σ compresses preservation, forgetting, technical obsolescence, and compatibility into one parameter.

(A, B, C) replaced by $(\widetilde{A}_1, \widetilde{B}_1, \widetilde{C}_1)$. To invoke Proposition 1, maintain

$$\widetilde{A}_1 > 0, \qquad \widetilde{B}_1 \geq 0, \qquad K > \chi^2 \widetilde{C}_1^2. \tag{103}$$

These effective-claim conditions are additional: separate inequalities $K > \chi^2 C_t^2$ in each period do not imply the last inequality because retention adds the two state returns before date-1 investment is chosen.[31] The date-1 firm hold-up wedge is therefore

$$\widetilde{H}_{F,1} = \widetilde{C}_1 - \widetilde{A}_1 = H_{F,1} + \delta\sigma H_{F,2}. \tag{104}$$

Under stationary rights and primitives it equals $(1 + \delta\sigma)H_F$. Retention scales the consequence of an incomplete claim; it does not eliminate that claim problem.

Write the stock inherited at the opening of period 2 as $S_{2,\text{in}} = \sigma \mathcal{L}_1$. Under Equation (103), applying Proposition 1 to the effective period-1 problem gives

$$S_{2,\text{in}}^{FB} - S_{2,\text{in}}^* = \sigma \left[\mathcal{L}_1(d_1^{FB}, e_1^{FB}) - \mathcal{L}_1(d_1^*, e_1^*) \right] \geq 0. \tag{105}$$

The inequality is strict whenever retention is positive and either firm investment is strictly below first best or complementarity makes the larger first-best provider input raise retained learning. Thus underinvestment is not only a lower current flow: it can enter the next period as a smaller installed state.

Now let the relationship break down before period 2 with probability $p \in [0, 1]$.[32] If cooperation survives, the period-2 firm claim is A_2 and bilateral value is C_2. After breakdown, the parties realize disagreement values $F_{0,2}$ and $P_{0,2}$. The full expected effective claims are

$$\begin{aligned}
\widetilde{A}_1^{(p)} &= A_1 + \delta\sigma[(1 - p)A_2 + pF_{0,2}], \\
\widetilde{B}_1^{(p)} &= B_1 + \delta\sigma[(1 - p)B_2 + pP_{0,2}], \\
\widetilde{C}_1^{(p)} &= C_1 + \delta\sigma[(1 - p)C_2 + p(F_{0,2} + P_{0,2})].
\end{aligned} \tag{106}$$

They satisfy $\widetilde{A}_1^{(p)} + \widetilde{B}_1^{(p)} = \widetilde{C}_1^{(p)}$. The expected future wedge between realized bilateral value and the firm's realized claim is

$$(1 - p)(C_2 - A_2) + pP_{0,2}. \tag{107}$$

[31] For example, let $k_F = k_P = 1$, $\chi = 0.75$, $A_t = B_t = 0.5$, $C_t = 1$, and $\delta = \sigma = 1$. Each period separately satisfies $1 > 0.75^2$, but $\widetilde{C}_1 = 2$ and $1 < 0.75^2 \widetilde{C}_1^2$. The effective bilateral quadratic then has a positive-curvature direction and no finite first best.

[32] Breakdown is exogenous and common across the architectures compared here; relational enforcement remains outside the extension (Baker et al., 2002).

Accordingly, the effective date-1 wedge becomes

$$\widetilde{H}_{F,1}^{(p)} = H_{F,1} + \delta\sigma \left[(1-p)H_{F,2} + pP_{0,2}\right]. \tag{108}$$

Any investment or inherited-stock ordering under breakdown additionally requires $\widetilde{A}_1^{(p)} > 0$, $\widetilde{B}_1^{(p)} \geq 0$, and $K > \chi^2[\widetilde{C}_1^{(p)}]^2$; the one-period conditions alone do not guarantee these inequalities. Breakdown does not make the provider's independent outside use disappear; it makes that use the relevant part of the inherited-state wedge.

The same calculation identifies when an accumulated stock changes the value of an architecture. Under a stationary rights bundle j, the expected bilateral return from one inherited unit is

$$R_j = (1-p)C(x_j) + p[F_0(x_j, y_j) + P_0(x_j)] \tag{109}$$

$$= C(x_j) - pv(1-y_j). \tag{110}$$

For a split arrangement s and a bundled arrangement u,

$$\Delta R \equiv R_s - R_u = (\lambda - v)(x_s - x_u) + pv(y_s - y_u). \tag{111}$$

Let $\Delta\phi_0$ denote their stock-independent flow-value difference and let $\kappa \geq 0$ be the incremental cost of implementing or moving to the split arrangement. Such implementation and switching costs are central to comparative governance (Argyres & Mayer, 2007; Williamson, 1991). At inherited stock S, its net advantage is

$$\text{Gain}_{s,u}(S) = \Delta\phi_0 - \kappa + S\Delta R. \tag{112}$$

This is a comparison conditional on an inherited stock measured in common date-value units; it does not yet let the architecture alter the law of motion for S or the breakdown probability. If $\Delta R > 0$ and $\kappa > \Delta\phi_0$, accumulation creates the threshold

$$S > \bar{S} \equiv \frac{\kappa - \Delta\phi_0}{\Delta R}. \tag{113}$$

If $\Delta R > 0$ but $\kappa \leq \Delta\phi_0$, the split arrangement already has higher net value at zero inherited stock. If $\Delta R \leq 0$, stock accumulation alone cannot create an eventual split advantage, although the stock-independent term may still favor it.

