

The Headless Firm: How AI Reshapes Enterprise Boundaries

Tassilo Klein

Mantix

tk@mantix.cloud

Sebastian Wieczorek

Mantix

sw@mantix.cloud

Abstract

The boundary of the firm is determined by coordination cost. We argue that agentic AI induces a structural change in how coordination costs scale: in prior modular systems, integration cost grew with interaction topology ($O(n^2)$ in the number of components); in protocol-mediated agentic systems, integration cost collapses to $O(n)$ while verification scales with task throughput rather than interaction count. This shift selects for a specific organizational equilibrium—the Headless Firm—structured as an hourglass: a personalized generative interface at the top, a standardized protocol waist in the middle, and a competitive market of micro-specialized execution agents at the bottom. We formalize this claim as a coordination cost model with two falsifiable empirical predictions: (1) the marginal cost of adding an execution provider should be approximately constant in a mature hourglass ecosystem; (2) the ratio of total coordination cost to task throughput should remain stable as ecosystem size grows. We derive conditions for hourglass stability versus re-centralization and analyze implications for firm size distributions, labor markets, and software economics. The analysis predicts a domain-conditional Great Unbundling: in high knowledge-velocity domains, firm size distributions shift mass from large integrated incumbents toward micro-specialized agents and thin protocol orchestrators.

Keywords: Agentic AI, Generative UI, Headless Architecture.

1 Introduction

The boundary of the firm is set by the cost of coordination. As Coase argued, firms expand when internal direction is cheaper than market exchange, and contract when market transaction costs fall. [19] The critical variable is how coordination costs scale: whether they grow primarily with the number of interaction edges between components, or primarily with task throughput. This paper advances a specific, testable claim: agentic AI can change the scaling regime of coordination. In prior modular systems (SOA, microservices), coordination cost grew with interaction topology—interfaces, release coupling, governance dependencies—often approaching a combinatorial surface. In agentic systems, by contrast, a larger share of integration work shifts from bilateral negotiation toward protocol-mediated tool access and outcome-based verification. If that shift holds, coordination can scale primarily with task throughput rather than with the number of pairwise integration edges. When coordination scales with throughput, modular delegation becomes economically sustainable. Execution can fragment into micro-specialized agents without triggering the superlinear coordination overhead that historically made modular systems economically unviable at scale. At the same time, fragmentation at the execution layer creates a usability problem: humans cannot navigate an exploding interface surface area. The architectural response is an hourglass equilibrium: a

Conflict of Interest: The authors are founders of Mantix (mantix.cloud), a company building infrastructure in the agentic enterprise space described in this paper. This paper represents independent academic work; no commercial claims are made herein.

personalized intent layer (Generative UI, or GenUI) at the top, a thin protocol-and-policy waist that compiles intent into constrained actions, and a competitive market of vertical execution agents at the bottom.

Definition 1 (Coordination Scaling Regime). *We define a coordination regime as topology-dominated when coordination cost scales with the number of interaction edges between components, and throughput-dominated when cost scales primarily with the number of tasks executed. The Headless Firm emerges when agentic systems shift coordination from the former regime to the latter.*

We call the organizational analogue of this equilibrium the *Protocol-Coordinated Firm (Headless Firm)*: a structure in which execution is decoupled from centralized management and coordinated via standardized protocols and verification rather than hierarchical supervision. The headless firm is “headless” in the specific sense that the user-facing intent surface and the execution surface are separated by a thin coordination layer; it is not a claim that governance disappears, but that it is compressed into standard contracts, policy gates, and evaluation. Our contributions are as follows. (1) We propose a coordination cost model that formalizes the distinction between topology-dominated and throughput-dominated coordination regimes (Section 5). (2) We derive two falsifiable empirical predictions from this model. (3) We characterize stability conditions for the hourglass equilibrium and identify four distinct re-centralization pathways (Section 11). (4) We derive labor market and firm size implications conditional on knowledge velocity (Section 9). The central claim is conditional: the hourglass is the stable organizational form if and only if per-task verification cost $v(k)$ grows sublinearly in workflow width k . This decomposition is not frictionless. The history of service-oriented architectures demonstrates that modular systems fail when coordination costs scale superlinearly with interaction edges. Latency, partial failures, schema drift, identity and compliance layers, and system-of-record gravity all reintroduce integration burdens. The burden of proof therefore lies in whether agentic modularity meaningfully changes the coordination scaling law: does verification and governance remain a throughput problem, or does it re-expand into a topology problem (blame allocation, cascading failures, bespoke policy)? If the new scaling regime holds, value concentrates at the architectural edges. Protocol coordination commoditizes into a thin waist, while differentiation shifts upward into personalized intent and downward into deep vertical execution. If it fails—because protocols thicken into proprietary gatekeepers, verification becomes topology-dependent, or liability centralizes—re-centralization follows. The argument of this paper is therefore conditional but concrete: *a change in the scaling law of coordination selects for an hourglass architecture, and that architecture shifts enterprise boundaries toward a headless firm.*

2 Related Work

Theory of the firm and transaction costs. Our work builds on Coase’s [19] transaction cost theory of the firm and Williamson’s [84] asset specificity extension. Prior work has examined how digital technologies affect firm boundaries: Brynjolfsson and Hitt [11] document IT-driven decentralization; Garicano and Rossi-Hansberg [31] model knowledge hierarchies and show that communication cost reductions lead to flatter organizations. We extend this line by introducing the knowledge decay rate as a second dimension that interacts with coordination cost to determine optimal firm structure.

Modular architectures and platform economics. Baldwin and Clark [8] establish the modularity thesis; Langlois [44] extends it to dynamic capabilities. Our hourglass equilibrium is related to the layered protocol model of Akhshabi and Dovrolis [3] and to the end-to-end principle [60]. The “thick waist” risk we identify in Section 11.1 is related to platform envelopment [23] and analysis of multi-sided markets [58].

Multi-agent AI systems. Recent work on tool-augmented language models [61] and agent benchmarks [56, 46] establishes empirical baselines for agent capability. The MCP protocol [5] provides the technical substrate our protocol waists assumes. Concurrent work on agentic delegation [76] and AI operating systems [50] addresses related coordination problems at the systems level. Our contribution is to bring these systems-level developments into contact with organizational economics.

Labor and automation. Autor et al. [6] on superstar firms and Acemoglu and Restrepo [1] on automation and labor demand provide the empirical backdrop for our Section 9 analysis. Song et al. [69] on between-firm wage dispersion is directly relevant to our scalable boutique prediction.

3 The Economic and Temporal Imperative

The hourglass thesis depends on two independent but interacting forces. First, AI reduces the spatial cost of coordinating across firm boundaries, altering the traditional transaction-cost calculus. Second, accelerating knowledge decay increases the temporal cost of maintaining broad internal capabilities. Together, these forces create the economic pressure that selects for modular, protocol-mediated firms.

3.1 Transaction-cost theory

Coase’s theory of the firm explains vertical integration as a response to the costs of using markets (discovering prices, negotiating contracts, and enforcing agreements) [19]. If AI agents reduce these costs, then market-based coordination becomes more attractive, and the efficient firm boundary may shrink. Importantly, this reduction is unlikely to be uniform: stateful dependencies, governance requirements, and SoR switching costs can keep large parts of the stack effectively “integrated” even when execution is modularized. This aligns with Williamson’s analysis of asset specificity, where highly specific investments (like complex SoR state) encourage integration to prevent hold-up problems [84].

3.2 The temporal constraint: knowledge decay

While Coase explains why firms shrink when spatial coordination costs fall, he does not account for the temporal stability of the firm’s internal processes. Historically, firms integrated vertical capabilities because best practices were stable; standard operating procedures could be codified and amortized over years. Today, high-stakes domains (medicine, law, compliance, software engineering) face an accelerating burden of knowledge. As Jones (2009) demonstrates, the expanding knowledge frontier forces individuals to narrow their expertise to reach the cutting edge [41]. This necessitates a shift from individual generalists to specialized teams—or, in our framework, specialized agents. This induces a Red Queen effect [79]: the firm must run ever faster simply to maintain its current relative standing. In this high-velocity regime, the integrated firm faces a “complexity wall” [74]: the cost of constantly rewriting internal tooling to match the frontier of knowledge exceeds the cost of purchasing execution from specialized external agents who amortize that maintenance across many clients. On this view, headless structure is attractive when two gradients steepen at once: market coordination becomes cheaper to invoke, while maintaining broad internal capabilities becomes more expensive as domains churn. However, the same logic cuts both ways: AI can also reduce *internal* coordination costs—automating monitoring, reporting, and information routing—which in isolation would support *larger* firms. The key question is therefore not whether AI reduces coordination, but *which kind*: transaction costs across firm boundaries, or complexity costs within them. On the internal side, even with AI-assisted management, organizations face diseconomies of scale rooted in complexity: innovation slows, adaptation lags, and incentive alignment degrades as firm size grows [82]. This mirrors Simon’s “bounded rationality”—the observation that no central planner, human

or AI, can process all the distributed information needed to coordinate a large, fast-moving system [66]. The principal–agent problem [39] persists: ensuring that delegated agents act in the principal’s interest still requires monitoring, and those monitoring costs can offset the benefits of delegation. More fundamentally, AI does not eliminate what Hayek called the “knowledge calculation problem” [36]: decentralized systems adapt to local signals faster than any optimized monolith. As domain knowledge accelerates—medical guidelines, legal precedent, software tooling—the cost of keeping internal processes current rises superlinearly [85], while external specialists can amortize that cost across many clients.

3.3 Empirical and theoretical constraints on firm growth

Empirically, firm sizes follow heavy-tailed distributions; for the United States, Axtell reports a Zipf-like distribution over firm sizes [7].

