The Organizational Panopticon

Operational Traces, Higher-Categorical Integration, and Concentration in AI Ecosystems

Anonymous research manuscript

5 November 2026

Prepared for consideration by *Higher Categorical Semantics*

Abstract

AI-mediated work can produce reusable records of its own performance: prompts, tool calls, corrections, evaluations, and histories of interaction. This creates a possible feedback from production to observation, from observation to learning, and from learning to the redesign of production. We develop a categorical and economic account of that feedback while separating informational integration from corporate ownership. An execution category, trace semantics, schema migrations, policy-indexed systems, and coherent interfaces distinguish several senses in which organizations may become integrated. Crucially, a faithful trace functor does not by itself imply effective reconstruction, causal identification, learnability, or a right to reuse the recorded activity. A learning-aware model of organizational scope is therefore conditioned on these additional requirements. Shared infrastructure, distribution, and agent interoperability can favor concentration, but specialization, open interfaces, privacy, organizational costs, and correlated failure can favor federation or plural ecosystems. Dated regulatory findings and primary technical work support individual mechanisms; they do not establish an inevitable transition to a single corporation. The resulting thesis concerns the concentration of effective control over operational information and executable workflows. We give an empirical research program, illustrative enterprise architectures, and extensive notes locating the argument within prior work on firms, surveillance, learning, databases, higher categories, and governance. This is a conceptual and analytical manuscript, without machine-checked proofs or finite model checks.

Keywords: operational traces; organizational learning; AI agents; information integration; higher categories; topos semantics; platform economics; infrastructural concentration.

Status and authorship. Prepared anonymously at the initiating researcher's direction, with OpenAI ChatGPT assistance in drafting, literature discovery, mathematical exposition, and typesetting. This manuscript requires human scholarly review; no institutional affiliation or editorial acceptance is claimed. References were checked on 5 November 2026; access and verification limits appear in the notes and source audit.

Contents

1 The problem: production that records how to produce 4

2 Intellectual genealogy and the limits of the metaphor 6

3 Evidence, provenance, and the status of the industrial claim 9

4 Objects of analysis: workflows, records, capabilities, and control 11

5 Execution categories and trace semantics 14

6 Why faithful observation does not imply reproducible skill 16

7 Learning from production: a constrained recursive model 19

8 Economies of scope with an operational-learning channel 22

9 Feedback, saturation, and concentration dynamics 24

10 Schema integration, data migration, and the meaning of a colimit 27

11 Higher-categorical integration and coherent interfaces 29

12 Topos semantics for context, visibility, and permission 31

13 Ownership, dependence, and effective informational concentration 34

14 An illustrative enterprise: incident response across organizational boundaries 36

15 Contrasting task families and the boundary of the thesis 39

16 Labor, expertise, and the incentives to reveal a method 43

17 Privacy, security, and the governance of derived capability 45

18 Competition, open systems, and systemic risk 48

19 The value of trace information is decision-specific 50

20 Operational knowledge as an asset and a contested option **53**

21 Versioning, provenance, correction, and forgetting **56**

22 Objections, counterarguments, and discriminating replies **58**

23 An empirical program that could support or weaken the thesis **61**

24 Conditional scenarios for the 2030s **65**

25 A governance pattern for jointly produced operational knowledge **67**

26 Conclusion: institutions that learn descriptions of their own operation **70**

A Technical conventions and extended analytical notes **71**

Notes and commentary **78**

References **85**

1 The problem: production that records how to produce

An organization uses knowledge to transform resources, answer questions, repair failures, and coordinate people. Increasingly, it also records those transformations in systems that another computational process can inspect. A debugging session can leave commands, patches, test results, comments, and rejected alternatives. A procurement decision can leave correspondence, revised specifications, approval histories, and exceptions. An interaction with an assistant can leave a request, a sequence of tool calls, an assessment, and a human correction. Such records have possible value both as evidence about completed work and as material for changing how future work is performed.

The central question is therefore not simply whether AI firms possess large datasets. It is whether an organization can make its own productive activity an input to the improvement of its productive machinery. This possibility creates a relationship between the organization as a producer and the organization as an object of observation. The same infrastructure that enables coordinated activity can make that activity legible, reusable, and, under further conditions, partly executable by a learned system. The distinction between a service and an experiment may become institutionally unstable even when it remains analytically indispensable.

We call this arrangement the *organizational panopticon*. The term designates an asymmetry of visibility: a controller can aggregate operational records while the people generating those records may possess only local knowledge of collection, retention, reuse, and consequential decisions. It does not imply perfect observation, universal collection, or a demonstrated policy of using every employee interaction for training. Those are separate empirical questions. The metaphor becomes useful only after these differences are stated. [1]

1.1 Three claims that should not be collapsed

The first claim is technological: instrumented activity can yield data useful for evaluation, retrieval, imitation, preference learning, or process redesign. The second is organizational: combining otherwise separated systems can reduce interface costs and increase the range of workflows an agent can complete. The third is industrial: firms that control useful data, compute, identity, and distribution may acquire durable advantages. These claims can reinforce one another, but none is equivalent to the others.

In particular, they do not jointly entail that the number of legal corporations must converge to one. Information can be integrated by common ownership, contractual partnerships, standards, federated protocols, or a jointly governed platform. Conversely, a single legal corporation may contain information systems that remain incompatible, inaccessible, or organizationally separated. Counting corporations is therefore an inadequate measure of the phenomenon under investigation.

Let N_L denote the number of legal entities and let N_E denote a specified measure of

effective informational independence. These variables need not move together:

$$\Delta N_L = 0 \quad \text{and} \quad \Delta N_E < 0 \tag{1}$$

can describe a partnership that leaves ownership unchanged while increasing common dependence on an identity provider, a model service, or a compute platform. The notation is intentionally conditional: there is no unique natural number of information systems without a definition of boundaries, tasks, and substitutability.

1.2 The thesis in a qualified form

Our thesis is that AI can add a *learning channel* to familiar incentives for integration. An acquisition, partnership, or deployment can give a controller access to operational patterns that improve the performance of future services. This is a mechanism to investigate, not a universal fact about the marginal value of corporate data. It depends on whether records are informative, usable, lawful to reuse, and technically connected to a learning or adaptation pipeline.

The basic dependency is

$$\text{activity} \longrightarrow \text{record} \longrightarrow \text{usable signal} \longrightarrow \text{changed capability}. \tag{2}$$

Each arrow can fail. Activity may occur outside the observed interface. A recorded success may omit the tacit judgment that made it possible. A legally restricted trace may be unavailable for training. A useful dataset may yield no affordable improvement. A better model may remain unable to execute a task because it lacks permission or access to a necessary service. The paper's mathematical distinctions are designed to make those failures visible.

Analytical claim. Operational integration can increase the opportunity to learn from production without guaranteeing either successful automation or a particular ownership structure.

The claim is stronger than an account in which data are an incidental by-product, but weaker than a theory of inevitable technological monopoly. Its explanatory value lies in identifying conditions under which one channel becomes economically consequential.

1.3 What the categorical language contributes

Category theory separates objects from relations and emphasizes the preservation of structure under translation. This is well suited to organizations whose tasks span heterogeneous services. An execution category records composable operations. A trace semantics records what instrumentation preserves or forgets. Schema mappings specify which kinds of records can be translated. Higher cells can record compatibility between alternative translations, while a policy-indexed structure distinguishes an available operation from an authorized one.

These constructions are not proposed as a mathematical replacement for political economy. They are a disciplined vocabulary for asking whether two systems are integrated in the same sense. A schema colimit, an identity federation, and a corporate merger are different constructions. Their interaction can be studied only if they are first kept distinct. Standard categorical foundations are used throughout, with applications developed here rather than attributed to the foundational texts [23, 30].

1.4 Scope, method, and contribution

The paper combines conceptual analysis, mathematical modeling, and a dated reading of selected evidence. Its contributions are a layered account of operational traces; a correction to the inference from faithful observation to reproducible skill; a learning-aware account of organizational scope; and a research program that measures dependence on operational systems rather than relying on legal entity counts. No claim of a comprehensive survey or a newly proved general theorem is made.

The framework applies most directly to work already mediated by software. Physical production, care, maintenance, and social negotiation introduce forms of context that may resist the proposed representations. Even in software organizations, the observed artifact may be the result of collaboration extending far beyond the instrumented interface. The analytical unit is consequently a *specified workflow under a specified observation regime*, not an entire person or an exhaustive theory of expertise. [2]

The following sections move from intellectual background and evidence to formal vocabulary, learning dynamics, economic incentives, and governance. The final research agenda is deliberately falsifiable. An account of concentration that cannot describe successful federation, failed integration, or declining returns to operational data would be a rhetorical prediction rather than an explanatory framework.

2 Intellectual genealogy and the limits of the metaphor

2.1 Visibility, discipline, and accumulation

Foucault's analysis of disciplinary institutions supplies a vocabulary for visibility linked to asymmetrical power [17]. Its relevance here is structural, not an identification of contemporary workplaces with prisons. An employee may know that work is recorded while remaining uncertain about which features are evaluated, how records are combined, and which future decisions will depend on them. Uncertainty about the observer can influence conduct even when observation is incomplete.

The additional issue introduced by learned systems is the reuse of records for capability production. An observer can do more than assess compliance or allocate rewards. It may con-

struct a service that performs some of the recorded operations. Surveillance and automation therefore overlap, but they remain separable: an audit log can support accountability without training, and a model can be trained on voluntarily contributed demonstrations without a disciplinary workplace. Treating them as identical would obscure both the legitimate uses of logging and the distinctive risks of its conversion into an asset.

Zuboff's account of surveillance capitalism connects data extraction to an accumulation process [18]. We extend the question to operational traces within production itself. The object of extraction may be an expert's sequence of corrections or an organization's handling of exceptions, rather than only a customer's externally visible behavior. This extension is our interpretation; it should not be attributed as a completed result of the cited article. [3]

2.2 Legibility is a transformation, not a transparent mirror

Scott's analysis of administrative legibility warns against confusing standardized representations with the activities they organize [21]. A schema selects names, units, categories, and relations. Those selections make some questions answerable and render other distinctions invisible. A system that records time-to-close may distinguish quick and slow cases while overlooking whether the underlying problem was resolved. A classifier trained on that record can reproduce the institutional measurement error with considerable precision.

This gives a productive reversal of the panopticon metaphor. The controller may see a great deal and still misunderstand what it sees. High-resolution recording is compatible with poor semantics. An organization can possess millions of trace events while lacking a stable interpretation of success, responsibility, or harm. The relevant inquiry is therefore not merely how much activity is recorded but what equivalence relation the record imposes on activity.

Let $e_1 \sim_T e_2$ mean that the observation system cannot distinguish two events. Let $e_1 \sim_Q e_2$ mean that they are equivalent for a substantive task criterion Q. Observational adequacy requires attention to the relationship between these two relations. When $\sim_T$ is coarser than $\sim_Q$, consequential distinctions are lost. When it is finer, the system may collect identifying detail irrelevant to the task. More information is not always better alignment.

2.3 Tacit skill and situated action

Polanyi's account of tacit knowing supplies a reason to resist equating explicit descriptions with expertise [20]. The operational consequence is not a proof that some tasks can never be automated. It is a warning that a successful action may depend on competencies that the actor neither articulates nor encodes in the visible record. A demonstration can be an important source of information while still underdetermining the conditions of successful reproduction.

Suchman's analysis of situated action further directs attention to the local circumstances in which instructions become meaningful [22]. An operational trace may preserve an approved

plan while failing to preserve the ongoing adjustments that made the plan workable. For our purposes, this suggests including context and exceptions in the semantic model rather than treating them as noise around a complete specification. [4]

The extension to AI-mediated work is consequently a tension between two transformations. Instrumentation can make previously inaccessible activity easier to represent. At the same time, the organization may redesign activity so that it fits what instrumentation can represent. The second transformation can look like progress in observability even if it narrows the range of practices the institution is willing to recognize.

2.4 Firms, contracts, and complementary assets

The economics of organizational boundaries already provides a substantial account of why firms integrate. Coase frames the boundary in terms of comparative coordination costs [1]. Williamson emphasizes transactional difficulties, including adaptation under contractual incompleteness [2]. Grossman and Hart analyze ownership through residual rights of control and investment incentives [4]. A new vocabulary of AI should not erase these antecedents.

The proposal here is to add a learning return to this comparative analysis. The organization may gain from an interface not only because it reduces the cost of a present transaction but because it supplies observations that change future production possibilities. Conversely, the prospect of that appropriation can change investment and cooperation incentives. The same trace that is valuable to a platform may make an expert less willing to reveal a method or an enterprise less willing to connect a sensitive workflow.

Complementarity theory already explains why organizational practices and technologies may be adopted in bundles [5]. Our use of complementarity is therefore an application rather than an original discovery. The specifically AI-related question is whether improvements in models, usable data, interfaces, and distribution reinforce one another strongly enough to change the efficient or privately preferred boundary of a system.

2.5 Contextual integrity and the normative boundary

Nissenbaum's contextual integrity treats appropriate information flow as dependent on social context, roles, information types, and transmission principles [19]. This provides a more useful starting point than the binary distinction between collected and uncollected data. A record created to debug a service may support one legitimate purpose without automatically supporting employment ranking, commercial training, or cross-customer inference.

The categorical formulation will reflect this by indexing availability on a policy context. An operation that exists in the engineering model may be absent from the authorized execution category. A transfer can preserve a schema while violating a purpose limitation. The formal system must therefore distinguish a well-typed data transformation from an institutionally acceptable one. [5]

Taken together, these traditions caution against two errors. One is technological determinism: assuming that a new information-processing capacity dictates a single industrial outcome. The other is conceptual conservatism: assuming that existing categories already exhaust what changes when productive activity can help train systems that act on the same environment. A useful account must preserve historical continuity while identifying the additional feedback precisely.

3 Evidence, provenance, and the status of the industrial claim

3.1 A hierarchy of evidential support

The empirical thesis should be decomposed before evidence is attached to it. Public material can establish that particular partnerships, interfaces, or learning methods exist. It can sometimes establish their contractual terms or measured performance in a specified setting. It rarely establishes the full internal flow of operational data across an enterprise. Public statements about partnerships should therefore not be transformed into assertions about undisclosed employee-training practices.

We distinguish four evidential categories: institutional findings within a stated jurisdiction and time window; technical demonstrations under specified evaluation conditions; interested-party statements about products or partnerships; and secondary reporting. None is worthless. Their inferential roles differ. A technical demonstration supports a possibility claim more directly than a market-wide prevalence claim. A regulatory decision supports a conclusion about the market investigated more directly than a prediction about all global AI ecosystems.

3.2 Cloud and model partnerships

The FTC's January 2025 staff report examines Microsoft–OpenAI, Amazon–Anthropic, and Google–Anthropic. It describes investment-related rights, compute spending commitments, and access to technical or business information. It identifies potential implications for input access, switching costs, and information advantages [13]. These observations support examining integration across capital, infrastructure, models, and distribution. They do not by themselves demonstrate that a single undertaking controls all participants, or that the partnerships are unlawful. The report explicitly presents a limited staff study rather than a comprehensive legal or economic analysis. [6]

The CMA's final cloud-services decision, dated 31 July 2025, identifies adverse effects on competition in UK cloud services, including concentration and barriers to switching or using multiple providers. It recommends prioritizing consideration of strategic-market-status investigations concerning Microsoft and AWS [14]. A recommendation at that stage is

Evidence	Supports	Does not establish
FTC staff study	Specified partnership structures and possible competitive implications	Universal illegality, complete market definition, or inevitable monopoly
CMA cloud decision	Findings about competition in UK cloud services in the investigated period	A global conclusion about every AI market or later enforcement outcome
Agent-learning research	Feasible connections between interaction records and adaptation in studied settings	Reliable reconstruction of an arbitrary company
Company announcements	Publicly asserted commitments and product relationships	Independently audited welfare gains
Reuters reporting	A dated reported account of a confidential filing	Independent inspection of the underlying filing by this manuscript

Table 1: Evidence is assigned to a claim of matching scope.

not equivalent to a later designation. The distinction is retained here because legal stage, geographic market, and market definition are part of the evidence, not incidental qualifications.

Anthropic's November 2024 AWS announcement is a primary company account of an infrastructure partnership [16]. It illustrates how model development can be linked to specialized compute arrangements. As a company statement, it is relevant evidence of publicly described commitments, with an obvious institutional interest in presenting the arrangement favorably.

Reuters reported on 29 September 2026 that a confidential Anthropic IPO filing described dependence on major technology partners. The report is available in a Reuters syndication [15]. We did not independently inspect the filing and use no precise financial estimate from it in the analytical model. This item is corroborative reporting, not the foundation of the paper's concentration thesis. [7]

3.3 What technical work actually demonstrates

Human-feedback training and preference optimization demonstrate ways in which selected human judgments can influence model behavior [44, 45]. They do not establish that raw logs automatically contain appropriate labels. Tool-using agents and agent-computer interfaces demonstrate that model behavior depends on how the environment is presented and acted upon [51, 52, 54]. The value of an integrated stack may therefore reside partly in the surrounding interface, rather than in the model weights alone.

Agent Lightning explicitly connects agent execution and learning infrastructure in a research framework [55]. Voyager provides a different example of accumulating reusable skills within a particular environment [56]. These are useful existence demonstrations of

components of the proposed feedback. Neither is an empirical demonstration of unrestricted enterprise self-automation. The present argument extrapolates a mechanism while explicitly withholding that stronger conclusion.

SWE-bench illustrates another distinction: a task benchmark can make aspects of software work measurable without exhausting software engineering [53]. Benchmark success, deployment reliability, organizational adoption, and labor substitution are different dependent variables. The manuscript deliberately avoids quoting current leaderboard values, which would add temporal fragility without resolving the main conceptual issues.

3.4 Absence of evidence and the testable residual

The publicly accessible evidence does not determine how much employee activity any named company actually trains on, the marginal value of its internal traces, or the fraction of integration gains attributable to those traces. These are open empirical quantities. A skeptical reader can accept the reported partnerships and technical methods while remaining unconvinced that the learning channel materially changes industrial structure.

That disagreement is productive. The proposed model asks what observations would discriminate it from an account based entirely on scale, subsidies, distribution, and conventional switching costs. If operational traces provide no transferable advantage after these variables are controlled, the stronger learning-based explanation loses support. If federated interfaces capture the same learning benefits at lower social cost, common ownership loses its privileged status in the explanation.

The appropriate evidential conclusion is thus conditional. There are documented structures compatible with informational integration and demonstrated techniques for learning from selected interactions. The claim that their combination systematically favors a small number of controlling ecosystems is a research hypothesis. Its strength depends on measurable conditions developed below.

4 Objects of analysis: workflows, records, capabilities, and control

4.1 Five layers and two boundary maps

A useful representation separates five layers. The *execution layer* contains productive actions and their environments. The *record layer* contains observations, provenance, and schemas. The *learning layer* contains transformations from selected records to models, policies, or retrieval resources. The *service layer* contains executable interfaces and operational permissions. The *control layer* contains authority to set terms, revoke access, change prices, and determine reuse.

These layers interact but should not be identified. A vendor can operate a model while another organization controls its retrieval corpus. An enterprise can hold a complete audit trail without permission to reuse it for model training. A cloud provider can possess a bottleneck in compute without access to the plaintext of a client's encrypted records. A platform can control the distribution interface without owning the model used underneath it.

Let L be a set of legal entities, S a set of service components, and R a set of enforceable control rights. A simplified ownership map is $o : S \to L$. Control is usually better represented by a relation

$$\Gamma \subseteq L \times S \times R, \tag{3}$$

because different entities may hold different rights over one component. There need be no function assigning every service to a unique effective controller. This is one reason why ownership concentration and infrastructural concentration can diverge. [8]

4.2 A workflow includes its environment

Define a task instance by

$$\tau = (X, Y, \mathsf{Env}, Q, \mathsf{Auth}, b), \tag{4}$$

where X is an input specification, Y an outcome specification, Env the relevant environment, Q a quality criterion, Auth an authorization regime, and b a resource budget. A workflow is a procedure for attempting such tasks. The same sequence of commands can succeed in one environment and fail in another. It can also be permissible for one actor and impermissible for another.

This definition prevents an overly narrow equation of automation with producing a plausible textual answer. An automated procurement agent may generate an accurate recommendation but lack authority to place an order. A coding agent may generate a valid patch but lack access to an internal service required for deployment. The capability boundary includes action, observation, and institutional permission.

For an actor a, time t, and tolerance ε, write

$$\mathcal{A}_{a,t}(\varepsilon, b) = \{\tau : \exists \pi \text{ authorized for } a,\ \mathbb{E}[\ell_\tau(\pi)] \le \varepsilon,\ \mathsf{cost}(\pi) \le b\}. \tag{5}$$

The expectation is with respect to a declared task distribution and deployment conditions. A larger model does not automatically enlarge this set if permissions shrink, costs rise, or the task distribution changes. Capability is relational rather than an intrinsic scalar attached to weights.

4.3 A trace is a selected representation

A trace record may contain timestamps, identities, requests, observations, actions, outcomes, and assessments. Each field is produced by an observation mechanism. A timestamp can

identify logging time rather than event time. A user identifier can refer to a delegated service account. A success label can record user acceptance rather than actual correctness. A trace semantics must therefore include an interpretation of fields, not only their serialization.

Represent a trace event as

$$r = (v, c, a, o, u, p, s), \tag{6}$$

where v identifies a schema version, c context, a action, o observed outcome, u assessed utility, p provenance, and s a policy status. Missing or uncertain entries are allowed. The notation does not prescribe a production logging format. It specifies distinctions that a model of reuse should not erase.

W3C provenance and distributed tracing provide established ways to represent parts of this structure [39, 38]. They support causal or operational bookkeeping at the level their specifications address. A trace format alone does not certify the truth of the event, the validity of its label, or permission for secondary use. Those require additional institutional and technical mechanisms.

4.4 Data stock versus usable evidence

The cardinality of a record collection is an especially weak proxy for learning value. Repeated boilerplate, duplicated trajectories, systematic errors, and inaccessible records may add bytes without adding useful distinctions. Conversely, a small collection of well-contextualized exceptions can be valuable for diagnosing a failure mode. We therefore distinguish stored data D_t from selected usable evidence Z_t:

$$Z_t = \mathsf{Select}(D_t; Q_t, \mathsf{rights}_t, \mathsf{schema}_t, \mathsf{budget}_t). \tag{7}$$

The selection operation may be a query, an evaluation protocol, a curation process, or a privacy-preserving transformation. Its output need not be a training dataset. Retrieval, process redesign, and access-policy changes are alternative routes from observations to improved service.