Equation (111) also supplies a useful zero result. If $p = 0$ and the two arrangements differ only in continuity, $\Delta R = 0$: once a state is sunk and cooperation is certain, transferring continuity alone does not raise its bilateral return. Continuity can still change ex ante investment by reallocating marginal claims, as Proposition 2 shows. With breakdown risk, however, continuity also preserves realized bilateral value, and an accumulated stock can make that preservation consequential. The extension therefore distinguishes an investment effect from an ex post stock-preservation effect; it does not by itself make ownership or organizational boundaries endogenous.

C. Contract Evidence and Coding

C.1. Purpose

The contract evidence establishes that cross-client reuse, rights over customer-specific adaptations, learned-state continuity, and assurance have distinct observable counterparts. The model treats a written provision as an economic right only when it changes effective authority, a disagreement position, or a permitted use (Aghion & Tirole, 1997; Boudreau, 2010). The evidence is descriptive; its limits are set out below.

The first evidence layer contains 12 public product-and-term bundles for listed enterprise-technology providers. Each bundle combines public legal terms, privacy or trust materials, retention documentation, and technical export documentation for a defined enterprise-AI product family. The unit is the reviewed product bundle, not the provider as a whole. A provider may offer other tiers, regions, subprocessors, models, or negotiated amendments with different rights. The source manifest records one row for each official page, with its title, URL, displayed date where available, access date, retrieval hash where automated retrieval succeeded, and any redirect or interpretation warning.

The second layer contains 13 AI-related agreements filed as SEC exhibits. Candidate agreements were located through targeted searches combining artificial intelligence or machine learning with customer data, training, improvement, source code, escrow, migration, transition, ownership, and exclusivity. We retained agreements with sufficiently specific operative language to identify an enterprise or domain-specific knowledge contribution and at least two learning-loop rights. The transactions span procurement, licensing, joint development, data collaboration, cloud infrastructure, and commercialization. They document variation in feasible contractual arrangements.

C.2. Coding rules and construct mapping

All clause codes use three levels:

- 0: no affirmative provision was located in the reviewed public source;
- 1: a partial, conditional, narrow, product-specific, or platform-internal provision was located;
- 2: an explicit and comparatively broad provision was located.

A zero does not establish legal absence. The relevant term may appear in a private order form, data-processing addendum, technical feature, regional agreement, or unreviewed document. The corpus is purposive and single-coded.[33]

The codebook separates four construct families. Reuse restrictions cover training, service improvement, telemetry, retention, and cross-client diffusion and inform effective isolation x. Adaptation allocation covers foreground intellectual property, customer-specific models, configurations, and derivative state. Continuity covers raw-data exit, learned-state delivery, independent operation, escrow, and transition and informs y. Assurance covers auditability, segregation, monitoring, retention controls, and expected sanctions and informs whether written restrictions become effective control. The released machine-readable codebook gives the full definitions.

C.3. Full descriptive matrices

Figure 5 displays the clause-level matrix. Tables 4 and 5 reproduce the numeric codes. Raw rows, the complete codebook, a source-level manifest, and coding notes are provided in machine-readable form. The bracketed row markers in both tables map directly to the manifest.

[33] The 0/1/2 values are ordinal descriptions of located provisions, not cardinal measures and not direct observations of the continuous model variables x and y. Structured coding of IT-outsourcing contracts motivates the clause-level approach (Chen & Bharadwaj, 2009; Susarla et al., 2010); research distinguishing access, control, and portability motivates keeping the construct families separate (Boudreau, 2010; Vijairaghavan et al., 2025).

