Endogenous Firm Scope via Organizational Complexity*

Sebastian Dyrda
University of Toronto

Guangbin Hong
Michigan State University

Joseph B. Steinberg
University of Toronto & NBER

March 26, 2026

Link to the current version

Abstract

Firms can scale across vast numbers of similar activities at negligible marginal cost, yet face sharp limits when expanding into even a few materially distinct environments. To rationalize this, we put forward and formalize the concept of organizational complexity—the cumulative heterogeneity of a firm's activities—and show that this single object disciplines expansion, shapes internal organization, and determines the cross-sectional distributions of firm scope and market concentration. Firms deploy a non-rival organizational blueprint across heterogeneous markets, but adapting the blueprint to each market requires scarce managerial attention. As the firm expands into more diverse activities, organizational complexity accumulates and coordination costs rise throughout the organization, making expansion decisions fundamentally non-separable: serving one market raises the cost of serving every other. The model predicts that firms grow by absorbing nearby, related activities first—generating a core–periphery internal structure—and that sufficiently advanced organizational technology gives rise to platform-type firms that can scale across many similar markets at low cost. The same mechanism compresses the firm size distribution: a firm with ten times the underlying capability achieves only about twice the scope. It also drives a wedge between local and aggregate concentration, because organizational complexity prevents any single firm from reaching across the full breadth of the economy. Calibrated to the cross-sectional distribution of firm scope among U.S. public firms, the model is quantitatively consistent with the observed patterns of related diversification, intangible-driven expansion, and stable aggregate concentration despite rising local market power.

Keywords: Firm scope; organizational complexity; intangible capital; firm size distribution; market concentration; platform firms

JEL Codes: D23; E23; L11; L22; L25; O33

*Dyrda: sebastian.dyrda@utoronto.ca; Hong: honggua2@msu.edu; Steinberg: joseph.steinberg@utoronto.ca. Address: University of Toronto, Max Gluskin House, 150 St. George Street, Toronto, ON M5S 3G7, Canada.

Large modern firms operate across an extraordinarily wide range of markets, products, and activities. Digital platforms such as Amazon, Google, or Uber simultaneously serve millions of distinct customer–product–location combinations, apparently at negligible marginal organizational cost. Yet many firms of comparable productivity and resources remain narrowly focused, and even the largest platforms exhibit sharp boundaries: Amazon dominates e-commerce logistics but has not replicated this success in healthcare; Google spans digital advertising and cloud computing but struggles in hardware. A central challenge for the theory of the firm is therefore not merely to explain why some firms become large, but to explain how firms select which activities to span when potential markets differ in their organizational requirements—and why scope appears cheaply expandable in some directions but prohibitively costly in others.

Existing theories of firm size and scope typically impose limits through fixed managerial spans (Lucas, 1978), discrete hierarchies of knowledge and control (Garicano, 2000; Garicano and Rossi-Hansberg, 2006), or per-market entry costs (Melitz, 2003). A common feature of these frameworks is that expansion decisions are effectively separable across markets: serving one activity does not alter the organizational cost of serving another. As a result, these models cannot account for the central empirical pattern—that firms expand cheaply across similar activities but face sharp limits across heterogeneous ones—because the cost of the marginal market does not depend on the composition of the firm's existing footprint. This paper proposes a different view. We argue that the binding constraint on firm scope is not the count of markets served, but the cumulative organizational complexity generated by heterogeneity across activities. Because organizational complexity aggregates heterogeneity across the entire footprint, expansion decisions become inherently non-separable: every new market alters the shadow cost of headquarters attention for all incumbent activities. Serving a new market is cheap when it is similar to what the firm already does, and expensive when it is not—even if the market itself is intrinsically profitable. Organizational complexity thus endogenously determines which markets the firm serves, not merely how many.

To formalize these ideas, we develop a model in which firms operate a centralized, non-rival blueprint that can be deployed across a continuum of heterogeneous markets. The blueprint—representing codified organizational knowledge, operational procedures, or proprietary technology—is created at headquarters and can be used simultaneously in any market the firm chooses to serve. However, effective deployment of the blueprint in a given market requires market-specific adaptation using scarce headquarters attention. Markets differ in their organizational distance from the firm's core: some are nearly identical to the environment in which the

blueprint was developed, while others require substantial customization. We summarize this variation through a heterogeneity index for each market and define organizational complexity as the integral of heterogeneity over the firm's footprint. Complexity enters the firm's problem through an organizational congestion function that raises the shadow price of headquarters attention as the firm's footprint becomes more heterogeneous. The firm therefore faces a fundamental tradeoff: the blueprint is non-rival and can be deployed at zero marginal cost in terms of replication, but each market served adds to organizational complexity and tightens the constraint on scarce headquarters resources. Scope is thus neither free nor mechanically costly—it depends on which markets the firm selects. This framework delivers four sets of results.

The first set of results establishes that optimal firm scope is bounded in organizational complexity rather than market count. Even though firms may operate across an arbitrarily large collection of activities, the total heterogeneity of those activities must remain finite at the optimum—otherwise, organizational congestion would drive the shadow cost of headquarters attention to infinity. This result reconciles two observations that appear contradictory under count-based theories: firms can scale across vast numbers of highly similar markets (as platforms do), yet the mass of materially distinct environments any firm serves remains sharply limited. Endogenous limits to scope emerge without imposing fixed per-market costs or discrete managerial constraints.

The second set of results characterizes a scope–technology tradeoff that governs expansion decisions. Holding blueprint quality fixed, adding markets raises organizational complexity and tightens internal congestion, reducing the value of incumbent activities. At the same time, because the blueprint is non-rival, serving more markets increases the marginal return to improving blueprint quality. These two forces operate in opposite directions: scope expansion simultaneously taxes existing operations and strengthens incentives to invest in organizational technology. Firm scope is therefore determined by the balance of these forces rather than by separable market-level profitability, and the profitability of any given market depends on the composition of the firm's entire footprint.

The third set of results concerns market selection and internal organization. When markets can be ordered by heterogeneity, optimal scope exhibits a cutoff structure: firms serve all markets sufficiently close to their organizational core and exclude those beyond a critical distance. This selection rule provides a theoretical foundation for the empirical finding that scope expansion occurs predominantly along related dimensions. Within the chosen footprint, scarce headquarters attention is allocated unevenly—concentrated in low-heterogeneity core activities where adaptation is most productive, while peripheral activities are operated through more standardized processes.

A core–periphery structure thus emerges endogenously from organizational scarcity, even in the absence of technological asymmetries or fixed costs.

The fourth set of results characterizes platform emergence as a limiting scalability regime. Allowing firms to invest in standardization technology does not, by itself, eliminate organizational limits: when blueprint quality is low, the firm remains in a bounded-scope regime where complexity constraints bind. However, when blueprint quality is sufficiently high and near-core market opportunities are sufficiently abundant, firms can replicate a common organizational recipe across arbitrarily many markets while keeping total complexity finite. Platform emergence is therefore a statement about the divergence of a bounded-complexity scale frontier—the largest footprint that can be operated profitably while keeping organizational complexity below a fixed cap—rather than a literal finite-parameter jump to infinite scope. Platform firms emerge as a distinct organizational form: not merely large firms, but firms whose organizational technology makes scope cheaply expandable along near-core margins.

Fifth, we trace out implications of the organizational complexity mechanism for the cross-sectional distribution of firm scope. Because the marginal market for a larger firm is more heterogeneous than the marginal market for a smaller firm, the mapping from underlying productivity to scope is concave: differences in capability translate into progressively smaller differences in scope as firms grow. When firm productivity follows a Pareto distribution—as standard models assume and evidence supports—the distribution of firm scope exhibits a thinner right tail than the productivity distribution that generates it. Organizational complexity thus provides a structural explanation for the relative scarcity of very large firms: the "missing giants" in the right tail of the firm size distribution reflect not the absence of highly capable firms, but the organizational constraints that prevent capability from translating fully into scope.

We then ask how organizational complexity shapes market concentration. Aggregate concentration depends not only on how large firms are, but on how widely any single firm spans distant parts of the market space. We introduce an overlap kernel that unifies local and aggregate concentration into a single object: its value at zero is average local concentration, and its integral over characteristic distance is aggregate concentration. Organizational complexity steepens the decay of this kernel by compressing the scope radii that generate distant same-firm overlap. As a result, organizational complexity reduces aggregate concentration even when local concentration remains high—a distinction between localized dominance and economy-wide reach that standard models of firm size cannot capture.

Finally, we calibrate the benchmark specification to match the distribution of firm scope doc-

umented by Hoberg and Phillips (2025) using text-based measures of product market presence for U.S. public firms. The calibrated model reproduces the mean, median, and dispersion of the empirical scope distribution and generates quantitatively significant tail compression: a firm with ten times the organizational capability of another achieves only twice the scope. Embedding the calibrated scope function in the multi-firm concentration framework, we find that organizational complexity reduces aggregate concentration by 56% relative to a frictionless benchmark in which scope scales linearly with capability. The model's predictions also align with specific empirical patterns: scope expansion concentrated in related markets, driven by R&D and acquisitions rather than capital expenditure, and scope-adjusted concentration measures that have remained stable despite rising traditional industry HHI.

Related literature This paper contributes to three distinct literatures on firm organization, scope, and scale. First, we build on the span-of-control and organizational economics tradition initiated by Lucas (1978) and developed by Garicano (2000) and Garicano and Rossi-Hansberg (2006). In these models, firm scope is mechanically bounded by discrete managerial or knowledge constraints: each additional activity requires a fixed or irreducible quantum of supervisory attention, so the number of tasks a firm can undertake is finite. Related work on coordination and information-processing costs (Radner, 1993; Zandt, 1999) emphasizes congestion in decentralized decision-making. Our approach departs from this tradition by modeling organizational complexity as a continuous object that aggregates heterogeneity across activities rather than counting them. Scope is not limited by a fixed managerial span but by an endogenous organizational constraint that firms can relax through investment in blueprint quality and adaptation technology. This formulation allows firms to operate across an unbounded measure of sufficiently similar markets while serving only a finite mass of materially distinct ones—a distinction that count-based models cannot capture.