4.5 Independence is defined relative to a task family

Two systems are operationally independent for a task family if either can complete the relevant tasks without requiring the other's continuing discretionary cooperation, apart from declared commodity inputs. This definition must specify a horizon, switching budget, and quality threshold. Otherwise a nominal alternative that requires years of migration can be counted as readily substitutable even when it offers no practical current option.

We will use independence as a graded property. Let $d_{ij}^{\mathcal{T}}$ denote the cost or loss incurred when service j becomes unavailable to actor i for tasks $\mathcal{T}$. This dependence matrix is directed and context-dependent. A common provider may be replaceable for one organization

and indispensable to another. Any global concentration statistic should aggregate such relationships transparently, rather than treating every observed connection as an identical unit of dependence.

These definitions establish the paper's central methodological rule: specify what is integrated, for which tasks, under whose authority, and with what consequences for exit. Without these parameters, a claim about a global information system risks describing a diagram rather than an institution.

5 Execution categories and trace semantics

5.1 The execution category

Fix an organizational domain and a version of its operational interfaces. Let $\mathcal{E}$ be a category whose objects are typed boundary states and whose morphisms are admissible productive processes. A boundary state includes the information required to compose processes: an artifact type, a relevant environment description, and the authority needed for the next step. Composition describes performing one process after another when their boundaries agree.

One can begin syntactically. Objects are workflow types and morphisms are well-typed programs built from declared operations. Quotienting by a chosen congruence identifies processes judged operationally equivalent. The congruence must be compatible with composition. If $f \sim g$ but placing them inside the same larger workflow can distinguish them, then $\sim$ is not an appropriate categorical quotient for that workflow language.

Identity morphisms represent doing no additional productive work at a boundary. This does not require a real implementation to consume zero time. The category abstracts from selected implementation details. A model concerned with elapsed time can enrich morphisms with costs or use a timed process category. The abstraction should be selected for the question, rather than treated as an exhaustive image of an enterprise. [9]

5.2 The trace category and compositional instrumentation

Let $\mathcal{D}$ have observable boundary descriptions as objects and composable trace fragments as morphisms. A trace functor

$$T : \mathcal{E} \longrightarrow \mathcal{D} \tag{8}$$

satisfies $T(g \circ f) = T(g) \circ T(f)$ and preserves identities. This equation says that the record of a composite operation is recoverable by the declared composition of its records. It is an assumption about instrumentation, not a universal property of logging.

For example, traces can carry explicit start and end boundaries, stable identifiers, and a concatenation rule. An operation that mutates hidden state may invalidate the abstraction unless that state is included at the boundary or its effect is represented. A sampler that

drops events may fail strict preservation if its choices depend on global history. A changing schema may require indexed trace categories rather than one fixed $\mathcal{D}$.

The following square expresses preservation of an operational composition:

$$\begin{array}{ccccc} x & \xrightarrow{f} & y & \xrightarrow{g} & z \\ {\scriptstyle T}\downarrow & & \downarrow{\scriptstyle T} & & \downarrow{\scriptstyle T} \\ Tx & \xrightarrow[Tf]{} & Ty & \xrightarrow[Tg]{} & Tz. \end{array} \tag{9}$$

Here the vertical mapsto notation is an explanatory diagram of a functor's action, not additional arrows inside $\mathcal{E}$ or $\mathcal{D}$. The distinction matters when a diagram mixes mathematical levels.

If only partial records are available, a weaker structure may be more appropriate. For records ordered by informativeness, one might specify a comparison

$$T(g) \circ T(f) \preceq T(g \circ f), \tag{10}$$

meaning that observing the whole composite provides at least the information retained by combining its separately observed fragments. A lax interpretation requires coherent comparison maps; writing an inequality alone is not a full construction of a lax functor.

5.3 What faithfulness means

For each pair of objects, a functor induces

$$T_{x,y} : \mathrm{Hom}_{\mathcal{E}}(x, y) \longrightarrow \mathrm{Hom}_{\mathcal{D}}(Tx, Ty). \tag{11}$$

Faithfulness means that each such map is injective. It distinguishes parallel operations in the category already chosen. It does not assert that every possible operation is represented by an object or morphism of $\mathcal{E}$, that traces distinguish different source or target objects, or that every trace in $\mathcal{D}$ is produced by an execution.

If $\bar{\mathcal{E}} = \mathcal{E}/\!\sim_Q$ is an operational quotient, faithfulness of $\bar{T} : \bar{\mathcal{E}} \to \mathcal{D}$ means that no two distinct Q-classes of parallel processes have the same trace. This is the precise categorical reading of the initiating proposal that economically relevant differences should remain visible. It is relative to the quotient and domain.

Fullness is another property: every morphism between Tx and Ty in $\mathcal{D}$ arises from a morphism of $\mathcal{E}$. A fully faithful functor is an equivalence onto its essential image, with the usual categorical qualifications [23]. Even that statement is an abstract relationship between categories, not a computational complexity bound for obtaining an executable process from an observed record. [10]

5.4 Parallel work and organizational structure

Many productive processes run concurrently. A symmetric monoidal category can represent this through $f \otimes g$, provided the chosen semantics accounts for resource sharing and interference. A lax monoidal trace semantics permits records from separate processes to combine into a joint record:

$$Tx \otimes Ty \longrightarrow T(x \otimes y). \tag{12}$$

The comparison need not be invertible. Joint observation can reveal correlations absent from separately stored records, while shared resources can introduce interference that an independent-product model omits.

Open-system approaches using structured or decorated cospans are especially useful when the boundaries of a process are central [27, 28]. An internal workflow can be represented with external input and output interfaces. Composing systems then requires a specified common interface and an ambient category with the necessary colimits. In this paper, the cospan perspective is a design language for interfaces, not a claim that every enterprise database is literally a decorated-cospan category.

5.5 Observation is a design decision

The functorial picture makes a political point precise. Choosing T chooses which operational differences survive recording. An organization can improve fault diagnosis by refining T, but it can also expose unnecessary personal information. It can design traces around worker participation or around managerial convenience. A more faithful trace for one equivalence relation may be less respectful of a deliberately protected distinction between professional performance and private behavior.

Consequently, faithfulness should never be an unconditional governance objective. The relevant design problem is to preserve task-relevant distinctions subject to justified restrictions. The next section shows why even an exceptionally informative trace semantics leaves major gaps between observation and reproducibility.

6 Why faithful observation does not imply reproducible skill

6.1 Six distinct requirements

The inference from detailed records to executable expertise passes through at least six requirements. *Discriminability* concerns whether relevant operations yield distinguishable observations. *Semantic adequacy* concerns whether recorded distinctions match the substantive task. *Reconstruction* concerns whether a procedure can recover an executable description.

Identification concerns whether the observations determine the behavior needed under new conditions. *Learnability* concerns finite resources, sample efficiency, and generalization. *Authorization* concerns whether the resulting action is permissible and operationally enabled.

Faithfulness addresses only the first of these within a specified category. It says nothing about how the functor is represented to the learner, whether an inverse can be computed efficiently, or whether the learner has the actuator needed to realize a recovered description. A perfect transcript of a negotiation can still omit the social authority that made the agreement binding. [11]

6.2 A categorical non-implication

Consider the inclusion of additive monoids $\mathbb{N} \hookrightarrow \mathbb{Z}$, regarded as one-object categories. The associated functor is faithful. Yet there is no functor $R : \mathbb{Z} \to \mathbb{N}$ whose composite with the inclusion is the identity on $\mathbb{N}$. A monoid homomorphism would have to send 1 and -1 to nonnegative elements summing to zero, so it cannot send 1 to 1.

This elementary symbolic example is not a model check of an organization. It isolates a logical point: a faithful embedding need not possess a functorial retraction defined on the entire observation category. An inverse on an appropriately restricted image is a different question. Even where such an inverse exists abstractly, a learning system still needs access to an effective representation and an implementation of its output.

For operational purposes, a reconstruction requirement might instead be stated as the existence of a compiler or controller C satisfying

$$\mathsf{Beh}(C(Tf)) \approx_Q \mathsf{Beh}(f), \tag{13}$$

under a declared environment, resource bound, and authorization regime. This is an extra assumption or an empirical target. It is not a consequence of injectivity of $T_{x,y}$. Moreover, C need not be a functor on all recorded material, since malformed, adversarial, or irrelevant records may have no admissible execution.

6.3 Stochastic observations and statistical experiments

Real instrumentation is often stochastic or incomplete. Let Θ represent relevant latent states and let $O : \Theta \rightsquigarrow Z$ be an observation channel. Distinct latent states may induce distinct distributions over traces while remaining hard to distinguish from a single sample. Injectivity of $\theta \mapsto O(\cdot \mid \theta)$ therefore does not give low sample complexity or robust classification.

Blackwell's comparison of experiments motivates an ordering by post-processing [34]. We use the explicit condition

$$O_2 = K \circ O_1 \tag{14}$$

to mean that O_2 can be simulated by applying a channel K to O_1. Under the usual decision-theoretic assumptions, this captures a sense in which O_1 is at least as informative for

decision problems. The manuscript does not invoke an unrestricted equivalence theorem for arbitrary measurable spaces. Markov-category language supplies a compositional treatment of probabilistic channels and conditional independence [33].

For approximate comparisons, one can specify an error directly:

$$\sup_{\theta} \mathrm{TV}\Big(O_2(\cdot \mid \theta), (K \circ O_1)(\cdot \mid \theta)\Big) \leq \delta. \tag{15}$$

For a fixed decision rule and losses in $[0, 1]$, this bounds the associated expectation discrepancy by δ. The bound concerns the compared distributions and bounded losses; it says nothing about the value of unmodeled objectives or the cost of implementing K.

6.4 Observational equivalence can hide causal differences

Suppose historical work records only cases in which an expert chooses a particular action under favorable conditions. Several environment models can agree on those records while disagreeing about what happens when the action is taken under an unobserved condition. Learning the observed association does not identify the missing intervention. Causal inference distinguishes observational from interventional distributions for exactly this reason [35].

The organizational issue is acute when failures are selectively logged, successful workers receive harder cases, or human approval removes risky proposals before outcomes occur. The record then reflects both a productive process and a selection mechanism. A model may learn that approved actions succeed while failing to learn the expert's reasons for refusing other actions. Those refusals may be among the most valuable parts of the skill.

One should therefore distinguish

$$P(Y \mid A = a, Z = z) \quad \text{from} \quad P(Y \mid \mathrm{do}(A = a), Z = z). \tag{16}$$

The equality of these expressions requires assumptions not supplied by detailed logging alone. Randomization, natural experiments, domain knowledge, or additional causal structure may help. In some domains, ethically acceptable interventions are limited, making uncertainty irreducible within the available observation regime. [12]

6.5 Skill can reside in the interface

The behavior of an agent depends on observations, available tools, latency, permissions, and recovery procedures as well as weights. A trace can be highly informative but generated by an environment the recipient cannot reproduce. The learner may know that a command was effective without possessing its private API, contract, or human escalation route.

This suggests a decomposition of reproduction cost:

$$C_{\text{reproduce}} = C_{\text{decode}} + C_{\text{learn}} + C_{\text{environment}} + C_{\text{authority}} + C_{\text{assurance}}. \tag{17}$$

The terms are an accounting device, not independently identified market estimates. They clarify why an infrastructure owner can retain a durable advantage even when competitors receive a similar corpus. They also clarify how portable interfaces can reduce that advantage without copying the entire organization.

6.6 An appropriately strong conclusion

The defensible claim is that richer, contextually appropriate traces can reduce uncertainty about selected workflows and support capability improvement when a usable learning and execution path exists. It is not that sufficiently faithful observation contains everything needed to reproduce every skill. This qualification strengthens the industrial argument by locating the scarce complementary assets: interpretation, deployment access, evaluation, and authority.

The panopticon is thus potentially powerful without being omniscient. Its economic significance can arise from a modest but systematic informational advantage over a large number of repeated tasks. Perfect reconstruction is neither established by the evidence nor required for that mechanism.

7 Learning from production: a constrained recursive model

7.1 From accumulated traces to usable updates

The initiating model writes an expanding dataset and a sequence of improved models. A more realistic formulation includes deletion, curation, rights, and deployment decisions. Let W_t describe employee workflows, U_t customer interactions, K_t available compute, and P_t the policy regime. Write

$$D_{t+1} = \mathsf{Retain}_{P_t}(D_t) \cup \Phi_{P_t}(W_t, U_t), \tag{18}$$

$$Z_{t+1} = \mathsf{Curate}(D_{t+1}, Q_t, P_t), \tag{19}$$

$$M_{t+1} = G(M_t, Z_{t+1}, K_t), \tag{20}$$

$$\pi_{t+1} = \mathsf{Deploy}(M_{t+1}, S_t, P_t, b_t). \tag{21}$$

Here S_t includes tools and infrastructure. The union is schematic: implementations may use versioned records, weighted examples, or probabilistic summaries. Neither $D_t \subseteq D_{t+1}$ nor a monotone improvement in M_t is assumed.

The deployed policy influences future work and future observations. This closes the feedback loop but also creates an identification problem: improvements in recorded performance may result from a better policy, easier cases, narrower service boundaries, or a changed assessment procedure. A feedback model should not silently identify all four. [13]

7.2 Three channels of improvement

Production traces can influence service through at least three distinct channels. In a *parametric channel*, selected observations change model parameters. In a *retrieval channel*, a fixed model gains access to a changing corpus or skill library. In an *organizational channel*, records change procedures, interfaces, or task allocation. A measured service improvement does not by itself identify which channel was responsible.

Retrieval-augmented generation provides a technical antecedent for separating model parameters from a retrievable knowledge resource [50]. The organizational extension is to treat a runbook, resolved incident, or validated tool sequence as an update to the service environment. Such updates may be easier to inspect or revoke than a distributed parameter change, although revocation can remain difficult if the material has already been copied or summarized elsewhere.

Human-feedback methods and direct preference optimization illustrate selected parametric update mechanisms [44, 45]. They require choices about demonstrations, comparisons, reward criteria, and optimization. A correction in a user interface is not automatically an unambiguous preference label: it can express a local factual repair, a change of goal, an accommodation to a bad interface, or a constraint the learner cannot observe.

7.3 Quality, diversity, and marginal learning value

Let $V(D; \mathcal{T})$ be the value of a record collection for a target task family. Its marginal value need not increase with size:

$$\Delta V(z \mid D) = V(D \cup \{z\}; \mathcal{T}) - V(D; \mathcal{T}) \tag{22}$$

can be positive, negligible, or negative under a particular learning pipeline. A corrupted or misleading record can reduce performance. A redundant record can provide little benefit. A rare exception can matter disproportionately when the task loss assigns substantial weight to catastrophic failure.

Empirical scaling work examines relationships between compute, data, and model size in specified language-model settings [48, 49]. It should not be promoted into a general law that every operational trace improves an enterprise model. For the present thesis, the relevant quantity is a task-specific marginal benefit net of curation, rights, and deployment costs.

We can write a private learning return from admitting dataset D_j as

$$L_j = \Delta\mathsf{ServiceValue}(D_j) - C_{\text{curation},j} - C_{\text{compute},j} - C_{\text{integration},j} - C_{\text{rights},j}. \tag{23}$$

This quantity may be positive for one controller and negative for another because their models, interfaces, and task distributions differ. A dataset has no context-free industrial value.

7.4 Sequential feedback and distribution shift

When an agent's actions determine its future observations, a static supervised-learning picture becomes inadequate. DAgger addresses distributional issues in imitation learning by collecting data under learner-induced states and obtaining expert supervision [46]. The institutional analogue is that an automation system may create cases that the historical workforce rarely encountered. Training only on earlier workflows can leave those states poorly covered.

Performative prediction explicitly studies settings in which deployment changes the distribution being predicted [47]. In organizations, workers adapt to evaluation, customers adapt to service behavior, and the task allocation changes as simple cases are automated. The sequence D_t is consequently endogenous to M_t and to the institution's incentives. A favorable training curve can coexist with a deteriorating understanding of the remaining difficult tasks.

The risk is not only ordinary distribution shift. If automated outputs are recycled without reliable provenance, the apparent diversity of operational knowledge can decline. Work on recursive training with generated data supplies a warning under its studied assumptions [57]. It does not imply that all synthetic data are harmful. The relevant design issue is preserving informative contact with the environment and distinguishing human judgments, model outputs, and verified outcomes. [14]

7.5 An evaluable capability frontier

Instead of assuming that the set of automated tasks grows monotonically, define a deployment frontier by Equation (5). Changes in this frontier can reflect quality, budget, or authorization. A useful longitudinal evaluation therefore holds some dimensions fixed. For example, compare performance on a frozen task distribution under a common budget, while separately measuring performance on newly arising tasks.

An organization should also record the distribution of human intervention. If an agent completes a task only because an expert repeatedly corrects it, the service may still be useful, but the automation claim should include that labor. Intervention is both a cost and a potential source of learning. Erasing it from the evaluation makes the recursive model look more successful than the actual production process.

The learning loop can thus generate a productive capability, an increasingly narrow imitation, or an unstable measurement regime. Its direction is an empirical property of the entire system. The economics of integration depends on which of these regimes actually prevails.

8 Economies of scope with an operational-learning channel

8.1 Use a vector of outputs

An economies-of-scope argument should not add heterogeneous quantities as though they were identical units. Let $q = (q_1, \ldots, q_m)$ describe outputs such as compute service, storage, model inference, identity service, and payment processing. For a partition $\mathcal{P}$ of the output coordinates, let q^B retain the coordinates in block B and set the rest to zero. A scope comparison is

$$C(q) < \sum_{B \in \mathcal{P}} C(q^B), \tag{24}$$

under matched quality, input prices, and feasibility conditions. This is the multiproduct form appropriate to the classical scope literature [3].

The comparison concerns a feasible production arrangement. It does not say that one legal owner is necessary to realize $C(q)$. A shared protocol or joint infrastructure organization may deliver some scope economies. Nor does it compare welfare unless externalities, market power, and the allocation of gains are included. Lower engineering cost can coexist with higher prices to dependent customers. [15]

8.2 A comparative institutional objective

Let a denote an arrangement: common ownership, a contract, a federation, or separated providers. Let $C_a(q)$ be production cost, J_a interface and adaptation cost, L_a the net learning benefit, and R_a expected privately borne operational loss. A private controller evaluates a discounted objective of the form

$$\Pi_a = \mathsf{Revenue}_a - C_a(q) - J_a + L_a - R_a. \tag{25}$$

An additional social objective accounts for consumer surplus, worker effects, privacy, external failure costs, and barriers to entry. There is no reason for the privately maximizing arrangement to maximize the social objective.

The distinctive channel is L_a. Bringing systems into an arrangement may make records jointly interpretable, enlarge evaluation coverage, or reduce the cost of learning a cross-system procedure. But L_a is net of the costs defined in Section 7. Counting every acquired dataset as pure benefit would predetermine the conclusion.

One can decompose an integration surplus as

$$\Delta_a = B_{\text{fixed}} + B_{\text{utilization}} + B_{\text{reuse}} + B_{\text{interface}} + L_a - C_{\text{organization}} - C_{\text{risk}}. \tag{26}$$

The terms are conceptual accounts that can overlap unless empirically identified with care. For instance, model reuse and interface savings can describe the same operational change from different perspectives. An empirical application must prevent double counting.

8.3 Complementarity without a monopoly theorem

Let $F(K, D, M, S, U)$ represent a chosen measure of useful service output. Positive cross-partials such as $F_{KD} > 0$ or $F_{MS} > 0$ express local complementarity under differentiability. They do not imply globally increasing returns, an everywhere profitable merger, or a unique market equilibrium. A bounded production function can have positive cross-partials over a region while exhibiting diminishing returns and substantial coordination costs.

For discrete design choices, increasing differences or supermodularity may be more appropriate than derivatives [5]. The economic interpretation is that the incremental return to one capability can increase when another is present. An improved model can make a reliable tool interface more valuable; a reliable interface can make model improvement easier to deploy. This supports studying bundles, not assuming that every complementary asset must share an owner.

Data-enabled learning and ordinary network effects should also be distinguished [6]. A user can benefit from a service's accumulated learning without directly interacting with other users. The industrial advantage depends on how learning transfers, whether it saturates, and whether rivals can obtain substitutes. Repeating the phrase network effect does not resolve these questions.

8.4 The value of a new workflow corpus

Suppose a platform acquires or connects to organization j. A useful accounting for its operational corpus is

$$\mathsf{Value}_j = \mathsf{RevenueAccess}_j + \mathsf{Assets}_j + \mathsf{Complementarity}_j \tag{27}$$

$$+ \mathsf{TransferableLearning}_j - \mathsf{Migration}_j - \mathsf{Restrictions}_j - \mathsf{Liabilities}_j. \tag{28}$$

Only the fourth term is the additional mechanism emphasized here. It may depend on retaining knowledgeable workers, preserving context, and respecting restrictions that prevent some records from being pooled. An acquisition that disrupts the practices generating high-quality observations can destroy the anticipated learning return.

This suggests a less obvious form of organizational diseconomy. A central controller may reduce duplicative infrastructure while also reducing the diversity of methods that made cross-organizational learning informative. Standardizing every workflow can improve immediate comparability at the expense of exploratory variety. The long-run benefit of integration may therefore be lower than the value inferred from the pre-integration corpus. [16]

8.5 A transparent interior optimum

To illustrate why one ecosystem need not minimize cost, consider a continuous relaxation in the number $n > 0$ of viable ecosystems:

$$C(n) = an + \frac{b}{n} + \frac{c}{n^2}, \qquad a, b, c > 0. \tag{29}$$

The first term represents duplicated fixed capacity; the inverse terms represent losses that intensify with concentration. The derivative is $a - b/n^2 - 2c/n^3$, and the second derivative is positive. Thus the relaxed minimum is uniquely characterized by the zero of the first derivative; the feasible integer choice compares neighboring admissible values. Whether that choice exceeds one depends on the parameters.

This is a sensitivity illustration, not an estimated industry cost curve. It does not assume that coordination always becomes harder under concentration: internal coordination can improve while governance or systemic losses increase. Keeping those components separate avoids the misleading claim that all forms of coordination cost move in the same direction.

The resulting economic thesis is comparative. The learning channel can shift the preferred boundary toward broader integration, but its magnitude must be weighed against organizational, competitive, and systemic costs. It provides a reason for concentration pressure, not a law of corporate convergence.

9 Feedback, saturation, and concentration dynamics

9.1 Two loops and their missing conditions

The proposed industrial mechanism contains a data loop and a capital loop. Service use can produce usable records; records can improve service; improved service can attract use. Separately, useful service can generate revenue; revenue can finance compute; compute can support further capability. Neither loop is automatic. Records may be unavailable, improvement may saturate, customers may multi-home, and revenue may fail to cover investment.