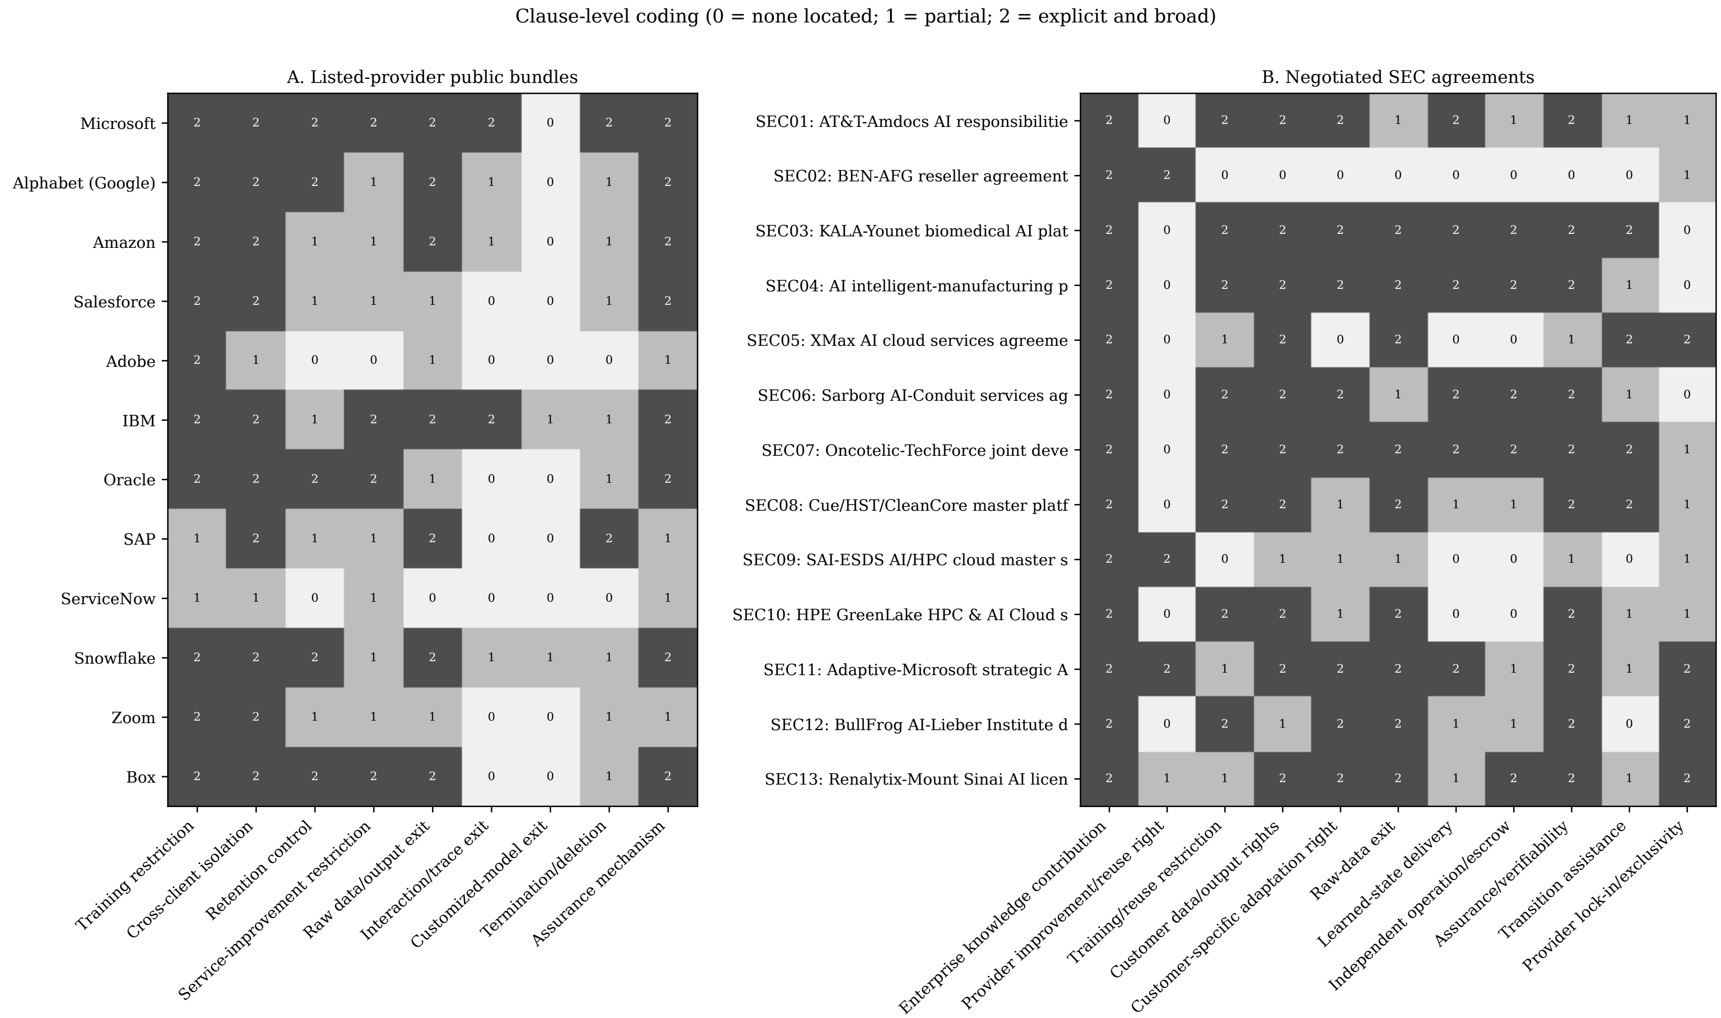

Figure 5. Full clause-level coding. A score of 0 means that no affirmative provision was located in the reviewed source; 1 denotes a partial or conditional provision; and 2 denotes an explicit and comparatively broad provision. The two panels have different dimensions because standardized public bundles and negotiated transactions reveal different portions of the learning-loop architecture.