Second, we contribute to the literature on intangible capital and firm scope. This literature builds on the insight that organizational knowledge is non-rival within the firm. McGrattan and Prescott (2010) formalize this property in the context of multinational production, showing that intangible capital can be deployed across locations without being depleted. Subsequent work extends this logic: Argente, Moreira, Oberfield, and Venkateswaran (2025) show that standardization lowers marginal costs and facilitates smooth expansion, while Crouzet, Eberly, Eisfeldt, and Papanikolaou (2024) emphasize replicability and external contracting frictions. In the trade literature, Bernard, Redding, and Schott (2010, 2011) develop models of multi-product firms with a core competence, where productivity declines with distance from that core. A common feature

of these frameworks is that scope decisions remain separable: each market or product is evaluated independently, and serving one activity does not alter the organizational cost of serving another. Our paper shares the non-rivalry property of intangible capital but introduces a countervailing force: deploying the blueprint across heterogeneous markets requires adaptation using scarce headquarters attention, and the organizational burden aggregates across the firm's footprint. This non-separability—absent from the existing literature—is what generates the selection problem, the core–periphery structure, the phase transition to platform-type organization, and the compression of the firm size distribution.

Third, our framework complements the empirical literature on management practices and firm performance (Bloom and Van Reenen, 2007; Bloom, Sadun, and Van Reenen, 2019) by providing a structural link between management technology and firm scope. In our model, improved management practices operate through more effective blueprint creation and adaptation efficiency, allowing firms to compress heterogeneity across markets and expand scope discontinuously. This connects to the platform literature (Bresnahan and Greenstein, 1999; Rochet and Tirole, 2003; Gawer and Cusumano, 2014), but through a distinct channel: platform emergence in our framework arises from the endogenous collapse of internal organizational constraints rather than from external network effects. The core–periphery structure that emerges within the firm resonates with evidence that large organizations concentrate managerial oversight in core activities while operating peripheral lines through standardized processes (Antràs and Rossi-Hansberg, 2009).

1 Environment

We model a single firm that operates across a set of heterogeneous markets. The firm is organized around a centralized, non-rival blueprint created at headquarters and deployed across markets. Effective deployment of the blueprint in a given market requires market-specific adaptation. Serving markets generates an organizational burden at headquarters, summarized by a scalar state variable called organizational complexity, which depends on the composition of the firm's footprint rather than on the number of markets served. The remainder of this section formalizes these elements.

1.1 Markets and Heterogeneity

We formalize markets and heterogeneity through four distinct objects, each building on the previous:

1. A market space (S, Σ, μ) providing the measurable structure over which firms optimize.
2. Market characteristics $\xi(s) \in \mathbb{R}^k$ capturing the primitive, multidimensional sources of het-

erogeneity at each market $s \in S$.

3. A heterogeneity index $d(s) \in \mathbb{R}_+$ aggregating characteristics into a scalar measure of organizational difficulty.
4. Organizational complexity $C(M) \in \mathbb{R}_+$ aggregating heterogeneity across the firm's footprint into a firm-level state variable.

This layered structure separates the abstract space of markets from the economic content of heterogeneity, and separates market-level heterogeneity from its firm-level consequences.

Layer 1: Market space. Let (S, Σ, μ) be a measure space of markets, indexed by $s \in S$. This space is the primitive object carrying all measurable structure; the σ-algebra Σ and measure μ determine which collections of markets are measurable and how they are weighted. A distinguished origination location H ("headquarters") lies outside S; no goods production takes place at H, and all production occurs in markets $s \in S$.

We use the term **market** broadly. An index s may represent:

- a geographic market,
- a product line,
- a business unit,
- a customer segment,
- or any operational activity requiring blueprint deployment.

Assumption 1.1 (Non-atomic market space)**.** *The measure space (S, Σ, μ) is non-atomic: for every $E \in \Sigma$ with $\mu(E) > 0$ and every $\varepsilon \in (0, \mu(E))$, there exists $E' \subseteq E$ with $E' \in \Sigma$ and $\mu(E') = \varepsilon$.*

Assumption 1.1 rules out indivisible markets of positive measure and ensures that any collection of markets can be subdivided arbitrarily finely. This is natural in environments where markets differ continuously along multiple characteristics and is standard in models with continua of activities.

Layer 2: Market characteristics. Each market $s \in S$ is characterized by a vector of heterogeneity-relevant attributes

$$\xi(s) \in \mathbb{R}^k,$$

where $\xi : S \to \mathbb{R}^k$ is measurable with respect to Σ and the Borel σ-algebra on $\mathbb{R}^k$. The components of $\xi(s)$ capture distinct primitive sources of market heterogeneity—geography, technology, regulation, customer composition, operational requirements, and so forth. These characteristics are exogenous to the firm.

In addition, each market s has local rival factor capacity $\bar{X}(s) > 0$, also exogenous.

Layer 3: Heterogeneity index. We summarize the multidimensional characteristics $\xi(s)$ through a scalar heterogeneity index

$$d(s) = \delta\big(\xi(s)\big),$$

where $\delta : \mathbb{R}^k \to \mathbb{R}_+$ is Borel measurable. The index $d(s)$ measures organizational difficulty: markets with higher $d(s)$ require weakly more adaptation effort to achieve a given level of effective blueprint deployment.

Assumption 1.2 (Nonnegativity)**.** *The heterogeneity index satisfies $d(s) \geq 0$ for μ-a.e. $s \in S$.*

Example 1.3 (Organizational distance)**.** *A useful interpretation of $d(s)$ is as organizational distance from the firm's core. For instance, $d(s)$ may take the form of a weighted norm,*

$$d(s) = \Big(\sum_{j=1}^{k} \omega_j \, |\xi_j(s)|^p\Big)^{1/p},$$

or an additive index $d(s) = \sum_{j=1}^{k} d_j(\xi_j(s))$. Different specifications correspond to different organizational technologies and substitution patterns across dimensions of heterogeneity. None of the results require $d(\cdot)$ to satisfy the axioms of a norm.

Layer 4: Organizational complexity. Serving markets imposes an organizational burden at headquarters. We summarize this burden by a scalar state variable called organizational complexity, which aggregates market-level heterogeneity across the firm's entire footprint.

Definition 1 (Organizational complexity)**.** *For any measurable set of served markets $M \subseteq S$, organizational complexity is*

$$C(M) \;\equiv\; \int_M d(s)\, \mu(ds). \tag{1}$$

Organizational complexity $C(M)$ is well defined in $[0, \infty]$ for any measurable $M \subseteq S$. By construction, $C(\cdot)$ satisfies the following properties for any measurable sets $M_1, M_2 \subseteq S$:

- *Nonnegativity:* $C(M) \geq 0$.
- *Monotonicity:* if $M_1 \subseteq M_2$, then $C(M_1) \leq C(M_2)$.
- *Additivity:* if $M_1 \cap M_2 = \emptyset$, then $C(M_1 \cup M_2) = C(M_1) + C(M_2)$.

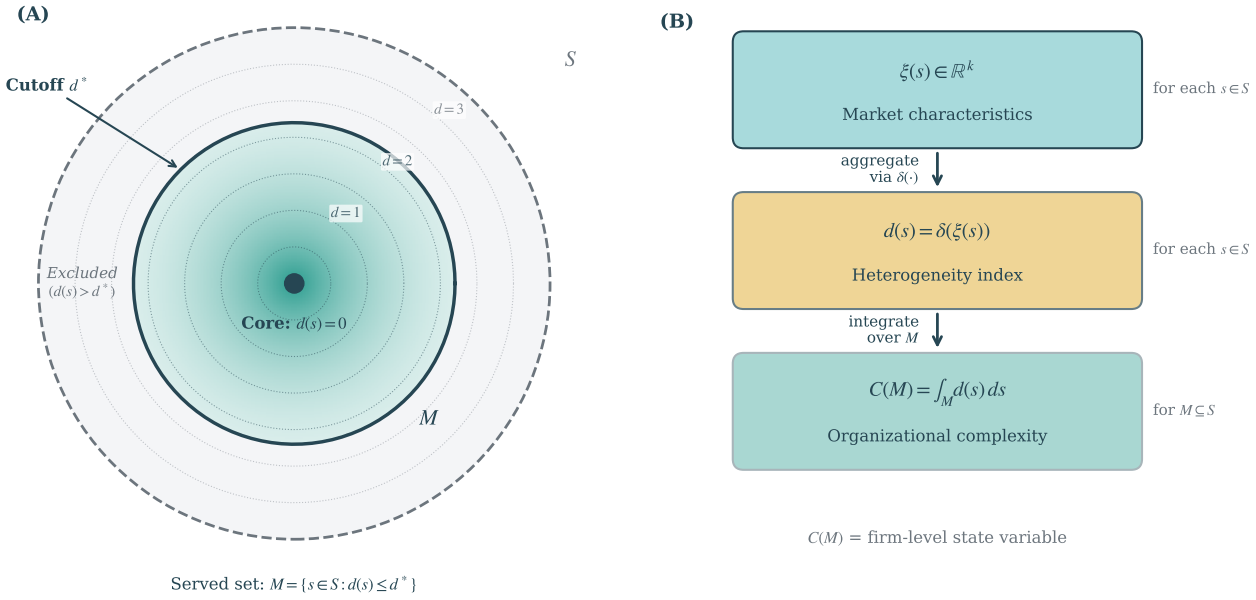


Figure 1: Market heterogeneity and endogenous firm scope.

Notes. Panel (a) depicts the market space S ordered by heterogeneity and the firm's endogenous footprint $M \subset S$. Under optimal organization, the firm serves all markets with heterogeneity below a cutoff d^*, so that $M = \{s \in S : d(s) \leq d^*\}$, while more heterogeneous markets are excluded. Panel (b) illustrates the layered structure of heterogeneity. Each market $s \in S$ is characterized by multidimensional attributes $\xi(s) \in \mathbb{R}^k$, which are aggregated into a scalar heterogeneity index $d(s) = \delta(\xi(s))$. Organizational complexity is defined as the integral of heterogeneity over the served set, $C(M) = \int_M d(s)\,\mu(ds)$.