Hypothesis 1 (Domain-Conditional Unbundling). *In sectors where (i) knowledge velocity is high, (ii) system-of-record switching costs are low, and (iii) protocol standardization is achievable, lower coordination friction shifts the firm size distribution toward smaller, specialized producers. This prediction does not apply uniformly: sectors with stable knowledge functions, high capital intensity, or strong network effects may experience consolidation rather than fragmentation. As AI agents reduce the transaction costs of external coordination, probability mass shifts from the “head” (large integrated firms) to the “tail” (micro-specialized entities). In a simple power-law view of firm sizes, this corresponds to a steepening of the exponent [29], such that the aggregate output of the long tail begins to rival that of the monolithic head.*




Figure 1: The hypothesized deformation of firm sizes (The Great Unbundling). Two distinct effects reshape the distribution: (1) **Unbundling**—falling transaction costs enable proliferation of micro-specialized verticals; (2) **Efficiency**—AI-enabled automation allows large firms to operate with fewer employees and reduced service requirements. Both mechanisms contribute to a flattened distribution, consistent with West’s analysis of hierarchical scaling constraints in organizations [82]. We provide the formal economic derivation of this shift—demonstrating why knowledge decay overrides AI’s span-of-control benefits—in Section 9.1.

The term “unbundling” has been used in enterprise software contexts to describe the shift from monolithic suites to specialized vertical solutions, a digital manifestation of Stigler’s lifecycle of vertical disintegration [71, 15]. We adopt this term but extend it to the *boundary of the firm itself*: we argue that unbundling is driven not only by the proliferation of micro-specialists (the standard narrative), but also by efficiency-driven downsizing of large incumbents—two distinct

forces that together reshape the firm size distribution. Several countervailing forces may constrain this unbundling in practice.

Platform concentration. AI infrastructure providers (foundation model vendors, cloud hyperscalers) may exhibit winner-take-most dynamics due to scale economies in training and network effects in ecosystems, potentially creating new large firms even as application-layer firms fragment.

Regulatory and capital requirements. Industries with high compliance burdens (finance, healthcare) or significant capital intensity (manufacturing, infrastructure) impose minimum viable firm sizes that resist full disaggregation.

Network effects and marketplace dynamics. Certain verticals may consolidate around dominant platforms that aggregate demand or supply, leading to bimodal rather than uniformly flattened distributions. These boundary conditions suggest that the Great Unbundling will be uneven across sectors, with the most pronounced effects in knowledge work, low-capital services, and domains where SoR lock-in is weakest. This long tail is also consistent with the rise of “gig economy” and marketplace coordination as a pre-GenAI precursor. As Sundararajan [72] observed, these platforms lowered matching costs and blurred traditional boundaries, yet they did not entirely dissolve the firm. A plausible reason is that many transaction costs remained stubbornly human (specification, monitoring, dispute resolution). If agentic systems reduce these frictions, they can unlock a more general version of marketplace coordination—the “Coasean singularity” argument [62]. Theoretically, Lucas models limits to firm growth via diminishing returns to managerial talent and increasing coordination costs [47]. These results motivate a “complexity wall”: beyond some scale, additional headcount adds disproportionate overhead rather than productive capacity.

4 Lessons from SOA and the Limits of Prior Modularity

Enterprise software has already lived through a large-scale attempt at modularization via service-oriented architectures (SOA) and microservices. That experience is a useful prior: exposing capabilities as services made ownership boundaries clearer, but it also surfaced new coordination costs that often offset the benefits, a phenomenon often described as the “microservices premium” [28].

4.1 Practical frictions in SOA

SOA deployments repeatedly surfaced the same family of integration burdens. Distributed systems costs arise because network latency, partial failures, and cascading outages shift failure modes from deterministic to probabilistic [22]. Interface drift compounds this: schemas and API contracts evolve continuously, turning backward compatibility into a permanent maintenance tax. More insidiously, services can agree on a field name while disagreeing on its meaning—a “semantic mismatch” that reveals the limits of pure syntactic coordination [25]. Finally, cross-service invariants require distributed transactions or compensating workflows, and identity, logging, and tracing tend to re-centralize into a platform layer, recreating coordination overhead under a new name. A particularly instructive failure is the “Smart Pipe” trap. Previous orchestration efforts—most notably the Enterprise Service Bus (ESB)—collapsed because they attempted to resolve coordination *syntactically*: rigid canonical schemas and fragile XSLT transformations broke whenever any endpoint changed [38]. Agentic orchestration avoids this failure mode because it coordinates *semantically*. Large language models act as fuzzy universal adapters that negotiate intent and map data across disparate schemas without brittle, hand-coded translators [16]. The hourglass architecture converts these historical lessons into design rules: the orchestrator owns versioned tool contracts and compatibility shims; shared evaluation harnesses push semantic alignment to the protocol layer; writes against the system of record are mediated through reversible, policy-gated workflows. Governance and observability remain concentrated in the thin waist rather than scattered across agents.

4.2 What agentic orchestration can overcome

Agentic systems do not repeal distributed-systems constraints, but they can substantially reduce the *human* costs of managing them. Three mechanisms are central.

Lower specification cost. Intent-first interaction turns under-specified requests into executable plans through iterative clarification, realizing Postel’s Robustness Principle [54]: agents accept ambiguous natural-language inputs while emitting conservative, structured tool calls. **Generated integration glue and model-driven adapter maintenance.** Agents can synthesize connectors, schema mappings, and migration scripts, and update them autonomously as APIs drift. Benchmarks such as ToolBench [56] confirm that tool selection and schema mapping can be learned and evaluated at scale, reframing integration as an engineering problem rather than an open-ended negotiation. **Continuous evaluation and routing.** Orchestrators can detect breakage, propose patches, and comparatively score services using outcome-based tests—supporting a market of interchangeable executors without manual re-integration.

4.3 Sublinear Coordination Hypothesis

The central claim of this paper can be stated as a testable economic hypothesis:

Hypothesis 2 (Sublinear Coordination Efforts). *In systems where execution is delegated to AI-mediated agents, total coordination cost grows more slowly than system complexity, provided that: (i) tool interaction is standardized at a shared protocol boundary, and (ii) verification and integration are shifted from bilateral negotiation to model-driven inference evaluated through scalable, outcome-based tests.*

If integration edges are collapsed at the protocol layer ($E = \mathcal{O}(n)$) and per-task verification cost satisfies $v(k) = \mathcal{O}(k^\gamma)$ for $\gamma < 1$, then total coordination cost grows sublinearly with system complexity for fixed task distribution. This condition fails—and hourglass architecture loses its advantage—when cross-provider coupling intensity q_t grows with ecosystem size, forcing the quadratic term in equation (3) to dominate. Empirically, $\gamma < 1$ requires that evaluation harnesses, policy templates, and outcome tests are reusable across agent compositions rather than constructed independently per task.

5 A Coordination Cost Model of the Hourglass Equilibrium

The preceding sections argued qualitatively that agentic systems can alter the scaling law of coordination. We now formalize that claim in a minimal economic model. The purpose of this model is not to capture all organizational complexity, but to isolate the structural condition under which the hourglass architecture becomes economically stable.

5.1 Coordination scaling: a minimal model

We make the scaling claim explicit with a minimal cost model.

Objects. Let n denote the number of independent execution providers (agents/services/components). Over a period, let T be the number of tasks executed. Each task t touches a subset of providers; let k_t be the number of providers involved in task t (the *workflow width*). Let E denote the number of maintained integration edges—adapters, tool contracts, schema mappings, and compatibility shims—required for the system to function.

Coordination cost decomposition. Define total coordination cost over the period as

$$C = C_{\text{integration}}(E) + C_{\text{verification}}(T, \{k_t\}) + C_{\text{governance}}(T), \quad (1)$$

where $C_{\text{integration}}$ is the cost of maintaining and versioning integration edges (API contracts, schema mappings, compatibility shims); $C_{\text{verification}}$ is the cost of checking that task outputs satisfy

correctness and policy requirements across the k_t involved providers; and $C_{\text{governance}}$, which we treat as linear in throughput by assumption (audit and routing-policy checks per task, not per integration pair), scales identically in both regimes and therefore does not determine which is preferable—the key question is whether $C_{\text{integration}}$ and $C_{\text{verification}}$ escape superlinear growth.

Two coordination regimes. In topology-dominated modularity (a common SOA/microservices failure mode), integration cost is edge-driven: $C_{\text{integration}}(E) = aE$, where, with weak protocol discipline, $E = \Theta(n^2)$ due to the combinatorial pairwise-edge growth famously formalized by Brooks [10], implying a superlinear integration term. In a protocol-mediated hourglass, pairwise integration is collapsed into a shared waist so each provider integrates once to the protocol boundary: $E = \Theta(n)$, provided protocol evolution remains backward-compatible and compatibility shims are centralized at the waist. Consequently, $C_{\text{integration}}$ scales approximately linearly in ecosystem size.