Table 4. Listed-provider public-bundle codes and source markers

Source/provider	Train	Isol.	Ret.	Svc.	Raw	Trace	Model	Term.	Assur.
`[MSFT]` Microsoft	2	2	2	2	2	2	0	2	2
`[GOOGL]` Alphabet (Google)	2	2	2	1	2	1	0	1	2
`[AMZN]` Amazon	2	2	1	1	2	1	0	1	2
`[CRM]` Salesforce	2	2	1	1	1	0	0	1	2
`[ADBE]` Adobe	2	1	0	0	1	0	0	0	1
`[IBM]` IBM	2	2	1	2	2	2	1	1	2
`[ORCL]` Oracle	2	2	2	2	1	0	0	1	2
`[SAP]` SAP	1	2	1	1	2	0	0	2	1
`[NOW]` ServiceNow	1	1	0	1	0	0	0	0	1
`[SNOW]` Snowflake	2	2	2	1	2	1	1	1	2
`[ZM]` Zoom	2	2	1	1	1	0	0	1	1
`[BOX]` Box	2	2	2	2	2	0	0	1	2

Primary-source note: Train = training restriction; Isol. = cross-client isolation; Ret. = retention control; Svc. = service-improvement restriction; Raw = raw-data/output portability; Trace = interaction-trace portability; Model = customized-model portability; Term. = termination/deletion; Assur. = assurance mechanism. Scores are 0, 1, or 2 under the protocol in Appendix C. Each bracketed marker identifies all official source rows with the same `case_id` in the released source manifest, which reports the full title, date, URL, access date, retrieval hash, and source warning; it is not an author–date citation.

Table 5. Negotiated AI-related SEC-agreement codes and source markers

Source/agreement	Know.	Reuse	Restr.	Result	Adapt.	Raw	State	Oper.	Assur.	Trans.	Lock
`[SEC01]` AT&T-Amdocs AI responsibilities amendment	2	0	2	2	2	1	2	1	2	1	1
`[SEC02]` BEN-AFG reseller agreement	2	2	0	0	0	0	0	0	0	0	1
`[SEC03]` KALA-Younet biomedical AI platform agreement	2	0	2	2	2	2	2	2	2	2	0
`[SEC04]` AI intelligent-manufacturing platform purchase agreement	2	0	2	2	2	2	2	2	2	1	0
`[SEC05]` XMax AI cloud services agreement	2	0	1	2	0	2	0	0	1	2	2
`[SEC06]` Sarborg AI-Conduit services agreement	2	0	2	2	2	1	2	2	2	1	0
`[SEC07]` Oncotelic-TechForce joint development agreement	2	0	2	2	2	2	2	2	2	2	1
`[SEC08]` Cue/HST/CleanCore master platform agreement	2	0	2	2	1	2	1	1	2	2	1
`[SEC09]` SAI-ESDS AI/HPC cloud master services agreement	2	2	0	1	1	1	0	0	1	0	1
`[SEC10]` HPE GreenLake HPC & AI Cloud statement of work	2	0	2	2	1	2	0	0	2	1	1
`[SEC11]` Adaptive-Microsoft strategic AI collaboration	2	2	1	2	2	2	2	1	2	1	2
`[SEC12]` BullFrog AI-Lieber Institute data use and commercialization agreement	2	0	2	1	2	2	1	1	2	0	2
`[SEC13]` Renalytix-Mount Sinai AI license and collaboration agreement	2	1	1	2	2	2	1	2	2	1	2

Primary-source note: Know. = enterprise knowledge contribution; Reuse = provider improvement/reuse right; Restr. = training/reuse restriction; Result = customer data/output rights; Adapt. = customer-specific adaptation right; Raw = raw-data exit; State = learned-state delivery; Oper. = independent operation/escrow; Assur. = assurance/verifiability; Trans. = transition assistance; Lock = provider lock-in/exclusivity. Scores are 0, 1, or 2 under the protocol in Appendix C. Each bracketed marker identifies the official source row with the same `case_id` in the released source manifest, which reports the full title, date, URL, access date, retrieval hash, and source warning; the agreements' full primary-document citations also appear in Section C.4. The marker is not an author–date citation.