Remark 1.4 (Role of the layered structure)**.** *The four-layer structure separates the abstract mathematical framework from its empirical implementation. Layers 1 and 4 define the theoretical objects—the market space (S, Σ, μ) and organizational complexity $C(M)$—while Layers 2 and 3 provide the bridge to data. In applications, the econometrician observes market characteristics $\xi(s)$ and must specify or estimate the aggregator $\delta(\cdot)$ that maps these characteristics into organizational heterogeneity $d(s) = \delta(\xi(s))$. The theoretical results hold for any measurable $d : S \to \mathbb{R}_+$, providing flexibility in how this mapping is constructed. This separation ensures that the theoretical propositions remain valid across a wide range of empirical settings, while the layered structure provides clear guidance on where modeling choices and data requirements enter.*

1.2 Blueprint Technology

Blueprint. The firm operates a centralized blueprint, denoted by $z \geq 0$. The blueprint is a non-rival organizational asset created at headquarters with the following properties:

- *Non-rivalry:* a given blueprint z can be deployed simultaneously across all served markets.

- *Centralization:* the blueprint is created and maintained at headquarters.
- *Genericity:* effective deployment of the blueprint in market s requires market-specific adaptation.

The blueprint does not directly produce output. Instead, it enters market-level payoffs through its effective deployment, which depends on local conditions and organizational inputs defined below.

Blueprint creation. The firm produces a blueprint at headquarters using headquarters attention. Let $h_z \geq 0$ denote the amount of headquarters attention allocated to blueprint creation. Blueprint quality is given by

$$z = \Gamma(h_z), \tag{2}$$

where $\Gamma : \mathbb{R}_+ \to \mathbb{R}_+$ is strictly increasing and continuous with $\Gamma(0) = 0$. Since Γ is strictly increasing, it admits an inverse $\chi \equiv \Gamma^{-1}$, so that producing a blueprint of quality z requires

$$h_z = \chi(z). \tag{3}$$

Blueprint deployment and adaptation. To deploy the blueprint in market s, the firm adapts it to local conditions using both headquarters and local organizational inputs. Let $h(s) \geq 0$ denote headquarters attention allocated to market s, and let $\ell(s) \geq 0$ denote local adaptation effort in market s. We restrict attention to measurable allocations $h(\cdot)$ and $\ell(\cdot)$.

Effective adaptation in market s is given by

$$a(s) = A\big(h(s), \ell(s)\big), \tag{4}$$

and effective blueprint deployment in market s is

$$z_s = \theta\big(a(s), d(s)\big)\, z, \tag{5}$$

where $\theta(\cdot,\cdot) \in (0,1]$ is a portability function. The heterogeneity index $d(s)$ enters (5) as the sufficient statistic for how the primitive characteristics $\xi(s)$ affect blueprint portability.

Assumption 1.5 (Adaptation and portability)**.** *The function $A : \mathbb{R}_+^2 \to \mathbb{R}_+$ is continuous and weakly increasing in each argument. The portability function $\theta : \mathbb{R}_+ \times \mathbb{R}_+ \to [0,1]$ is continuous, weakly increasing in a, weakly decreasing in d, and satisfies*

$$\lim_{a \to \infty} \theta(a, d) = 1 \quad \textit{for all } d \geq 0.$$

Assumption 1.5 captures that blueprint deployment requires market-specific adaptation using both headquarters and local inputs, and that adaptation improves the effectiveness with which the blueprint can be applied. Greater heterogeneity reduces portability, but sufficiently intensive adaptation can fully overcome these frictions.

1.3 Organizational Constraints and Payoffs

Local rival capacity. Each market s has limited local organizational resources $\bar{X}(s) > 0$. These resources must be allocated between production and local adaptation. Let $x(s) \geq 0$ denote production input in market s. Local feasibility requires

$$x(s) + \ell(s) \leq \bar{X}(s). \tag{6}$$

Scarcity of headquarters attention. Headquarters attention is rival within the firm and must be allocated across competing uses: blueprint creation and market-specific adaptation. We model organizational scarcity through a complexity-dependent coordination wedge: as the firm serves a more heterogeneous set of markets, each unit of headquarters attention becomes more costly to deploy. Deploying headquarters attention then incurs the total cost

$$q\big(C(M)\big)\left(\chi(z) + \int_M h(s)\,\mu(ds)\right). \tag{7}$$

Assumption 1.6 (Organizational congestion)**.** *The organizational congestion function $q : \mathbb{R}_+ \to \mathbb{R}_+$ satisfies:*

1. $q(0) = 1$*;*
2. q *is continuously differentiable and strictly increasing, with* $q'(C) > 0$ *for all* $C \geq 0$*;*
3. $\lim_{C\to\infty} q(C) = \infty$*.*

We refer to $q(C)$ as organizational congestion. It captures the idea that broader and more heterogeneous scope raises coordination burdens, thereby increasing the marginal cost of headquarters attention anywhere in the firm. We normalize the base price of headquarters attention to 1 (in units of the numeraire), so $q(C)$ represents the complexity-adjusted shadow price of attention.

Note that the congestion multiplier $q(C)$ applies to all uses of headquarters attention, including blueprint creation $\chi(z)$. This reflects the fact that improving the organizational blueprint—refining operational procedures, upgrading proprietary technology, redesigning coordination protocols—requires sustained managerial focus that competes with the day-to-day demands of adapting the

blueprint across heterogeneous markets. Blueprint creation cannot be fully insulated from the coordination burden of the firm's existing footprint.

Market-level payoff. Operating in market s generates payoff

$$R\big(z_s, x(s)\big), \tag{8}$$

where $R(\cdot,\cdot)$ is weakly increasing in both arguments. The function R summarizes demand, pricing, and production conditions in market s. We assume $s \mapsto R(z_s, x(s))$ is measurable for all feasible allocations and impose Assumption 1.7 to ensure that the firm's objective is well defined.

Assumption 1.7 (Conditional integrability)**.** *For any fixed blueprint quality $z \geq 0$, there exists an integrable function $\bar{R}_z \in L^1(S,\mu)$ such that for all measurable $M \subseteq S$ and all feasible allocations $\{x(s), \ell(s), h(s)\}_{s\in M}$ consistent with the deployment and feasibility constraints (10)–(11) stated in Section 1.4 below,*

$$R\big(z_s, x(s)\big) \leq \bar{R}_z(s) \quad \textit{for } \mu\textit{-a.e. } s \in S.$$

Assumption 1.7 requires that, for any fixed blueprint quality z, the firm's total potential gross payoff from serving markets is finite. The assumption does not impose a uniform bound on returns across different blueprint qualities and therefore does not restrict the scale or profitability of technological improvements. Instead, it rules out pathological cases in which, holding the firm's technology fixed, payoffs could be made arbitrarily large by serving an unbounded set of markets or by choosing extreme local allocations.

1.4 Firm Problem

The firm chooses a blueprint quality $z \geq 0$, a measurable set of markets $M \subseteq S$, and measurable allocations $\{x(s), \ell(s), h(s)\}_{s\in M}$ to maximize net payoff:

$$\max_{z,\, M,\, \{x,\ell,h\}} \left\{ \int_M R\big(z_s, x(s)\big)\, \mu(ds) \;-\; q\big(C(M)\big) \left(\chi(z) + \int_M h(s)\, \mu(ds) \right) \right\}, \tag{9}$$

subject to, for all $s \in M$,

$$z_s = \theta\big(A\big(h(s), \ell(s)\big),\, d(s)\big)\; z, \tag{10}$$

$$x(s) + \ell(s) \leq \bar{X}(s). \tag{11}$$

Served markets. We interpret the firm's footprint M as the support of market-level activity. Formally, a market is served if the firm allocates positive resources to it:

$$s \in M \quad \Longleftrightarrow \quad x(s) + \ell(s) + h(s) > 0.$$

Accordingly, whenever we refer to a feasible plan $(M, z, \{x, \ell, h\})$, we take $M = \{s \in S : x(s) + \ell(s) + h(s) > 0\}$ up to μ-null sets.

2 Baseline Results

The results characterize how firms organize their scope and internal operations in the presence of organizational complexity. We distinguish between two margins. The extensive margin determines which markets are served and governs the firm's overall footprint. The intensive margin determines how organizational resources are allocated across served markets. Both margins are disciplined by the same underlying mechanism—organizational complexity and the scarcity of headquarters attention—but they operate in distinct ways. The results below first establish limits and ordering on the extensive margin, and then characterize how firms optimally allocate headquarters effort across markets on the intensive margin.

2.1 Organizational Complexity and Endogenous Limits to Firm Scope

We need several assumptions first.

Assumption 2.1 (Positive heterogeneity)**.** $d(s) > 0$ *for* μ*-a.e.* $s \in S$.

Assumption 2.1 rules out a positive-measure set of perfectly frictionless markets with zero heterogeneity. It ensures that any non-negligible expansion of scope entails some organizational burden, while still allowing the firm to scale across many nearly homogeneous markets with arbitrarily small $d(s)$.

Assumption 2.2 (No free deployment)**.** *For any market* s *with* $d(s) > 0$*, blueprint portability satisfies*

$$\theta\big(A(0, \ell), d(s)\big) = 0 \quad \textit{for all } \ell \geq 0.$$

Assumption 2.2 rules out costless deployment of the blueprint across heterogeneous markets, even when the firm invests arbitrarily in blueprint quality. As a result, serving any market with positive heterogeneity requires strictly positive organizational resources, preventing free expansion

along the extensive margin. Local expertise cannot perfectly substitute for headquarters attention; some HQ involvement is always required to deploy the blueprint.

Assumption 2.3 (No positive payoff without effective deployment)**.** *For any market s and any feasible production input x satisfying $x \leq \bar{X}(s)$, the payoff satisfies*

$$R(0, x) \leq 0.$$

Assumption 2.3 rules out free-standing market rents: favorable local market conditions alone cannot generate positive payoffs for the firm in the absence of effective blueprint deployment.

We are in a position to state the first formal result of the paper.