A feedback diagram is therefore a list of possible dependencies, not a dynamical conclusion. To establish self-reinforcement one must specify signs, delays, attenuation, and external shocks. To establish concentration one must additionally specify how a firm's improvement changes its position relative to competitors. Aggregate industry learning can be rapid even when no firm captures a durable advantage.

Let d_t, m_t, u_t be normalized deviations in usable evidence, deployed capability, and service

use around an operating point. A deliberately simple cyclic linearization is

$$\begin{pmatrix} d_{t+1} \\ m_{t+1} \\ u_{t+1} \end{pmatrix} = \begin{pmatrix} 1-\delta_D & 0 & \alpha \\ \beta & 1-\delta_M & 0 \\ 0 & \gamma & 1-\delta_U \end{pmatrix} \begin{pmatrix} d_t \\ m_t \\ u_t \end{pmatrix}, \tag{30}$$

with $0 < \delta_D, \delta_M, \delta_U \leq 1$ and nonnegative cross-effects. The attenuation terms can represent obsolescence, depreciation, or loss of attention. These parameters are not measurements of an actual platform.

In this positive cyclic model the threshold at which the leading eigenvalue passes through one is

$$\alpha\beta\gamma = \delta_D\delta_M\delta_U. \tag{31}$$

The ratio of the two sides is a useful local gain indicator. It concerns a particular linearization. It is not a threshold for global monopoly, and a bounded nonlinear system can stabilize even when its small-state feedback is amplifying.

9.2 Saturation changes the interpretation

Usable learning value can saturate before the number of customers does. Let $h(d)$ be increasing and bounded, and suppose the model update depends on $h(d_t)$ rather than d_t. At large d, a small derivative $h'(d)$ weakens the data-to-capability link. Additional users may then contribute revenue or distribution advantages without contributing much new information.

The reverse can also occur. A narrow service may have few users but encounter rare high-value cases. Their records can be more useful than a much larger repetitive corpus. A concentration theory based on user count alone can consequently misidentify the relevant advantage. Data-enabled learning is task-specific, and its competitive significance depends on transfer and saturation [6].

Compute constraints introduce another nonlinear limit. If a company cannot afford the experiments needed to exploit its records, the usable data stock can grow without a corresponding capability change. Conversely, a provider can buy model capability from elsewhere while contributing little proprietary learning. The loops should be measured separately rather than inferred from their coexistence in a diagram.

9.3 Appropriability determines whether learning concentrates

Let g_i be firm i's capability improvement and let η_{ij} be the share of its useful effect that becomes available to firm j through publication, standards, labor mobility, licensing, or shared models. When off-diagonal spillovers are large, learning can improve the industry without strongly favoring the original collector. When improvements remain exclusive and affect customer choice, the same technological process can increase asymmetry.

A schematic competitive advantage is

$$a_i = g_i - \sum_{j \neq i} \omega_{ij} \eta_{ji} g_j, \tag{32}$$

where the weights describe competitive overlap. This is not a structural estimation equation. It emphasizes that the relevant variable is relative, appropriable improvement. An integrated system that learns quickly can still lose share if rivals learn faster or customers value independence more strongly. [17]

The platform literature also distinguishes interactions among participant groups [7]. An AI service may serve developers, enterprise buyers, workers, and end users, with prices and access terms set differently across them. A subsidy to one group can change data collection and distribution without showing that production costs have fallen. Financially sustained growth should not be mistaken for evidence of an engineering inevitability.

9.4 Compatibility can either disperse or concentrate power

Network-effects research shows why compatibility choices can matter for adoption [8]. For our setting, a shared interface can lower the cost of trying another provider and therefore reduce dependence. The same interface can concentrate power if a single actor controls certification, identity, discovery, or access to the high-value capabilities behind it.

Consider two services that accept the same request schema. Their practical substitutability may still differ because one relies on a proprietary retrieval corpus, a nonportable user history, or a privileged action endpoint. Syntactic compatibility reduces one migration cost while leaving others intact. It can even increase dependence on the shared identity or orchestration layer.

This motivates tracking which layer becomes a bottleneck after an interoperability improvement. A market can become less concentrated in model weights and more concentrated in distribution or identity. A single aggregate concentration index may conceal that relocation.

9.5 Dynamic counterforces

Several processes weaken concentration. Improvements in model efficiency can make smaller deployments viable. Portable records can reduce the incumbency advantage from accumulated history. Openly available models can lower entry costs in selected layers. Enterprise demand for independence can support multi-provider architectures. Organizational complexity can make the incumbent's broad integration slower than a specialist's focused adaptation.

Counterforces can coexist with substantial infrastructural dependence. A specialist may compete vigorously while relying on a common cloud or distribution platform. The right empirical question is which rights remain contestable and which become difficult to substitute. The likely endpoint, if concentration pressure remains strong, may be a small set of overlapping control systems rather than a single corporate tree.

The model therefore predicts conditional pathways, not a calendar. It gives no basis for assigning a numerical probability to one corporation controlling the industrial system by a particular year. The scenarios in Section 24 retain the initiating interest in the 2030s while treating it as a horizon for observation rather than a foregone outcome.

10 Schema integration, data migration, and the meaning of a colimit

10.1 Schemas and instances have different types

The expression $D_{\text{global}} \simeq \operatorname{colim}_i D_i$ is suggestive but incomplete until its ambient category and diagram are specified. A schema, a database instance, a permission policy, and a service interface are not interchangeable objects. Gluing schemas does not by itself resolve the identity of records, preserve access restrictions, or create a globally executable agent.

In a basic categorical database model, a small category $\mathcal{C}$ serves as a schema. Objects describe entity types or attribute domains, arrows describe functional dependencies, and equations constrain paths. An instance is a functor $X : \mathcal{C} \to \mathsf{Set}$ in a chosen universe. This model is an established starting point for functorial data migration [25]; richer algebraic database frameworks handle additional structure [26].

Suppose organizations share an overlap schema $\mathcal{C}_0$ and schema maps

$$\mathcal{C}_1 \longleftarrow \mathcal{C}_0 \longrightarrow \mathcal{C}_2. \tag{33}$$

When the pushout is formed in Cat, the result

$$\mathcal{C} = \mathcal{C}_1 \sqcup_{\mathcal{C}_0} \mathcal{C}_2 \tag{34}$$

integrates declared schema structure. It does not say that all employee records or customer identities have been pooled. The overlap can be empty, restrictive, or based on a deliberately limited interface. [18]

10.2 Compatibility of data requires actual witnesses

Let $X_i : \mathcal{C}_i \to \mathsf{Set}$ be instances. If their restrictions along $\mathcal{C}_0$ agree strictly, the ordinary universal property gives an extending functor on the pushout schema. If they agree only up to an isomorphism, that isomorphism is part of the integration data. One must choose an appropriate bicategorical or homotopical construction, or first make a coherent strict presentation. The symbol $\simeq$ cannot do this work by itself.

In practice, the restrictions may not agree at all. An employee identifier in a source-code system may refer to a service account, while the corresponding human-resources identifier

refers to a person. Matching them by string equality can assign a machine action to a human or combine records belonging to different people. This is not a small defect in the colimit. It is an error in the diagram supplied to the construction.

The relevant compatibility data include entity resolution rules, units, clock conventions, version histories, and uncertainty. A system that cannot justify a match should retain the uncertainty rather than manufacture a global identity. Category theory can organize justified identifications; it does not discover their truth from syntax alone.

10.3 Restriction and the two Kan extensions

A schema functor $F : \mathcal{C} \to \mathcal{C}'$ induces restriction

$$\Delta_F : [\mathcal{C}', \mathsf{Set}] \longrightarrow [\mathcal{C}, \mathsf{Set}], \qquad \Delta_F(Y) = Y \circ F. \tag{35}$$

For small schemas and appropriate completeness conditions on the target, left and right Kan extensions supply adjoints:

$$\Sigma_F \dashv \Delta_F \dashv \Pi_F. \tag{36}$$

These are different data transformations. Informally, the left extension constructs information through colimit-shaped combinations, while the right extension uses limit-shaped compatibility. Their actual behavior is determined by the schema map and the instance; neither is simply a universal command to merge databases.

For example, integrating a ticket system with a deployment system can use a schema map that relates tickets to changes. A left extension may create an integrated presentation with necessary identifications. A query involving simultaneous compatibility constraints may instead require a limit. The choice should follow the semantics of the question. Confusing these constructions can turn missing information into an apparent fact.

10.4 A colimit has a universal property, not a welfare guarantee

The universal property of a colimit organizes compatible maps out of a diagram. It says nothing directly about cost, justice, security, or ownership. If a diagram omits a protected distinction, its colimit will not reintroduce that distinction as an ethical correction. If the maps encode a bad identity match, their formal consistency does not make the match true.

It is nevertheless economically significant that an integrated presentation can reduce repeated translation. A common task interface may allow one agent to query several systems with stable meanings. The saving arises from an implementation of the presentation, an indexing strategy, and an authorization protocol. A categorical existence statement does not supply a scalable query engine or a tractable migration algorithm. [19]

10.5 Federated integration as a first-class alternative

A common schema does not require a common datastore. Each organization can expose a policy-filtered view, retain local records, and execute local queries. A federation can compose authorized answers or exchange aggregates. This arrangement can capture interoperability gains while preserving important boundaries. Its feasibility depends on latency, failures, trust, query expressiveness, and accountability.

Conversely, a single owner may prefer federation within its own organization because business units have different data obligations or failure domains. The technical distinction is therefore not centralized corporation versus distributed corporation. It is the location and revocability of particular information flows and control rights.

The colimit metaphor is strongest when it helps expose what must be agreed: interfaces, identifications, compatibility witnesses, and permitted maps. It is weakest when it is used to imply that all data naturally belong in one total space. The mathematical construction is conditional on a diagram; the institutional question is who gets to determine that diagram.

11 Higher-categorical integration and coherent interfaces

11.1 Why a category of components is sometimes insufficient

A first-order diagram records that a connector translates one schema into another. An organization also needs to know whether two translation routes agree, whether a policy is preserved, and how these agreements behave when another system is added. These are relations between relations. They motivate a two-dimensional or higher-categorical treatment.

Consider systems A, B, C and connectors $F : A \to B$, $G : B \to C$, and $H : A \to C$. A natural isomorphism $\alpha : GF \Rightarrow H$ can record a specified agreement between the indirect and direct translations. It is not merely an assertion that someone has tested both once. In a categorical model, the components of α and their naturality equations are part of the data.

In a strict category one can demand equality $GF = H$. This is often too rigid for versioned systems with different internal representations. In a bicategory, associativity and units can hold through coherent invertible two-cells. Higher-category language is justified when these coherence witnesses are actually needed; it is unnecessary decoration when no such structure is modeled. [20]

11.2 Double categories separate two kinds of movement

Open-system composition naturally distinguishes changes of components from connections between components. A double category can represent systems as objects, schema or

implementation changes as vertical arrows, interfaces as horizontal arrows, and compatibility squares as two-cells. Structured-cospan constructions provide an established mathematical framework of this kind [28, 29].

An illustrative square is

$$\begin{array}{ccc} A & \overset{J}{\rightsquigarrow} & B \\ {\scriptstyle u}\downarrow & & \downarrow{\scriptstyle v} \\ A' & \underset{J'}{\rightsquigarrow} & B'. \end{array} \tag{37}$$

The horizontal arrows denote interfaces, not ordinary functors in the same category as the vertical arrows. A two-cell specifies how the interface is transported by the changes u and v. In an enterprise setting, this can model a version migration that must preserve a contract between two services.

The economics of integration includes the cost of maintaining such squares. Pairwise API compatibility is insufficient if independently updated connectors cease to agree on identity, units, or policy. Common ownership may reduce negotiation cost, but it does not abolish semantic maintenance. A federation may perform the same work through shared governance and versioned standards.

11.3 What higher information can preserve

Let $\mathcal{E}_\infty$ be an $(\infty, 1)$-category of workflows in a domain where higher paths are meaningful operational equivalences. Let $T_\infty : \mathcal{E}_\infty \to \mathcal{D}_\infty$ be a trace semantics. Full faithfulness means equivalences of mapping spaces:

$$\mathrm{Map}_{\mathcal{E}_\infty}(x, y) \simeq \mathrm{Map}_{\mathcal{D}_\infty}(T_\infty x, T_\infty y). \tag{38}$$

This preserves not only classes of processes but the higher structure represented in those spaces. Merely injecting the sets of connected components is much weaker. It can distinguish process classes while losing how equivalent implementations are related.

The foundational setting is standard higher-category theory [31]. Its enterprise interpretation is proposed here. Most real compliance evidence is not invertible, so an $(\infty, 1)$-category alone may be the wrong ambient structure for that evidence. Noninvertible refinements and obligations can call for a bicategory, an $(\infty, 2)$-category, or a separate ordered policy structure. One should not force all institutional relations into homotopies. [21]

11.4 Homotopy colimits do not repair bad semantics

An ordinary colimit can impose strict identifications. A homotopy-coherent colimit retains specified equivalences and their coherence in an appropriate higher setting. This is useful when overlapping representations agree up to a meaningful equivalence rather than literal equality. It does not convert incompatible claims into consistent ones.

For a diagram $D : I \to \mathcal{C}_\infty$, a colimit object L is characterized by

$$\mathrm{Map}(L, X) \simeq \lim_{i \in I^{\mathrm{op}}} \mathrm{Map}(D_i, X) \tag{39}$$

with the induced contravariant diagram understood. The right side includes coherent compatibility data. Existence depends on the ambient category and the diagram. The expression is not a reason to merge records that are legally or semantically incompatible.

An organization might retain both a source-code identity and a finance identity together with a revocable, evidenced relation between them, rather than replacing them by one permanent identifier. A higher or relational representation can preserve that distinction. Its actual benefit depends on a concrete data model and governance arrangement, not on the adjective higher.

11.5 Interoperability has higher-order costs

For n systems, a fully connected undirected interface graph has $\binom{n}{2}$ edges. This is a maximum pair count for that graph, not an unavoidable engineering cost. A hub, a common schema, or a sparse task graph can require fewer connectors. Furthermore, consistency around triples and longer cycles can be costly even when the pairwise interfaces all exist.

Let $G = (V, E)$ describe the interfaces actually needed. A more informative cost is

$$J(G) = \sum_{e \in E} c_e + \sum_{\sigma \in \mathcal{K}} c_{\mathrm{coh}}(\sigma), \tag{40}$$

where $\mathcal{K}$ indexes selected coherence obligations. It might include schema triangles, permission-preservation squares, or version-migration diagrams. The expression is an accounting model; there is no claim that every simplex of an arbitrary nerve imposes an independent cost.

The higher-categorical perspective therefore sharpens the integration argument in two directions. It explains why a shared semantic substrate can be valuable, and it explains why integration remains difficult after ownership boundaries disappear. Coherence is work. A concentrated system can fail at it, and a distributed system can organize it successfully.

12 Topos semantics for context, visibility, and permission

12.1 Context-indexed information

An enterprise does not have one undifferentiated store of available facts. Information is visible to particular roles, for particular purposes, in particular environments and periods. A topos-theoretic perspective can represent this dependence without identifying global availability with local possession.

Let $\mathcal{C}$ be a small category of information contexts. We choose arrows $c' \to c$ to mean that information at c can be restricted to the more limited context c'. A presheaf $X : \mathcal{C}^{\mathrm{op}} \to \mathsf{Set}$ assigns a set of admissible local records to each context and a restriction operation to each arrow. This convention must be fixed explicitly: reversing the interpretation of the arrows reverses the intended information flow.

The presheaf category $\widehat{\mathcal{C}} = [\mathcal{C}^{\mathrm{op}}, \mathsf{Set}]$ is a topos. Standard topos theory supplies finite limits, exponentials, and a subobject classifier [36]. The application does not claim that these abstract structures automatically implement access control. It uses them to formulate a representation in which context dependence is explicit. [22]

12.2 Permission as a subobject

Let A represent possible actions and let $\mathsf{Allowed} \hookrightarrow A$ be a subobject selecting those permitted in each context. A characteristic map $A \to \Omega$ represents the predicate internally. The interpretation is local: an action can be allowed in one context and unavailable in another. The subobject is stable under the restrictions encoded by the model.

That stability is a substantive condition. If a policy becomes more permissive after a purported restriction, the chosen context order may be wrong, or the policy may require additional attributes. Real authorization can depend on time, delegation, purpose, location, and the action's history. Attribute-based access-control guidance gives an established technical account of such distinctions [42].

One should not infer from topos logic that all policy decisions are intrinsically undecidable or subjective. Nor should one impose classical yes/no semantics where the model intentionally represents partial evidence. The internal logic describes the selected semantic universe. Legal and institutional meaning still enters through the interpretation of its objects and predicates.

12.3 Gluing compatible local records

Add a Grothendieck topology J to $\mathcal{C}$ only when a meaningful notion of covering has been supplied. A sheaf then requires uniquely compatible gluing for covering families. This can model local views that jointly determine an admissible global record. Applied sheaf theory supplies tools for reasoning about local consistency and global compatibility [32].

The topology is not free. A family of departmental views should not be called a cover merely because a central controller wants their union. It must have a semantic justification in the modeled information system. Some collections of local records have no authorized global counterpart. Others admit multiple unresolved interpretations. A presheaf model may be more appropriate than a sheaf when unique gluing is not justified.

Sheafification is therefore not an operational command to create missing consent, resolve disputed identities, or eliminate meaningful ambiguity. It constructs a sheaf relative to a chosen topology and ambient semantics. If that changes the interpretation of a record, the

change must be explained institutionally. The mathematical universality of the construction cannot supply the missing justification. [23]

12.4 Purpose-sensitive availability

A useful context can be a tuple $c = (a, p, t, e)$ consisting of actor, purpose, time, and environment. The same underlying record can support a debugging purpose and remain unavailable for employment assessment. Its availability is therefore not a property of its bytes alone. This aligns with the contextual-integrity motivation without reducing that normative theory to a static access matrix [19].

Suppose an operational trace is admitted to a dataset by a map $i : Z \to D$. A model-training transformation should factor through the subobject of permitted examples:

$$Z \longrightarrow D_{\text{train}} \hookrightarrow D. \tag{41}$$

The factorization records a policy requirement. It does not show that every derived artifact can later be removed if permission changes. Training, caching, summarization, and downstream sharing can create additional dependencies requiring provenance and revocation mechanisms.

This distinction is central to the integration thesis. A global identity service may make records technically linkable across contexts while the policy model forbids the linkage. Technical linkability should not be silently converted into permitted integration. A system that respects that separation can deliberately preserve informational plurality within a common infrastructure.

12.5 The higher-topos refinement

When records include meaningful equivalences and coherence witnesses, one can consider space-valued presheaves and appropriate sheaf or descent conditions in an ∞-topos [31]. The value of the refinement is to retain structured evidence of equivalence rather than collapse it to a truth value. For example, alternative identity correspondences might carry evidence, provenance, and compatibility across contexts.

This is a proposed modeling direction. The paper does not assert that a deployed AI company's information architecture is an ∞-topos, or that all its permissions satisfy higher descent. A satisfactory implementation would have to specify the site, the data values, the equivalences, the topology, and the interpretation of descent. Without those ingredients, the terminology would add prestige without explanatory content.

12.6 The institutional consequence

The topos perspective replaces the image of one complete database with a structured family of local views and admissible transitions. Integration can then mean improving the composition

of authorized local information rather than maximizing global exposure. This provides a mathematically motivated alternative to the idea that an optimal agent should see everything.

The central governance choice becomes the structure of contexts and restrictions. Who defines a purpose? Who can change a role? Which local views are allowed to determine a joint record? Who can contest a linkage? These are institutional questions that the model makes explicit. It does not settle them by internal logic alone.

13 Ownership, dependence, and effective informational concentration

13.1 Legal independence and practical dependence

An enterprise can retain separate management, investors, and branding while depending on another actor for compute, identity, model access, or customer distribution. Dependence does not erase its legal identity. It changes its feasible responses to a price increase, access restriction, failure, or change in terms. A useful theory of industrial integration must describe both the legal boundary and the operational dependence.

The property-rights perspective is relevant because ownership concerns control over decisions not exhaustively specified in advance [4]. In digital systems, however, important control can also be exercised through a service contract or technical interface. Revocation of an API, withdrawal of a model endpoint, or change in identity terms may alter another firm's capabilities without a transfer of its shares.

This motivates a multilayer dependency graph. Nodes represent actors and services. Edges record a specified dependence: a compute requirement, an authentication dependency, a data-use right, or a distribution gateway. Edge labels include the task family, expected duration, substitution cost, and holder of the relevant discretionary authority. An unlabeled graph would confuse ordinary exchange with a bottleneck. [24]

13.2 Concentration is a vector before it is a scalar

Let ℓ index layers such as compute, identity, models, orchestration, and distribution. If a defensible allocation of task-weighted activity shares $s_{i\ell}$ is available, with $\sum_i s_{i\ell} = 1$, define

$$H_\ell = \sum_i s_{i\ell}^2, \qquad N_{\mathrm{eff},\ell} = H_\ell^{-1}. \tag{42}$$

These are descriptive concentration measures, not a legal market definition. Their validity depends on the chosen activity denominator and the allocation of shared control. A cloud-revenue share and a share of identity-mediated task executions measure different things.

A vector $(H_{\mathrm{compute}}, H_{\mathrm{identity}}, \ldots)$ is often more informative than a weighted sum. If aggregation is necessary, the weights should correspond to an explicit task population or loss function.

Otherwise a decline in concentration at an easily replaceable layer can mask increasing concentration at a critical bottleneck.

Rights can overlap. If two companies can each veto a workflow, attributing one half of its control to each would obscure the fact that either can stop it. A veto-dependence matrix should therefore be reported separately from normalized activity shares. There is no universal scalar that faithfully summarizes all dimensions of control.

13.3 Exit as an operational capability

Define an exit problem for organization i and provider j by a task family $\mathcal{T}$, quality threshold ε, time horizon h, and switching budget b. Let

$$E_{ij}(h, b, \varepsilon) = 1 \tag{43}$$

when i can replace the relevant services of j under those constraints while meeting the declared threshold. In realistic applications this can be graded by the fraction of tasks preserved. A contractual statement that data can be exported does not establish $E_{ij} = 1$.

Exit can require a portable model configuration, historical context, policy mappings, credentials, evaluation artifacts, and staff expertise. An export containing only raw records may leave the organization unable to reconstruct derived indexes, permissions, or operational assumptions. Switching-cost research provides a classical antecedent for understanding how such dependencies can create ex post market power [9].