C.4. Selected primary-source interpretations

Isolation and customer-controlled learning. The AT&T–Amdocs amendment restricts storage or training on AT&T prompts and data absent written agreement, allocates specified improvements, and requires documentation and repository integration of AI-generated work product.[34] The KALA–Younet agreement combines data sovereignty, client-controlled infrastructure, restrictions on third-party model benefit, customer rights in custom agents and training-data models, audit rights, and an architecture designed for on-premises operation.[35] The intelligent-manufacturing platform agreement allocates specified customer-exclusive models, business logic, configurations, data, and documentation and includes conditional source-code escrow for self-maintenance.[36] The Sarborg AI–Conduit services agreement likewise conditions training or fine-tuning on permission and requires delivery of specified source code, object code, and maintenance documentation.[37] These examples make isolation, adaptation allocation, and continuity separately visible.

Provider-reusable learning. The BEN–AFG agreement requires recurring customer datasets, feedback, validation, and domain expertise and grants the provider rights to use specified customer and reseller data for service improvement and model training.[38] These provisions correspond to the model's provider-reuse margin.

Data exit without complete continuity. The XMax cloud agreement recognizes rights in stored customer data and provides a migration window, while service restrictions limit extraction, replication, or competing-model use.[39] The Cue agreement provides conditional object-code or container escrow and transition assistance but restricts the released material to maintenance of existing projects and preserves limits on modification and reverse engineering.[40] These provisions illustrate why raw-data

---

[34]Amdocs Development Limited, *Amendment No. 11 to Agreement No. 53258.C*, Agreement No. 53258.A.011 (Nov. 17, 2023), filed as SEC Exhibit 4.a(6), SEC EDGAR exhibit (accessed Aug. 1, 2026).

[35]KALA BIO, Inc., *Platform Development and Exclusive License Agreement* with 2624465 Ontario Inc. o/a Younet AI (Mar. 3, 2026), filed as SEC Exhibit 10.1, SEC EDGAR exhibit (accessed Aug. 1, 2026).

[36]CCSC Technology International Holdings Limited, *AI-Enabled Smart Manufacturing Platform Technology Purchase Agreement* with Asia Resource Holdings Limited (Apr. 29, 2026), filed as SEC Exhibit 10.1, SEC EDGAR exhibit (accessed Aug. 1, 2026).

[37]SARBORG Limited and Conduit Pharmaceuticals Inc., *Services Agreement* (Dec. 12, 2024), filed as SEC Exhibit 10.1, SEC EDGAR exhibit (accessed Aug. 1, 2026).

[38]Brand Engagement Network, Inc., *Exclusive Reseller Agreement (SaaS)* with AFG Companies, Inc. (Aug. 19, 2023), filed as SEC Exhibit 10.15, SEC EDGAR exhibit (accessed Aug. 1, 2026).

[39]XMax AI Inc. and SuperX AI Technology USA, *Cloud Services Agreement* (Apr. 22, 2026), filed by Datavault AI Inc. as SEC Exhibit 10.1, SEC EDGAR exhibit (accessed Aug. 1, 2026).

[40]CleanCore Solutions, Inc., *Master Platform Agreement* among HST Technologies, Inc., the customer/SPV, and CleanCore Solutions, Inc. as a limited-joinder party (July 2, 2026), filed as SEC Exhibit 10.3, SEC EDGAR exhibit (accessed Aug. 1, 2026).

delivery, learned-state delivery, and independent operation can be separate bottlenecks.

Foreground rights and category carve-outs. The Oncotelic–TechForce agreement allocates specified AI-related foreground intellectual property, including algorithms, models, training data, inference engines, pipelines, and improvements, and combines that allocation with transition duties.[41] The SAI–ESDS agreement distinguishes customer data from system logs, usage data, telemetry, and operational data, illustrating how customer-data ownership can leave provider-generated learning outside the protected category.[42] The HPE GreenLake statement of work limits specified uses of customer-provided data and supplies certain return rights without establishing complete portability of learned state.[43]

Strategic collaborations. The Adaptive–Microsoft, BullFrog–Lieber, and Renalytix–Mount Sinai agreements partition domain data, know-how, machine-learning effort, foreground outputs, field-of-use rights, exclusivity, commercialization, and surviving rights.[44] These collaborations differ from standardized cloud subscriptions, but they reveal how parties allocate claims when enterprise knowledge and provider adaptation are both central inputs.

C.5. Limitations and confirmatory priorities

Public terms are mutable and may omit negotiated order forms, data-processing addenda, model-provider flow-down, regional differences, and technical feature availability. Retrieval hashes can reveal that a page has changed; they do not preserve the legal text, and two official pages could not be retrieved automatically. Redirects and unstable source identities are flagged rather than silently treated as current support. SEC exhibits are selected by filing materiality and public availability and may overrepresent bespoke collaborations or transactions with unusual continuity provisions. Agreement roles differ, so the corpus cannot support a common treatment comparison.