Proposition 2.1 (Bounded organizational complexity)**.** *Suppose that organizational congestion is unbounded (Assumption 1.6), that gross payoffs are conditionally integrable (Assumption 1.7), and that positive payoffs require effective blueprint deployment (Assumptions 2.2 and 2.3). Then any optimal firm choice $(M^\star, z^\star, \{x^\star, \ell^\star, h^\star\})$ satisfies*

$$C(M^\star) \;=\; \int_{M^\star} d(s)\, \mu(ds) \;<\; \infty.$$

Proof. See Appendix A.1. □

Proposition 2.1 formalizes a bounded organizational capacity view of scope: even though the firm may operate across a very large collection of markets, the total organizational burden generated by heterogeneity, $C(M) = \int_M d(s)\, \mu(ds)$, must remain finite at the optimum because escalating complexity raises the marginal cost of headquarters resources through $q(C)$.

Corollary 2.1 (Finite mass of materially distinct markets)**.** *Fix any heterogeneity threshold $\tau > 0$ and define the set of materially distinct markets*

$$M^\star_\tau \;\equiv\; \{\, s \in M^\star : d(s) \geq \tau \,\}.$$

Then

$$\mu(M^\star_\tau) \;\leq\; \frac{1}{\tau} \int_{M^\star} d(s)\, \mu(ds) \;=\; \frac{C(M^\star)}{\tau} \;<\; \infty,$$

so the optimal organization serves only a finite (in measure) set of markets whose heterogeneity exceeds τ.

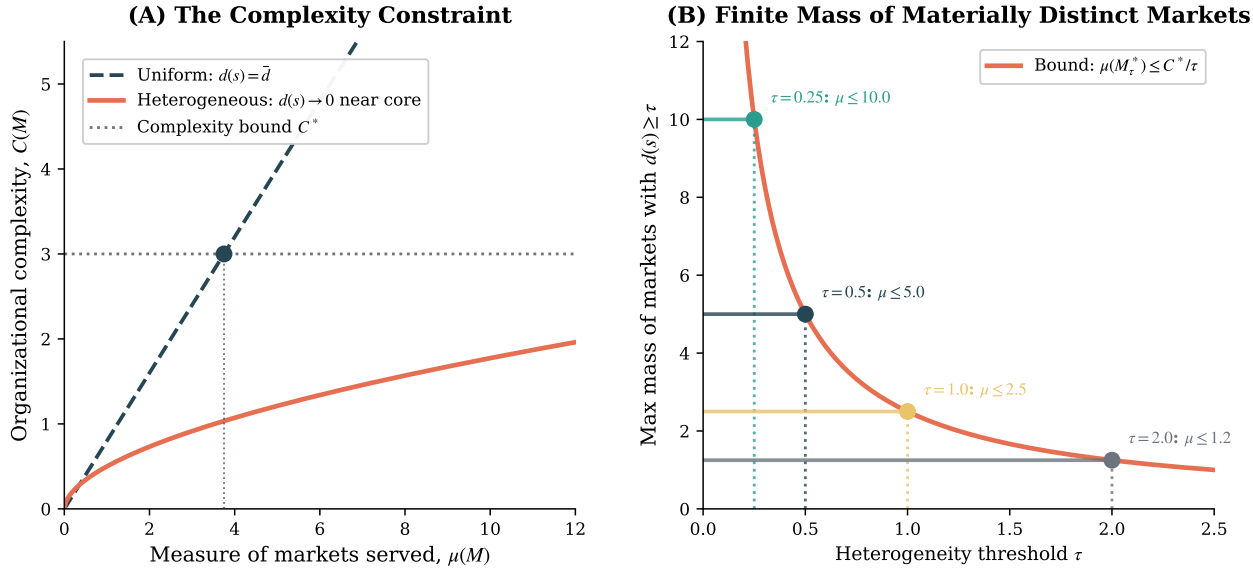


Figure 2: Organizational complexity and bounded scope

Proof. See Appendix A.1. □

Corollary 2.1 translates Proposition 2.1 into an economically meaningful statement about the mass of materially distinct markets. For any resolution level $\tau > 0$, the firm can serve only a finite (in measure) set of markets with heterogeneity at least τ. At the same time, the model allows the firm to cover a vast—and potentially unbounded—measure of highly similar markets with $d(s) < \tau$, and in particular with arbitrarily small $d(s)$.

This distinction captures the scalability of platform-like firms. For example, Amazon can serve millions of product–location niches because most of them are handled through standardized logistics, modular IT systems, and templated compliance processes—in the model, these niches correspond to markets with very small heterogeneity $d(s)$. The binding constraint is not the raw count of niches, but the aggregate heterogeneity of the footprint: the mass of materially different environments (large-d markets) that require bespoke adaptation and managerial attention.

Relation to span-of-control theories. Classic theories of firm size and scope based on span of control, such as Lucas (1978) and Garicano (2000), impose an exogenous limit on expansion through a fixed or irreducible managerial burden associated with each additional activity or task. In these models, managerial attention is scarce and expansion requires discrete allocations of supervisory or problem-solving resources. As a result, firm scope is mechanically bounded: the number of markets or activities a firm can operate in is finite and determined by technological parameters

governing managerial capacity.

The mechanism in our model is fundamentally different. We do not impose per-market fixed costs, discrete entry costs, or uniform lower bounds on organizational burden. Instead, the constraint on expansion operates through total organizational complexity,

$$C(M) = \int_M d(s)\,\mu(ds),$$

which aggregates heterogeneity across markets in a smooth and continuous manner. This formulation allows firms to serve a large—potentially unbounded—set of sufficiently similar markets, while sharply limiting the mass of materially distinct environments that require bespoke adaptation and coordination.

As a consequence, firm scope is genuinely endogenous. Firms expand not by exhausting a fixed span of managerial control, but by endogenously reshaping the heterogeneity associated with additional markets through standardization, modular organizational design, and scalable operational blueprints. In this sense, large platform-type firms—such as Amazon—operate across vast numbers of markets not because managerial constraints are absent, but because organizational heterogeneity has been compressed, allowing scope to scale without proportionate increases in coordination costs. This distinction separates our approach from traditional span-of-control models and highlights a novel channel through which organizational technology governs firm scope.

2.2 Organizational Complexity, Scale, and the Scope Trade-off

Definition 2 (Market-level value function)**.** *Fix a market $s \in S$, a blueprint quality $z \geq 0$, and an organizational complexity level $C \geq 0$. The market-level value function $\Pi_s(z; C)$ is defined as the value of the inner optimization problem*

$$\begin{aligned} \Pi_s(z;C) \quad &\equiv \quad \sup_{\substack{x\geq 0,\ \ell\geq 0,\ h\geq 0 \\ x+\ell\leq \bar{X}(s)}} \Big\{ R(z_s, x) \;-\; q(C)\,h \Big\}, \\ \textit{s.t.} \quad z_s \;&=\; \theta\big(A(h,\ell),\ d(s)\big)\, z. \end{aligned} \tag{12}$$

The blueprint creation cost $q(C)\chi(z)$ is common across markets and is handled at the outer problem.

The function $\Pi_s(z; C)$ represents the value of operating market s given the organizational prices induced by complexity C and a fixed blueprint quality z, taking the firm's overall scope and blueprint investment as given.

Assumption 2.4 (Blueprint effectiveness). *For each market s, the market-level payoff function $R(z_s, x)$ is weakly increasing in effective blueprint deployment z_s.*

Assumption 2.4 formalizes that greater effective deployment of the blueprint weakly improves performance in any given market. The assumption imposes no curvature or diminishing returns on market payoffs and serves only to ensure that blueprint investment is economically meaningful rather than a purely organizational artifact.

Proposition 2.2 (Organizational congestion). *Fix a market $s \in S$ and a blueprint level $z \geq 0$. The market-level value function $\Pi_s(z; C)$ is weakly decreasing in organizational complexity C. In particular, for any served set $M \subseteq S$ and any measurable $\Delta \subseteq S \setminus M$ with $\mu(\Delta) > 0$,*

$$\Pi_s\big(z; C(M \cup \Delta)\big) \;\leq\; \Pi_s\big(z; C(M)\big).$$

Proof. See Appendix A.2. □

Remark 2.5 (Strictness). *The inequality in Proposition 2.2 is strict for an incumbent market $s \in M$ whenever the optimal inner allocation at $C(M)$ uses strictly positive headquarters attention, $h^*(s; z, C(M)) > 0$, so that the shadow price $q(C)$ is locally relevant. Markets operated purely through the standardized blueprint ($h^*(s) = 0$) are locally insulated from changes in organizational complexity.*

Proposition 2.2 isolates a negative internal organizational externality that operates conditional on blueprint quality. Holding the blueprint z fixed, expanding the firm's footprint raises organizational complexity and thereby increases the shadow price of headquarters attention. Because headquarters attention is a shared and rival input used for market-level adaptation, this increase in organizational scarcity weakly lowers the value of all incumbent markets.

This mechanism implies that market selection is inherently non-separable. The profitability of operating in a given market depends not only on local conditions, but also on the composition of the firm's overall footprint through its effect on organizational complexity. Importantly, this does not imply that expansion is unprofitable in aggregate: new markets may generate sufficient value to offset the organizational tax imposed on incumbent activities. Rather, Proposition 2.2 establishes that expansion decisions are interdependent across markets.

It is crucial to distinguish this congestion effect from the return to the blueprint itself. Because the blueprint is non-rival, expanding scope increases the marginal benefit of investing in

blueprint quality. For example, consider a firm headquartered in the United States that initially operates in the EU and subsequently expands to China. Entering China increases the scale over which the blueprint can be deployed, strengthening the incentive to improve z. At the same time, serving China raises organizational complexity, tightening the firm's organizational constraint and increasing the cost of adaptation in all markets, including the EU. Proposition 2.2 captures this latter force holding z fixed, while Proposition 2.3 captures the countervailing effect: the return to blueprint investment rises with the firm's scope. Whether incumbent markets become more or less profitable in equilibrium depends on the balance between these two forces.