Verification as the decisive term. With a protocol waist, the dominant question becomes how verification scales with workflow width. We write per-task verification as We separate per-task verification into two components:

$$v(k) = v_{\text{local}}(k) + v_{\text{coupling}}(k) \quad (2)$$

where v_{local} represents the cost of checking constraints against reusable harnesses, and v_{coupling} represents the cost of checking cross-provider invariants. The hourglass architecture requires that $v_{\text{local}}(k)$ grows slowly in workflow width k . We formalize this relying on the following condition:

Lemma 1 (Verification Reuse Condition). *Suppose tasks are evaluated using a library of reusable outcome tests and policy templates drawn from a fixed domain vocabulary. If the marginal cost of extending this library grows sublinearly with workflow width, then the local per-task verification cost satisfies $v_{\text{local}}(k) = \mathcal{O}(k^\gamma)$ for some $\gamma < 1$.*

This condition emerges when evaluation artifacts are shared across workflows, so extending the library amortizes across many compositions rather than scaling with each additional provider. The exponent γ can be interpreted as the inverse reuse factor of verification artifacts: when $\gamma \approx 0$, a single reusable test harness suffices; when $\gamma \rightarrow 1$, each additional provider requires proportionally more verification work. The condition $\gamma < 1$ is an essential driver of the hourglass architecture’s efficiency. Intuitively, the hourglass architecture is efficient when most correctness checks can be applied independently to individual agents using reusable evaluation artifacts, and only a minority of tasks require enforcing tight consistency across multiple providers. Verification becomes topology-like again when correctness depends on cross-provider invariants (v_{coupling}). Let $q_t \in [0, 1]$ denote *coupling intensity* for task t (e.g., the probability that task t triggers cross-provider reconciliation). A minimal way to express total verification cost over T tasks is:

$$C_{\text{verification}} = \sum_{t=1}^T (bv_{\text{local}}(k_t) + dv_{\text{coupling}}(k_t)) = \sum_{t=1}^T (bk_t^\gamma + dq_t k_t^2), \quad (3)$$

The quadratic form of the coupling term follows from the assumption that cross-provider invariant checking requires pairwise consistency verification between each pair of coupled providers involved in a task. If task t triggers coupling between $m_t \leq k_t$ providers, and each pair requires independent reconciliation, verification work scales as $\mathcal{O}(m_t^2)$. We simplify by letting $q_t = m_t/k_t$ denote the fraction of involved providers that are coupled, giving a coupling term proportional to $q_t \cdot k_t^2$. The quadratic form represents a worst-case bound when invariants require pairwise reconciliation; hierarchical or sampled verification schemes reduce the exponent but do not eliminate topology dependence.

Proposition 1 (Hourglass Instability). *If coupling intensity grows proportionally with ecosystem size, i.e., $q_t = \Theta(n)$, then verification cost becomes quadratic and the hourglass architecture loses its coordination advantage.*

Scope condition. This instability proposition sharply bounds the scope of the theory. We predict hourglass stability in read-heavy or reversible-write workflows where coupling intensity remains low. Conversely, the hourglass model fails in domains requiring tightly coupled state machines, safety-critical control systems, or high-invariant core ledgers, where coupling forces topology-dependent verification.

5.2 Observability as a Precondition

The coordination cost model in Section 5.1 treats workflows as machine-addressable objects: it assumes that n , T , and k_t are measurable, that integration edges E can be maintained, and that verification harnesses can be applied to task outputs. These assumptions are not automatically satisfied. In practice, the majority of enterprise processes exist only implicitly — encoded in UI interaction sequences, human memory, and undocumented branching logic that no API exposes. Process observability is therefore not a downstream engineering concern but a logical precondition of the model itself: without a machine-readable representation of how workflows actually operate, the coordination cost decomposition in equation (1) cannot be instantiated. This gap is structural rather than incidental. Enterprise applications were designed to expose data models, not process topology. The reachable state space as traversed in practice — including exception paths, permission-conditioned branches, and context-dependent defaults — is systematically absent from any API-level representation. For the hourglass model to be instantiated, any process observability primitive must satisfy three abstract conditions: (i) topological completeness — the reachable state space as actually traversed must be represented, not only the nominal path; (ii) semantic interpretability — transitions must carry intent, not merely syntactic descriptors, so the orchestration layer can reason about what a step means rather than only what it does; and (iii) distributional awareness — traversal frequencies must be captured, so that normal-flow paths are distinguishable from exception cases. This last condition is necessary for the verification reuse condition in Lemma 1: reusable harnesses presuppose sufficient knowledge of workflow structure to identify which compositions are common and which are exceptional. These conditions cannot be satisfied at the API layer alone; the delta between what an application exposes programmatically and how users actually traverse it is systematically absent from any server-side representation. The gap these conditions expose defines a necessary infrastructure layer that is not yet standardized — analogous to the role distributed tracing played as a precondition for microservices observability. Until such a layer exists, the coordination cost decomposition in equation (1) remains a theoretical construct rather than an operationalizable framework.

5.3 Falsifiable predictions

The model yields two empirical predictions that can be evaluated with production telemetry and engineering effort accounting.

Prediction 1 (Edge-addition marginal cost). *Let $\Delta C_{\text{integration}}$ be the incremental integration work required to add one new execution provider. Without a protocol waist, adding a provider induces new pairwise edges ($\Delta E = \Theta(n)$), so $\Delta C_{\text{integration}} \propto n$. With a protocol waist, the new provider integrates once to the protocol boundary ($\Delta E = \Theta(1)$), so $\Delta C_{\text{integration}} \approx \text{constant}$.*

Prediction 2 (Verification stability under ecosystem growth). *In a stable hourglass, C/T should remain approximately stable or grow slowly with n —not proportionally to n or n^2 . Conversely, if coupling intensity q rises with ecosystem size (e.g., due to stateful cross-provider invariants), then*

(3) predicts that C/T will increase sharply with n via the k_t^2 term, signaling hourglass collapse and re-centralization pressure.

In the agentic setting, n is the number of vertical agents, k_t is the workflow width of a single user task, E represents maintained tool contracts and schema mappings, and q_t reflects the density of cross-agent invariants and write-authority constraints. The model provides a selection condition: the hourglass is stable when $E = \Theta(n)$ and q_t remains low; it collapses toward re-centralization when coupling forces the quadratic term in (3) to dominate. This selection condition is the economic logic behind Sections 6 and 7: vertical specialization achieves low γ by staying within a narrow domain; generative intent keeps k_t manageable by compiling user goals rather than exposing raw agent interfaces.

5.4 What agents do not overcome (and may worsen)

Several frictions remain binding:

- **System-of-record lock-in:** Switching SoR is costly due to migration, reconciliation, downstream dependencies, and audit risk. This mirrors Moore’s distinction between systems of record (which resist change) and systems of engagement (which require fluidity) [51]. However, the monolithic nature of the SoR is increasingly contested by decentralized data architectures:
 - **Composable ERP:** Frameworks emphasizing the decoupling of core business logic into modular, best-of-breed microservices rather than relying on vendor-locked suites [43].
 - **Zero-copy Federation:** Open table formats (e.g., Apache Iceberg) and data lakehouses [14] attempt to decouple canonical data state from proprietary application logic.

Historically, these unbundling initiatives have stalled due to severe integration complexity, network latency during distributed joins, and failures in scaling manual data governance [32]. In the context of the headless firm, agentic orchestration offers a novel mechanism to overcome these historic frictions: agents can automate semantic alignment across fragmented microservices and autonomously monitor data quality across distributed environments, potentially enabling the modularization of the SoR itself.

- **Write authority and liability:** delegating writes requires robust authorization, approvals, provenance, and replayability. This creates a classic principal-agent problem where the agent’s objective function may not perfectly align with the principal’s intent, a challenge framed in economics as incomplete contracting [34, 33].
- **Tail risks and debugging:** stochastic behavior shifts failure from “wrong response code” to “wrong outcome,” raising the burden of monitoring and dispute resolution. These failures often manifest as negative side effects or reward hacking [4, 87, 46, 21], requiring scalable oversight mechanisms beyond simple unit tests.

To manage these stochastic tail risks programmatically, the orchestration layer requires robust evaluation harnesses. Current benchmarks rely on fixed app environments, obscuring how reliability degrades under real-world variation: empirical work shows task success fluctuating by more than 50% across app design and content variations, with failure modes such as looping and hallucinating actions varying strongly by environment [78].

5.5 Three illustrative configurations

The stability condition derived above depends critically on the intensity of cross-provider coupling and the rigidity of the underlying system of record. To make this dependency concrete, we outline three stylized configurations that vary in state density and invariant strictness. These cases illustrate where the hourglass is robust, conditionally stable, or structurally constrained.

Hourglass with a dominant SoR (swappable execution). A CRM or ERP remains the canonical record, but many specialized agents (e.g., enrichment, outreach drafting, ticket triage) can be swapped because they mainly *read* from the SoR and propose changes that are mediated by the orchestrator via human-in-the-loop or policy-gated writes.

Hourglass in low-SoR-density work (high substitutability). In research, analytics, and content workflows, the “record” is mostly documents and versioned artifacts. Execution providers are easily substituted because state is shallow and outputs are highly reviewable.

Limited impact (SoR and invariants dominate). For core ledgers and high-integrity transaction processing (e.g., settlement), correctness is defined by strict invariants, regulated retention, and deterministic reconciliation. Agents may assist at the margins, but the architecture remains SoR-centered unless the underlying data infrastructure transitions to a federated, zero-copy model [14] where agents maintain cross-ledger semantic consistency.

5.6 Empirical protocol

The predictions in Section 5.3 can be evaluated using production telemetry from multi-agent deployments. For Prediction 1, the relevant measurement is the engineering time required to integrate a new execution provider as a function of current ecosystem size n . In a topology-dominated system, this should grow proportionally with n ; in a protocol-mediated system, it should be approximately constant. This can be estimated via developer effort logs (e.g., commit histories, ticket cycle times) in organizations running both legacy SOA and agentic middleware stacks. For Prediction 2, the relevant measurement is C/T —total coordination-related engineering cost divided by task throughput—tracked longitudinally as the agent ecosystem grows. Confounds include changes in task complexity distribution $\{k_t\}$ over time and organizational learning effects. A controlled comparison between organizations adopting protocol-mediated architectures versus those extending bespoke integrations would provide the cleanest test. We leave full empirical validation to future work; the purpose of these predictions is to establish that the model is falsifiable in principle.