A confirmatory contract study should archive dated snapshots, use at least two independent coders

[41] Oncotelic Therapeutics, Inc., *Joint Development, Manufacturing, and Licensing Agreement* with TechForce Robotics, Inc. (Mar. 31, 2026), filed as SEC Exhibit 10.1, SEC EDGAR exhibit (accessed Aug. 1, 2026).

[42] Sustainable AI Infrastructure Ltd., *Master Services Agreement* among SAI AU No. 2 Pty Ltd., SpoChub Solutions Private Limited, ESDS Cloud FZ-LLC, and ESDS Software Solutions Limited (Mar. 31, 2026), filed as SEC Exhibit 10.17, SEC EDGAR exhibit (accessed Aug. 1, 2026).

[43] Soluna Holdings, Inc., *HPE GreenLake Services Custom Statement of Work* between Hewlett Packard Enterprise Company and Soluna AL CloudCo, LLC (June 18, 2024), filed as SEC Exhibit 10.105, SEC EDGAR exhibit (accessed Aug. 1, 2026).

[44] Primary agreements: Adaptive Biotechnologies Corporation, *Microsoft Artificial Intelligence & Research Group Strategic Collaboration Agreement* (Dec. 11, 2017), SEC Exhibit 10.2, EDGAR exhibit; BullFrog AI Management, LLC, *Data Use and Technology Partnership Agreement* with Lieber Institute, Inc. (Sept. 7, 2023), SEC Exhibit 10.1, EDGAR exhibit; and Renalytix AI plc, *Exclusive License and Collaboration Agreement* with Icahn School of Medicine at Mount Sinai (May 30, 2018), SEC Exhibit 10.2, EDGAR exhibit. All three were accessed Aug. 1, 2026.

blinded to theoretical predictions, and report Krippendorff's alpha for the ordinal categories (Hayes & Krippendorff, 2007). Interviews with procurement, legal, security, and enterprise-architecture professionals would be needed to distinguish formal rights from operational practice. The released corpus should therefore be read as a transparent institutional map and construct-validation instrument.

D. Architecture and Notation Tables

The main text contains the construct and implication tables needed to follow the argument. This appendix collects the fuller architecture and notation references.

Table 6. Illustrative rights bundles at the four corners

Architecture	(x, y)	**Rights pattern**	**Primary incentive effect**	**Primary limitation**
Hosted H	$(0, 0)$	Reuse permitted; post-exit operation unavailable	Maintains provider claims from reuse and continued focal operation	Exposes the firm to diffusion and continuity hold-up
Protect but trap T	$(1, 0)$	Isolated; post-exit operation unavailable	Raises the firm's claim by removing diffusion exposure	Leaves the firm dependent on incumbent operation
Mobile but reusable M	$(0, 1)$	Reuse permitted; post-exit operation available	Raises the firm's continuity claim and proprietary-knowledge investment	Permits diffusion and reduces the provider's claim from focal operation
Controlled C	$(1, 1)$	Isolated; post-exit operation available	Combines diffusion protection with firm continuity	Can weaken provider adaptation and require costly governance

Notes. The corners are descriptive reference bundles. Proposition 4 compares the best outcomes over the full split feasible set and a pre-specified bundled path.

Table 7. Extended notation for the appendices and evidence materials

Symbol	Name	Meaning
$f(d, e)$	General learning technology	Differentiable learning function used to state the general local comparative statics in Appendix A.
c_F, c_P	General costs	Strictly convex firm and provider investment costs.
D_F, D_P, m, Ω	Local response terms	Own-action curvature terms, technological cross-partial, and determinant of the general private first-order system.
W, Φ	Governance value	Bilateral value before governance cost and net equilibrium value $\Phi = W - G$.
f_x, f_y	Linear governance costs	Linear costs of isolation and continuity in the continuous governance problem.
c_x, c_y, c_{xy}	Governance-cost curvature	Quadratic curvatures and the interaction term in the continuous governance problem.
$\gamma(z)$	Bundled rule	Pre-specified Lipschitz, piecewise-smooth path tying isolation and continuity to one governance choice.
y_D, y_S, y_O	Continuity components	Data/output delivery, learned-state delivery, and independent operation in the weakest-link extension.
x_R, x_D	Isolation components	Productive-reuse restriction and focal diffusion protection in the decomposed-isolation extension.
$\bar{x}, h, a$	Promised isolation, hidden reuse, assurance	Variables in the hidden-reuse extension.
$q, \psi, c_a(h)$	Enforcement primitives	Level, curvature, and provider-borne net cost in the hidden-reuse subgame.
ρ_e	Nonrival adaptation use	Additional outside return from direct adaptation that coexists with focal operation, unlike the provider's disagreement value ω.
u, η, k_M	Contractor extension	Contractor implementation investment, complementarity with proprietary-knowledge investment, and cost curvature.
S_t, σ, δ	Learning-asset stock	AI learning asset stock, interperiod retention rate, and discount factor in the two-period extension.
$\widetilde{A}_1, \widetilde{B}_1, \widetilde{C}_1$	Effective dynamic claims	Date-1 claims including the discounted retained-state return in period 2.
p, R_j	Breakdown and state return	Probability of breakdown before period 2 and expected bilateral return on inherited state under architecture j.
$\Delta\phi_0, \kappa, \Delta R$	Architecture comparison	Stock-independent value difference, incremental implementation cost, and per-unit inherited-state return difference.
$\bar{S}$	Stock threshold	Inherited state above which the higher-return architecture overcomes its net fixed disadvantage when $\Delta R > 0$.