Proposition 2.3 (Returns to blueprint scale)**.** *Fix an organizational complexity level $C \geq 0$. Suppose Assumption 2.4 holds and that sufficient regularity obtains for differentiation under the integral sign (see Lemma A.2 in the Appendix). Then for any served set $M \subseteq S$, the marginal return to blueprint quality*

$$\frac{\partial}{\partial z} \int_{s \in M} \Pi_s(z; C)\, \mu(ds)$$

is weakly increasing in M with respect to set inclusion. In particular, if $M_1 \subseteq M_2$, then

$$\frac{\partial}{\partial z} \int_{s \in M_2} \Pi_s(z; C)\, \mu(ds) \;\geq\; \frac{\partial}{\partial z} \int_{s \in M_1} \Pi_s(z; C)\, \mu(ds).$$

Proof. See Appendix A.2. □

Because the blueprint is non-rival, a marginal improvement in blueprint quality raises effective productivity in every market where it is deployed. Serving a larger set of markets therefore increases the aggregate return to blueprint investment, holding organizational complexity fixed. In our earlier quantitative work (Dyrda, Hong, Sajid, and Steinberg, 2025), this mechanism appeared in a highly stylized and tightly parametrized environment with explicit production technologies, discrete multinational expansion, and general equilibrium feedbacks. In that setting, scalable complementarity between intangible investment and multinational expansion emerged from specific functional forms and calibration choices. Proposition 2.3 shows that this logic is considerably more general: it follows directly from non-rival deployment at the organizational level and does not rely on any particular market structure, production technology, or parametric specification.

Net effect of scope expansion. Proposition 2.2 shows that, holding blueprint quality fixed, expanding the firm's footprint raises organizational complexity and thereby weakly lowers the value of all incumbent markets. Proposition 2.3 shows that expanding scope raises the marginal

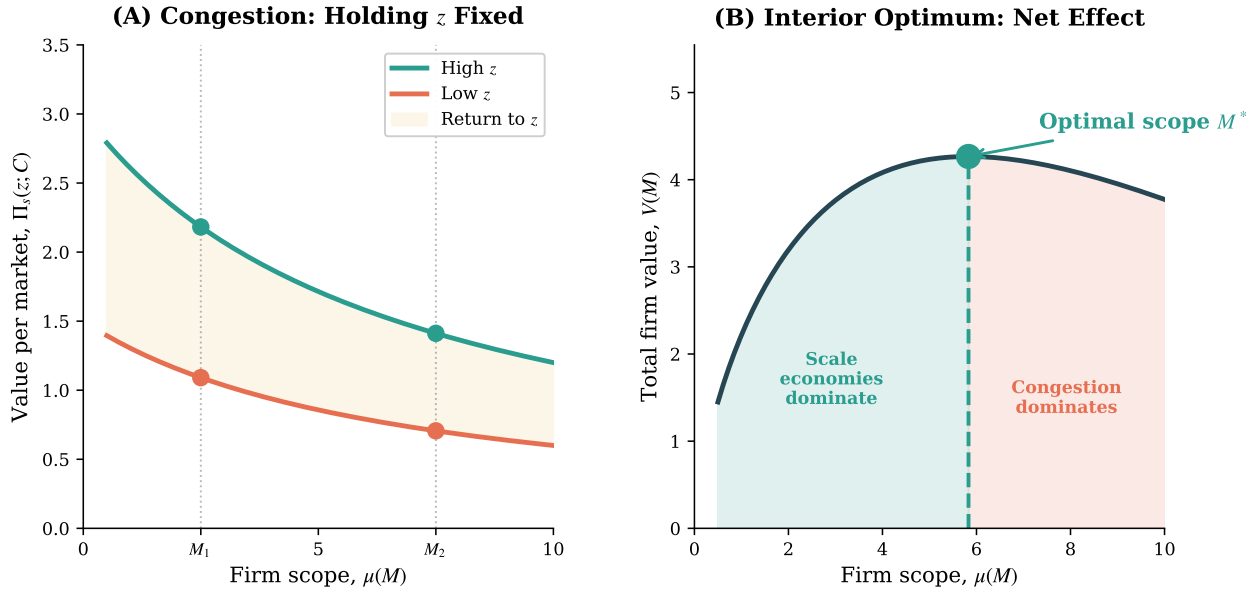


Figure 3: The scope–technology trade-off

return to investing in non-rival blueprint quality. When blueprint quality is chosen optimally, these forces operate in opposite directions.

For any incumbent market $s \in M$, the change in its value following an expansion $M' = M \cup \Delta$ admits the decomposition

$$\Delta_\Pi = \underbrace{\left[\Pi_s\big(z^*(M'); C(M')\big) - \Pi_s\big(z^*(M); C(M')\big)\right]}_{\text{blueprint adjustment}} + \underbrace{\left[\Pi_s\big(z^*(M); C(M')\big) - \Pi_s\big(z^*(M); C(M)\big)\right]}_{\text{organizational congestion}}. \tag{13}$$

The second term is weakly negative by Proposition 2.2. The first term captures the effect of adjusting blueprint quality at the new complexity level and is weakly positive whenever scope expansion induces an increase in the optimal blueprint. As a result, the net effect of scope expansion on incumbent market performance is a priori ambiguous and depends on the balance between these two forces.

Taken together, Propositions 2.2 and 2.3 characterize the opposing effects of scope expansion on incumbent market values and on the incentives to invest in organizational technology. Given these forces, the firm's footprint choice reduces to a selection problem over markets: which markets to serve, taking blueprint quality and organizational complexity as given. The next result shows that, under a mild ordering assumption on market heterogeneity, this selection problem admits a simple

cutoff structure.

Assumption 2.6 (Heterogeneity ordering)**.** *Markets can be ranked by heterogeneity* $d(s)$*. Fix* (z, C)*. The market-level value* $\Pi_s(z; C)$ *is weakly decreasing in* $d(s)$*: if* $d(s_1) \leq d(s_2)$ *then* $\Pi_{s_1}(z; C) \geq \Pi_{s_2}(z; C)$*.*

Assumption 2.6 ensures that markets can be ordered by organizational difficulty, so that more heterogeneous markets are weakly less profitable to serve for a given blueprint and organizational capacity. This assumption provides a monotone structure to the extensive margin of expansion without imposing any restrictions on returns to scale or within-market technologies.

Proposition 2.4 (Cutoff selection in heterogeneity)**.** *Fix blueprint quality* $z \geq 0$*. Suppose Assumption 2.6 holds. Then the firm's optimal footprint* $M^\star$ *is monotone in heterogeneity: if* $s_2 \in M^\star$ *and* $d(s_1) \leq d(s_2)$*, then* $s_1 \in M^\star$*. Equivalently, there exists a threshold* $d^\star \in [0, \infty]$ *such that*

$$s \in M^\star \quad \Longleftrightarrow \quad d(s) \leq d^\star,$$

up to μ*-null sets.*

Proof. See Appendix A.2. □

Proposition 2.4 shows that optimal market selection is ordered by heterogeneity. Because organizational complexity increases with both the scope and the composition of activities, serving more heterogeneous markets imposes a disproportionately larger organizational burden on the firm. At the same time, holding blueprint quality fixed, the value of operating a market is lower when that market is more distant from the blueprint's origination environment. The firm therefore expands selectively, serving activities that are sufficiently similar to its core operations and excluding those that are too heterogeneous.

Importantly, the cutoff structure does not rely on per-market fixed costs or exogenous decreasing returns. It follows from organizational scarcity interacting with the heterogeneity ordering (Assumption 2.6), which provides the monotone structure needed for market selection: expansion tightens a shared internal constraint, and markets are selected based on their contribution to overall organizational complexity. As a result, firm growth follows an ordered pattern—near or closely related activities are served before distant or heterogeneous ones—even though the blueprint itself is non-rival.

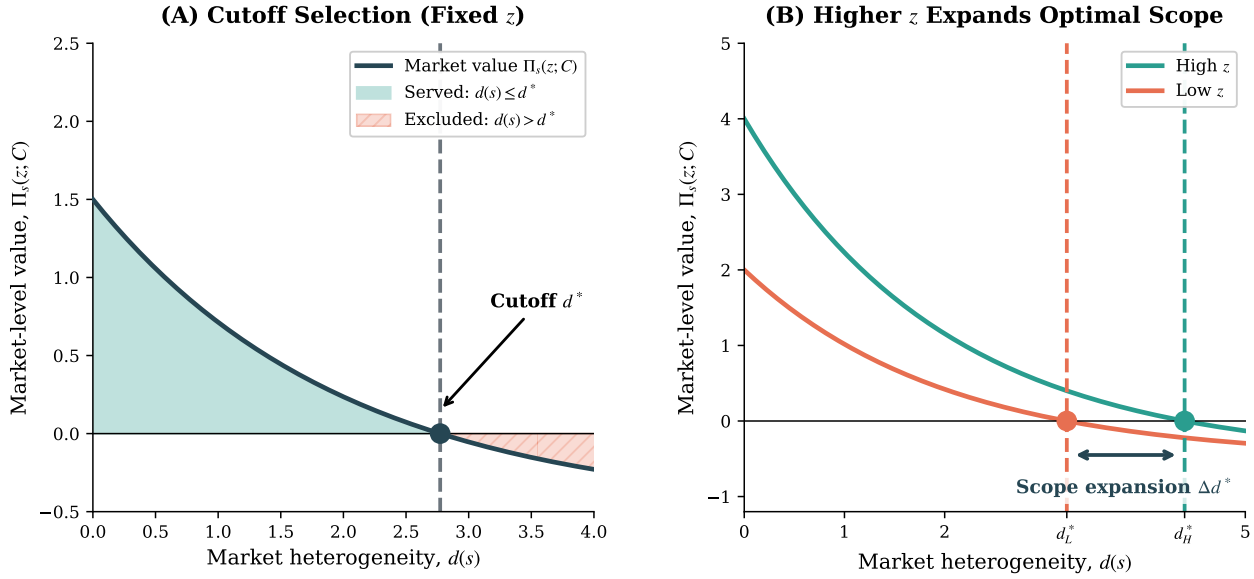


Figure 4: Cutoff selection in heterogeneity

This mechanism aligns closely with recent empirical evidence documenting that modern firms expand scope primarily by entering highly related product markets rather than by diversifying into unrelated activities. Using text-based measures of product market scope, Hoberg and Phillips (2025) show that scope expansion over the past several decades has been concentrated in related markets and has occurred largely without increases in traditional operating segments. Proposition 2.4 provides a theoretical foundation for this pattern: when organizational capacity is scarce, firms optimally prioritize related expansions that add less organizational complexity per unit of scope.