5.7 Early Empirical Signals

While a formal, large-sample test of the coordination cost model awaits longer-term enterprise data, early signals from production agentic deployments align with the sublinear scaling hypothesis. For instance, integration times for new, highly specific tools using the Model Context Protocol (MCP) frequently approach $\mathcal{O}(1)$ developer effort compared to building bespoke API clients [5]. Furthermore, the ability to swap specialized agent backends (e.g., changing a search provider or a drafting model) without altering the user-facing intent surface provides anecdotal confirmation that the hourglass waist effectively isolates the topology of the execution layer from the throughput of user requests.

6 The Execution Layer: Vertical Intelligence

When coordination scales sublinearly, the boundary of the firm shifts. As external coordination becomes cheaper than maintaining broad internal capabilities, execution rationally migrates outward into specialized agents. Modularity ceases to be an architectural preference and becomes an economic sorting mechanism.

Historically, SOA decomposed software without decomposing ownership. Agentic systems differ because inference-driven coordination lowers the cost of integrating independent execution units. As integration frictions decline, the efficient unit of production becomes narrower.

6.1 Vertical Agents and Workflow Ownership

Vertical agents are AI-enabled services that specialize deeply within a single workflow. They do not span many business functions; they integrate tightly within one. Unlike horizontal suites, which

generalize across domains, vertical agents optimize for local objectives, proprietary context, and execution policies. Their advantage is not the base model. It is the coupling of domain data, evaluation metrics, and tool execution inside a narrow operational loop. As coordination costs fall, specialization deepens. This is the execution-layer analogue of modular architecture: specialization without integration penalty.

6.2 Defensibility Through Contextual Accumulation

A common objection is that commoditized foundation models eliminate durable advantage. This assumes advantage resides in model weights. In vertical systems, it resides in accumulated context. Enterprise work generates large volumes of unstructured “dark data” — emails, PDFs, chat threads, negotiation traces [30]. The strategic value of accumulated operational context follows the logic of the Resource-Based View [9] but at the data rather than capability level—a distinction developed in the literature on data network effects [49] and proprietary training data moats [77]. A vertical agent embedded in a workflow can capture this data, convert it into evaluation sets and feedback loops, and continuously refine execution policies. Over time, the system learns why specific contracts fail, which clauses delay negotiation, which client behaviors predict churn. This produces a contextual data network effect: performance improves with embedded use, and replication requires access to proprietary operational traces [17]. In the language of the Resource-Based View (RBV), the agent’s accumulated context becomes a VRIN (Valuable, Rare, Inimitable, Non-substitutable) asset, whereas the underlying model weights remain a commodity [9]. Crucially, this dynamic reinforces micro-specialization: dark data is highly context-specific (jargon, exception handling, local norms). Generalist systems tend to regress to the mean because they cannot justify the context windows, tuning effort, and workflow instrumentation required to decode a niche vertical’s unstructured traces. Recent work on “task interference” offers a mechanistic complement: Wang et al. show that training a single model across disparate tasks can induce negative transfer, while mixture-of-experts (MoE) architectures [64] limit interference by routing inputs to specialized subnetworks [65]. The Headless Firm can be read as the organizational analogue of this MoE pattern. Vertical agents can therefore over-optimize for local context and convert dark data into reliable execution. Finally, this architecture can drive a systemic increase in software quality. In a monolithic suite, features compete for internal resources, often resulting in “good enough” mediocrity across the board. In a market of vertical agents, each component must compete on the quality of its specific output to survive. As argued in analyses of modular architectures by Baldwin and Clark, unbundled designs allow individual modules to innovate at their own distinct rates, avoiding the regression-to-the-mean dynamics typical of integrated systems [8].

7 The Interaction Layer: Generative Intent

If execution unbundles into vertical agents, interaction must rebundle at the point of use. Specialization lowers coordination cost between machines, but it raises coordination cost for humans: as the execution layer fragments, the user-facing surface area expands. Few users can productively orchestrate dozens of agent-specific interfaces without reintroducing friction at the boundary.

Generative UI is the architectural response to this coordination overhead induced by UI surface fragmentation. The intent layer does not merely “chat”; it utilizes structured primitives to compile user intent into a constrained interaction plan, orchestrate minimal surfaces to resolve missing parameters, and render auditable confirmations for irreversible actions. The objective is not maximal flexibility but *bounded malleability*: a stable interaction frame within which task-specific micro-interfaces can be composed and discarded.

7.1 Bounded Malleability and Abstract Primitives

An engineered intent layer must preserve coherence under heterogeneity. As tasks vary across domains and agents, the interface adapts—but not arbitrarily. To achieve this, Generative UI relies on four abstract primitives:

1. **Intent Compiler:** Parses natural language into a structured, parameterized interaction plan.
2. **Constraint Validator:** Checks proposed actions against predefined policy rails and enterprise governance rules before execution.
3. **Reversible Preview:** Renders the expected state transition (e.g., a "dry run" of a database write) in a legible format.
4. **Policy Surface:** Exposes the underlying permissions and auditing logs contextually to the user.

Persistent anchors provide continuity: visible state, permissions, activity history, and a legible preview of "what will happen." Around these anchors, the system synthesizes ephemeral widgets only when needed to clarify ambiguity or collect missing inputs. Intent, in this architecture, is compiled rather than loosely interpreted. Natural-language goals are transformed into structured execution artifacts that downstream orchestration can validate and constrain. Where intent is underspecified, the system requests only the information necessary to produce a safe plan. Where actions are irreversible, it presents a confirmation surface that exposes scope, affected systems, and expected state transitions before execution proceeds. The interaction layer therefore becomes a coordination interface rather than a conversational veneer. It rebundles experience while preserving modular execution.

7.2 Trust Semantics and Coordination Scalability

Because agentic execution is stochastic and toolchains can degrade, the interaction layer must make delegation inspectable. Users must be able to see what systems will be touched, what records will change, and which policy context governs the action. Intermediate states, partial completions, and exceptions must be rendered explicitly rather than hidden behind conversational abstraction. In this sense, Generative UI functions as a contract boundary between intent and execution. It binds user goals to downstream tool invocation under explicit constraints, transforming open-ended language into legible state transitions. Crucially, this layer must itself be evaluated as a coordination mechanism. The relevant question is not aesthetic quality but structural scalability: does it reduce clarification cycles as execution complexity grows? Does it keep irreversible actions legible under variation? Does it remain stable under tool-schema change and policy drift? If interface variability scales with ecosystem diversity, cognitive load becomes topology-dependent and the hourglass collapses from above. Recent work on generative interfaces and agentic HCI [45, 81] motivates these architectural primitives, while advances in dynamic interface synthesis and test-time adaptation suggest that interaction can be evaluated systematically rather than treated as an aesthetic layer [18, 86]. The interaction layer thus mirrors the logic of the execution layer. A complete evaluation framework for the generative intent layer would require at minimum three metrics: (1) clarification cycle reduction rate—the ratio of task completions requiring zero or one clarification exchange versus baseline unaugmented interface; (2) irreversible action legibility—user accuracy in predicting post-action state under time pressure; (3) cognitive load stability under tool-schema drift—whether error rates remain bounded as the underlying agent ecosystem evolves. We leave empirical measurement of these metrics to future work but propose them as the minimal sufficient evaluation protocol for any production deployment of generative intent architectures. Just as protocol standardization compresses integration edges among agents, generative intent compresses interaction edges for users—provided that bounded malleability and explicit confirmation semantics prevent cognitive load from reintroducing topology-driven coordination costs.

8 The Hourglass Architecture of the “Headless” Firm

The preceding sections derived two structural shifts. First, execution unbundles into a competitive market of vertical agents once coordination scales sublinearly. Second, interaction must rebundle through generative intent to prevent fragmentation at the user boundary.

Taken together, these forces imply a specific architectural equilibrium. Execution fragments downward; interaction coheres upward; coordination compresses in between. The resulting structure is an hourglass, shown in Figure 2.

Wide specialization at the bottom and wide personalization at the top are connected by a thin protocol waist that standardizes coordination.




Figure 2: Convergence toward the hourglass equilibrium. Transitional configurations concentrate value in a single architectural layer (top-heavy: brand and UX; middle-heavy: orchestration logic; bottom-heavy: domain expertise). The stable end-state emerges when protocol standardization commoditizes orchestration, pushing differentiation toward the high-context edges: personalized intent above and deep domain execution below. This architectural evolution mirrors the layered convergence of Internet protocol stacks documented by Akhshabi and Dovrolis [3]. The stability condition — that the protocol waist remain thin — is formalized as the condition $\gamma < 1$ in Section 5.

Adapting the layered protocol model of [3], the hourglass consists of three layers:

Top (Generative UI) — Solves the coordination overhead induced by UI surface fragmentation. A dynamic interaction layer captures user intent and synthesizes ephemeral interfaces. Rather than navigating dozens of specialized agents, users interact with a unified intent surface. Interaction is rebundled even as execution is unbundled.

Middle (Protocol Orchestrator) — Solves the verification and governance overhead of delegated execution. A thin protocol layer compiles intent into structured commands and enforces safety, provenance, and policy-gated writes. Unlike a passive gateway, the orchestrator maintains a consistent view of state despite stochastic agent behavior [50]. Coordination is compressed into standardized contracts.