The learning channel adds a further asymmetry. The outgoing customer may recover its local records while the provider retains capabilities improved through permitted aggregate learning. Whether this is acceptable depends on the agreement, governance, and nature of the derived information. It also means that data portability alone need not equalize future service capability.

13.4 Bottlenecks and common failure domains

An operational bottleneck can be defined by the disruption of tasks when a service or controlling actor is removed. This is a counterfactual dependence question. It differs from counting outgoing graph edges, since a frequently used service may have cheap substitutes while an infrequently used authorization service is indispensable.

For a declared task distribution, let $v(B)$ be the quality loss when a set of components B becomes unavailable, after allowed adaptations. A control set with large $v(B)$ and small practical substitutability is significant even if no member owns a large share of the legal entities in the ecosystem. Overlapping bottlenecks can produce correlated dependence across nominal competitors.

Common infrastructure does not establish common corporate intention. A shared outage, a common model family, and a coordinated commercial strategy are different hypotheses. The

model can identify correlated exposure without alleging collusion. This distinction is especially important when describing a small number of substrates underlying many competing services. [25]

13.5 The target of the industrial hypothesis

The strongest version of the thesis concerns a reduction in the number of *independently governable operational alternatives*. This is more demanding than observing that systems share a schema and less demanding than asserting common ownership. It asks whether affected actors can change providers, refuse data reuse, preserve critical workflows, and maintain meaningful institutional discretion.

One can observe many firms alongside few compute substrates, few identity systems, concentrated model access, and closely linked data flows. Such a configuration is consistent with the paper's mechanism. It is not its only possible outcome. Open standards and distributed governance can produce extensive informational integration while keeping the control vector substantially more plural.

An adequate empirical account should therefore report legal structure, task dependence, rights, and exit together. Any one of these dimensions can be misleading in isolation. The next section makes the distinctions concrete in an enterprise workflow.

14 An illustrative enterprise: incident response across organizational boundaries

14.1 The task and the relevant systems

Consider a software service that coordinates orders with an external logistics provider. An incident begins when deliveries fail after a deployment. Resolving it can require customer support, telemetry, source code, an order database, deployment history, supplier communication, and a payment adjustment. This is a symbolic architecture example, not a finite model check or a claim about an actual company.

The substantive task is not merely to answer why deliveries failed. It is to diagnose the cause, restore safe operation, communicate accurately, and make authorized adjustments. A candidate agent needs read access to selected logs, authority to propose or apply a patch, a constrained supplier interface, and an escalation path for irreversible or disputed decisions. The correct boundary state includes these rights.

The process can be decomposed into typed stages:

$$\mathsf{Alert} \xrightarrow{d} \mathsf{Diagnosis} \xrightarrow{p} \mathsf{Proposal} \xrightarrow{a} \mathsf{AuthorizedPlan} \xrightarrow{e} \mathsf{Outcome}. \tag{44}$$

This sequential presentation is only one view. Several investigations can proceed in parallel,

and a failed intervention can return the workflow to diagnosis. An execution category should retain those branches rather than treating the successful path as the whole skill.

14.2 What integration contributes

A shared identity and schema can reduce friction when the agent relates an alert to a deployment, a deployment to a code change, and a code change to affected orders. A common orchestration layer can carry provenance and permissions across those steps. Integration can also make it easier to evaluate the complete workflow, rather than separately optimizing response time in support and patch speed in engineering.

Let V_i be the value of each isolated service and let I_{ij} describe a specified cross-service benefit. A heuristic decomposition is

$$V_{\text{joint}} = \sum_i V_i + \sum_{(i,j)\in E} I_{ij} + I_{\text{higher}} - C_{\text{integration}}. \tag{45}$$

The higher-order term allows benefits or conflicts that cannot be decomposed into pairs. Its inclusion prevents the pairwise interface count from being mistaken for a complete model of value. None of these terms is presumed positive.

For example, connecting support and deployment histories may improve diagnosis, while joining employee identities to every customer record may create little task benefit and substantial privacy cost. The architecture should expose the minimum useful relation. Integration is a design choice about sufficient access, not a requirement for unlimited visibility. [26]

14.3 The operational-learning opportunity

A well-contextualized incident trace can contain the initial hypothesis, observations that refuted it, a rejected patch, a human explanation, the authorized intervention, and the measured outcome. Such a trace may be useful for future retrieval, evaluation, or training. A transcript containing only the final successful command can be much less informative because it omits the conditional reasoning and rejected paths.

Process mining is an established approach to extracting process information from event data [37]. The extension proposed here combines that organizational perspective with agent learning. The task is to determine which parts of the recorded process transfer, not to assume that a discovered event pattern is already an executable policy.

The learning target should be stated explicitly. It might be improved classification of incidents, better selection of diagnostic tools, fewer unnecessary changes, or more accurate escalation. Each target requires a different evaluation. A generic label such as resolved is too coarse to distinguish a correct intervention from an incident that disappeared for an unrelated reason.

14.4 Two architectures with different control implications

In an integrated architecture, one controller can operate the identity service, orchestration layer, data platform, and model service. It can reduce negotiation costs and preserve common provenance. It can also combine a broad range of operational information and become difficult to replace.

In a federated architecture, each organization retains local records and authority. The agent receives purpose-limited views, sends signed requests, and composes permitted results. A shared schema and connector contract can provide much of the same task-level integration. The federation incurs additional coordination and failure-handling costs, but may preserve independent governance and local accountability.

The comparison cannot be settled by counting systems. A centralized datastore can contain strict purpose partitions; a federation can rely on one dominant identity or orchestration provider. The relevant measurements are task quality, total cost, information exposure, revocation behavior, and feasible exit. A hybrid arrangement may dominate either idealized endpoint.

14.5 Reliability is compositional only under stated assumptions

Suppose a workflow contains k necessary stages with failure events $F_1, \ldots, F_k$. Without independence, a basic upper bound is

$$P\left(\bigcup_{j=1}^{k} F_j\right) \leq \sum_{j=1}^{k} P(F_j). \tag{46}$$

Under a suitable independence assumption, successful completion probabilities multiply. The latter expression should not be used when stages share a model, identity service, or infrastructure fault. Common components can make failures strongly correlated.

An integrated platform can improve local reliability by reducing inconsistent interfaces while increasing the impact of a common failure. A federation can reduce common exposure while adding network and agreement failures. The architecture problem is therefore to manage the joint distribution of failure, not simply to maximize the quality of individual tools. [27]

14.6 What the example establishes

The example identifies a plausible source of scope economies: reliable composition across data, tools, identity, and evaluation. It also identifies the extra learning channel: a coherent incident history can be reused to improve selected future decisions. Both benefits can exist under several ownership arrangements.

The panopticon concern arises when the same infrastructure gives a controller broad visibility and unilateral reuse authority without equivalent accountability to the participants.

The technical alternative is not necessarily to eliminate records. It is to make collection, interpretation, access, and derived use explicit and contestable while preserving the observations needed for safe operation.

15 Contrasting task families and the boundary of the thesis

15.1 A comparative method without fabricated evidence

The incident-response example gives a setting favorable to the mechanism: several relevant artifacts are digital, their relationships can be represented, and some outcomes can be evaluated through operational tests. It would be a mistake to infer that all productive activity has the same structure. This section contrasts three task families to identify which assumptions do the work. The cases are analytical constructions, not field observations or measured estimates.

For each family, ask four questions. What part of the relevant state is observed? Which actions are available to the learner? What makes an outcome good? How much of the learned procedure transfers to another environment? These questions separate a trace-rich setting from one in which the trace is adequate for the intended decision. They also reveal why a platform's broad reach can be valuable in one domain and insufficient in another.

The comparison should hold the proposed learning channel fixed in view. A system may improve performance by providing a better interface while learning little from local activity. Another may learn a useful diagnostic pattern while remaining unable to execute the required intervention. The industrial argument changes depending on which of these situations obtains.[28]

15.2 Software repair with an incomplete oracle

Consider a task that asks for a repair to a software defect in a maintained repository. Available records include an issue description, code, dependency versions, test failures, proposed patches, review comments, and later incidents. A learner can potentially observe both intermediate actions and some outcomes. This makes the setting attractive for operational learning, as software-agent research illustrates in specified benchmark environments [53, 54].

The important qualification is the oracle. Passing the available tests does not establish that the patch satisfies every requirement, preserves all security properties, or remains maintainable. A trace labeled successful can record only that a particular test suite accepted it. If human reviewers supplied unstated knowledge, the record may not explain why one passing patch was preferred to another.

The environment is also part of the skill. A patch can rely on a private library, an

undocumented deployment assumption, or a service contract. A competitor receiving the patch history may reproduce its syntax without reproducing its operational value. Integration can help by connecting code, tests, deployment, and incident feedback, but the gain depends on the quality of those connections.

This family therefore supports a relatively strong version of the mechanism when the task has a stable environment, informative tests, preserved review context, and a usable execution interface. It supports a weaker version when tests are poor proxies, dependencies change rapidly, or deployment authority remains unavailable. A code trace is unusually amenable to composition, yet it still does not contain an unconditional proof of production correctness.

15.3 Enterprise approval with an authority gap

Consider a request to adjust a supplier order after a service failure. The organization has correspondence, contract terms, inventory records, an approval history, and payment interfaces. An agent can summarize evidence and propose an action. The relevant quality criterion includes not only factual accuracy but consistency with the organization's obligations and the decision-maker's authority.

A historical record can show that a manager approved an adjustment. It may not show whether approval depended on a temporary delegation, a negotiated relationship, or an exception that should not be generalized. The model can imitate the visible decision while misunderstanding its institutional basis. Here the authority gap can be more consequential than the prediction gap.

The appropriate automation boundary may therefore place the agent upstream of a human or institutional authorization. The agent prepares a proposal, assembles evidence, and checks declared constraints. An authorized actor decides whether to proceed. This can be a valuable service without being a complete replacement for the actor. The trace of the proposal can support learning even when the consequential decision remains outside the automated policy.

Integration can still have a substantial return. Linking order records, identities, contract versions, and prior communications can reduce information-search costs. Yet pooling more historical approvals does not necessarily create a transferable decision rule. If authority is highly context-specific, the principal asset may be a well-governed interface that makes the relevant context explicit.

This case also shows why a common identity provider can become strategically important. It may determine which actor's instruction is recognized and which delegation is valid. A model supplier can be replaceable while the authorization layer remains a bottleneck. Measuring model concentration alone would miss the decisive dependency.

15.4 Physical maintenance with an observation gap

Consider maintenance of a piece of equipment whose operator records a diagnosis, a repair action, and a later outcome. The digital record may omit sound, vibration, resistance, environmental conditions, or the operator's embodied inspection. A textual account can be precise about the final action while remaining incomplete about the observation that made the action appropriate.

Additional sensors and structured observations can reduce this gap, but their usefulness depends on placement, calibration, reliability, and interpretation. A camera is not automatically an adequate observation channel for every physical variable. A rich multimodal trace can also be expensive to collect, store, and label. The industrial mechanism therefore depends on the cost of instrumenting the relevant state, not only on model capability.

Action introduces a second gap. A learner can infer a likely repair procedure without possessing a robot, the required tools, access to the site, or the authority to intervene. Simulation can support experimentation while leaving uncertainty about real operating conditions. A transferable description of an action is different from a transferable capacity to perform it.

The case need not defeat operational learning. Records can improve fault triage, parts selection, documentation, or escalation without automating the physical intervention. The appropriate capability set is modular. An organization can gain from learning a narrow diagnostic component while retaining a large role for local human expertise.

For concentration, the crucial question is whether the scarce complement is broadly shareable infrastructure or highly local access and skill. If the latter dominates, a universal trace corpus may have limited cross-domain value. The same mechanism that encourages broad digital integration can coexist with geographically and institutionally distributed maintenance capacity.

15.5 Three kinds of missing complement

The cases can be summarized by the gap between what the trace records and what successful execution requires. Software repair often foregrounds the quality of the evaluation oracle. Enterprise approval foregrounds authority and interpretation of institutional context. Physical maintenance foregrounds observation and action access. These are emphases rather than exclusive classifications: every family can contain all three gaps.

One can express the distinction through a task-specific feasibility condition

$$\mathsf{Feasible}(\tau) = \mathsf{Observe}(\tau) \wedge \mathsf{Interpret}(\tau) \wedge \mathsf{Act}(\tau) \wedge \mathsf{Authorize}(\tau) \wedge \mathsf{Evaluate}(\tau), \qquad (47)$$

where each predicate stands for a declared threshold rather than an absolute property. The conjunction is a way to organize the requirements. It is not a theorem that these five predicates exhaust every condition of safe deployment. [29]

Task family	Potentially useful trace	Complement that can remain missing
Software repair	Patches, tests, reviews, deployment outcomes	Adequate oracle, stable dependencies, production context
Enterprise approval	Evidence, proposals, decisions, delegation records	Valid authority, current obligations, exceptional context
Physical maintenance	Sensor records, diagnoses, actions, outcomes	Relevant physical observation, actuators, local access

Table 2: The same trace-to-capability chain encounters different limiting conditions. The table is analytical, not an empirical ranking of sectors.

15.6 Implications for the value of integration

Integration is valuable when it supplies the complement that is limiting the task. Joining more code histories helps little if the missing element is a deployment approval. A better identity service helps little if the relevant physical state is unobserved. A larger model helps little if the evaluation criterion rewards the wrong outcome. The value of scope depends on which constraint is binding.

This gives a possible explanation for specialization within an integrated industry. A platform can provide broad infrastructure while a specialist supplies the difficult local observation, institutional interpretation, or action environment. The specialist's position is stronger when its complement is hard to codify and substitute. It is weaker when continued interaction makes that complement readily transferable to the platform.

The dynamic relationship is therefore a bargaining problem over evolving capabilities. A platform can help a specialist produce better records, which can improve the service and also change future dependence. A federation can attempt to share the improvement while retaining local authority. The outcome depends on rights, transferability, and the alternatives available at each stage.

The contrasting cases constrain the thesis in a useful way. They identify where a broad operational-learning system should be expected to perform well and where its gains may remain narrow. A theory of concentration becomes more credible when it predicts these limits rather than treating every unrecorded condition as a temporary detail to be eliminated.

16 Labor, expertise, and the incentives to reveal a method

16.1 Workers as producers of present and future capability

An employee who solves a problem can produce two outputs: a present result and a record useful for future automation. The second output is often difficult to separate from ordinary work when collaboration occurs through an instrumented interface. Its existence does not determine who should own it, how it should be compensated, or which uses are appropriate. It does make those questions part of the organization of production.

The labor consequence depends on tasks rather than a simple count of occupations. Automation can displace particular activities, complement others, or create new work. Acemoglu and Restrepo's task-based framework provides a relevant antecedent [11]. The operational-learning channel adds the possibility that an expert's present contribution changes the future allocation of the very tasks through which that expertise is exercised.

This can produce an incentive conflict. A worker may benefit from a useful assistant and from a less repetitive workflow while anticipating that detailed demonstrations will reduce future bargaining power. An enterprise customer can face an analogous concern when its distinctive operating method improves a supplier's general service. These reactions are not irrational resistance to efficiency. They reflect the distribution of future gains and losses.

16.2 A dynamic participation condition

Let w_t be current compensation, a_t the worker's benefit from assistance, c_t the effort cost of producing a useful demonstration, and r_{t+1} the expected future effect on the worker's position. A schematic participation value is

$$U_t = w_t + a_t - c_t + \beta\, \mathbb{E}[r_{t+1} \mid \textsf{reuse rules}]. \tag{48}$$

The term r_{t+1} can be positive or negative. It can include reduced drudgery, new tasks, greater productivity, lost discretion, or displacement risk. The formula is not a wage theory; it identifies a missing term in accounts that value traces only from the collector's perspective.

Clear reuse rules, worker participation, training opportunities, or gain-sharing can alter the participation condition. A requirement to produce high-quality demonstrations without a credible arrangement for future benefits can instead encourage minimal disclosure or strategic adaptation. The quality of the learning corpus is therefore endogenous to institutional trust. [30]

16.3 Measurement changes behavior

If the system rewards a proxy such as speed, closure rate, or apparent acceptance, workers and agents can adapt to that proxy. Multi-task incentive analysis explains why rewarding measured activities can distort effort away from less measurable ones [10]. Operational traces can amplify the problem when the proxy becomes both a performance metric and a training target.

For example, a system that learns from accepted suggestions may favor suggestions that are easy to accept rather than those that solve the hardest problems. A worker may avoid recording uncertainty if uncertainty is interpreted as poor performance. The resulting dataset can become more polished and less informative. Increasing telemetry density can then reduce the honesty of the observed signal.

The same effect can occur in customer interactions. Users learn the phrases that make an assistant comply or the shortcuts that bypass a bad workflow. Those actions express adaptation to the system, not necessarily stable preferences or expert methods. A model trained indiscriminately on them may internalize interface failures as desirable conduct.

16.4 Skill is distributed across a collective

A successful operational trace can depend on documentation maintained by others, a colleague's earlier correction, an inherited codebase, and an institution's accumulated conventions. Attributing the trace to the visible individual can overstate individual authorship while understating collective work. Conversely, treating every record as a corporate asset can erase the contributions of identifiable workers and external communities.

This is another reason to preserve provenance. A record should identify relevant sources and transformations where feasible, while avoiding unnecessary personal exposure. Documentation methods such as datasheets and model cards provide useful antecedents for structured accounts of origin, intended use, and limitations [62, 63]. Their application to operational traces would require additional attention to labor relations and organizational context.

The categorical picture can represent a workflow as a composite assembled from many contributors. It does not determine a unique allocation of economic credit. A composite's value is not generally the sum of the isolated values of its morphisms, and complementary contributions can make marginal attribution unstable. Any compensation rule therefore introduces normative and institutional choices beyond composition itself.

16.5 A possible deskilling feedback

If a system automates routine cases, workers may encounter fewer opportunities to maintain foundational skills while being assigned more unusual failures. Their workload becomes harder to evaluate from historical averages. The organization's ability to recover from automation

failure can then decline even as ordinary-case performance improves.

A second effect is loss of methodological diversity. If everyone is guided through the same assistant interface, fewer alternative approaches may be explored. The learning system receives a narrower distribution and may become less able to respond to unfamiliar conditions. The apparent self-description of the organization can become a description of its own constrained interface.

These are plausible mechanisms, not empirical findings about every AI deployment. They imply measurements: frequency of independent problem solving, recovery performance during model outages, diversity of successful methods, and the quality of escalation. A responsible evaluation would include these quantities alongside productivity gains.

16.6 From extraction to negotiated learning

The operational-learning channel can be organized through explicit negotiation rather than unilateral capture. Participants can distinguish logging required for safety, optional contribution to a shared corpus, organization-specific adaptation, and broader commercial reuse. They can define access to evaluations, correction procedures, and benefit allocation.

Such an arrangement need not eliminate asymmetry, but it changes the institutional meaning of recording. The relevant question becomes whether the people whose activity makes the system useful have credible influence over its reuse. The panopticon thesis is strongest where that influence is weak and the future capability generated from traces is concentrated elsewhere.

17 Privacy, security, and the governance of derived capability

17.1 Separate operational purposes

An organization can require records for reliability, incident investigation, or accounting without thereby justifying every downstream use. Collection, retention, access, training, evaluation, retrieval, and employment assessment should be represented as distinct operations. A permission to perform one is not automatically a permission to perform the others.

The original GDPR text provides a relevant legal reference for purpose limitation, data minimization, and grounds for processing [67]. The present discussion is conceptual and does not determine current obligations for any particular deployment or jurisdiction. Its institutional point is narrower: a design that treats all technically accessible traces as one reusable corpus erases distinctions that governance systems may need to preserve. [31]

17.2 An authorization architecture

Authentication identifies a principal; authorization governs what that principal may do. OAuth and OpenID Connect are foundational specifications for distinct parts of this architecture [40, 41]. They should not be collapsed into a generic notion of identity. An authenticated agent can still lack permission to read a record, combine it with another source, train on it, or execute a consequential action.

The architecture should associate a request with a principal, delegated authority, declared purpose, resource scope, and expiry. It should also preserve the difference between proposing an action and executing it. Classical protection principles, including least privilege and complete mediation, remain relevant [43]. The integration of AI does not remove the need for ordinary security engineering.

Authorization needs to follow derived artifacts. A summary can reveal restricted material even when the raw record is not returned. A retrieval index can retain a deleted source. A learned model can expose information about its training data. Operational governance therefore needs a provenance graph linking records, transformations, and deployed resources, together with realistic limits on deletion and revocation.

17.3 Learning without centralizing all raw records

Federated learning is an established technical alternative in which training data remain distributed while model updates are coordinated [58]. It can reduce raw-data movement and support local control. It does not by itself remove central orchestration, participation bias, communication cost, or leakage from updates.

Secure aggregation can protect individual contributions within an aggregate under its stated protocol assumptions [71]. Differential privacy constrains how much an output distribution can depend on an individual's data under a specified adjacency relation and privacy parameters [59]. These tools address different threats. Neither establishes that an entire organizational system is private or fairly governed.

A representative differential-privacy condition is

$$P[\mathcal{M}(D) \in B] \leq e^{\varepsilon} P[\mathcal{M}(D') \in B] + \delta \tag{49}$$

for adjacent datasets D, D' and measurable output events B. The meaning of adjacency is crucial. Protecting the presence of one event may provide little protection for a worker whose activity contributes thousands of related events. Privacy composition and repeated releases also matter. [32]

17.4 Derived data are not automatically safe

Gradient-sharing attacks show that some update representations can leak private examples under studied conditions [60]. Training-data extraction research similarly demonstrates that learned language models can sometimes reveal memorized material [61]. These results do not imply that every deployed model leaks every record. They show why the assertion that only gradients or only weights are shared is not a complete privacy argument.

The operational setting introduces particularly sensitive combinations: internal source code, customer incidents, employee corrections, and business decisions can appear in one workflow. A controller may infer information from their relationships even when no single field looks sensitive. Data minimization should therefore consider joins and derived inferences, not merely the list of fields retained in isolation.

Security also creates a counterforce to unrestricted integration. Connecting every tool to one broadly authorized agent can enlarge the consequences of a compromised identity, malicious instruction, or incorrect plan. Purpose partitions and independent approvals can reduce the reachable set of harmful actions. The associated friction may be a deliberate safety feature rather than wasteful fragmentation.

17.5 Governance over the learning channel

An effective arrangement can specify which traces enter which channel of improvement, what evaluations must precede deployment, how participants contest misleading labels, and how benefits and risks are reviewed. It can preserve a distinction between organization-specific adaptation and cross-customer model improvement. It can also limit the retention of raw material while preserving auditable aggregate evidence.