2.3 Adaptation and Core–Periphery Organization

Assumption 2.7 (Adaptation regularity). *The functions $R(z_s, x)$, $\theta(a, d)$, and $A(h, \ell)$ are continuously differentiable in their relevant arguments, with $R_1(z_s, x) \geq 0$, $\partial_a \theta(a, d) \geq 0$, and $\partial_h A(h, \ell) > 0$.*

Assumption 2.7 is a regularity condition ensuring that market-level payoffs and effective blueprint deployment respond smoothly to changes in headquarters attention and local adaptation effort. The assumption guarantees that organizational effort is productive at the margin and allows us to characterize how scarce headquarters attention is allocated across markets, without imposing curvature or restrictions on returns to scale.

Assumption 2.8 (Single-crossing in heterogeneity). *For any fixed $a \geq 0$, the marginal effectiveness*

of adaptation is weakly decreasing in heterogeneity: if $d_1 \leq d_2$*, then*

$$\partial_a \theta(a, d_1) \;\geq\; \partial_a \theta(a, d_2).$$

Assumption 2.8 imposes a single-crossing property in market heterogeneity, so that headquarters-driven adaptation is relatively more effective in less heterogeneous markets. This structure delivers monotone allocation of organizational resources within the firm's footprint while remaining agnostic about within-market technologies.

Proposition 2.5 (Core–periphery allocation of headquarters attention)**.** *Fix blueprint quality* $z \geq 0$*, organizational complexity* C*, and a served set* M*. Consider the market-level inner problems defining* $\Pi_s(z; C)$ *for* $s \in M$*.*

Suppose Assumptions 2.7–2.8 hold, that local rival capacity $\bar{X}(s)$ *is constant across served markets, and that, for each* $s \in M$*, the optimal choice of headquarters attention* $h^\star(s)$ *is interior. Then headquarters attention allocated across markets is weakly decreasing in heterogeneity: for any* $s_1, s_2 \in M$*,*

$$d(s_1) \leq d(s_2) \quad \Longrightarrow \quad h^\star(s_1) \;\geq\; h^\star(s_2).$$

Proof. See Appendix A.3. □

Proposition 2.5 characterizes how a firm allocates scarce organizational resources within its optimal footprint. Because headquarters attention is a shared and costly input, the firm cannot adapt all markets equally. Markets that are closer to the blueprint's origination environment have higher marginal returns to headquarters-driven adaptation, as organizational effort translates more effectively into usable blueprint deployment. By contrast, in more heterogeneous markets, the same organizational effort yields smaller gains. The firm therefore concentrates headquarters attention in low-heterogeneity core markets and operates high-heterogeneity peripheral markets with less intensive organizational involvement.

This core–periphery structure does not arise from fixed costs, technological constraints, or differences in market size. Rather, it is an endogenous response to organizational scarcity: as the firm expands and organizational complexity rises, headquarters attention is rationed toward activities where it is most productive. As a result, firms appear centralized in their core operations while maintaining broader, more standardized peripheral activities. This pattern is consistent with empirical evidence showing that large firms tend to concentrate managerial oversight and customization in core activities while relying on standardized, less intensive organizational processes

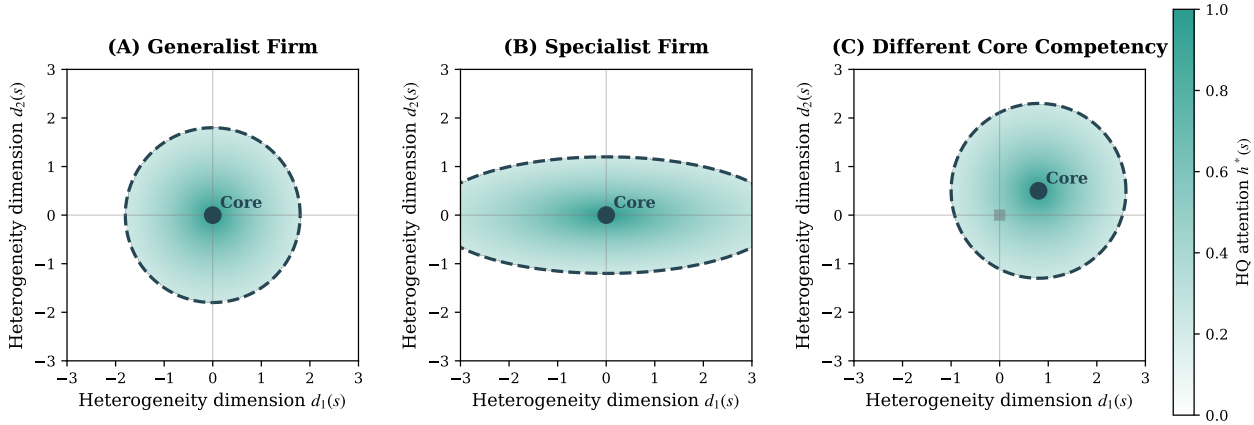


Figure 5: Core–periphery organization within the firm

in more peripheral lines of business.

Corollary 2.2 (Effectiveness gradient)**.** *Under the conditions of Proposition 2.5, effective blueprint deployment is weakly decreasing in heterogeneity:*

$$d(s_1) \leq d(s_2) \quad \Longrightarrow \quad z^\star_{s_1} \geq z^\star_{s_2}.$$

The core–periphery allocation of headquarters attention implies an effectiveness gradient across markets. Activities that are closer to the firm's organizational core operate with higher effective blueprint deployment, while more heterogeneous peripheral activities are served in a more standardized and less effective manner. Organizational complexity therefore shapes not only which markets are served, but also how intensively and effectively they are operated.

Remark 2.9 (Standardization as a substitute for managerial attention)**.** *The production function for effective deployment, $z_s = \theta(A(h(s), \ell(s)), d(s))\, z$, implies a substitution between blueprint quality z and headquarters attention $h(s)$. To achieve a given level of effective deployment in a heterogeneous market, the firm can either allocate rival organizational effort to adapt the blueprint (raising the portability term θ) or invest in the non-rival quality of the blueprint itself (raising z).*

As a result, improvements in organizational technology reduce the marginal need for bespoke managerial adaptation in heterogeneous markets. This substitution underlies the core–periphery allocation of headquarters attention and foreshadows how investments in non-rival organizational capital relax internal capacity constraints.

3 Platform Emergence as a Limiting Regime

The results in Section 2 characterize firm scope and internal organization when organizational complexity imposes a binding constraint on expansion. In the baseline environment, firms may operate across many highly similar activities, but the aggregate heterogeneity of the footprint must remain finite, so organizational scarcity ultimately limits scope through congestion.

This section asks whether that trade-off can be endogenously relaxed. We introduce an organizational technology—standardization—that reduces the effective heterogeneity associated with serving markets. Unlike the within-footprint adaptation choices studied in Section 2.3, which allocate scarce headquarters attention across served activities, standardization operates at a distinct margin: it reshapes how baseline market heterogeneity aggregates into organizational complexity. As a result, standardization does not merely lower the level of organizational burden; it can change how that burden scales with scope. We show below that platform emergence can arise from a pure replication mechanism under abundant near-core opportunities, even when standardization is inactive ($\sigma \equiv 0$). The extended environment with endogenous standardization provides a richer framework for applications but is not required for the limiting result (Proposition 3.2). Figure 7, Panel (A), illustrates this distinction: without standardization, organizational complexity grows without bound as the footprint expands, whereas sufficiently strong standardization keeps complexity uniformly below a fixed cap $\bar{C}$ even as the firm's footprint grows arbitrarily large.

The key point is that platform-type organization should be interpreted here as a limiting scalability regime, not as a literal finite-parameter jump to infinite scope. In any finite environment, footprint remains finite. What strong organizational technology can do is make the bounded-complexity scale frontier arbitrarily large: for any target footprint, sufficiently strong blueprint technology together with sufficiently fine near-core market granularity can support that footprint while keeping total organizational complexity uniformly bounded. This delivers a sharp and economically natural notion of platform emergence without relying on a knife-edge phase transition.

3.1 Extended model: endogenous heterogeneity reduction

We extend the baseline environment by allowing the firm to invest in standardization that reduces the effective heterogeneity of served markets. To distinguish pre- and post-standardization objects, we write $d_0(s)$ for the baseline heterogeneity index previously denoted $d(s)$.

Effective heterogeneity and deployment. Let $g : \mathbb{R}_+ \to (0, 1]$ be measurable and weakly decreasing, with $g(0) = 1$. Given market-specific standardization intensity $\sigma(s) \geq 0$, define effective

heterogeneity by

$$\tilde{d}(s) \;\equiv\; d_0(s)\, g(\sigma(s)).$$

Effective blueprint deployment in market s is then

$$z_s \;=\; \theta\big(A(h(s),\ell(s)),\, \tilde{d}(s)\big)\, z \;=\; \theta\big(A(h(s),\ell(s)),\, d_0(s)g(\sigma(s))\big)\, z. \tag{14}$$

Standardization costs. Standardization is costly and interacts with blueprint quality. We assume separability:

$$\kappa(\sigma, z) \;=\; \psi(z)\, c(\sigma),$$

where c is increasing and convex with $c(0) = 0$, and ψ is decreasing with

$$\lim_{z\to\infty} \psi(z) = 0.$$

Intuitively, higher blueprint quality reduces the unit cost of implementing standardization.

Complexity with standardization. Organizational complexity of a served set M is now

$$C(M,\sigma) \;\equiv\; \int_{s\in M} \tilde{d}(s)\, \mu(ds) \;=\; \int_{s\in M} d_0(s) g(\sigma(s))\, \mu(ds).$$

Extended market-level value. Fix a market $s \in S$, blueprint quality $z \geq 0$, and organizational complexity level $C \geq 0$. Define the extended market-level value function

$$\hat{\Pi}_s(z; C) \;\equiv\; \sup_{\substack{x,\ell,h,\sigma\geq 0 \\ x+\ell\leq \overline{X}(s)}} \left\{ R\big(\theta\big(A(h,\ell), d_0(s)g(\sigma)\big)z,\; x\big) - q(C)\, h - \psi(z)c(\sigma) \right\}. \tag{15}$$

As in Section 2.2, the blueprint creation cost $q(C)\chi(z)$ is common across markets and is handled at the outer problem.