Bottom (Vertical Agent Market) — Solves the overhead of maintaining internal capabilities in high-velocity knowledge domains. A competitive ecosystem of micro-specialized agents executes tasks. Because these agents are decoupled from UI and orchestration logic, they can evolve at the speed of their domain without destabilizing the broader system. The architecture is structurally asymmetric. Differentiation concentrates at the edges—personalized interaction above and deep domain expertise below—while the middle collapses into a thin, standardized protocol waist. As Figure 2 illustrates, transitional configurations concentrate value in a single layer (top-heavy, middle-heavy, or bottom-heavy). The stable end-state emerges when protocol standardization com-

moditizes orchestration and competitive agent markets commoditize execution, forcing durable advantage toward the high-context edges. The “Headless” Firm is therefore not defined by the absence of structure, but by the separation of interaction and execution through a thin coordination layer.

8.1 Transitional Patterns and Convergence

Current systems exhibit unstable intermediate configurations in which value concentrates in a single layer. *Top-heavy* systems (e.g., AI wrappers) concentrate value in brand and UX, while execution relies on commodity APIs. *Middle-heavy* systems (e.g., proprietary orchestrators such as Zapier) concentrate value in coordination logic. *Bottom-heavy* systems (e.g., vertical AI specialists) concentrate value in domain expertise and proprietary data. These are transitional equilibria. As protocol standards mature and agent markets deepen, coordination logic commoditizes and execution becomes interchangeable. Differentiation migrates to the high-context edges: personalized interaction at the top and deep specialization at the bottom. The hourglass is therefore not a design preference but a convergence outcome.

8.2 Coordination scalability: the manager-to-managed inversion

The hourglass architecture is not merely a technical design pattern—it represents a fundamental inversion of the traditional management pyramid. The hourglass inverts the relationship between hierarchical control and transaction costs (as discussed in Section 2). Instead of imperative micro-management (“execute task X, then Y, then report status”), the orchestrator operates declaratively: it routes intent to capable agents, enforces policy constraints, and verifies outcomes. The orchestrator does not coordinate *how* agents execute—only *what* they must satisfy (safety, provenance, state consistency). This shifts the coordination model from hierarchical command-and-control to market-based outcome verification [48]. First, standardized protocols eliminate bespoke integration work. Unlike microservices that require custom API contracts and versioning coordination, agents in a mature hourglass ecosystem communicate through shared protocols (e.g., Model Context Protocol, standardized tool schemas), reducing the orchestrator’s role to routing and validation rather than negotiation. Second, outcome-based evaluation enables autonomy. The orchestrator does not monitor *process*—it verifies *results*. An agent that consistently produces correct outputs requires no supervision, analogous to the shift from time-based employment to task-based completion. Third, market mechanisms replace top-down allocation. In a competitive agent marketplace, tasks are allocated via bidding, reputation, or capability matching rather than centralized assignment [59, 62]—the orchestrator becomes an auctioneer, not a planner, aligning with Hayek’s argument that decentralized signals aggregate information more efficiently than central planning [36]. Finally, coordination shifts from static routing to intelligent delegation [76]: the dynamic transfer of authority based on real-time context, where the orchestrator negotiates delegation contracts that enforce auditability and the right to revoke authority if safety constraints are violated. This inversion has profound implications for scalability. If coordination overhead grows sublinearly (or even logarithmically) with the number of agents, then the economic ceiling on firm size shifts dramatically. A single orchestrator could theoretically manage millions of micro-specialized agents—not through heroic management, but through *protocol-mediated indifference*. The orchestrator does not care *which* legal research agent drafts a memo, only that the output meets specified quality thresholds and provenance requirements. Empirically, this predicts a bifurcation in the firm size distribution (recall Figure 1): a small number of massive orchestrators (protocol-setters like AWS, Stripe, or future MCP-compliant platforms) and a long tail of hyper-specialized execution agents. The middle—integrated firms that bundle orchestration with execution—becomes economically unstable, squeezed between the efficiency of thin protocols and the depth of vertical specialists. This mirrors the platform economics literature’s prediction that two-sided markets concentrate around

dominant intermediaries while enabling a long tail of niche suppliers [53]. This is why the hourglass is not just architecturally elegant but economically inevitable: it is the only structure that scales coordination costs sublinearly while preserving the adaptability benefits of decentralized specialization. The transitional patterns (top-heavy, middle-heavy, bottom-heavy) represent temporary market inefficiencies that will erode as protocol standardization and agent marketplaces mature.

9 Economic and Labor Implications

If the hourglass architecture is the equilibrium form under sublinear coordination scaling, then firm boundaries and software economics must adjust accordingly. We analyze three specific shifts: the dual effect of AI on firm size, the rise of the “scalable boutique,” and the inversion of the make-or-buy decision.

9.1 The Dual Effect on Firm Size

The preceding analysis raises an apparent paradox: if AI reduces coordination costs, why predict smaller firms rather than larger ones? Williamson’s 1967 model of hierarchical control suggests that reducing “loss of control” through better monitoring should enable firms to grow [83]. In his framework, firm size is limited by the degradation of information as it passes through hierarchical layers—formalized as a control transmission factor $\beta < 1$, where each additional layer multiplies the control loss. If AI increases β (by reducing information distortion through real-time dashboards and automated reporting) and expands the span of control s (by automating routine supervision), then the optimal number of hierarchical levels n could increase, supporting larger integrated firms.

Calvo and Wellisz (1978) refined this model, arguing that the nature of supervision matters: if employees cannot predict when they are being monitored (e.g., through stochastic AI-driven audits), the loss of control may be even lower than Williamson’s deterministic model suggests [13]. This implies AI could enable *dramatically* larger hierarchical firms by solving the principal-agent problem at scale. Empirical evidence supports this mechanism: Brynjolfsson et al. (2023) find that generative AI tools increase worker productivity by 14% on average (and 34% for novice workers) by disseminating best practices and reducing information asymmetry [12]. Murray et al. (2021) document how algorithmic management expands the span of control by automating routine supervision, enabling managers to oversee larger teams without degradation in oversight quality [52]. Puranam et al. (2014) observe that digital communication technologies enable flatter organizational structures by reducing the need for middle management layers [55]. In the extreme case, decentralized autonomous organizations (DAOs) replace hierarchical control entirely with smart contracts, theoretically achieving $\beta \approx 1$ through protocol-enforced transparency [35].

However, this analysis assumes a critical constraint: that the *knowledge embodied in the firm’s processes is stable*. Williamson’s model treats the firm’s production function as fixed; the challenge is purely one of coordination and monitoring. In this regime, AI’s reduction of internal coordination costs unambiguously favors larger integrated firms, as the benefits of economies of scale and scope compound without the offsetting cost of constant adaptation.

The hourglass architecture emerges when we relax this assumption. In high-velocity knowledge domains—medicine, law, compliance, software engineering—the firm faces a *temporal constraint* that Williamson’s model does not capture: the decay rate of codified knowledge [41]. Internal tooling, standard operating procedures, and domain-specific models become obsolete faster than they can be maintained. This introduces a “complexity wall”: the cost of continuously rewriting internal systems to track the frontier of knowledge scales superlinearly with firm size, because larger firms have more legacy systems, more interdependencies, and more organizational inertia [37, 74]. AI’s effect on firm size is therefore *asymmetric*. It reduces *both* internal coordination costs (higher β , larger s , fewer n) *and* external transaction costs (cheaper market-based coordination via agent marketplaces and

standardized protocols). The net effect depends on which dominates. We argue that in high-velocity domains, the external transaction cost reduction dominates for four reasons. First, the firm’s internal AI investment faces a compounding disadvantage: foundation model capabilities advance faster than any single firm can fine-tune against, meaning internal AI tooling depreciates toward the frontier at an accelerating rate. External specialists who serve many clients can amortize continuous retraining and evaluation costs; integrated firms cannot. This is a specific instance of the general principle that adaptation costs scale with the breadth of the capability portfolio maintained internally [37]. Second, a vertical specialist (e.g., a legal research agent) can amortize the cost of tracking regulatory changes, updating domain models, and refining evaluation harnesses across hundreds of clients—an integrated firm must bear this cost internally for each domain it internalizes, creating diseconomies of scope in adaptation. Third, external specialists can update their models and processes independently, without coordinating release cycles or regression-testing against the broader system; even an AI-optimized monolith must centrally plan updates, creating bottlenecks [36]. Fourth, in a market of vertical agents, poor performers are replaced instantly thanks to low switching costs—in an integrated firm, underperforming internal teams face organizational friction (politics, sunk costs, career concerns) that slows the feedback loop. Empirically, this predicts a *conditional* effect: AI will enable larger firms in domains with stable knowledge (e.g., logistics, manufacturing, routine transaction processing), where Williamson’s mechanism dominates. But in high-velocity domains, AI will accelerate unbundling, as the transaction cost collapse outpaces the control loss reduction. The hourglass is the equilibrium architecture for the latter regime. This resolves the paradox: AI does not uniformly shrink or grow firms. It *bifurcates* the firm size distribution based on the knowledge decay rate of the domain. Slow-moving domains consolidate; fast-moving domains fragment. The Great Unbundling is domain-specific, not universal.

9.2 The Scalable Boutique: Return on Talent vs. Return on Scale

This bifurcation carries an asymmetric wage implication that operates primarily in fast-moving domains: precisely where unbundling dominates, small AI-augmented firms can pay *higher* salaries than large integrated incumbents. However, this dynamic is strictly conditional on surplus capture: scalable boutiques can only pay a wage premium if they capture the surplus generated by AI leverage at the orchestration and specialization layers, rather than having it extracted via platform rents or inference pricing by foundation model providers. Assuming this condition holds, while in slow-moving domains where consolidation prevails the traditional large-firm wage premium likely persists, empirical evidence suggests this premium is broadly eroding, and in some high-tech domains, reversing entirely.