The NIST generative-AI profile is a primary institutional reference for organizing risk management [70]. The contribution of this manuscript is to connect such governance to the economics of integration. Privacy restrictions can reduce a collector's private learning return, but they can also improve trust and the quality of contributions. Security separation can increase interface cost while reducing correlated operational loss.

The correct objective is not maximum collection. It is useful, justified, and accountable learning under a declared threat model and institutional arrangement. This changes the integration problem from removing every boundary to selecting which boundaries should remain and how permitted composition should occur.

18 Competition, open systems, and systemic risk

18.1 Integration can be beneficial without being contestable

An integrated agent can save users time, reduce repeated authentication, improve consistency, and make complex services accessible. These are substantive benefits. A critical account loses credibility if it treats every reduction in interface friction as an exercise of power with no productive content. The question is how those benefits are produced, distributed, and made compatible with continued choice.

The private controller can have an incentive to preserve a bottleneck after useful integration has been achieved. It may make some interfaces portable while retaining exclusive access to context, permissions, or distribution. The same system can therefore combine real efficiency gains with an increase in dependence. There is no contradiction: production efficiency and contestability are different dimensions of performance.

The learning channel can deepen this ambiguity. A shared operational corpus can improve service while making it harder for new entrants to match the incumbent's experience. Alternatively, transferable methods and accessible models can distribute the gain. The empirical issue is the degree of exclusivity, not merely the existence of learning.

18.2 Open models challenge a simple ownership narrative

Technical reports for Llama 3, DeepSeek-V3, and Qwen3 document distinct model-development programs and model families [64, 65, 66]. Their relevance here is the existence of alternative sources of model capability and differing deployment possibilities. The paper does not infer equivalent openness, licenses, safety properties, or performance from their being discussed together.

Availability of model weights can lower dependence on a particular inference endpoint for some tasks. It does not eliminate the need for compute, operational expertise, data rights, distribution, or a maintained tool environment. A market can become more plural at the model layer while remaining concentrated at the infrastructure layer. Conversely, local deployment can preserve important independence even when the model was developed by a large organization. [33]

This counterexample matters for the industrial thesis. If capability can be acquired through multiple channels, an incumbent's operational dataset may be a narrower advantage than its marketing suggests. A small firm might combine a widely available model with a specialized corpus and a superior interface. The tendency toward integration can then coexist with entry and specialization.

18.3 Interoperability as a remedy and as a new bottleneck

The Digital Markets Act provides a legislative reference for certain interoperability and access obligations in its defined settings [68]. The Data Act separately addresses switching between data-processing services [69]. These texts illustrate institutional efforts to make boundaries more contestable. They do not establish that any particular AI system is covered by every obligation or that interoperability can be imposed without technical trade-offs.

An effective portability arrangement should ask what must move for a task to remain feasible: records, embeddings, schemas, policy mappings, model configuration, credentials, and evaluation evidence. A file export can satisfy a narrow technical requirement while leaving the operational dependency intact. The framework's exit criterion offers a way to evaluate practical consequences without equating every export with independence.

A shared standard can also become a governance bottleneck if one actor controls its evolution or certification. A neutral schema is not institutionally neutral merely because several firms use it. Decisions about version changes, extensions, and conformance can favor particular implementations. The governance of the interface is part of the competitive structure. [34]

18.4 Polycentric governance is a serious alternative

Ostrom's work supplies a substantive antecedent for governance through multiple centers rather than an exclusive choice between a central state and a private market [12]. In this setting, a consortium, cooperative, public-interest infrastructure provider, or sectoral federation could organize shared interfaces while preserving local authority. This is an institutional possibility, not a guarantee of success.

Polycentric arrangements can suffer from slow decisions, uneven participation, underinvestment, and capture. A dominant member can exercise substantial power without formal ownership. Nonetheless, they provide a countermodel to the claim that informational integration requires corporate unification. The relevant comparison is their actual cost, accountability, adaptability, and ability to sustain alternatives.

One practical design is to distribute different rights: a common schema governed by participants, local custody of raw records, portable identity assertions, interchangeable model providers, and shared evaluations. Such separation can preserve a useful task interface while preventing one controller from accumulating every complementary asset. Its feasibility depends on the task and on the quality of coordination.

18.5 Concentration changes the distribution of failure

Suppose losses from several services are random variables $L_1, \ldots, L_m$. Then

$$\operatorname{Var}\left(\sum_i L_i\right) = \sum_i \operatorname{Var}(L_i) + 2\sum_{i<j} \operatorname{Cov}(L_i, L_j). \tag{50}$$

Shared identity, models, infrastructure, or update procedures can increase common exposure. The formula does not prove that integration always raises covariance; integration can also improve protection and recovery. It shows why evaluating each service's average reliability separately misses an important dimension of systemic risk.

Concentration can also increase the impact of semantic failure. If many organizations rely on similar model-generated interpretations of an ambiguous policy or event, an error can propagate through otherwise independent institutions. Diverse models, local expertise, and independent decision procedures may provide resilience even when their individual average accuracy is lower. [35]

18.6 The social integration frontier

The socially preferred arrangement balances useful interoperability against privacy, market power, fragility, and the preservation of independent judgment. It can involve extensive technical integration with strong institutional separation. It can also involve deliberately duplicated infrastructure where independence has substantial option value.

The economic pressure toward concentration is therefore real only as a comparison among feasible arrangements under particular conditions. It is not an entitlement to remove boundaries. The strongest practical implication is to study which interfaces can be shared without concentrating every consequential control right in the same actor.

19 The value of trace information is decision-specific

19.1 From distinguishability to a decision problem

An information system can distinguish many events without making an economically relevant decision easier. To connect traces to productive value, specify a latent condition θ, an observation channel O, an action space A, and a loss $\ell(\theta, a)$. With prior μ and a permitted class of decision rules Δ, define

$$\mathcal{R}(O; \mu, \ell, \Delta) = \inf_{\delta \in \Delta} \mathbb{E}_{\theta \sim \mu,\, z \sim O(\cdot|\theta)}[\ell(\theta, \delta(z))]. \tag{51}$$

The value of an observation is a comparison with an appropriate baseline, such as an observation that carries no task information. A loss function for diagnosing an incident will value different distinctions from one for assigning responsibility or forecasting demand.

This formalization also clarifies a limitation of the fidelity objective. A record can retain an employee's typing rhythm, personal style, and incidental browsing without improving the declared task. Such details may increase statistical identifiability while contributing little decision value. Conversely, a coarse record of the precise error condition can be highly useful. The task, not raw detail, determines the relevant comparison. [36]

19.2 The free-information qualification

If an additional observation is free and the decision rule can always ignore it, the best achievable risk cannot increase merely because that observation is available. This elementary observation should constrain loose claims that more data intrinsically make rational decisions worse. The negative marginal values discussed earlier concern a specified learning pipeline, costs, constraints, or a changed environment, not an unconstrained optimum with free access and the option to ignore information.

In practice the qualifications are decisive. Processing more records consumes compute and attention. A fixed training procedure may overweight redundant examples. An interface may force a worker to respond to noisy alerts. A privacy requirement may make certain joins impermissible. A record can also change behavior through monitoring. These mechanisms can make an expanded collection privately or socially harmful without contradicting the free-information comparison.

The economic model should therefore write net value as

$$\mathsf{NetVol}(O) = \mathcal{R}(O_0) - \mathcal{R}(O) - C_{\text{acquire}} - C_{\text{interpret}} - C_{\text{use}} - C_{\text{external}}, \tag{52}$$

with compatible units or a clearly stated normalization. Some external costs may be governed by constraints rather than a monetary term. The expression is an accounting framework, not a claim that every cost has an agreed price.

19.3 Task relevance and statistical information

Measures such as entropy or mutual information quantify statistical relationships under a specified distribution. A high information quantity about an identity variable does not imply high value for a maintenance task. A record may be rich in information about who performed an action and poor in information about why it worked. The variable whose uncertainty is reduced matters.

The post-processing order introduced in Section 6 is useful because it compares channels through explicit simulation. But practical decisions often require a narrower comparison: whether one representation supports a particular task within an acceptable loss. Two channels can be incomparable in a general information order while one is clearly preferable for the declared workflow.

This suggests a task-indexed preorder $O_1 \succeq_{\mathcal{T}} O_2$ when the achievable risks for a specified family of task losses are no worse under O_1. The definition must state the allowed policies, resource constraints, and distributions. It need not coincide with a universal comparison over all decision problems. [37]

19.4 Cross-system value can be synergistic

Two records can each have low value in isolation and substantial value jointly. A deployment identifier without an incident history may reveal little about a failure; the incident history without the deployment link may likewise remain ambiguous. Their join can identify a useful relation. This provides a decision-theoretic interpretation of interoperability benefits.

Let $V(Z)$ be the improvement in a declared decision objective using observation Z. Define a joint surplus

$$I(Z_1, Z_2) = V(Z_1, Z_2) - V(Z_1) - V(Z_2) + V(\varnothing). \tag{53}$$

The surplus can be positive, zero, or negative under net-value accounting. It should not be presumed positive for every pair of datasets. Redundancy can make separate values overlap, while joint processing can introduce costs or expose sensitive relationships.

The same mathematical possibility creates a governance problem. A pair of individually innocuous views can reveal protected information when combined. The productive and privacy consequences of a join are linked by the information it creates. The right design may permit the useful decision while restricting the general joined dataset or providing a limited answer service.

19.5 Robustness across environments

A trace corpus can support a decision under its historical distribution and fail under a changed one. A robust evaluation can replace a single distribution with a declared uncertainty set $\mathcal{Q}$:

$$\mathcal{R}_{\mathrm{rob}}(O) = \inf_{\delta \in \Delta} \sup_{Q \in \mathcal{Q}} \mathbb{E}_Q[\ell(\theta, \delta(Z))]. \tag{54}$$

The usefulness of this objective depends on whether $\mathcal{Q}$ captures plausible deployment changes. A mathematically convenient uncertainty set can omit the rare organizational transition that matters most. Robust notation is not a substitute for understanding the environment.

This perspective gives a possible benefit to diversity. Different organizations can generate different operational contexts, broadening coverage. Yet indiscriminate pooling can also mix incompatible labels and procedures. The net learning gain depends on whether the corpus preserves contextual distinctions or erases them in pursuit of scale. [38]

19.6 Implications for concentration

The decision model identifies when trace accumulation can be an industrial advantage: the records must improve important decisions, the improvement must be usable at acceptable cost, and the collector must capture enough of the resulting benefit. Volume alone supplies none of these conditions. A task-specific specialist can outperform a broad platform when it has better labels, a better environment, or a more appropriate action interface.

Conversely, a platform can gain from joining records that specialists cannot legally or technically combine. Its advantage can then reside in the relationship between datasets rather than in the individual datasets. This is an additional reason to examine identity and permission layers. They determine which joint observations are feasible and under whose control.

The industrial thesis is strengthened by this narrower statement. It no longer requires the claim that all information naturally wants to pool. It requires that certain high-value decision problems exhibit a net gain from particular forms of authorized combination, and that the resulting gains are difficult to reproduce under alternative arrangements.

20 Operational knowledge as an asset and a contested option

20.1 A corpus is not a self-contained machine

It is tempting to describe operational traces as capital because they can support future production. The analogy is useful only if the supporting conditions are included. A corpus can depend on a maintained schema, functioning software, contextual interpretation, and continuing legal permission. If any of these disappears, the apparent stock of knowledge can become costly archival material rather than productive capacity.

Let an operational asset be a tuple

$$\mathcal{K} = (D, \mathsf{Interp}, \mathsf{Rights}, \mathsf{Tools}, \mathsf{Eval}), \tag{55}$$

where the components identify records, interpretation, permitted uses, executable tools, and evaluation resources. The tuple is a conceptual productive bundle, not an accounting classification. Two firms holding identical bytes can possess different operational assets because their rights or deployment environments differ.

This formulation explains why an acquisition can fail to transfer the expected knowledge. The acquiring firm may receive records and software while losing workers who understood exceptions, local relationships, and undocumented meanings. A standardized migration can preserve fields while destroying their operational interpretation. The lost component is not visible in a simple data-volume measure. [39]

20.2 Codification investment is endogenous

High-quality traces often require deliberate effort. Someone must label a failure, explain a correction, preserve a rejected alternative, or maintain a schema. If the organization treats this work as a costless by-product, it can underinvest in the very evidence that the learning mechanism requires. Instrumentation does not eliminate the labor of interpretation.

Let e be effort invested in codification and let $q(e)$ measure the resulting usefulness of a record resource. The controller's benefit depends on $B(q(e))$, while the contributor bears a cost $c(e)$. If the contributor captures only a share of the future benefit, privately chosen effort can differ from the joint optimum. The exact relationship depends on contracts, trust, incentives, and the possibility of withholding or distorting information.

This connects the learning channel to incomplete-contract and incentive analysis [2, 4, 10]. The new observation is that instrumented work can make codification appear automatic while leaving its quality highly dependent on human cooperation. A large low-trust corpus may be less useful than a smaller resource produced under explicit, credible arrangements.

20.3 Acquisition value and the timing of learning

An acquirer may value a company for a current operational corpus, for its capacity to generate future records, or for the tools and personnel needed to interpret them. These are distinct assets. A historical corpus can depreciate quickly if the environment changes. A continuing workflow can remain valuable even if old records become obsolete.

The discounted learning value can be represented as

$$\mathcal{L} = \mathbb{E}\left[\sum_{t\geq 0} \beta^t \Big(B_t(Z_t) - C_t(Z_t)\Big)\right], \qquad 0 < \beta < 1, \tag{56}$$

subject to continuing rights, task relevance, and deployment access. This is a generic valuation structure, not an estimate. It makes clear why future data-generating capacity can matter more than the nominal size of an acquired archive.

It also creates a feedback into acquisition strategy. A platform may prefer companies whose workflows fill gaps in its coverage, rather than those with the largest record count. Alternatively, it may find that the cost of interpreting a specialized corpus exceeds its transfer value. The model permits both outcomes. [40]

20.4 The option value of preserving alternatives

Operational independence can have value even when it is not exercised immediately. A portable workflow, a maintained local model, or a second viable supplier can preserve the option to respond to a future price increase, outage, policy change, or political restriction. A current-cost comparison that ignores this option can overstate the attraction of integration.

The option is task-specific. A duplicate system that cannot perform essential actions provides little resilience. A modestly more expensive alternative that preserves a critical workflow can be valuable. The cost of maintaining it should be compared with the distribution of future losses, not only with ordinary-period utilization.

The same reasoning applies to workers' independent competence. Maintaining the ability to diagnose a system without the dominant assistant can be costly in normal operation and valuable during a failure. Treating all such duplication as waste overlooks an insurance function. The empirical question is whether the maintained alternative is effective under the relevant stress.

20.5 Learning rents and collective contributions

A controller can earn a return from a capability assembled from many sources: employee corrections, customer interactions, open documentation, public software, and purchased compute. The fact that the final model or service is controlled by one firm does not imply that the underlying contribution was individual or exclusively proprietary. Nor does collective contribution imply that every use should be unrestricted.

The distributional question concerns who receives the return from a jointly generated resource and who bears the risks of reuse. A worker may supply valuable corrections without sharing in the resulting productivity gain. A customer may gain better service while losing a distinctive operating advantage. An external community may contribute methods that are incorporated into a closed orchestration layer.

These issues cannot be settled by marginal-contribution arithmetic alone. With complements, the value of one contribution depends on the presence of others. Different counterfactual baselines produce different allocations. A governance arrangement must therefore state its principles explicitly rather than imply that a unique fair division is encoded in the production function. [41]

20.6 What becomes a strategic bottleneck

If model capability becomes more accessible, the scarce asset can move to evaluation, contextual interpretation, or authorized action. A platform need not own every record to control the route from records to useful service. It can hold the only accepted identity assertion, the necessary payment interface, or the deployment approval channel.

This relocation of scarcity reinforces the multilayer perspective. The value of the operational asset is the value of a bundle. Policies that address one component can shift advantage to another. Data portability can coexist with nonportable permissions; accessible weights can coexist with a closed tool environment; open schemas can coexist with exclusive distribution.

The concentration thesis is therefore best stated in terms of control over the bundle's indispensable relations. An institution that can combine records, interpret them, evaluate

changes, and execute approved actions may possess a powerful position even when it owns only some of the underlying components.

21 Versioning, provenance, correction, and forgetting

21.1 A self-description must remain interpretable over time

An operational record is generated under a particular version of software, schema, model, policy, and organizational practice. Replaying it later can fail because the environment has changed. A command can have a different effect, an identifier can have been reassigned, or an approval rule can have expired. A useful organizational memory must preserve enough temporal context to interpret the record.

We can represent the relevant version state by

$$v = (v_{\text{schema}}, v_{\text{tools}}, v_{\text{model}}, v_{\text{policy}}, v_{\text{environment}}). \tag{57}$$

The tuple need not enumerate every physical detail. It should include the distinctions necessary for the intended reuse. A compact version identifier can refer to a maintained description, but it is useful only if that description remains available and trustworthy.

This introduces an organizational maintenance cost into the learning channel. Accumulation without semantic maintenance can produce a growing archive with declining executable value. The institution may know that something worked while no longer knowing under which conditions it worked. [42]

21.2 Time and provenance are different relations

Event time, observation time, ingestion time, and correction time can differ. A delayed incident report can be recorded after the intervention it explains. A model-generated summary can appear earlier in a search index than the corrected primary source. Sorting records by one timestamp does not recover the causal structure of the workflow.

Provenance can instead record that one artifact was derived from another, that an action used a particular resource, and that a decision was attributed to a principal. W3C PROV supplies an established vocabulary for parts of this structure [39]. The present application adds an emphasis on the policy and version conditions under which a derivation may be reused.

Let G_P be a graph of records, transformations, evaluations, and deployed artifacts. A path can indicate a chain of influence. The graph is not necessarily a complete causal model: unrecorded human decisions and environmental events can remain outside it. Its value is to support traceability and targeted inquiry, not to guarantee that every effect has been captured. [43]

21.3 Correcting a record is not the same as correcting a capability

Suppose an incident is initially labeled as a network failure and later diagnosed as an authorization defect. The corrected label may need to propagate to a retrieval index, an evaluation set, a runbook, and a training resource. Updating the original row alone can leave several derived artifacts with the old interpretation.

This can be modeled by a correction request attached to a node of G_P, followed by a traversal of affected descendants. The traversal identifies candidates for review. It does not determine automatically whether each descendant remains invalid: some derived artifact may be unaffected by the corrected field, while another may require a full rebuild.

A simple dependency flag therefore has limited meaning. It can support a conservative review policy, but an exact impact analysis requires semantics of the transformations. In a categorical model, the relevant diagrams should express how the corrected data map into later resources. If those maps are unknown or opaque, the organization cannot claim complete correction merely because it has a dependency graph.

21.4 Deletion and machine unlearning

Deleting a raw record, removing it from retrieval, and removing its influence from a trained model are different operations. Work on machine unlearning studies methods and definitions for changing learned systems after data removal requests [72, 73]. The literature provides important technical approaches under specified training structures and guarantees; it does not make arbitrary influence removal costless or universally exact.

An operational architecture should distinguish at least raw-data deletion, index deletion, cache invalidation, evaluation-set correction, and model-update handling. The appropriate response depends on the use, the applicable requirement, and the available technical guarantee. A system that treats all five as one delete button can create a misleading account of control. [44]

These costs affect the economics of integration. A broad operational-learning system may benefit from combining many sources while incurring large obligations to track correction and withdrawal. A modular or locally adapted system can sometimes make those obligations easier to isolate. The relevant comparison includes lifecycle maintenance, not just initial training or migration cost.

21.5 Semantic migrations can fail to commute

Let $F : \mathcal{C} \to \mathcal{C}'$ be a schema migration and let Q be a query or decision procedure. Ideally, a corresponding migrated procedure Q' preserves the relevant result under the data transformation. This can be expressed by a commuting or coherently comparable diagram. A change in field meaning, units, or missing-value treatment can make the diagram fail.

The failure need not be a software bug. It may reveal that the new schema intentionally changes the institution's categories. For example, merging several incident types can simplify reporting while eliminating distinctions needed for a safety evaluation. The migration can be technically successful and analytically inappropriate for that later use.

Higher-categorical comparisons can record the intended relation between old and new procedures. They cannot supply an equivalence where the new representation has genuinely forgotten a necessary distinction. In that case the correct response may be to retain an older view, revise the task, or acknowledge that historical comparisons are not valid.

21.6 Forgetting can be productive

The accumulation metaphor encourages treating every deletion as a loss. Yet obsolete, misleading, or unjustifiably retained records can reduce the quality and legitimacy of organizational learning. Forgetting can improve a system by removing stale assumptions, limiting exposure, or forcing a reassessment of a procedure that no longer applies.

The appropriate question is not whether the institution remembers everything. It is whether it retains the information needed for justified purposes and can explain the lifecycle of the capability derived from it. A memory system that cannot correct or forget is not necessarily more intelligent; it may simply be less governable.

The organizational panopticon thus has a temporal dimension. Power can reside in the ability to preserve and recombine old traces after their original context has disappeared. Accountability requires meaningful limits on that ability, together with realistic statements about which derived effects can and cannot be removed.

22 Objections, counterarguments, and discriminating replies

22.1 Objection: this is ordinary organizational learning

Organizations have long documented procedures, trained replacements, and improved machinery through experience. Calling the process an AI panopticon can overstate novelty and understate continuity. The objection is correct if the thesis is presented as the invention of learning from work.

The narrower claim concerns a change in the cost and reach of a particular channel. Software-mediated activity can produce machine-readable records at the point of execution, and learning systems can use selected records to change behavior across a growing range of interfaces. Whether this is quantitatively transformative must be measured. The paper's novelty is not the existence of feedback but its integration with trace semantics, action environments, and infrastructural control.

An empirical reply would therefore compare the incremental cost of codification and transfer with earlier systems. If AI adds little beyond conventional documentation and process mining, the stronger claim should be reduced accordingly. Historical continuity is a constraint on interpretation, not a reason to ignore a potentially large change in scale or transferability. [45]

22.2 Objection: data are nonrival, so concentration should weaken

A dataset can be copied without destroying the original. This can support sharing, entry, and competition. If useful records and methods diffuse readily, the collector's advantage may be temporary. The objection is especially important where models and interfaces are accessible to many providers.

Nonrivalry in use does not imply nonexcludability, equal rights, equal interpretation costs, or equal access to the action environment. A copied trace can remain unusable without context, tools, or permission. The paper's concentration mechanism concerns the control of complementary relations, not the physical scarcity of bytes.

The discriminating evidence is therefore whether independent actors can obtain and exploit substitutes at reasonable cost. Where they can, the operational-learning channel may improve an industry without strongly concentrating it. Where they cannot, nonrival technical content can coexist with exclusive institutional control.