Restricted-economy problem. To state a limiting result without claiming literal infinite scale at finite parameter values, we introduce a family of finite opportunity sets. For each $\rho \geq 1$, the feasible market set is $S^\rho \subseteq S$, and conditional on blueprint quality z, the firm chooses $M \subseteq S^\rho$

together with measurable allocations $\{x(s), \ell(s), h(s), \sigma(s)\}_{s\in M}$ to maximize

$$\hat{V}^{\rho}(M, z, x, \ell, h, \sigma) \equiv \int_{s\in M} R(z_s, x(s))\,\mu(ds) - q\big(C(M,\sigma)\big)\left[\chi(z) + \int_{s\in M} h(s)\,\mu(ds)\right] - \int_{s\in M} \psi(z)\,c(\sigma(s))\,\mu(ds), \quad (16)$$

subject to (14) and the baseline feasibility constraints.

Remark 3.1 (Relation to standardization-based scale models)**.** *Recent work, most notably Argente et al. (2025), emphasizes standardization as a force that facilitates expansion by increasing replicability and flattening cost schedules. Our framework also features organizational costs, captured by the complexity-dependent shadow price $q(C)$ of headquarters attention, but standardization operates through a different channel. Rather than directly lowering marginal production costs, standardization lowers effective heterogeneity, which governs how organizational complexity aggregates across activities. When sufficiently many near-core opportunities are available, scope can expand along margins that add little to organizational burden even though headquarters attention remains scarce. In this sense, standardization relaxes internal organizational constraints by changing how burden scales with scope, not merely by shifting marginal operating costs.*

3.2 Results

Assumption 3.2 (Finite-granularity approximation)**.** *There exists a family of measurable subsets $\{S^{\rho}\}_{\rho\geq 1}$ such that:*

(i) $S^{\rho} \subseteq S^{\rho'}$ *whenever* $\rho \leq \rho'$*;*
(ii) $0 < \mu(S^{\rho}) < \infty$ *for every* ρ*;*
(iii) $\bigcup_{\rho\geq 1} S^{\rho} = S$*;*
(iv) for every $\delta > 0$*,*

$$\mu(S^{\rho}_{\delta}) \xrightarrow[\rho\to\infty]{} \infty, \qquad S^{\rho}_{\delta} \equiv S^{\rho} \cap \{s \in S : d_0(s) \leq \delta\}.$$

Assumption 3.2 says that each finite economy has a large but finite set of opportunities, while the mass of sufficiently near-core opportunities expands without bound as ρ rises. Economically, ρ is a granularity index: larger ρ corresponds to finer product varieties, customer niches, geographic subdivisions, or standardized service requests.

Assumption 3.3 (Near-core deployment efficiency). *The deployment technology satisfies:*

(i) ***Perfect portability at zero heterogeneity:*** $\theta(a, 0) = 1$ *for all* $a > 0$.

(ii) ***Continuity and monotonicity:*** *For each* $a > 0$*, the function* $\theta(a, d)$ *is continuous and non-increasing in* d.

(iii) ***Uniform feasible recipe:*** *There exist* $(\bar{x}, \bar{h}, \bar{\ell})$ *with* $\bar{x}, \bar{h} > 0$, $\bar{\ell} \geq 0$, *such that*

$$\bar{x} + \bar{\ell} \leq \inf_{s \in S} \overline{X}(s), \qquad A(\bar{h}, \bar{\ell}) > 0.$$

(iv) ***Zero standardization cost at zero:*** $c(0) = 0$.

Assumption 3.3 ensures that sufficiently similar markets can be served using a common, well-behaved organizational recipe. Part (i) normalizes deployment in markets identical to headquarters. Part (ii) implies that portability deteriorates smoothly as heterogeneity rises. Part (iii) guarantees the existence of a uniform feasible recipe that can be replicated across markets. Part (iv) ensures that abstaining from standardization carries no direct cost.

Assumption 3.4 (Increasing returns to blueprint quality). *There exists* $\bar{x} > 0$ *such that:*

(i) ***Unbounded revenues:***

$$\lim_{z \to \infty} R(z, \bar{x}) = \infty.$$

(ii) ***Revenue growth dominates blueprint creation cost:*** *for every* $\alpha \in (0, 1]$,

$$\lim_{z \to \infty} \frac{R(\alpha z, \bar{x})}{\chi(z)} = \infty.$$

Interpretation. Assumption 3.4 is the source of scalability in the model. Blueprint quality is a non-rival organizational asset; improving it raises payoff in every market where it is deployed, while creation costs are paid only once. Part (ii) requires these scalable benefits eventually to dominate the common fixed cost of blueprint creation.

Proofs of the formal results below are collected in Appendix A.5.

Lemma 3.1 (Uniform near-core surplus under bounded complexity). *Fix any* $\bar{C} < \infty$. *Suppose Assumptions 3.3 and 3.4 hold. Then there exist constants* $\varepsilon(\bar{C}) > 0$, $z_1(\bar{C}) < \infty$, *and a function* $\underline{\pi}(\cdot; \bar{C}) : [z_1(\bar{C}), \infty) \to (0, \infty)$ *such that:*

(i) for all $z \geq z_1(\bar{C})$, *all* $C \in [0, \bar{C}]$, *and all markets* s *with* $d_0(s) \leq \varepsilon(\bar{C})$,

$$\hat{\Pi}_s(z; C) \;\geq\; \underline{\pi}(z; \bar{C}) \;>\; 0;$$

(ii) the lower bound satisfies

$$\lim_{z\to\infty} \underline{\pi}(z;\bar{C}) = \infty, \qquad \lim_{z\to\infty} \frac{\underline{\pi}(z;\bar{C})}{\chi(z)} = \infty.$$

Lemma 3.1 isolates the mechanism. Once blueprint quality is high enough, every sufficiently near-core market can be served using a common organizational recipe that yields a uniformly positive surplus, and this surplus dominates blueprint creation costs asymptotically.

Definition 3 (Bounded-complexity scale frontier)**.** *Fix $\bar{C} < \infty$, $\rho \geq 1$, and $z \geq 0$. The bounded-complexity scale frontier is*

$$\mathcal{K}^{\rho}(z;\bar{C}) \;\equiv\; \sup\left\{ \mu(M) : \begin{array}{l} M \subseteq S^{\rho},\ \exists \text{ measurable feasible allocations } (x,\ell,h,\sigma) \\ \text{such that } C(M,\sigma) \leq \bar{C} \text{ and } \hat{V}^{\rho}(M,z,x,\ell,h,\sigma) > 0 \end{array} \right\},$$

with the convention $\mathcal{K}^{\rho}(z;\bar{C}) = 0$ if the set is empty.

The object $\mathcal{K}^{\rho}(z;\bar{C})$ is the largest footprint that can be operated profitably while keeping total organizational complexity below the fixed cap $\bar{C}$.

Proposition 3.1 (Finite-economy boundedness of platform scale)**.** *Suppose Assumption 3.2 holds. Then for every fixed $\rho \geq 1$, $z \geq 0$, and $\bar{C} < \infty$,*

$$\mathcal{K}^{\rho}(z;\bar{C}) \;\leq\; \mu(S^{\rho}) \;<\; \infty.$$

In particular, no finite-granularity economy exhibits literally unbounded feasible or optimal scope.

Proposition 3.2 (Limiting platform emergence)**.** *Fix any complexity cap $\bar{C} < \infty$. Suppose Assumptions 3.2, 3.3, and 3.4 hold, and let Lemma 3.1 apply at that $\bar{C}$. Then for every target scale $K < \infty$, there exist $\rho_K < \infty$ and $z_K < \infty$ such that for all $\rho \geq \rho_K$ and all $z \geq z_K$,*

$$\mathcal{K}^{\rho}(z;\bar{C}) \;\geq\; K.$$

Equivalently, the bounded-complexity scale frontier diverges in the joint limit: for every $\bar{C} < \infty$,

$$\mathcal{K}^{\rho}(z;\bar{C}) \xrightarrow[(\rho,z)\to(\infty,\infty)]{} \infty$$

in the cofinal sense that

$$\forall K < \infty \; \exists(\rho_K, z_K) \text{ such that } \mathcal{K}^\rho(z; \bar{C}) \geq K \quad \forall \rho \geq \rho_K, \; \forall z \geq z_K.$$

Remark 3.5 (Economic interpretation)**.** *Proposition 3.2 is the precise sense in which platform-type organizations emerge in the model. For every finite economy, footprint remains finite. But as near-core opportunities become finer and more abundant and as blueprint quality improves, the firm can profitably replicate a common blueprint across arbitrarily many markets while keeping total organizational complexity below a fixed cap $\bar{C}$. Platform emergence is therefore a statement about the divergence of a bounded-complexity scale frontier, not a literal finite-parameter jump to infinite optimal scope.*

Remark 3.6 (Role of standardization)**.** *The proof of Proposition 3.2 uses the feasible policy $\sigma \equiv 0$. Thus the proposition isolates a pure replication mechanism under abundant near-core opportunities. Additional structure on g and c could make standardization an active driver of the result, but that requires stronger assumptions than those imposed here.*

Figure 6 illustrates Propositions 3.1 and 3.2 jointly. Panel (A) plots the bounded-complexity scale frontier $\mathcal{K}^\rho(z; \bar{C})$ against blueprint quality z for a sequence of granularity levels ρ. For every finite ρ, the frontier saturates at $\mu(S^\rho)$—the dotted horizontal lines—confirming Proposition 3.1. As ρ rises, the ceiling lifts and the frontier shifts outward: with finer near-core granularity, greater scale becomes feasible at each level of blueprint quality. The dashed horizontal line marks an arbitrary target scale K; one can read off the minimum blueprint quality required to reach K at each granularity level, and observe that this threshold falls as ρ increases. Panel (B) displays the same object as a contour map in (ρ, z)-space. The iso-scale curves $\mathcal{K}^\rho = K$ shift leftward as the target K falls, confirming that for any target footprint, sufficiently large (ρ, z) makes it attainable—the cofinal divergence stated in Proposition 3.2.