The mechanism is *marginal productivity of labor*. In a traditional large firm, an individual employee’s contribution is diluted across thousands of colleagues and layers of bureaucracy. In a small, AI-augmented firm, each employee controls a disproportionate “leverage” of capital and technology. A lean startup using generative AI and cloud infrastructure can serve millions of users (analogous to pre-agentic high-leverage cases such as WhatsApp, which served 450 million users with 55 engineers at acquisition—a ratio that agentic architectures may extend further by replacing human coordination with protocol-mediated agent delegation), generating revenue per employee that rivals or exceeds that of 100,000-person conglomerates. Autor et al. (2020) document the rise of “superstar firms” that achieve extraordinary productivity with lean headcounts, capturing market share through technological advantage rather than scale [6].

WhatsApp had approximately 55 employees at the time of its \$19B acquisition by Facebook in February 2014, serving over 450 million monthly active users. See contemporaneous reporting: e.g., Olga Khazan, *The Atlantic*, February 19, 2014.

Song et al. (2019) provide complementary evidence: two-thirds of the rise in U.S. earnings inequality from 1978 to 2013 occurred *between* firms, not within them [69]. High-wage workers increasingly sort into high-wage firms, and these high-wage firms are not necessarily the largest—they are the most productive. This sorting is enabled by technology: cloud platforms, AI tools, and global logistics allow small firms to access the same infrastructure once reserved for giants, leveling the playing field in operational efficiency while preserving advantages in adaptability.

The hourglass architecture amplifies this dynamic. In a headless firm, a small team can orchestrate a vast network of vertical agents without bearing the overhead of a large internal workforce—creating a “scalable boutique” model that captures enterprise revenue while maintaining startup agility. Three mechanisms drive the higher compensation potential:

1. **Hyper-leveraged productivity.** A lean firm (e.g., Midjourney generating an estimated \$200M in revenue with fewer than 40 employees) achieves revenue per employee an order of magnitude above a traditional software company; even a modest marginal product payout vastly exceeds a large-firm median salary.
2. **Catastrophic retention risk.** Small firms face a “must-win” scenario for talent: losing a single key engineer is catastrophic at 10 people but negligible at 10,000, forcing above-market retention wages.
3. **Open market competition.** Unlike large firms that exercise monopsony power over employees in similar roles, small firms must compete in the open market for every hire, bidding wages upward.

Cobb and Lin (2017) document this shift empirically, showing that the large-firm wage premium has declined significantly since the 1980s, particularly in technology-intensive industries [20]—the same premium that historically justified sorting into large firms has been structurally eroded by the technology leverage available to small teams. This trend accelerates in the agentic era: as orchestration becomes commoditized and execution is delegated to external agents, the optimal firm size for maximizing *return on talent* (not return on scale) shrinks dramatically.

This has profound implications for labor markets. The Great Unbundling does not merely redistribute work from large firms to small ones—it redistributes *value capture* from shareholders to high-skilled workers. In a world where a 10-person firm can generate \$100M in revenue, the bargaining power of each employee is immense. This predicts a bifurcation in compensation: a small number of highly skilled workers in scalable boutiques will capture extraordinary wages, while workers in commoditized roles (whether in large firms or as external agents) face downward pressure from automation and competition.

9.3 The Inversion of the Make-or-Buy Decision

The economic consequences of the hourglass equilibrium are summarized in Figure 3. Panel (a) shows how inference-based systems undermine the traditional zero-marginal-cost logic of software, while Panel (b) illustrates how modular agentic ecosystems outpace tightly coupled monoliths through independent, compounding innovation.

Historically, buying SaaS was cheaper than hiring engineers to build and maintain bespoke internal tools. Agentic AI can invert this equation.

First, many enterprise offerings package general-purpose models behind seat-based pricing. In practice, customers may pay for integration, governance, and procurement convenience—but they may also face perceived “per-token” markups when model usage is embedded in premium stock-keeping units (SKUs).

Second, AI can weaken the classic “zero marginal cost” intuition of software [57]. While traditional SaaS distributes a fixed codebase at low incremental cost, AI features incur ongoing inference cost *per interaction*, pushing vendors toward usage-based pricing and making high-frequency workflows




Figure 3: Structural economic shifts in the agentic era. (a) Stylized representation of marginal cost dynamics: traditional SaaS approaches zero marginal cost at scale while agentic AI inference incurs per-token cost that scales with usage volume, following empirical inference cost curves documented in [88]. (b) Conceptual illustration of innovation rate divergence between monolithic and modular architectures, consistent with modularity theory [8] but not derived from empirical data. Both panels are schematic; empirical calibration is left to future work.

expensive at scale [27]. This pressure is one reason seat-based pricing feels brittle once a product becomes a thin wrapper around metered model calls.

Third, AI coding assistants and agentic development workflows reduce the time-to-build for narrow tools. This enables *disposable software* [75]: lightweight, purpose-built scripts or apps created quickly for a temporary need and discarded or rewritten as requirements change, rather than committing to a complex long-lived platform. Crucially, this lowers the barrier to exit as much as the barrier to entry. Because the sunk cost of these ephemeral agents is near-zero, firms can switch providers instantly without the “lock-in” friction and high switching costs that classically define legacy enterprise software markets [63], forcing agents to compete continuously on performance.

Recent analyses of inference economics [88] indicate that while training costs are one-time capital expenditures (CapEx), inference costs behave as variable operating expenses (OpEx) that scale linearly with usage. Unlike legacy SaaS, where the marginal cost of serving the N^{th} user approaches zero, the marginal cost of an agentic workflow remains significant. This pushes pricing away from flat-rate subscriptions (which bleed margin under heavy use) toward outcome-based or consumption-based models.

This shift marks an economic transition from “SaaS” (selling tools to humans) to “Service-as-Software” (selling outcomes to firms). While traditional SaaS competes for a relatively constrained IT budget, vertical agents compete for the much larger payroll and outsourced-services budget. As argued by HFS Research [80], this creates a non-linear growth opportunity by expanding the addressable market from software spend to services spend.

9.4 The Democratization of the Stack

The hourglass architecture creates two symmetric entry points (Figure 2). At the **frontend entry**, a small team can build a compelling generative UI by orchestrating best-in-class agents rather than recreating backend capabilities. At the **backend entry**, a domain expert can build a vertical agent by encoding specialized process knowledge, relying on the ecosystem for distribution and interaction.

Because coordination is standardized at the waist, neither side must internalize the entire stack. This operationalizes a strategic form of Inverse Conway’s Law [67]: firms can adopt a headless architecture to induce modular, high-velocity organizational structure.

This aligns with the “Software 2.0” paradigm [42]. Scarcity shifts from writing explicit code toward curating data, defining objectives, and constructing evaluation harnesses. In a headless ecosystem, builders compete on semantics rather than syntax.

10 Open Problems

The central open question is whether the relocated coordination layer—verification, governance, and market allocation—can scale with throughput rather than topology in practice. Each item below represents a mechanism that could tip the balance toward re-centralization if left unsolved.

Incentive design and market structure. Agent marketplaces require mechanisms for reputation, routing, pricing, and capability discovery. The open question is how to prevent race-to-the-bottom equilibria while preserving meaningful competition and specialization.

Verification, liability, and auditability. As execution is delegated to autonomous agents, the primary friction shifts from vendor selection to behavioral verification. This motivates research in evaluation harnesses, persistent agent identity, decentralized identity (DID) systems [70], provenance tracking, and liability allocation under stochastic execution.

Protocol standardization. The thin waist depends on shared tool contracts, semantic interoperability, and agent-interface specifications (e.g., Model Context Protocol [5]). Without standards, coordination costs re-fragment into bespoke integrations.

Data governance and vertical integration. Vertical agents accumulate high-context operational traces. Ensuring privacy, regulatory compliance, and controlled reuse of proprietary workflow data is foundational for defensible specialization without unacceptable leakage.

State integrity and composable systems of record. Multi-agent ecosystems must preserve invariants across composable ERPs and zero-copy architectures. Whether agents can safely mediate irreversible writes and high-integrity state transitions remains an open empirical question.

Robustness under distribution shift. Agent performance can degrade under changing tools, policies, or adversarial inputs. Understanding cascading failures in multi-agent toolchains and designing graceful degradation remains open.

11 Critical Discussion: The Risks of Re-Centralization

If the hourglass architecture is the equilibrium under sublinear coordination scaling, then its stability depends on preventing coordination costs from re-accumulating elsewhere in the system. Each of the following risks represents a distinct mechanism through which coordination could revert from throughput-based scaling back to topology-based scaling, reintroducing centralization pressures.

11.1 The Aggregator Paradox: The Risk of a “Thick Waist”

A central critique of the hourglass thesis is that protocol layers in digital markets rarely capture value; instead, aggregators built atop them do, a phenomenon formalized as platform envelopment [23]. If the Orchestrator controls identity, routing, reputation, or pricing, it may cease to function as a neutral protocol and instead become a rent-extracting gatekeeper — a “thick waist” that re-centralizes coordination.

This is not merely a value-capture concern but a scaling concern. A protocol remains thin only if it minimizes discretionary policy and pricing power. Once routing decisions, identity verification, or reputation scoring become proprietary and vertically integrated, coordination cost begins scaling with platform governance rather than standardized contracts.

See also Thompson (2015) for an influential practitioner account of aggregation dynamics in digital markets: <https://stratechery.com/aggregation-theory/>

Open standards such as the Model Context Protocol (MCP) [5] reduce this risk by commoditizing tool interaction before monopoly control can emerge. Historically, attempts to build “smart pipes” (e.g., ESBs) re-centralized integration logic and ultimately failed because they internalized too much coordination intelligence. As formalized by the classic end-to-end principle in system design [60], the stable design is the minimal pipe that allows maximal intelligence at the edges.