22.3 Objection: diseconomies of organization will dominate

Broad integration can slow decisions, increase internal politics, produce incompatible goals, and weaken specialist expertise. A large organization can collect extensive records while learning poorly from them. The cost of coordinating a universal stack may exceed the savings from sharing infrastructure.

This objection is incorporated directly through organizational costs and task specificity. The framework does not assume that integration surplus is always positive. It asks whether the learning channel changes the comparison at the margin. A failed acquisition, a successful specialist, or a high-performing federation can therefore count as evidence about the mechanism rather than as an exception dismissed by definition.

The crucial measurement is the net return, including migration, semantic maintenance, worker cooperation, and lifecycle governance. If those costs routinely outweigh the learning benefit, the industrial prediction weakens. A theory that describes all failures as temporary implementation issues would not be falsifiable.

22.4 Objection: privacy and security will prevent integration

Legal restrictions, contractual confidentiality, encryption, and institutional preferences can limit data pooling. Purpose partitions can make a global trace corpus impermissible or

operationally undesirable. A strong version of this objection would treat such limits as sufficient to preserve plural systems.

The reply is conditional. Restrictions can indeed block the learning channel and favor local or federated arrangements. But some useful integration can occur through permitted views, aggregate updates, retrieval interfaces, or model reuse. The amount of concentration depends on which rights are preserved and which coordination functions remain centralized.

Privacy can also be used rhetorically to defend a closed ecosystem while the controller retains broad internal use. The correct comparison examines actual access and control, not the label privacy-preserving. A system can protect raw records from competitors while concentrating the power to derive and deploy capabilities from them. [46]

22.5 Objection: agents make silos obsolete without mergers

If agents can translate between schemas and operate existing interfaces, they may reduce the need for vertically integrated providers. A capable intermediary could compose heterogeneous services on behalf of a user. This is a serious technological counterforce to the thesis.

The paper agrees that translation can reduce some interface costs. It does not follow that agents can supply missing permissions, resolve disputed identities, guarantee semantics, or bypass a provider's discretionary control. The intermediary can itself become a new bottleneck through memory, identity, routing, or default tool selection.

The discriminating question is whether the agent improves practical exit and independent governance. If it makes providers interchangeable under a user's control, it supports federation. If it hides multiple suppliers behind one nonportable orchestration layer, it can relocate concentration. The same technical capability can support either arrangement.

22.6 Objection: the panopticon metaphor prejudges the conclusion

The metaphor can encourage treating all observation as coercive and all integration as a plan for replacement. It can also imply perfect visibility that no real system possesses. These risks are genuine. The paper therefore restricts the term to asymmetrical visibility and control over consequential reuse.

The framework can describe a participatory learning institution that logs extensively, gives contributors influence, preserves purpose limits, and supports exit. Such an institution would realize parts of the operational-learning mechanism with weaker panopticon characteristics. The distinction is an intended feature of the model, not a retreat from it.

If the metaphor obscures these differences in a particular application, operational-learning infrastructure is the more neutral description. The substantive claims survive without the metaphor. Their validity depends on the specified observation, learning, control, and market relationships.

22.7 Objection: no finite checks or formal proofs means no rigor

The paper is a conceptual and analytical manuscript. Its rigor should be assessed through clear types, stated assumptions, correct use of standard mathematics, explicit non-implications, and appropriately scoped evidence. Machine checking can validate a formalized statement, but it cannot substitute for choosing the right industrial interpretation or obtaining the missing empirical evidence.

The manuscript therefore avoids presenting the concentration thesis as a proved theorem. The mathematical examples clarify logical relationships, while the empirical program states what remains to be investigated. A future formalization could verify selected constructions, and an empirical study could test particular mechanisms. Neither is claimed to have been completed here. [47]

22.8 The cumulative reply

The objections show why the initiating idea should be formulated as a conditional mechanism rather than a universal trend line. Every major counterforce enters the framework through an identifiable quantity: information value, transferability, appropriation, coordination cost, rights, or exit. The framework is useful only if those quantities can make the proposed effect weak, absent, or reversed.

What remains is a substantive and testable possibility. When productive activity yields valuable operational evidence, when the evidence can be combined with complementary infrastructure, and when its benefits are exclusive and difficult to substitute, AI can strengthen incentives for broad integration and concentrated control. If one or more of those conditions fail, other industrial arrangements can remain viable or preferable.

23 An empirical program that could support or weaken the thesis

23.1 Separate the dependent variables

The empirical program begins by distinguishing four outcomes: task performance, learning from traces, operational dependence, and market concentration. A deployment can improve performance without learning from local traces. It can learn locally without increasing dependence. It can increase dependence without increasing legal concentration. A design that uses one outcome as a proxy for all four cannot test the proposed mechanism.

For task performance, record quality, cost, latency, intervention, and recovery under declared conditions. For trace learning, estimate the incremental effect of selected operational evidence relative to a suitable comparison. For dependence, measure the feasible substitutions

defined in Section 13. For industrial structure, report layer-specific shares and control rights with a transparent denominator.

The unit of observation should usually be an organization-task-time tuple, rather than a worker or a whole industry alone. This permits the same organization to be independent for one task and dependent for another. It also avoids interpreting a shift in the mix of tasks as an improvement in performance on a stable population.

23.2 Identify the learning channel

A useful comparison separates changes in model parameters, retrieval access, tool interfaces, and workflow design. If all change simultaneously, a performance gain cannot be attributed to operational traces alone. Where feasible and appropriate, compare versions with matched models and interfaces but different permitted access to a curated trace resource.

The estimand can be stated as

$$\tau_{\text{trace}} = \mathbb{E}[Y(\mathsf{trace\ access} = 1) - Y(\mathsf{trace\ access} = 0)], \tag{58}$$

for a specified task population and deployment regime. This notation defines a target, not an identification result. Selection, interference, learning spillovers, and changes in human behavior can prevent a simple comparison from estimating it.

Randomized access may be inappropriate when safety-relevant information would be withheld. Alternatives include phased deployment, historical version comparisons with explicit limitations, or matched evaluation environments that do not affect live decisions. An evaluation of the hypothesis should not itself justify indiscriminate recording of workers or customers. [48]

23.3 Measure integration independently of size

An integration index should distinguish schema compatibility, identity linkage, workflow reach, and discretionary control. A large enterprise can score highly on size and poorly on interoperability. A small federation can expose a coherent task interface. Using revenue or employee count as a direct proxy for integration would obscure these possibilities.

One candidate is a task-coverage measure

$$I_{\mathcal{T}} = \sum_{\tau \in \mathcal{T}} w_\tau \mathbf{1}_{\mathsf{Comp}(\tau)}, \tag{59}$$

where $\mathsf{Comp}(\tau)$ means that the task is composable across its required interfaces, weights sum to one, and composability includes semantic and authorization requirements. For an infinite task population, replace the sum by an integral. This is a descriptive index; it does not imply that maximizing coverage is the correct objective.

A separate visibility measure can report which classes of records a controller can access, for which purposes, and for how long. A separate control measure records veto rights and feasible exit. These dimensions can move in opposite directions. For example, a federated query system can improve task coverage while reducing raw-record visibility.

23.4 Test the additional return to scope

The learning-aware integration hypothesis predicts a larger integration surplus where records are informative, transferable, and costly for rivals to reproduce. It predicts weaker effects where data are redundant, rights restrict reuse, models are readily substitutable, or interfaces capture most gains without common ownership.

An empirical model might relate a measured integration outcome Y_{it} to usable trace access Z_{it}, interface compatibility I_{it}, and their interaction:

$$Y_{it} = \alpha_i + \lambda_t + \beta Z_{it} + \gamma I_{it} + \theta Z_{it} I_{it} + X'_{it}\delta + \epsilon_{it}. \tag{60}$$

The interaction coefficient is not automatically causal evidence of complementarity. Adoption is endogenous; successful firms may invest in both variables. Appropriate instruments, natural variation, or carefully designed comparisons would be needed, and none is asserted here. [49]

23.5 Falsification and competing explanations

Several findings would weaken the thesis. Operational traces might add little value beyond public documentation and task-specific examples. Improvements might transfer readily to rivals. Integration gains might be explained almost entirely by identity and workflow design while model learning contributes little. Federated systems might achieve equivalent performance with lower exit costs. Organizational complexity might outweigh the advantage of a broad corpus.

Other findings would support the mechanism: a robust incremental benefit from context-rich traces; gains concentrated in tasks spanning several systems; persistent advantages tied to exclusive access rather than model scale alone; and increasing dependence after trace-rich integrations even when legal ownership is unchanged. Support would still be local to the studied setting.

The distinction between public and private evidence also matters. A company can claim proprietary operational learning without disclosing enough to assess it. Researchers should report the resulting uncertainty rather than treat the claim as a measured intangible asset. Confidential access can improve evidence but may also limit independent replication.

23.6 Report missingness, censoring, and intervention

An evaluation should report which work is absent from the operational corpus. Difficult cases may move to an unrecorded conversation, a worker may repair an error outside the instrumented interface, or a failed service may stop emitting records. The observed population can then become easier precisely when the system performs poorly. A high completion rate among retained traces need not represent the full task population.

Let R indicate inclusion in the evaluated corpus. In general,

$$\mathbb{E}[Y \mid R = 1] \neq \mathbb{E}[Y] \tag{61}$$

is possible, especially when inclusion depends on task difficulty or the outcome. More detailed fields within the retained records do not resolve selection into the corpus. A faithful trace semantics on the observed execution category can coexist with a badly incomplete choice of that category.

This is a second limitation of the panopticon metaphor. The controller may possess dense records of the interface it operates while overlooking work that occurs around the interface. Workers' repairs, negotiations, and tacit coordination can be essential to the apparent success of the automated service. An evaluation that attributes their results entirely to the model confuses the visible component with the full production system.

A reporting protocol can distinguish automatically completed tasks, completed tasks with human intervention, escalations, abandoned attempts, and tasks diverted before an agent begins. It should state how each category enters cost and quality measures. Intervention should be counted even when it is beneficial: the point is accurate attribution, not treating human participation as failure.

Where possible, use independent task-intake or outcome records to estimate coverage, subject to the same privacy and governance constraints as the main study. The comparison should not rely exclusively on the system whose completeness is being evaluated. When coverage cannot be estimated, the study should state the resulting limit rather than imply that a complete self-description has been obtained. [50]

23.7 A longitudinal evidence register

A useful observational program would maintain a dated register of relevant changes: partnership terms when public, interface portability, model deployment options, data-use policies, switching experiences, and common-cause incidents. Every item should retain its source type and evidence window. This would prevent a regulatory finding from one period from becoming an unqualified description of a later one.

The register should also record successful exits, effective federations, and failed integrations. A database containing only acquisitions and large partnerships would be selected in favor of the thesis. Its purpose is to compare pathways, not to accumulate examples of a predetermined

industrial future.

The empirical program is thus demanding but tractable in parts. It replaces a universal prediction with linked claims whose conditions can be examined separately. That is the appropriate standard for an argument that combines technical possibility, institutional behavior, and market structure.

24 Conditional scenarios for the 2030s

24.1 Scenarios are not estimated forecasts

The initiating argument asks what industrial organization might look like in the 2030s. The available evidence does not justify a numerical forecast or a deadline for convergence. We instead describe conditional scenarios using variables already introduced: transferable learning value, interface cost, appropriation, practical exit, and systemic loss. The scenarios identify what would have to be true for each pathway to become plausible.

They are not mutually exhaustive. Different sectors can follow different pathways, and an ecosystem can move between them. A firm can centralize model access while federating data, or decentralize inference while retaining a common identity bottleneck. The goal is to avoid treating one visible dimension as the whole industrial system.

24.2 Scenario A: concentrated integrated ecosystems

In this pathway, useful operational learning remains substantially exclusive; large fixed costs persist; identity and distribution reinforce model deployment; and customers face significant switching costs. A small number of ecosystems integrate many formerly distinct services. Legal entities remain numerous, but their task-level independence declines.

The model predicts broad workflow coverage, dense internal trace reuse, and control concentrated at several complementary layers. This can yield useful services and rapid adaptation. It can also make exit difficult and common failures costly. Its signature is not merely a large cloud bill, but a combination of operational reach, exclusive learning, and discretionary bottlenecks.

Evidence favoring this scenario would include persistently low substitutability after nominal portability improvements, durable gains tied to exclusive operational context, and increasing reliance on common identity or orchestration services. Evidence against it would include successful provider substitution and rapid diffusion of comparable capability to independent arrangements. [51]

24.3 Scenario B: interoperable federation

In this pathway, shared schemas, portable identity assertions, accessible models, and local execution reduce the cost of composing services without transferring all records or control rights. Operational learning occurs partly through local adaptation, negotiated aggregates, and shared evaluations. Integration remains extensive, but the capacity to govern and replace components is more distributed.

The principal challenge is sustaining the shared infrastructure. Participants must fund maintenance, agree on versions, resolve disputes, and prevent capture by dominant members. The success of federation depends on institutions as well as protocol design. A diagram of interoperable services does not establish a viable governance model.

Evidence favoring this scenario would include comparable task performance across multiple interchangeable implementations, practical exits within declared budgets, and durable local authority over sensitive records. Evidence against it would include a nominal federation whose identity, certification, or orchestration remains controlled by one difficult-to-replace actor.

24.4 Scenario C: sovereign and sectoral fragmentation

In this pathway, national-security concerns, procurement rules, data sovereignty, sector-specific obligations, or resilience requirements support multiple partially separated ecosystems. Some duplication persists because independent control has substantial value. Integration occurs within bounded domains, with carefully negotiated interfaces between them.

Fragmentation can preserve autonomy and reduce some common-cause exposure. It can also increase costs, slow knowledge diffusion, and reduce the quality of services available to smaller participants. The social value of separation depends on whether the protected boundaries correspond to meaningful rights and failure domains or merely to entrenched local interests.

Evidence favoring this pathway would include persistent limits on cross-domain data reuse, investments in independent deployment capacity, and task interfaces designed around controlled exchange. Its failure mode is costly duplication without real independence, where apparently separate systems rely on the same inaccessible components underneath. [52]

24.5 Scenario D: modular specialization around common substrates

A fourth pathway combines accessible model capability with strong specialization. Many firms build narrow services and domain-specific operational corpora, while relying on a smaller set of compute and distribution infrastructures. The model layer becomes more plural than the substrate layer. Scope economies matter, but specialized task knowledge limits the value of a universal operational corpus.

This configuration can generate substantial competition at the service layer while preserving infrastructural concentration. It is neither a clean federation nor a conventional vertically integrated monopoly. Its key question is whether specialists can change substrates or negotiate favorable terms as their services grow.

Evidence favoring this scenario would include durable specialist advantages, limited cross-domain transfer of operational traces, and multiple viable model suppliers. Evidence against it would include systematic acquisition or displacement of specialists once their operational methods become legible to a broad platform.

24.6 Parameters that move the boundary

The scenarios can be compared along two conceptual axes: the net gain from broad integration and the cost of concentrated control. High gains combined with weak safeguards favor concentrated ecosystems. High gains combined with effective portability and shared governance favor federation. High separation value can sustain sovereign or sectoral systems. Strong task specificity can sustain modular specialization.

These axes are not observed numerical indexes in the present manuscript. They are a way to organize comparative inquiry. The sensitivity diagrams included in the companion source use only declared illustrative parameters and should not be read as market estimates.

24.7 A prediction that remains meaningful under plural outcomes

The paper's most defensible expectation is that the relation between operational traces and executable workflows will become an increasingly important object of bargaining, architecture, and competition. Even a plural industry can exhibit that tendency. Firms may compete over who may observe, interpret, train on, and execute the methods embodied in everyday work.

The appropriate monitoring target is therefore a vector of practical independence, trace reuse, and control across layers. A future with many logos can still be concentrated; a future with shared information standards can still be institutionally plural. The difference lies in the rights, substitutions, and governance represented by the connections.

25 A governance pattern for jointly produced operational knowledge

25.1 Define the relation before enabling the flow

The analysis suggests a concrete institutional pattern: a trace-to-capability arrangement that specifies the relation between productive participation and permitted improvement. This is a

research design pattern, not a model legal contract. Its purpose is to identify the information an organization would need to make its learning channel reviewable by the participants.

The arrangement should begin with a task family and the purposes for which recording is necessary. It should identify the actors generating records, the entity operating the service, the parties controlling derived artifacts, and the bodies able to review disputed uses. An undifferentiated statement that data improve the service is inadequate because it leaves the relevant service, beneficiaries, and modes of improvement unspecified.

Represent an admitted use by a record

$$g = (\mathsf{principal}, \mathsf{resource}, \mathsf{purpose}, \mathsf{transform}, \mathsf{scope}, \mathsf{expiry}, \mathsf{review}). \tag{62}$$

The tuple records a grant or other declared authority. It does not establish that the authority is valid or sufficient. A technical system must connect the tuple to enforceable checks and an institutional process capable of correcting mistakes. [53]

25.2 Make the improvement channel explicit

The arrangement should distinguish at least four channels. Operational diagnostics use records to investigate and repair the service. Organization-specific adaptation changes a resource used for that organization's tasks. Shared improvement contributes to capabilities offered to other participants. Consequential assessment uses records to make decisions about the people or organizations producing them.

These channels can require different restrictions even when they begin with the same trace. A debugging record may be essential for resolving a failure while being inappropriate for ranking an employee. A customer may accept local retrieval without agreeing to a broader training contribution. A shared corpus may be acceptable when governed collectively and unacceptable when subject to unilateral commercial reuse.

25.3 Preserve correction and disagreement

A participant should be able to contest a record's interpretation as well as its factual content. An expert's rejection of an agent proposal may be labeled as a preference when it was actually a response to a missing constraint. A customer complaint may indicate a changed goal rather than a failure. The correction mechanism should preserve these distinctions.

An auditable correction event can refer to the original record, state the disputed interpretation, identify the responsible reviewer, and record the effect on derived resources. Where disagreement remains unresolved, the system can retain a disputed status rather than silently convert one party's interpretation into a training label. This can improve both legitimacy and learning quality. [54]

Channel	Resource affected	Question for review
Diagnostics	Incident records and repair procedures	Is the collection proportionate to resolving the operational issue?
Local adaptation	Local index, policy, or model configuration	Can the organization inspect, correct, and replace the adapted resource?
Shared improvement	Cross-participant model or corpus	Who benefits, who controls reuse, and which contributions remain attributable?
Consequential assessment	Personnel or enterprise decision process	Are the criteria justified, contestable, and appropriate to the decision?

Table 3: One trace can enter several channels with different governance questions.

25.4 Evaluate exit and retained dependence

The arrangement should specify what an exiting participant receives and which capabilities remain with the provider. Portable records, schema documentation, local adaptations, and evaluation artifacts can support continuity. Some shared capability may be retained under the agreement, while other material may require deletion or restricted use. The distinctions should be disclosed before the relationship becomes difficult to leave.

A practical exit exercise can compare the declared task family before and after substitution under a fixed budget and horizon. This would be an empirical evaluation of the arrangement, not a proof that every future dependency has been removed. The paper does not perform such an exercise; it proposes the criterion as a way to make portability consequential.

25.5 Separate the reviewer from the immediate beneficiary

Where the same controller benefits from broader reuse and decides whether reuse is appropriate, a conflict of interest can arise. Independent review, participant representation, or shared governance can reduce that conflict. The appropriate institution will vary, and no single body is prescribed here.

This pattern connects the technical and economic arguments. Better-defined rights can reduce a platform's unilateral freedom while increasing trust, contribution quality, and sustainable participation. The choice is not exhausted by unrestricted extraction and complete informational isolation. A jointly governed learning relationship can preserve useful integration while limiting the concentration of authority over its future results.

26 Conclusion: institutions that learn descriptions of their own operation

AI-mediated organizations can turn productive activity into records that support future production. This is a significant organizational possibility because the records can include expert corrections, rejected alternatives, tool interactions, and responses to failure. The factory can become both a site of work and a source of evidence for constructing machinery that performs parts of that work.

The possibility does not establish an exhaustive self-description. A record is selected by an observation mechanism, interpreted through a schema, and admitted to use under a policy. Its value depends on the task, the environment, the learning procedure, and the available actions. Rich visibility can coexist with mistaken labels, missing context, and limited generalization.

The categorical analysis clarifies the distinctions. A trace functor preserves a specified composition. Faithfulness distinguishes parallel processes in a chosen category. Full faithfulness concerns the structure of mappings. Neither property alone supplies an efficient reconstruction procedure, causal identification, a reliable policy, or authority to act. The additional conditions are not minor technicalities; they identify the complementary assets on which the industrial mechanism depends.

The database and higher-categorical constructions clarify integration. A schema colimit depends on declared maps and identifications. Data compatibility requires witnesses, and policy compatibility requires more than a common representation. Higher cells can retain agreement between alternative translations, while context-indexed semantics can preserve justified boundaries. Informational integration can therefore be achieved through several architectures and ownership arrangements.

The economic analysis adds a learning return to established accounts of scope, contracts, and organizational boundaries. Shared infrastructure and coherent interfaces can lower cost; operational traces can improve selected future capabilities. Yet these benefits can saturate, fail to transfer, or be outweighed by organizational and systemic losses. Positive feedback and complementarity create concentration pressures only under additional conditions of appropriability, competition, and practical dependence.

The empirical sources support parts of this account at matching levels of specificity. Dated regulatory findings identify partnership structures and cloud-market problems. Technical research demonstrates components of learning and agent execution under specified conditions. None establishes the inevitable emergence of one legal corporation or the complete automation of organizational expertise. The industrial thesis remains a set of linked, testable hypotheses.

The sharper object of study is effective informational independence. Many corporations can depend on a few compute, identity, model, and orchestration systems. Conversely, shared schemas and federated protocols can integrate useful information while preserving local control. The distinction between these outcomes concerns who can set terms, refuse reuse,

replace a component, and maintain essential work after a provider withdraws.

This reframes the panopticon concern. The issue is not simply that organizations collect data. It is that productive participation can become a continuing contribution to a privately controlled representation of how production is done. Acquiring an organization may then provide a possible learning corpus as well as revenue, assets, and employees. That additional value exists only when the corpus remains interpretable, usable, and connected to deployment.

The corresponding governance question is equally concrete: who may transform the traces of collective work into future executable capability, under which restrictions, and with what accountability to those whose activity generated them? There is no categorical construction that answers this normative question automatically. There are, however, categorical distinctions that prevent it from being concealed inside the language of seamless integration.

Analytical claim. AI can transform organizations from institutions that use knowledge into institutions that also accumulate reusable operational representations. The concentration of control over those representations and their execution environments is a plausible industrial tendency, not a theorem of technological destiny.