Appendix References

Aghion, P., & Tirole, J. (1997). Formal and real authority in organizations. *Journal of Political Economy*, *105*(1), 1–29. https://doi.org/10.1086/262063

Argote, L., & Miron-Spektor, E. (2011). Organizational learning: From experience to knowledge. *Organization Science*, *22*(5), 1123–1137. https://doi.org/10.1287/orsc.1100.0621

Argyres, N. S., & Mayer, K. J. (2007). Contract design as a firm capability: An integration of learning and transaction cost perspectives. *Academy of Management Review*, *32*(4), 1060–1077. https://doi.org/10.5465/amr.2007.26585739

Arrow, K. J., Chenery, H. B., Minhas, B. S., & Solow, R. M. (1961). Capital–labor substitution and economic efficiency. *Review of Economics and Statistics*, *43*(3), 225–250. https://doi.org/10.2307/1927286

Baker, G., Gibbons, R., & Murphy, K. J. (2002). Relational contracts and the theory of the firm. *Quarterly Journal of Economics*, *117*(1), 39–84. https://doi.org/10.1162/003355302753399445

Benkard, C. L. (2000). Learning and forgetting: The dynamics of aircraft production. *American Economic Review*, *90*(4), 1034–1054. https://doi.org/10.1257/aer.90.4.1034

Besanko, D., Doraszelski, U., Kryukov, Y., & Satterthwaite, M. (2010). Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica*, *78*(2), 453–508. https://doi.org/10.3982/ECTA6994

Binmore, K., Rubinstein, A., & Wolinsky, A. (1986). The Nash bargaining solution in economic modelling. *RAND Journal of Economics*, *17*(2), 176–188. https://doi.org/10.2307/2555382

Boot, A. W. A. (2000). Relationship banking: What do we know? *Journal of Financial Intermediation*, *9*(1), 7–25. https://doi.org/10.1006/jfin.2000.0282

Boudreau, K. J. (2010). Open platform strategies and innovation: Granting access vs. devolving control. *Management Science*, *56*(10), 1849–1872. https://doi.org/10.1287/mnsc.1100.1215

Chatain, O., & Plaksenkova, E. (2026). Incentive alignment in ecosystems: The role of complementarity types and multi-sided bargaining. *Strategic Management Journal*, *47*(1), 228–256. https://doi.org/10.1002/smj.70012

Chen, Y., & Bharadwaj, A. (2009). An empirical analysis of contract structures in IT outsourcing. *Information Systems Research*, *20*(4), 484–506. https://doi.org/10.1287/isre.1070.0166

Choe, C., Cong, J., & Wang, C. (2023). Softening competition through unilateral sharing of customer data [Published online February 13, 2023]. *Management Science*, *70*(1), 526–543. https://doi.org/10.1287/mnsc.2023.4689

Dierickx, I., & Cool, K. (1989). Asset stock accumulation and sustainability of competitive advantage. *Management Science*, *35*(12), 1504–1511. https://doi.org/10.1287/mnsc.35.12.1504

Edlin, A. S., & Reichelstein, S. (1996). Holdups, standard breach remedies, and optimal investment. *American Economic Review*, *86*(3), 478–501.

Ethiraj, S. K., Kale, P., Krishnan, M. S., & Singh, J. V. (2005). Where do capabilities come from and how do they matter? A study in the software services industry. *Strategic Management Journal*, *26*(1), 25–45. https://doi.org/10.1002/smj.433

Grossman, S. J., & Hart, O. D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, *94*(4), 691–719. https://doi.org/10.1086/261404

Grout, P. A. (1984). Investment and wages in the absence of binding contracts: A Nash bargaining approach. *Econometrica*, *52*(2), 449–460. https://doi.org/10.2307/1911498

Hagiu, A., & Wright, J. (2023). Data-enabled learning, network effects, and competitive advantage. *RAND Journal of Economics*, *54*(4), 638–667. https://doi.org/10.1111/1756-2171.12453