However, this outcome is contingent on protocol governance. If standards remain vendor-controlled or allow proprietary extensions, the waist can re-thicken. Sublinear coordination scaling therefore depends not just on technical design, but on institutional neutrality.

11.2 The Cognitive Load Trap: Interaction Re-Thickening

The hourglass assumes that generative interfaces reduce human coordination costs by compressing navigation into intent. Yet HCI research on Cognitive Load Theory suggests that users rely on spatial consistency and stable affordances to construct mental models [73]. A fully malleable interface risks increasing cognitive coordination cost rather than reducing it. If interface variability scales with task diversity, then human coordination burden may grow with system complexity. In that case, interaction cost re-scales with topology rather than collapsing into an intent abstraction layer. The viable design path is *bounded malleability*. Generative interfaces must preserve stable anchors (navigation rails, persistent affordances) while synthesizing ephemeral micro-interfaces around tasks. The objective is to eliminate navigation friction without eliminating cognitive predictability. If this balance fails, the interaction layer becomes a new coordination bottleneck, undermining the hourglass equilibrium.

11.3 The Jevons Paradox of Inference: The Return of Marginal Cost

Economic theory suggests that as AI inference costs drop (following a Moore’s Law trajectory), the pressure to shift from “SaaS” (flat-rate) to “Service-as-Software” [26] (outcome-based) pricing may evaporate. If intelligence becomes “too cheap to meter,” why unbundle? This view ignores the Jevons Paradox [40]. As the cost per token drops, demand for reasoning depth increases. We are witnessing a structural shift toward test-time compute scaling [68], where models allocate additional inference cycles to verify and refine chains of thought. Early “Level 2” reasoners (like OpenAI’s o1) demonstrated that performance scales not just with model parameters, but with “thinking time”—ensuring that the computational cost of high-integrity judgments remains significant. More precisely, the marginal cost MC of executing a task via an autonomous agent is not merely its inference cost:

$$MC(a) = C_{\text{compute}}(a) + \mathbb{E}[r(a) \cdot p(a)] \quad (4)$$

where $r(a)$ denotes the liability risk (the cost of an error or safety violation) of an autonomous action a , and $p(a)$ denotes the probability of an incorrect execution. Even if $p(a)$ falls with continuous model improvement and $C_{\text{compute}}(a)$ approaches zero, $r(a)$ scales directly with the value of the underlying transaction. For high-value enterprise actions (contract execution, financial settlement, regulatory filings), the expected liability term $\mathbb{E}[r(a) \cdot p(a)]$ is orders of magnitude larger than the compute cost, creating a permanent floor on the zero-marginal-cost intuition that does not respond to Moore’s Law. This formalizes the intuition in Agrawal et al. [2] within the coordination cost framework developed in Section 5. Furthermore, the true marginal cost of an agent is not compute, but *liability*. As Agrawal et al. argued, while AI commoditizes prediction, it increases the value of judgment—specifically the liability for incorrect judgments [2]. The risk of an agent executing a trade or signing a contract scales linearly with volume, creating a permanent floor on zero-marginal-cost scaling. Economically, this is an insurance premium. As agents move from “drafting” (low risk) to “executing” (high risk), the vendor must price in the cost of errors. Unlike compute, which follows Moore’s Law, insurance premiums do not halve every 18 months; they scale with the value of the transaction.

11.4 The Trust Boutique: Survival of the Middle

Finally, we predict a “Missing Middle” where mid-sized firms die. Critics argue that mid-sized firms survive on high-trust relationships that agents cannot replicate. We refine our prediction: the “Missing Middle” will not disappear, but transform into “Trust Boutiques”—specialized liability wrappers. These firms will no longer sell execution (which is commoditized by agents) but accountability (the legal and moral liability for outcomes). They survive by becoming the “moral crumple zone” [24] for algorithmic systems, selling the human handshake that an agent cannot forge. However, effectively acting as an insurance wrapper requires far fewer humans than performing the execution itself. Thus, the firm size (headcount) still unbundles, even if the firm entity survives. (For a discussion on the wage implications of this shift, see the ‘Scalable Boutique’ analysis in Section 9.2).

Synthesis

These four pathways are not fully independent. Platform capture at the waist (Section 11.1) and liability concentration in trust boutiques (Section 11.4) are particularly likely to co-occur: a thick-waist aggregator has strong incentives to also aggregate liability in order to extract rents from the execution layer, creating a self-reinforcing centralization dynamic. The remaining two pathways—cognitive load accumulation (Section 11.2) and the liability floor (Section 11.3)—are more likely to operate independently and can be mitigated at the layer level without systemic restructuring: bounded malleability addresses the former; outcome-based pricing models address the latter. The durability of the hourglass equilibrium therefore depends most critically on preventing the co-occurrence of thick-waist capture and liability aggregation, since these two forces together would reconstitute a vertically integrated intermediary under a new name. If coordination remains compressed into standardized contracts and outcome verification, the hourglass persists as the stable architectural form of the agentic economy.

12 Conclusion

For most of economic history, large firms existed because coordination was expensive. Specialists needed managers; managers needed directors; directors needed strategy. Each layer added overhead but also reduced the cost of working across complex, interdependent tasks. AI agents challenge this logic at the root: if coordination can be compressed into a shared protocol and verification can scale with task volume rather than interaction complexity, the overhead that justified large integrated firms becomes optional. The headless firm is the organizational form that follows from this change. Execution migrates outward to specialized agents; interaction coheres inward through a generative intent layer; in between sits a thin protocol waist that keeps the whole system composable. The resulting hourglass is not a design preference—it is an economic selection pressure. Firms that internalize broad capabilities in high-velocity domains face a “complexity wall” that compounds as domains accelerate. Firms that route execution through competitive agent markets face only the cost of verification and governance, which the model predicts can remain throughput-scaled. The argument is conditional. If protocol layers thicken into proprietary gatekeepers, if verification becomes entangled with cross-agent liability, or if generative interfaces amplify rather than reduce cognitive load, the hourglass collapses and reintegration follows. Ultimately, the Headless Firm is not a prediction about specific technologies or the trajectory of foundation models; it is a structural observation about economic selection pressure. As the underlying laws of coordination scaling change, the architectures that dominate the coming decade will be those that push execution to the specialized edges while compressing interaction into a generative, protocol-mediated waist.

Limitations. Our coordination cost model is stylized: the functional forms in equations (1)–(3) are chosen for tractability rather than empirical fit, and the exponent γ is asserted to be less than one under the hourglass rather than derived from first principles. The Great Unbundling hypothesis (Hypothesis 1) is stated with domain-level scope conditions but has not been tested against longitudinal firm size data. The labor market predictions in Section 9 rest on the scalable boutique mechanism, which depends on the assumption that high-leverage small firms can maintain bargaining power as agent execution commoditizes; this assumption may not hold if model providers extract value through inference pricing rather than capability differentiation. Finally, the analysis focuses on the DACH/EU enterprise context; regulatory differences (GDPR, AI Act compliance requirements) may shift the coordination cost landscape in ways not captured by our model. The open research agenda is therefore not whether the Headless Firm is the right idea—it is whether the engineering and institutional conditions that stabilize it can be built fast enough to outpace the re-centralizing forces of platform capture, liability concentration, and governance complexity.

References

- [1] Daron Acemoglu and Pascual Restrepo. Tasks, automation, and the rise in us wage inequality. *Econometrica*, 90(5):1973–2016, 2022.
- [2] Ajay Agrawal, Joshua Gans, and Avi Goldfarb. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press, Boston, MA, 2018.
- [3] Saamer Akhshabi and Constantinos Dovrolis. The evolution of layered protocol stacks leads to an hourglass-shaped architecture. In *Proceedings of the ACM SIGCOMM 2011 Conference*, pages 206–217, 2011.
- [4] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [5] Anthropic. Model context protocol: Specification and architecture. Technical report, Anthropic, 2024. <https://modelcontextprotocol.io>.
- [6] David Autor, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, 135(2):645–709, 2020.
- [7] Robert L. Axtell. Zipf distribution of U.S. firm sizes. *Science*, 293(5536):1818–1820, 2001.
- [8] Carliss Y. Baldwin and Kim B. Clark. *Design Rules: The Power of Modularity*. MIT Press, 2000.
- [9] Jay Barney. Firm resources and sustained competitive advantage. *Journal of Management*, 17(1):99–120, 1991.
- [10] Frederick P. Brooks. *The Mythical Man-Month: Essays on Software Engineering*. Addison-Wesley, Boston, MA, 1975.
- [11] Erik Brynjolfsson and Lorin M Hitt. Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4):23–48, 2000.
- [12] Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond. Generative ai at work. Working Paper 31161, National Bureau of Economic Research, 2023.