This qualified thesis retains the force of the initiating argument while making it open to counterevidence. It predicts neither universal transparency nor a single final owner. It identifies a mechanism through which ordinary work, technical architecture, and industrial power can become increasingly connected, and supplies a vocabulary for comparing the arrangements through which that connection is governed.

A Technical conventions and extended analytical notes

A.1 A compact dictionary of the model

Symbol	Meaning	Qualification
$\mathcal{E}$	Category of productive processes	Domain, boundary states, and operational equivalence must be specified
$\mathcal{D}$	Category of records	Composition must represent the declared record semantics
T	Trace semantics	Strict functoriality is an assumption, not an automatic property of logs
$\sim_Q$	Task-relevant equivalence	Must be a congruence to form the intended quotient category
$\mathcal{A}_{a,t}$	Feasible automated tasks for an actor	Relative to quality, budget, environment, and permission

Symbol	Meaning	Qualification
D_t	Retained trace collection	Can shrink or change interpretation
Z_t	Curated usable evidence	Need not be a training dataset
M_t	Model state	Distinct from tools, retrieval, and deployment authority
π_t	Deployed policy	Includes its interface and action constraints
$\mathcal{C}$	Schema or context category	These two applications use different interpretations
$\Sigma_F, \Delta_F, \Pi_F$	Data migration functors	Left extension, restriction, and right extension respectively
Γ	Control-rights relation	Can contain overlapping, conditional rights
H_ℓ	Layer-specific concentration	Requires a declared denominator and allocation rule
E_{ij}	Practical exit condition	Relative to tasks, time, quality, and budget
L_a	Net operational-learning return	Includes curation and deployment costs

The dictionary is intended to prevent category errors across the paper. In particular, $\mathcal{C}$ as a schema and $\mathcal{C}$ as a category of visibility contexts are separate models even when the same typographic letter is used. A joint implementation would require a declared relationship between them. The paper does not identify data migration with permission restriction. [55]

A.2 Policy-indexed execution and the Grothendieck construction

Let B be a category of policy contexts and let

$$\mathsf{Exec} : B^{\mathrm{op}} \longrightarrow \mathsf{Cat} \tag{63}$$

be a pseudofunctor. Its fiber $\mathsf{Exec}(b)$ is an execution category under context b. A base arrow $u : b \to c$ induces reindexing $u^* : \mathsf{Exec}(c) \to \mathsf{Exec}(b)$. The pseudofunctor includes coherent isomorphisms comparing repeated reindexing with reindexing along a composite.

This model assumes that the selected reindexings preserve the admissibility of the operations they transport. If an operation loses all meaning under a restriction, one may need partial reindexing, a different base, or an explicit denial outcome. It is preferable to change the model rather than pretend that every permission change induces an ordinary functor.

The Grothendieck construction has objects (b, x) and arrows

$$(b, x) \longrightarrow (c, y) \quad \text{given by} \quad (u : b \to c,\ f : x \to u^* y). \tag{64}$$

This separates a change of context from work performed within a context. The construction is standard categorical machinery [23]; its use to organize purpose-sensitive workflows is an application proposed here. The resulting arrows are semantic comparisons or executions as specified by the interpretation, not automatically network requests. [56]

An indexed trace semantics adds $\mathsf{Trace} : B^{\mathrm{op}} \to \mathsf{Cat}$ and fiberwise functors T_b. Compatibility requires coherent comparisons between $T_b u^*_{\mathsf{Exec}}$ and $u^*_{\mathsf{Trace}} T_c$. These comparisons express what it means for tracing and changing context to agree. If redaction changes the record, the comparison may be noninvertible or require a weaker target structure.

A.3 Ordered records and the writer effect

For a monoid $(R, \cdot, e)$ of record fragments, the writer construction on sets is

$$W(X) = R \times X, \quad \eta_X(x) = (e, x), \quad \mu_X(r, (s, x)) = (r \cdot s, x). \tag{65}$$

It supplies a familiar algebraic account of computations that produce a result and a record. The monoid need not be commutative. For ordered logs, concatenating r before s can carry a different meaning from the reverse order. Lawvere's functorial semantics and standard monad theory provide foundational context for algebraic presentations of effects [24, 36].

This abstraction does not imply faithful tracing. A writer computation can emit the same record for different operations or emit nothing. Nor does the monad implement retention, identity verification, or data rights. Those are extra structures. The point is to separate the algebra of record accumulation from the informational and institutional properties of the record. [57]

A stateful learner can be represented at a different level by a state transformer $S \to S \times X$, where S includes a model and relevant resources. Repeatedly applying such a transformer is not a guarantee of improvement. The state can drift, the objective can be wrong, and the environment can change. Algebraic composability and beneficial recursion are different claims.

A.4 What an effective inverse would require

Suppose a trace semantics is represented by a total computable injection on a countable set of encoded processes. Exhaustive search can sometimes yield a partial inverse that terminates on its image. This observation prevents an overstatement of the reconstruction objection: injectivity does not inherently make recovery impossible under every effective presentation.

The missing requirements are still substantial. The learner must know the encoding and forward procedure; search may be infeasible; membership in the image may be difficult; the recovered syntax may require unavailable resources; and behavioral generalization can remain unresolved. The distinction is between existence, computability on a domain, complexity, and operational realization. [58]

The relevant industrial question is therefore comparative cost. A controller with privileged interfaces, contextual data, and evaluation infrastructure may reproduce a task more cheaply than a rival even when both have access to a trace encoding. That advantage is entirely compatible with imperfect observation and incomplete automation.

A.5 Observation granularity and minimal sufficient access

For a specified task distribution, let Z be the available record and $S(Z)$ a coarser representation. A decision-specific sufficiency condition can be stated by requiring that the best achievable task loss using $S(Z)$ equals that using Z, within an allowed tolerance. This is a task-relative criterion, not a claim that S is statistically sufficient for every parameter or decision problem.

The design question is to find a useful coarsening subject to privacy, interpretability, and implementation constraints. A low-dimensional summary may suffice for one diagnostic task while failing for a rare safety-critical exception. A representation that discards personal identity can reduce exposure but still retain sensitive business information through correlations.

This viewpoint changes the meaning of increasing trace faithfulness. The objective is not necessarily to distinguish every productive process. It can be to preserve the distinctions needed for a declared decision while refusing unnecessary linkage. Statistical channels and contextual policies can therefore support intentionally nonfaithful observations that are adequate for their purpose. [59]

A.6 A constrained optimization for integration

Let a range over feasible architectures. Let $Q(a)$ measure task quality, $C(a)$ operational cost, $P(a)$ a declared privacy loss measure, $R(a)$ systemic risk, and $E(a)$ an exit-capability measure. A policy problem can be written as

$$\min_{a} \quad C(a) + \lambda_R R(a) \tag{66}$$

$$\text{subject to} \quad Q(a) \geq q_0, \qquad P(a) \leq p_0, \qquad E(a) \geq e_0. \tag{67}$$

This formulation makes clear that integration is not a one-dimensional maximization. It also exposes the normative choices in metrics and thresholds. Some privacy or rights constraints may be noncompensable and should remain constraints rather than be monetized in the objective.

The feasible set is itself institutional. Contracts, technical standards, public infrastructure, and governance reforms can make architectures possible that a private firm would not create alone. An observed concentrated outcome therefore does not prove that alternatives are technically impossible. It can reflect the current distribution of investment and coordination capacity. [60]

A.7 Sensitivity illustrations and their interpretation

The companion source includes a plot of Equation (29) and a plot of the leading eigenvalue in Equation (30). All coefficients are declared illustrative constants. The plots visualize comparative statics; they are neither fitted to market data nor estimates of an optimal number of actual firms.

The first plot shows that the same formal cost family can favor a boundary solution or an interior one as the concentration-loss coefficients change. The second shows that a local feedback loop can be attenuating or amplifying. Neither plot adds empirical support to the industrial hypothesis. Their value is to make the model's conditionality visible. [61]

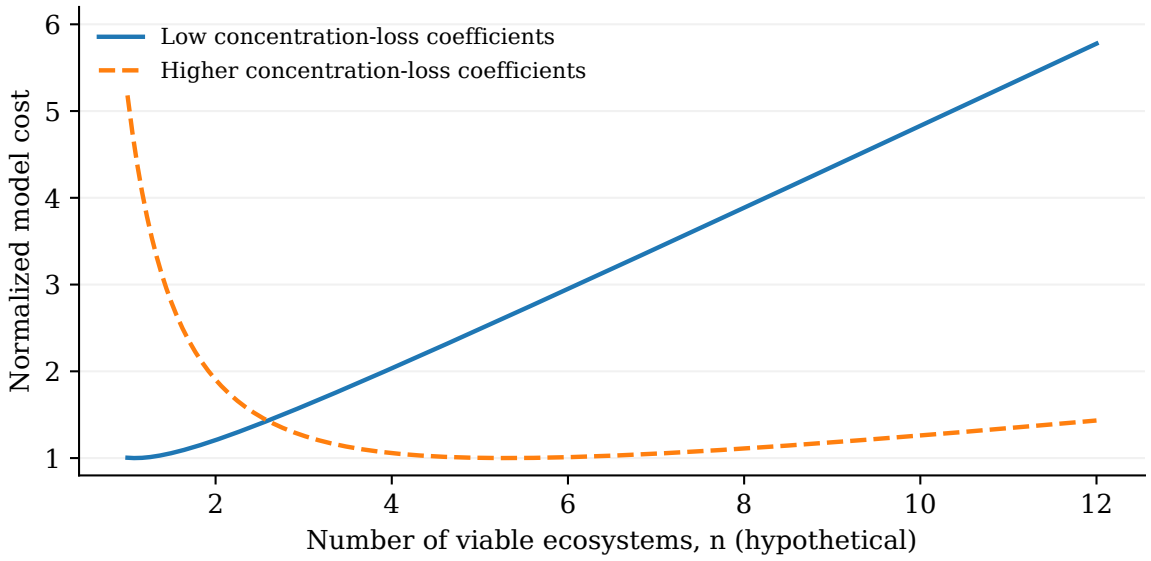


Figure 1: Illustrative cost curves $C(n) = an + b/n + c/n^2$ on $1 \leq n \leq 12$, normalized separately to compare shapes. The solid curve uses $(a, b, c) = (1, 1, 0.1)$; the dashed curve uses $(1, 16, 32)$. These are hypothetical parameters, not market estimates. The horizontal variable counts viable ecosystems in the model, not legal corporations observed in data.

A.8 A claim-status register

The paper uses several kinds of claim, and they should not be read at the same evidential level. The following register summarizes their status without replacing the detailed qualifications in the main text.

Claim	Status	Additional requirement
Some interactions can support model adaptation	Demonstrated in specified technical settings	Task-specific replication for a new deployment

Claim	Status	Additional requirement
Certain cloud/model partnerships link several assets	Dated institutional and public evidence	Current terms for a later-period claim
Faithfulness does not supply all reconstruction requirements	Conceptual and mathematical distinction	A specified effective compiler or learner for a positive result
Operational traces add a return to integration	Conditional economic mechanism	Measurement of incremental net learning value
Few legal firms must remain	Not established	No such prediction follows from this framework
Informational independence can decline without a merger	Model-compatible institutional possibility	Task, rights, and substitution evidence in an application
Federation can realize useful integration	Architecture possibility	Sustainable governance, reliable interfaces, and measured performance
A company's full operation is recoverable from logs	Not established	Adequate context, identification, implementation, resources, and authority

The register also limits the paper's novelty claim. Trace recording, organizational learning, economies of scope, contextual privacy, and categorical data integration have substantial prior literatures. The contribution is their synthesis around a particular institutional mechanism and a set of distinctions that constrain its interpretation. No priority claim is made for the individual ingredients. [62]

A.9 Reference verification and remaining uncertainty

The source audit records how each reference was checked. Most technical items were verified through author papers, official proceedings, standards, or publisher records. Some foundational items were verified at the metadata level rather than through full-text access. The Reuters item remains secondary reporting about a confidential document. Company statements retain their interested-party status.

Verified bibliographic identity is weaker than independent verification of every result in a source. This manuscript does not rerun experiments, reproduce private contracts, or establish current legal applicability. References are included to locate prior work and support claims of matching scope. Where the paper develops an application or inference beyond the source, it identifies that move as its own proposal. [63]

This standard of reference use matters for the argument itself. A paper about informational asymmetry should not conceal asymmetry in its own evidence. The distinction between what

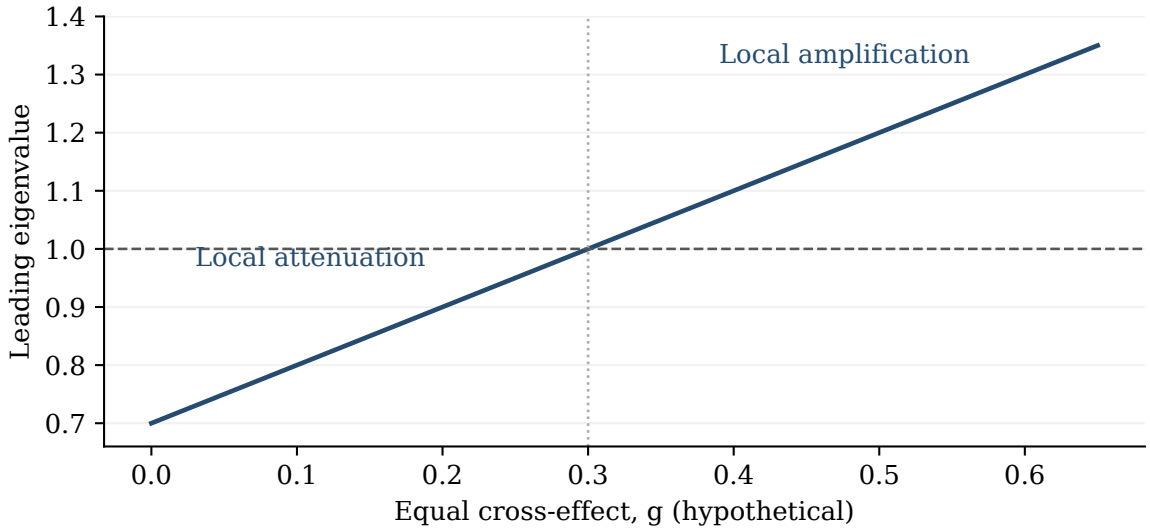


Figure 2: Illustrative leading eigenvalue of the symmetric cyclic linearization with attenuation $\delta_D = \delta_M = \delta_U = 0.3$ and $\alpha = \beta = \gamma = g$. Its leading eigenvalue is $0.7 + g$. Crossing one marks local amplification in this model, not a monopoly threshold.

was observed, what was inferred, and what remains unverified is part of the method rather than an afterthought.

Notes and commentary

[1]The term panopticon carries historical and political associations that a technical model cannot exhaust. Here it identifies a possible asymmetry in operational visibility and authority. It is not used as an established quantitative property of any named company's internal monitoring practices. Section 2 places this restricted usage alongside Foucault, Zuboff, and contextual-integrity theory.

[2]A workflow is not a person. Predicting a worker's next software action does not establish that the system represents the worker's understanding, intentions, or experience. The paper concerns observable task performance, authority over records, and economic consequences of reuse. It neither assumes a theory of consciousness nor treats behavioral imitation as personal equivalence.

[3]The historical comparison is selective. Foucault analyzes disciplinary power, while Zuboff analyzes an economic logic of extracting and commercializing behavioral information. Neither source establishes the technical possibility of reconstructing an enterprise from telemetry. The operational-learning mechanism is a separate, conditional proposal developed in this manuscript.

[4]Tacitness is not used as an immunity argument against learning systems. Some unarticulated regularities may be learned from rich observations. The limited conclusion is that explicit logs need not contain every causally relevant variable, and that any claim of reproducibility must specify what additional observations, interventions, or environmental assumptions are available.

[5]Contextual integrity is a normative framework, not merely a mechanism for asking users to click a consent dialog. The policy-indexed constructions below are a technical way to record selected norms. They do not determine which norms are justified or whether workplace consent is sufficiently free. Those remain substantive questions of governance and power.

[6]The FTC report's evidence window matters: respondent information extends through September 2024, supplemented by public information through January 2025. Its findings should not be silently relabeled as a complete account of partnerships in November 2026. The present paper uses the report to identify mechanisms, not to infer undisclosed subsequent contractual terms.

[7]The source supplied with the initiating text was a Reuters URL. Its direct page was inaccessible during verification, but a dated Reuters syndication with matching title and byline was located. The bibliography links that accessible version. This distinction prevents an accessible news report from being misrepresented as access to a confidential primary document.

[8]The relation Γ is a modeling choice, not a complete theory of corporate law. Rights

may be conditional, contested, delegated, time-limited, or jointly exercised. Where these distinctions affect the task, replace a static relation by a category of grants and revocations or by a temporal policy structure.

[9]The category $\mathcal{E}$ is specified, not discovered automatically by calling arbitrary business records objects and arrows. Associativity, identities, and the compatibility of the operational equivalence must be supplied by the workflow semantics. These are ordinary categorical requirements; the economic interpretation does not relax them.

[10]Faithfulness is a property on hom-sets. The condition $T(f) = T(g) \Rightarrow f \simeq g$ is meaningful only when the equivalence and the types of f, g have been specified. In higher settings, injectivity of a function is replaced by a requirement on mapping spaces; Section 11 distinguishes this from equivalence of mapping spaces.

[11]The six requirements are proposed as an analytical checklist for this problem. They are not six independent axioms of a universal learning theorem. Different tasks can combine or refine them. Their purpose is to block an invalid inference from a structural observation property to an unrestricted claim about skill transfer.

[12]The causal discussion does not claim that every workflow requires a full structural causal model. It identifies a failure mode in extrapolating observed action-outcome associations. Recording an expert's reasons, rejected actions, and local context can help, but none of these records automatically removes hidden confounding or establishes counterfactual validity.

[13]The original additive data equation is a useful first approximation when retention is monotone. It becomes misleading when data are deleted, permissions expire, schemas change, or derived records replace raw events. Equation (21) makes these possibilities explicit without prescribing a particular data platform.

[14]The cited recursion paper is referenced under its arXiv title. A related published version has a different title. We retain the exact preprint title to avoid mixing bibliographic identities. Its results concern particular recursive data-generation regimes; no universal impossibility claim about synthetic data is inferred.

[15]The initiating scalar expression $C_{\text{integrated}}(\sum_i q_i)$ is meaningful only when the q_i are commensurable or already embedded in a common output space. Equation (24) uses a vector to avoid adding, for example, storage capacity directly to payment volume. The classical reference's metadata was verified; its subscription full text was not inspected in this session.

[16]A learning return can be appropriable without being socially additional. A controller may capture knowledge that was previously shared informally or available to workers, rather than create new productive knowledge. Welfare analysis must distinguish redistribution of control from an increase in total capability, and it must include whose opportunities are reduced by exclusive reuse.

[17]Learning and concentration are different outcomes. A positive feedback within one service is insufficient to show increasing market share, since prices, capacity, quality differences, spillovers, and rival responses also matter. The industrial hypothesis requires a bridge from

internal capability gains to appropriable competitive advantage.

[18]A pushout is relative to specified maps. The same two schemas can produce different integrations under different choices of overlap and identification. Writing a gluing symbol without the diagram can conceal the most consequential design decision: which distinctions are declared to represent the same thing.

[19]The word colimit is used literally only for an identified diagram in an identified category. Broader phrases such as informational integration are sociotechnical descriptions. A technically valid schema colimit may realize only a small part of that broader integration, and may be implemented virtually without moving all raw data to one location.

[20]A two-cell can represent many things in applications: a transformation, a refinement, a simulation, or a compatibility witness. These interpretations are not automatically interchangeable. A two-cell called an audit does not become a legal attestation merely by appearing in a bicategory. Its intended meaning must be declared.

[21]The notation $\mathcal{E}_\infty$ does not assert that a complete higher category of an enterprise has been constructed. It specifies the structure required for a proposed refinement. If only ordinary workflows and transformations are available, the lower-dimensional model should be retained. Higher language earns its role by preserving a distinction the simpler model loses.

[22]A presheaf of records is a semantic model, not a security boundary. Its restriction maps describe intended behavior. Enforcement still requires authenticated identities, complete mediation, revocation, isolation, and operational assurance. A mathematical model of an authorization rule does not establish that a deployed service obeys it.

[23]It is tempting to call every failed enterprise integration a sheaf-cohomological obstruction. That requires a specified site, coefficients or suitable higher structure, and a demonstrated relationship between the obstruction and the practical problem. This paper uses local-to-global compatibility directly and makes no unsupported claim that arbitrary policy conflicts are measured by a cohomology group.

[24]A supplier relationship is not automatically a relation of domination. Dependence becomes consequential when substitution is costly, delayed, legally constrained, or technically infeasible for the task. The graph should distinguish revocable discretionary access from commodity supply under effective alternatives.

[25]The phrase distributed monopoly is retained only as an informal warning about concentrated practical dependence across legal boundaries. It is not a legal finding of monopoly, a claim of coordination, or a substitute for market definition. Distributed concentration or ecosystem dependence is usually the more precise description.

[26]The example intentionally spans read operations, proposals, authorizations, and actions. An agent that is excellent at diagnosis may still be unsuitable for deployment or payment decisions. Treating every stage as the same generic tool call conceals differences in reversibility, liability, and required evidence.

[27]The reliability expressions are analytical relationships, not experimental validation. They

are included to show where an independence assumption enters. A deployed service would need empirical measurement of common causes, recovery behavior, and human intervention before any numerical reliability claim was justified.

[28]The task-family comparison is a method for exposing assumptions. It is not evidence that software work is uniformly automatable, that enterprise approvals are uniformly simple, or that physical expertise is intrinsically inaccessible to computation. The relevant variables differ within each family as well as between families.

[29]The feasibility expression is intentionally modular. An organization can improve one predicate without satisfying the others. Reporting such partial progress is more informative than classifying an entire occupation as automated or unautomated. The capability definition in Equation (5) remains relative to a task distribution, quality threshold, and budget.

[30]The participation expression is deliberately descriptive. It does not claim that compensation alone resolves rights or power questions, nor that workers can always refuse collection freely. In settings with unequal bargaining power, formally agreed terms can coexist with limited meaningful choice.

[31]References to legislation identify original official texts and selected concepts. They are not a current legal opinion, a claim that every record is personal data, or a conclusion that consent is always the appropriate basis. Applicability, amendments, regulatory interpretation, and the facts of a particular deployment require separate analysis.

[32]Event-level, user-level, and organization-level privacy are different promises. A parameter is uninterpretable without the neighboring-dataset relation, mechanism, accounting method, and release policy. The equation states a mathematical condition; it does not certify the privacy of the illustrative architecture.