Hart, O., & Moore, J. (1990). Property rights and the nature of the firm. *Journal of Political Economy*, *98*(6), 1119–1158. https://doi.org/10.1086/261729

Hayes, A. F., & Krippendorff, K. (2007). Answering the call for a standard reliability measure for coding data. *Communication Methods and Measures*, *1*(1), 77–89. https://doi.org/10.1080/19312450709336664

Hoetker, G., & Mellewigt, T. (2009). Choice and performance of governance mechanisms: Matching alliance governance to asset type. *Strategic Management Journal*, *30*(10), 1025–1044. https://doi.org/10.1002/smj.775

Jones, C. I., & Tonetti, C. (2020). Nonrivalry and the economics of data. *American Economic Review*, *110*(9), 2819–2858. https://doi.org/10.1257/aer.20191330

Kang, M.-P., Mahoney, J. T., & Tan, D. (2009). Why firms make unilateral investments specific to other firms: The case of OEM suppliers. *Strategic Management Journal*, *30*(2), 117–135. https://doi.org/10.1002/smj.730

Kremer, M. (1993). The O-ring theory of economic development. *Quarterly Journal of Economics*, *108*(3), 551–575. https://doi.org/10.2307/2118400

Milgrom, P., & Roberts, J. (1990). The economics of modern manufacturing: Technology, strategy, and organization. *American Economic Review*, *80*(3), 511–528. https://doi.org/10.2307/2006681

Milgrom, P., & Shannon, C. (1994). Monotone comparative statics. *Econometrica*, *62*(1), 157–180. https://doi.org/10.2307/2951479

Nash, J. F. (1950). The bargaining problem. *Econometrica*, *18*(2), 155–162. https://doi.org/10.2307/1907266

Nöldeke, G., & Schmidt, K. M. (1995). Option contracts and renegotiation: A solution to the hold-up problem. *The RAND Journal of Economics*, *26*(2), 163–179. https://doi.org/10.2307/2555911

Oxley, J. E. (1997). Appropriability hazards and governance in strategic alliances: A transaction cost approach. *Journal of Law, Economics, and Organization*, *13*(2), 387–409. https://doi.org/10.1093/oxfordjournals.jleo.a023389

Pacheco-de-Almeida, G., & Zemsky, P. B. (2012). Some like it free: Innovators' strategic use of disclosure to slow down competition. *Strategic Management Journal*, *33*(7), 773–793. https://doi.org/10.1002/smj.971

Parmigiani, A., & Mitchell, W. (2009). Complementarity, capabilities, and the boundaries of the firm: The impact of within-firm and interfirm expertise on concurrent sourcing of complementary components. *Strategic Management Journal*, *30*(10), 1065–1091. https://doi.org/10.1002/smj.769

Poppo, L., & Zenger, T. (2002). Do formal contracts and relational governance function as substitutes or complements? *Strategic Management Journal*, *23*(8), 707–725. https://doi.org/10.1002/smj.249

Rajan, R. G. (1992). Insiders and outsiders: The choice between informed and arm's-length debt. *Journal of Finance*, *47*(4), 1367–1400. https://doi.org/10.1111/j.1540-6261.1992.tb04662.x

Rosen, J. B. (1965). Existence and uniqueness of equilibrium points for concave *N*-person games. *Econometrica*, *33*(3), 520–534. https://doi.org/10.2307/1911749

Sakhartov, A. V., & Helfat, C. E. (2026). Resource redeployment in multi-business firms: Centralized or decentralized? *Strategic Management Journal*, *47*(7), 1878–1906. https://doi.org/10.1002/smj.70077

Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, *28*, 2503–2511.

Sharpe, S. A. (1990). Asymmetric information, bank lending, and implicit contracts: A stylized model of customer relationships. *Journal of Finance*, *45*(4), 1069–1087. https://doi.org/10.1111/j.1540-6261.1990.tb02427.x

Susarla, A., Subramanyam, R., & Karhade, P. (2010). Contractual provisions to mitigate holdup: Evidence from information technology outsourcing. *Information Systems Research*, *21*(1), 37–55. https://doi.org/10.1287/isre.1080.0204

Vijairaghavan, V., Hidaji, H., & Nault, B. R. (2025). Consequences of resorting to fines and investments to regulate data portability [Published online August 26, 2025]. *Management Science*, *72*(5), 3761–3779. https://doi.org/10.1287/mnsc.2024.08404

Weigelt, C. (2013). Leveraging supplier capabilities: The role of locus of capability deployment. *Strategic Management Journal*, *34*(1), 1–21. https://doi.org/10.1002/smj.1998

Williamson, O. E. (1991). Comparative economic organization: The analysis of discrete structural alternatives. *Administrative Science Quarterly*, *36*(2), 269–296. https://doi.org/10.2307/2393356