- [13] Guillermo A. Calvo and Stanislaw Wellisz. Supervision, loss of control, and the optimum size of the firm. *Journal of Political Economy*, 86(5):943–952, 1978.
- [14] Ryan Carlson et al. Apache iceberg: An architectural look under the covers. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, 2024.
- [15] Martin Casado and Matt Bornstein. The great unbundling. Andreessen Horowitz Blog, 2024.
- [16] David Chappell. *Enterprise Service Bus*. O’Reilly Media, 2004.
- [17] Jerry Chen. Castles in the cloud: The new moats. Greylock Partners Blog, 2017.
- [18] Xiang “Anthony” Chen, Tiffany Kneare, and Yang Li. The GenUI study: Exploring the design of generative UI tools to support UX practitioners and beyond, 2025.
- [19] Ronald H. Coase. The nature of the firm. *Economica*, 4(16):386–405, 1937.
- [20] J. Adam Cobb and Ken-Hou Lin. Growing apart: The changing firm-size wage premium and its inequality consequences. *Organization Science*, 28(3):429–446, 2017.
- [21] Xiang Deng et al. Mind2web: Towards a generalist agent for the web. In *NeurIPS*, 2023.
- [22] L Peter Deutsch. The eight fallacies of distributed computing. Technical report, Sun Microsystems, 1994. Failures 1 through 7 were coined in 1994; the 8th was added later by James Gosling.
- [23] Thomas Eisenmann, Geoffrey Parker, and Marshall W. Van Alstyne. Strategies for two-sided markets. *Harvard Business Review*, 84(10):92, 2006.
- [24] Madeleine Clare Elish. Moral crumple zones: Cautionary tales in human-robot interaction. *Engaging Science, Technology, and Society*, 5:40–60, 2019.
- [25] Eric Evans. *Domain-Driven Design: Tackling Complexity in the Heart of Software*. Addison-Wesley Professional, 2004.
- [26] Phil Fersht. Service-as-software: The new paradigm. HFS Research Blog, 2025.
- [27] Stanford Institute for Human-Centered AI (HAI). Artificial intelligence index report 2025, 2025.
- [28] Martin Fowler and James Lewis. Microservices: a definition of this new architectural style. *MartinFowler.com*, 2014.
- [29] Xavier Gabaix. Power laws in economics and finance. *Annual Review of Economics*, 1(1):255–294, 2009.
- [30] John Gantz and David Reinsel. The digital universe in 2020: Big data, bigger digital shadows, and biggest growth in the far east, 2012. IDC Analytical Study.
- [31] Luis Garicano and Esteban Rossi-Hansberg. Organization and inequality in a knowledge economy. *The Quarterly Journal of Economics*, 121(4):1383–1435, 2006.
- [32] Gartner. Latest enterprise resource planning (erp) insights, 2025.
- [33] Dylan Hadfield-Menell and Gillian K Hadfield. Incomplete contracting and ai alignment. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 417–422, 2019.

- [34] Oliver Hart and John Moore. Property rights and the nature of the firm. *Journal of Political Economy*, 98(6):1119–1158, 1990.
- [35] Samer Hassan and Primavera De Filippi. Decentralized autonomous organization. *Internet Policy Review*, 10(2):1–10, 2021.
- [36] Friedrich A. Hayek. The use of knowledge in society. *The American Economic Review*, 35(4):519–530, 1945.
- [37] Rebecca M. Henderson and Kim B. Clark. Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Administrative Science Quarterly*, 35(1):9–30, 1990.
- [38] Gregor Hohpe and Bobby Woolf. *Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions*. Addison-Wesley Professional, 2003.
- [39] Michael C Jensen and William H Meckling. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4):305–360, 1976.
- [40] William Stanley Jevons. *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-Mines*. Macmillan and Company, 1865.
- [41] Benjamin F. Jones. The burden of knowledge and the ”death of the renaissance man”: Is innovation getting harder? *The Review of Economic Studies*, 76(1):283–317, 2009.
- [42] Andrej Karpathy. Software 2.0. *Medium*, 2017.
- [43] Alexander Kyte et al. The future of erp is composable. Technical report, Gartner, 2020.
- [44] Richard N Langlois. Modularity in technology and organization. *Journal of Economic Behavior & Organization*, 49(1):19–37, 2002.
- [45] Yaniv Leviathan et al. Generative ui: Llms are effective ui generators. Technical report, Google Research, 2025.
- [46] Xiao Liu et al. Agentbench: Evaluating llms as agents. In *ICLR*, 2024.
- [47] Robert E. Lucas. On the size distribution of business firms. *The Bell Journal of Economics*, 9(2):508–523, 1978.
- [48] Thomas W. Malone and Kevin Crowston. The interdisciplinary study of coordination. *ACM Computing Surveys*, 26(1):87–119, 1994.
- [49] Bertin Martens et al. The digital transformation of news media and the rise of disinformation and fake news. Technical report, JRC Digital Economy Working Paper 2018-02, 2020.
- [50] Kai Mei et al. Aios: Llm agent operating system. In *Conference on Language Modeling (COLM)*, 2025.
- [51] Geoffrey Moore. Systems of engagement and the future of enterprise IT. *AIIM White Paper*, 2011.
- [52] Alex Murray, Jen Rhymer, and David G. Sirmon. Algorithms as work designers: How algorithmic management influences the design of jobs. *Human Resource Management Review*, 31(2), 2021.

- [53] Geoffrey G. Parker, Marshall W. Van Alstyne, and Sangeet Paul Choudary. *Platform Revolution: How Networked Markets Are Transforming the Economy and How to Make Them Work for You*. W. W. Norton & Company, 2016.
- [54] Jon Postel. Transmission control protocol. Technical report, RFC 793, Internet Engineering Task Force, September 1981. The Robustness Principle is formulated in Section 2.10.
- [55] Phanish Puranam, Oliver Alexy, and Markus Reitzig. What’s “new” about new forms of organizing? *Academy of Management Review*, 39(2):162–180, 2014.
- [56] Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. Toolllm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*, 2023.
- [57] Jeremy Rifkin. *The Zero Marginal Cost Society: The Internet of Things, the Collaborative Commons, and the Eclipse of Capitalism*. Palgrave Macmillan, 2014.
- [58] Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003.
- [59] Alvin E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, 2002.
- [60] Jerome H. Saltzer, David P. Reed, and David D. Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems (TOCS)*, 2(4):277–288, 1984.
- [61] Timo Schick et al. Toolformer: Language models can teach themselves to use tools. *Advances in Neural Information Processing Systems*, 2024.
- [62] Peyman Shahidi, Gili Rusak, Benjamin S. Manning, Andrey Fradkin, and John J. Horton. The coasean singularity? demand, supply, and market design with AI agents. Working Paper 34468, National Bureau of Economic Research, 2025. Presented at the 2025 Economics of Transformative AI Workshop.
- [63] Carl Shapiro and Hal R. Varian. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business Press, Boston, MA, 1998.
- [64] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarsz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer, 2017.
- [65] Leyang Shen, Gongwei Chen, Rui Shao, Weili Guan, and Liqiang Nie. Mome: Mixture of multimodal experts for generalist multimodal large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [66] Herbert A Simon. A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1):99–118, 1955.
- [67] Matthew Skelton and Manuel Pais. *Team Topologies: Organizing Business and Technology Teams for Fast Flow*. IT Revolution, 2019.

- [68] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.
- [69] Jae Song, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. Firming up inequality. *The Quarterly Journal of Economics*, 134(1):1–50, 2019.
- [70] Manu Sporny, Dave Longley, Markus Sabadello, Drummond Reed, Orie Steele, and Christopher Allen. Decentralized identifiers (dids) v1.0. Recommendation, W3C, 2022.
- [71] George J. Stigler. The division of labor is limited by the extent of the market. *Journal of Political Economy*, 59(3):185–193, 1951.
- [72] Arun Sundararajan. *The Sharing Economy: The End of Employment and the Rise of Crowd-Based Capitalism*. MIT Press, Cambridge, MA, 2016.
- [73] John Sweller. Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2):257–285, 1988.
- [74] Joseph A. Tainter. *The Collapse of Complex Societies*. Cambridge University Press, 1988.
- [75] Dave Thomas. The case for disposable software: ‘big architecture’ not always essential. In Diomidis Spinellis and Georgios Gousios, editors, *Beautiful Architecture: Leading Thinkers Reveal the Hidden Beauty in Software Design*. O’Reilly Media, 2009.
- [76] Nenad Tomašev, Matija Franklin, and Simon Osindero. Intelligent ai delegation. *arXiv preprint arXiv:2602.11865*, 2026.
- [77] Andrew Trask et al. Data moats and proprietary training sets in foundation models. *Advances in Neural Information Processing Systems*, 2023.
- [78] Karen Ullrich, Jingtong Su, et al. Openapps: Simulating environment variations to measure ui agent reliability. *arXiv preprint arXiv:2511.20766*, 2025.
- [79] Leigh Van Valen. A new evolutionary law. *Evolutionary theory*, 1:1–30, 1973.
- [80] Ashish Venkatesan and Nimit Jhunjunwala. Stop waiting for proof: Services-as-Software is driving non-linear growth. Technical report, HFS Research, 2025.
- [81] Yu Wang and Y. Lu. Agentic media: Reimagining the future of communication. Technical Report MSR-TR-2025-62, Microsoft Research, 2025.
- [82] Geoffrey West. Scale: The universal laws of growth, innovation, sustainability, and the pace of life in organisms, cities, economies, and companies, 2017.
- [83] Oliver E. Williamson. Hierarchical control and optimum firm size. *Journal of Political Economy*, 75(2):123–138, 1967.
- [84] Oliver E Williamson. The economics of organization: The transaction cost approach. *American Journal of Sociology*, 87(3):548–577, 1981.
- [85] Stefan Wuchty, Benjamin F Jones, and Brian Uzzi. The increasing dominance of teams in production of knowledge. *Science*, 316(5827):1036–1039, 2007.

- [86] Haijun Xia. Generative and malleable user interfaces. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Workshop on GenAI for UI)*, 2024.
- [87] Shuyan Zhou et al. Webarena: A realistic web environment for building autonomous agents. In *ICLR*, 2024.
- [88] Boqin Zhuang, Jiacheng Qiao, Mingqian Liu, Mingxing Yu, Ping Hong, and Rui Li. Beyond benchmarks: The economics of AI inference. *arXiv preprint arXiv:2510.26136*, 2025.