[33]Open source, open weights, publicly documented, and commercially downloadable are not synonyms. The cited reports are technical sources for named model families. Any procurement or licensing conclusion would require the exact release terms and deployment facts, which this manuscript does not attempt to adjudicate.

[34]The legislative references are deliberately separated. The DMA is not a generic mandate that every AI model interoperate with every other model, and the Data Act is not a claim that all derived enterprise capability must be transferable. The paper draws a conceptual connection to contestability while retaining the statutes' distinct subjects and scope.

[35]Diversity is not automatically beneficial. Independent poor systems can fail often, and apparent provider diversity can conceal a common dependency. The relevant quantities are common-cause exposure, recovery capacity, and the distribution of losses. A resilience argument should measure those properties rather than count logos.

[36]The decision model uses an explicit prior and loss to make the value question definite. In a real organization these choices can be disputed. Different stakeholders may care about different errors, and a single expected-loss objective can hide distributional harms. A multi-objective or constrained model may therefore be preferable.

[37]The notation $\succeq_{\mathcal{T}}$ is introduced for this manuscript. It is not a new general characterization theorem. Restricting the decision class can weaken an information order and change which observations count as equivalent. That is a feature when the purpose is to preserve task utility while limiting unnecessary exposure.

[38]A broad corpus can improve coverage and still require local adaptation. The model does not assume that all workflows share one invariant task distribution. In some domains, the most valuable information is precisely the fact that two apparently similar actions have different meanings under different institutional conditions.

[39]Calling a record bundle an asset here makes no claim about financial reporting, ownership of employee knowledge, or an applicable valuation standard. It means only that the bundle can contribute to future productive activity under specified conditions. Legal rights and accounting treatment require separate analysis.

[40]The acquisition mechanism is a hypothesis about incentives. The paper supplies no estimate of how much any actual transaction price reflects operational-learning value. Evidence for that claim would require transaction-specific documents or credible variation separating learning motives from revenue, talent, distribution, and conventional strategic considerations.

[41]The paper does not propose a universal property right in every operational trace or a unique compensation formula. Its narrower claim is that a credible institutional arrangement for contribution and reuse can affect both distribution and data quality. Ethical legitimacy and productive performance can interact rather than form wholly separate questions.

[42]Versioning is not only a software release number. The same tool can behave differently under a changed permission, business rule, or external service. A trace intended for future learning should preserve the environmental distinctions relevant to its target use, subject to minimization and retention constraints.

[43]A provenance edge states a declared relationship. It can be incomplete or false. Cryptographic signatures can authenticate who made a statement or whether bytes changed, but they do not establish that the stated event occurred as described. Trust in provenance depends on the recording process and its accountability.

[44]The cited unlearning sources are primary research papers. Their relevance is the distinction between removing stored examples and changing a learned system. No particular method is claimed to satisfy every legal erasure requirement or every model architecture. The manuscript does not evaluate an implementation.

[45]The argument does not require an unprecedented industrial mechanism in every respect. A familiar mechanism can become more consequential when its costs, scope, or appropriability change. Conversely, new terminology does not establish that such a change has occurred. The empirical agenda is designed to separate these possibilities.

[46]This is a structural possibility, not an allegation about a named provider. Privacy and security are substantive values and can justify genuine boundaries. The analytical concern is that claims about these values should be evaluated against the architecture's actual

information flows, rights, and available alternatives.

[47]The companion contains editable sources and plotting code for exposition, without finite model checks, proof-assistant artifacts, or claims of machine-verified economic conclusions. These are different forms of scholarly support. Their absence does not relax the requirement that the mathematical statements themselves be accurate.

[48]No empirical study is performed in this manuscript. The proposed estimands and designs are a research agenda, not reported findings. Randomization is mentioned as one possible identification tool, subject to consent, safety, organizational constraints, and interference between units.

[49]A positive estimated interaction can arise from omitted variables, nonlinear measurement, or selection. Economic complementarity is a property of a specified objective or technology under conditions, not a synonym for any positive regression coefficient. The proposed equation is a way to state a question, not a completed econometric design.

[50]Missingness is a property of the observation regime, while censorship or redaction can also be a justified policy choice. The purpose of reporting coverage is not to require unlimited collection. It is to keep a performance claim proportional to the population actually observed and to disclose where protected or unavailable activity limits inference.

[51]The scenario does not require a single corporation or explicit coordination among dominant providers. It concerns a small number of controlling ecosystems, which may overlap or compete. Any allegation of illegal conduct would require a separate evidential and legal analysis.

[52]Sovereignty is used here as an institutional preference for independent authority over specified operations. It is not assumed to guarantee privacy, democratic accountability, or technical resilience. A locally controlled system can be highly surveillant or fragile; the actual rights and dependencies still need examination.

[53]The design pattern is derived from the paper's distinctions among records, learning, action, and control. It does not assume that individual consent alone can authorize every workplace use or that an expiry field can erase all derived influence. Those questions depend on the institutional setting and the specific technical transformation.

[54]Contestability does not require every participant to edit every record or model. It requires a meaningful route for relevant evidence and objections to reach an accountable decision process. The appropriate access level can vary with confidentiality, safety, and the rights of other participants.

[55]The schema model uses covariant instances $\mathcal{C} \to \mathsf{Set}$, following a common categorical database convention. The visibility model uses contravariant presheaves on a context category. Both conventions are valid in their respective settings. Reversing one without adjusting the interpretation would reverse the intended direction of information.

[56]A fibration or indexed-category presentation records coherent dependence on context. It does not guarantee that a real permission service enforces the fibers. A deployment still

needs a correspondence between policy decisions, executable checks, and recorded outcomes. This is the gap between a semantic specification and operational assurance.

[57]Writing a record and changing a model are different effects. A writer monad describes accumulation under a chosen monoid operation. A learning update generally changes state and depends on data selection, optimization, and deployment. Combining these effects requires a specified semantics; arbitrary monads do not automatically commute or admit a distributive law.

[58]No cryptographic hardness assumption is needed for the paper's main correction. The correction is logical: faithfulness alone does not supply a reconstruction procedure with the required domain, resource bound, environmental access, and authorization. Some special faithful encodings do have efficient inverses; that is additional useful structure, not a contradiction.

[59]Blackwell's experiment comparison and Fritz's Markov-category formulation motivate reasoning about information through decision and post-processing relations. The purpose-specific optimization proposed here is not quoted as a theorem of either source. Its feasibility and computational properties depend on the representation and loss function.

[60]The optimization is a conceptual formulation, not an implementable solver specification. No numerical weights, feasible architecture set, or validated privacy and risk metrics are supplied. Its purpose is to display choices that a one-variable cost comparison would omit.

[61]The plotting code is included for reproducible exposition only. It does not perform finite model checking, certify a theorem, or evaluate a real organization's architecture. The manuscript contains analytical examples and diagrams without a formal verification package.

[62]A valid synthesis should state both its antecedents and what it adds. The present manuscript adds a layered connection between operational learning and effective infrastructural control, together with corrections to overly strong conclusions. It does not claim to invent surveillance analysis, data-network economics, functorial databases, or higher-category theory.

[63]All URLs were checked against accessible retrieval results on 5 November 2026. A link resolving to a publisher record can remain paywalled or partly inaccessible; the audit states those limits rather than implying a full-text reading. Living documentation is cited with an access date. No search-result tokens or invented identifiers appear in the bibliography.

References

[1] R. H. Coase (1937). *The Nature of the Firm*. Economica 4(16), 386-405. https://doi.org/10.1111/j.1468-0335.1937.tb00002.x.

[2] O. E. Williamson (1971). *The Vertical Integration of Production: Market Failure Considerations*. American Economic Review 61(2), 112-123. https://competitionandappropriation.econ.ucla.edu/wp-content/uploads/sites/95/2017/08/Williamson1971.pdf.

[3] J. C. Panzar and R. D. Willig (1981). *Economies of Scope*. American Economic Review 71(2), 268-272. https://www.jstor.org/stable/1815729.

[4] S. J. Grossman and O. D. Hart (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.

[5] P. Milgrom and J. Roberts (1990). *The Economics of Modern Manufacturing: Technology, Strategy, and Organization*. American Economic Review 80(3), 511-528. https://web.stanford.edu/~milgrom/publishedarticles/The%20Economics%20of%20Modern%20Manufacturing.pdf.

[6] A. Hagiu and J. Wright (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.

[7] J.-C. Rochet and J. Tirole (2003). *Platform Competition in Two-Sided Markets*. Journal of the European Economic Association 1(4), 990-1029. https://doi.org/10.1162/154247603322493212.

[8] M. L. Katz and C. Shapiro (1986). *Technology Adoption in the Presence of Network Externalities*. Journal of Political Economy 94(4), 822-841. https://doi.org/10.1086/261409.

[9] J. Farrell and P. Klemperer (2006). *Coordination and Lock-In: Competition with Switching Costs and Network Effects*. CEPR Discussion Paper 5798; subsequently a Handbook of Industrial Organization chapter (2007). https://cepr.org/publications/dp5798.

[10] B. Holmstrom and P. Milgrom (1991). *Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design*. Journal of Law, Economics, and Organization 7, special issue, 24-52. https://web.stanford.edu/~milgrom/publishedarticles/Multitask%20Principal%20Agent.pdf.

[11] D. Acemoglu and P. Restrepo (2019). *Automation and New Tasks: How Technology Displaces and Reinstates Labor*. Journal of Economic Perspectives 33(2), 3-30. https://www.aeaweb.org/articles?id=10.1257%2Fjep.33.2.3.

[12] E. Ostrom (2010). *Beyond Markets and States: Polycentric Governance of Complex Economic Systems*. American Economic Review 100(3), 641-672. https://www.aeaweb.org/articles?id=10.1257%2Faer.100.3.641.

[13] Federal Trade Commission, Office of Technology (2025). *Partnerships Between Cloud Service Providers and AI Developers*. Staff report, January 2025; evidence through September 2024 and public information through January 2025. https://www.ftc.gov/reports/ftc-staff-report-ai-partnerships-investments-6b-study.

[14] Competition and Markets Authority (2025). *Cloud Services Market Investigation: Final Decision Report*. Decision dated 31 July 2025. https://assets.publishing.service.gov.uk/media/688b8891fdde2b8f73469544/final_decision_report_31.7.25.pdf.

[15] E. Wang and K. Hu, Reuters (2026). *Exclusive: Anthropic IPO Prospectus Lays Bare Deep Dependence on Big Tech Partners*. 29 September 2026; Reuters syndication on Investing.com. https://www.investing.com/news/stock-market-news/exclusiveanthropic-ipo-prospectus-lays-bare-deep-dependence-on-big-tech-partners-4923585.

[16] Anthropic (2024). *Powering the Next Generation of AI Development with AWS*. Company announcement, 22 November 2024. https://www.anthropic.com/news/anthropic-amazon-trainium.

[17] M. Foucault (1977). *Discipline and Punish: The Birth of the Prison*. Translated by Alan Sheridan; English edition, Vintage/Pantheon; French original 1975. https://www.penguinrandomhouse.com/books/55026/discipline-and-punish-by-michel-foucault-and-alan-sheridan/.

[18] S. Zuboff (2015). *Big Other: Surveillance Capitalism and the Prospects of an Information Civilization*. Journal of Information Technology 30, 75-89. https://doi.org/10.1057/jit.2015.5.

[19] H. Nissenbaum (2009/2010). *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press; electronic copyright record 2009, commonly cited print edition 2010. https://doi.org/10.1515/9780804772891.

[20] M. Polanyi (1966/2009). *The Tacit Dimension*. Original 1966; University of Chicago Press reprint 2009, foreword by Amartya Sen. https://press.uchicago.edu/ucp/books/book/chicago/T/bo6035368.

[21] J. C. Scott (1998). *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press; linked paperback edition 1999. https://yalebooks.yale.edu/book/9780300078152/seeing-like-a-state/.

[22] L. Suchman (2007). *Human-Machine Reconfigurations: Plans and Situated Actions*. Cambridge University Press, second edition. https://www.cambridge.org/core/books/humanmachine-reconfigurations/9D53E602BA9BB5209271460F92D00EFE.

[23] E. Riehl (2016). *Category Theory in Context*. Dover Publications; author-hosted text. https://math.jhu.edu/~eriehl/context.pdf.

[24] F. W. Lawvere (1963/2004). *Functorial Semantics of Algebraic Theories*. Doctoral dissertation, 1963; Reprints in Theory and Applications of Categories 5 (2004), 1-121. https://www.tac.mta.ca/tac/reprints/articles/5/tr5abs.html.

[25] D. I. Spivak (2012). *Functorial Data Migration*. Information and Computation 217, 31-51; arXiv:1009.1166. https://arxiv.org/abs/1009.1166.

[26] P. Schultz, D. I. Spivak, C. Vasilakopoulou, and R. Wisnesky (2016). *Algebraic Databases*. arXiv:1602.03501; cited as preprint. https://arxiv.org/abs/1602.03501.

[27] B. Fong (2015). *Decorated Cospans*. Theory and Applications of Categories 30, 1096-1120. https://www.tac.mta.ca/tac/volumes/30/33/30-33abs.html.

[28] J. C. Baez and K. Courser (2020). *Structured Cospans*. Theory and Applications of Categories 35, 1771-1822. https://www.tac.mta.ca/tac/volumes/35/48/35-48abs.html.

[29] J. C. Baez, K. Courser, and C. Vasilakopoulou (2022). *Structured versus Decorated Cospans*. Compositionality 4, article 3. https://compositionality.episciences.org/13517.

[30] B. Fong and D. I. Spivak (2019). *An Invitation to Applied Category Theory: Seven Sketches in Compositionality*. Cambridge University Press; linked preprint titled Seven Sketches in Compositionality. https://arxiv.org/abs/1803.05316.

[31] J. Lurie (2009). *Higher Topos Theory*. Princeton University Press; arXiv:math/0608040. https://arxiv.org/abs/math/0608040.

[32] J. Curry (2014). *Sheaves, Cosheaves and Applications*. Doctoral thesis, University of Pennsylvania; arXiv:1303.3255. https://arxiv.org/abs/1303.3255.

[33] T. Fritz (2020). *A Synthetic Approach to Markov Kernels, Conditional Independence and Theorems on Sufficient Statistics*. Advances in Mathematics 370, 107239; arXiv:1908.07021. https://arxiv.org/abs/1908.07021.

[34] D. Blackwell (1953). *Equivalent Comparisons of Experiments*. Annals of Mathematical Statistics 24(2), 265-272. https://doi.org/10.1214/aoms/1177729032.

[35] J. Pearl (2009). *Causality: Models, Reasoning, and Inference*. Cambridge University Press, second edition. https://doi.org/10.1017/CBO9780511803161.

[36] M. Barr and C. Wells (1985/2005). *Toposes, Triples and Theories*. Springer, 1985; Reprints in Theory and Applications of Categories 12 (2005), 1-287. https://www.tac.mta.ca/tac/reprints/articles/12/tr12abs.html.

[37] W. van der Aalst (2016). *Process Mining: Data Science in Action*. Springer, second edition. https://doi.org/10.1007/978-3-662-49851-4.

[38] OpenTelemetry Authors (2026 access). *Traces*. Official documentation; living document, accessed 5 November 2026. https://opentelemetry.io/docs/concepts/signals/traces/.

[39] L. Moreau and P. Missier, editors (2013). *PROV-DM: The PROV Data Model*. W3C Recommendation, 30 April 2013. https://www.w3.org/TR/prov-dm/.

[40] D. Hardt, editor (2012). *The OAuth 2.0 Authorization Framework*. IETF RFC 6749; cited as foundational specification, not complete current deployment guidance. https://www.rfc-editor.org/rfc/rfc6749.

[41] N. Sakimura, J. Bradley, M. Jones, B. de Medeiros, and C. Mortimore (2023). *OpenID Connect Core 1.0 Incorporating Errata Set 2*. OpenID Foundation specification, 15 December 2023. https://openid.net/specs/openid-connect-core-1_0.html.

[42] V. C. Hu et al. (2014). *Guide to Attribute Based Access Control (ABAC) Definition and Considerations*. NIST Special Publication 800-162; updated 2019. https://www.nist.gov/publications/guide-attribute-based-access-control-abac-definition-and-considerations.

[43] J. H. Saltzer and M. D. Schroeder (1975). *The Protection of Information in Computer Systems*. Proceedings of the IEEE 63(9), 1278-1308. https://web.mit.edu/Saltzer/www/publications/protection/.

[44] L. Ouyang et al. (2022). *Training Language Models to Follow Instructions with Human Feedback*. NeurIPS 2022; arXiv:2203.02155. https://arxiv.org/abs/2203.02155.

[45] R. Rafailov, A. Sharma, E. Mitchell, S. Ermon, C. D. Manning, and C. Finn (2023). *Direct Preference Optimization: Your Language Model Is Secretly a Reward Model*. NeurIPS 2023; arXiv:2305.18290. https://arxiv.org/abs/2305.18290.

[46] S. Ross, G. Gordon, and D. Bagnell (2011). *A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning*. AISTATS, PMLR 15, 627-635. https://proceedings.mlr.press/v15/ross11a.html.

[47] J. Perdomo, T. Zrnic, C. Mendler-Dunner, and M. Hardt (2020). *Performative Prediction*. ICML, PMLR 119, 7599-7609. https://proceedings.mlr.press/v119/perdomo20a.html.

[48] J. Kaplan et al. (2020). *Scaling Laws for Neural Language Models*. arXiv:2001.08361; cited as preprint. https://arxiv.org/abs/2001.08361.

[49] J. Hoffmann et al. (2022). *Training Compute-Optimal Large Language Models*. NeurIPS 2022; arXiv:2203.15556. https://arxiv.org/abs/2203.15556.

[50] P. Lewis et al. (2020). *Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks*. NeurIPS 2020; arXiv:2005.11401. https://arxiv.org/abs/2005.11401.

[51] S. Yao et al. (2023). *ReAct: Synergizing Reasoning and Acting in Language Models*. ICLR 2023; arXiv:2210.03629. https://arxiv.org/abs/2210.03629.

[52] T. Schick et al. (2023). *Toolformer: Language Models Can Teach Themselves to Use Tools*. NeurIPS 2023; arXiv:2302.04761. https://arxiv.org/abs/2302.04761.

[53] C. E. Jimenez et al. (2024). *SWE-bench: Can Language Models Resolve Real-World GitHub Issues?*. ICLR 2024; arXiv:2310.06770. https://arxiv.org/abs/2310.06770.

[54] J. Yang et al. (2024). *SWE-agent: Agent-Computer Interfaces Enable Automated Software Engineering*. NeurIPS 2024; arXiv:2405.15793. https://arxiv.org/abs/2405.15793.

[55] X. Luo et al. (2025). *Agent Lightning: Train ANY AI Agents with Reinforcement Learning*. arXiv:2508.03680; cited as preprint. https://arxiv.org/abs/2508.03680.

[56] G. Wang et al. (2023). *Voyager: An Open-Ended Embodied Agent with Large Language Models*. arXiv:2305.16291; cited as preprint. https://arxiv.org/abs/2305.16291.

[57] I. Shumailov et al. (2023/2024). *The Curse of Recursion: Training on Generated Data Makes Models Forget*. arXiv:2305.17493; related published version appears under a different title. https://arxiv.org/abs/2305.17493.

[58] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Aguera y Arcas (2017). *Communication-Efficient Learning of Deep Networks from Decentralized Data*. AISTATS, PMLR 54, 1273-1282. https://proceedings.mlr.press/v54/mcmahan17a.html.

[59] C. Dwork and A. Roth (2014). *The Algorithmic Foundations of Differential Privacy*. Foundations and Trends in Theoretical Computer Science 9(3-4), 211-407; DOI 10.1561/0400000042. https://www.cis.upenn.edu/~aaroth/Papers/privacybook.pdf.

[60] L. Zhu, Z. Liu, and S. Han (2019). *Deep Leakage from Gradients*. Advances in Neural Information Processing Systems 32. https://proceedings.neurips.cc/paper/2019/hash/60a6c4002cc7b29142def8871531281a-Abstract.html.

[61] N. Carlini et al. (2021). *Extracting Training Data from Large Language Models*. USENIX Security 2021; arXiv:2012.07805. https://arxiv.org/abs/2012.07805.

[62] T. Gebru et al. (2018/2021). *Datasheets for Datasets*. arXiv:1803.09010; Communications of the ACM 64(12), 86-92 (2021). https://arxiv.org/abs/1803.09010.

[63] M. Mitchell et al. (2019). *Model Cards for Model Reporting*. ACM FAT* 2019, 220-229; arXiv:1810.03993. https://arxiv.org/abs/1810.03993.

[64] A. Grattafiori et al. (2024). *The Llama 3 Herd of Models*. arXiv:2407.21783; cited as technical report. https://arxiv.org/abs/2407.21783.

[65] DeepSeek-AI et al. (2024/2025). *DeepSeek-V3 Technical Report*. arXiv:2412.19437; cited as technical report. https://arxiv.org/abs/2412.19437.

[66] A. Yang et al. (2025). *Qwen3 Technical Report*. arXiv:2505.09388; cited as technical report. https://arxiv.org/abs/2505.09388.

[67] European Parliament and Council (2016). *Regulation (EU) 2016/679 (General Data Protection Regulation)*. Original Official Journal text; especially Articles 5, 6, 15, 20, and 22. https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng.

[68] European Parliament and Council (2022). *Regulation (EU) 2022/1925 (Digital Markets Act)*. Original Official Journal text; especially Articles 6 and 7. https://eur-lex.europa.eu/eli/reg/2022/1925/oj/eng.

[69] European Parliament and Council (2023). *Regulation (EU) 2023/2854 (Data Act)*. Original Official Journal text; especially Chapter VI on switching between data processing services. https://eur-lex.europa.eu/eli/reg/2023/2854/oj/eng.

[70] National Institute of Standards and Technology (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. NIST AI 600-1, July 2024. https://doi.org/10.6028/NIST.AI.600-1.

[71] K. Bonawitz et al. (2017). *Practical Secure Aggregation for Privacy-Preserving Machine Learning*. ACM CCS 2017, 1175-1191. https://research.google/pubs/practical-secure-aggregation-for-privacy-preserving-machine-learning/.

[72] L. Bourtoule et al. (2019/2020). *Machine Unlearning*. arXiv:1912.03817; cited as preprint version. https://arxiv.org/abs/1912.03817.

[73] A. Ginart, M. Y. Guan, G. Valiant, and J. Zou (2019). *Making AI Forget You: Data Deletion in Machine Learning*. NeurIPS 2019; arXiv:1907.05012. https://arxiv.org/abs/1907.05012.