Figure 7 complements the formal results by illustrating the economic mechanism through which standardization enables platform emergence. Panel (A) plots organizational complexity $C(M, \sigma)$ against footprint $\mu(M)$ under four scenarios of increasing standardization intensity. Without standardization, complexity grows without bound as the firm expands—either linearly when baseline heterogeneity is uniform, or at an accelerating rate when the firm must eventually reach into more distant markets. Moderate standardization slows this growth but does not eliminate it: the complexity cap $\bar{C}$ still binds at some finite footprint. Under strong standardization—the platform regime—effective heterogeneity is reduced sufficiently that $C(M, \sigma)$ converges to a finite

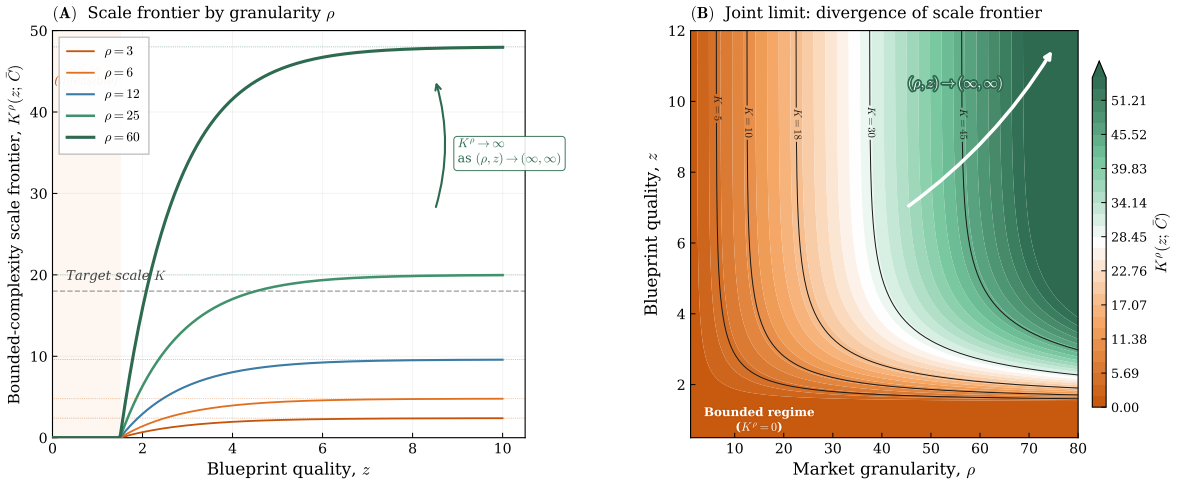


Figure: Platform emergence as asymptotic expansion of the bounded-complexity scale frontier.
Panel (A): Each curve plots $K^\rho(z;\bar{C})$ for a fixed granularity ρ; dashed lines mark $\mu(S^\rho)$. Panel (B): Contour map of K^ρ in (ρ, z)-space; iso-scale curves show how any target K is attainable for sufficiently large (ρ, z).

Figure 6: Platform emergence as an asymptotic expansion of the bounded-complexity scale frontier.

Notes: Panel (A) plots the bounded-complexity scale frontier $\mathcal{K}^\rho(z;\bar{C})$ against blueprint quality z for a sequence of granularity levels ρ. For every finite ρ, feasible scale is finite because the opportunity set S^ρ has finite measure (dotted horizontal lines). As ρ rises, the mass of near-core opportunities expands and the frontier shifts outward. Panel (B) displays $\mathcal{K}^\rho(z;\bar{C})$ as a contour map in (ρ, z)-space. Iso-scale curves show that for any fixed complexity cap $\bar{C}$ and any target scale K, sufficiently large ρ and sufficiently high z make K attainable with strictly positive payoff (Proposition 3.2).

limit, remaining below $\bar{C}$ for all $\mu(M)$. This is the mechanism underlying Proposition 3.2: not a reduction in organizational burden at the margin, but a qualitative change in how that burden scales with scope. Panel (B) translates this into value implications. Below a threshold z^*, the complexity constraint binds and the firm operates in the bounded regime; above z^*, the platform regime dominates and yields strictly higher value, with the gap widening in blueprint quality.

Throughout this section, blueprint quality z is treated as an exogenous technology parameter. Proposition 3.2 should therefore be read as describing how the bounded-complexity scale frontier varies across firms or environments with different organizational technologies and different degrees of near-core market granularity.

4 Implications for the Firm Scope Distribution

This section characterizes how the organizational complexity mechanism shapes the cross-sectional distribution of firm scope. The key object is the mapping from blueprint quality to the firm's optimal scope. To obtain a sharp distributional result, it is useful to specialize the general environment to a benchmark ordered economy in which markets are indexed directly by heterogeneity. The point of the section is not to prove the most general possible theorem, but to isolate

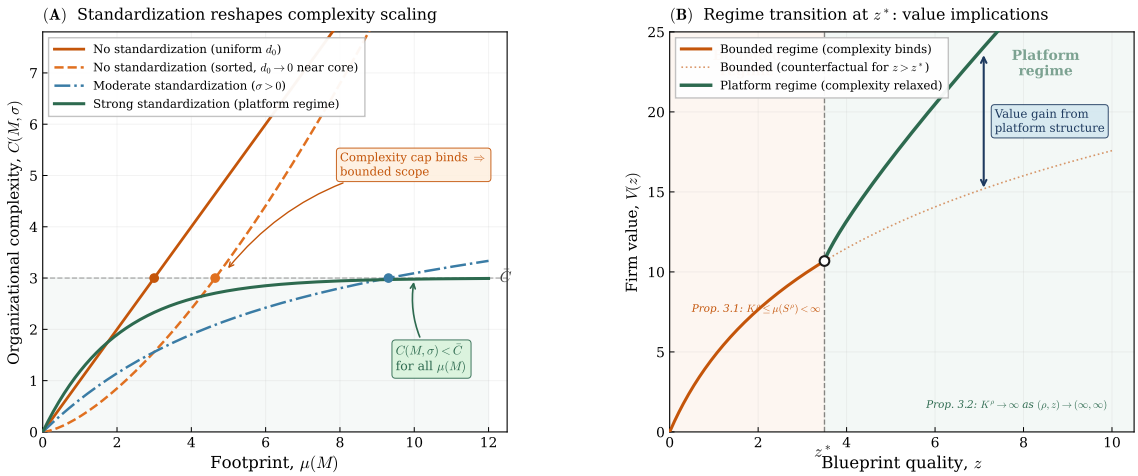


Figure: Standardization mechanism and platform emergence.
Panel (A): Organizational complexity $C(M,\sigma)$ vs. footprint $\mu(M)$ under varying standardization intensity. Strong standardization keeps C bounded as $\mu(M) \to \infty$. Panel (B): Firm value under bounded vs. platform regimes; above z^*, platform structure yields strictly higher value.

Figure 7: Standardization mechanism and platform emergence.

Notes: Panel (A) plots organizational complexity $C(M, \sigma)$ against footprint $\mu(M)$ under varying standardization intensity. Without standardization, complexity grows without bound and the complexity cap $\bar{C}$ binds at finite footprint. Under strong standardization (platform regime), complexity remains uniformly bounded below $\bar{C}$ as $\mu(M) \to \infty$. Panel (B) shows firm value under bounded versus platform regimes. Below z^*, the complexity constraint binds and the bounded regime is optimal. Above z^*, the platform regime yields strictly higher value; the dotted line shows the counterfactual bounded-regime value.

the mechanism that turns capability dispersion into scope dispersion.

4.1 A benchmark mapping from blueprint quality to scope

Consider the benchmark environment in which markets are indexed directly by heterogeneity $d \in [0, \bar{d})$, where $\bar{d} \in (0, \infty]$, and the market measure is Lebesgue measure. Under Proposition 2.4, the optimal footprint of a firm with blueprint quality z takes cutoff form,

$$M^*(z) = \{d \in [0, \bar{d}) : d \leq d^*(z)\},$$

so scope is summarized by the cutoff itself:

$$\Phi(z) \equiv \mu(M^*(z)) = d^*(z).$$

In this benchmark, organizational complexity generated by serving all markets up to heterogeneity d is

$$C(d) = \int_0^d u \, du = \frac{d^2}{2}, \tag{17}$$

where the second equality uses the fact that in this benchmark, markets are indexed directly by their heterogeneity level, so that $d(s) = s$ and the market measure is Lebesgue.

To characterize the scope mapping, it is convenient to work with a reduced-form representation of the marginal market condition.

Assumption 4.1 (Boundary representation)**.** *There exist positive C^1 functions $B, K : [0, \bar{d}) \to \mathbb{R}_+$ such that any interior optimal cutoff $d^*(z)$ satisfies*

$$z\, B(d^*(z)) = K(d^*(z)). \tag{18}$$

Moreover,

$$B'(d) < 0 \qquad \textit{and} \qquad K'(d) > 0 \qquad \textit{for all } d \in (0, \bar{d}).$$

The function $B(d)$ collects the marginal payoff from deploying the blueprint in the boundary market at heterogeneity d, while $K(d)$ collects the full marginal organizational cost of expanding the footprint to include that market. As heterogeneity rises, portability and boundary profitability fall, so $B(d)$ decreases. At the same time, adaptation and congestion burdens rise, so $K(d)$ increases. Equation (18) is the reduced-form cutoff condition: the firm expands until the blueprint-scaled benefit of the boundary market equals its full organizational cost.

Define the required-blueprint function

$$H(d) \equiv \frac{K(d)}{B(d)}. \tag{19}$$

The object $H(d)$ is the blueprint quality required to profitably absorb markets at heterogeneity d.

Proposition 4.1 (Benchmark scope mapping)**.** *Suppose Assumption 4.1 holds and that H defined in (19) is strictly increasing on $[0, \bar{d})$. Then:*

(i) For every z in the range of H, the interior optimal cutoff is unique and satisfies

$$d^*(z) = H^{-1}(z).$$

(ii) The scope function is

$$\Phi(z) = H^{-1}(z),$$

and is therefore strictly increasing in blueprint quality.

(iii) If $\bar{d} < \infty$ *and*

$$\lim_{d \uparrow \bar{d}} H(d) = \infty,$$

then

$$\lim_{z \to \infty} \Phi(z) = \bar{d}.$$

Proposition 4.1 makes the central object of this section explicit. The entire distribution problem reduces to the shape of $H(d)$. Standard models effectively impose a linear or log-linear relationship between capability and scale. Here, the mapping is mediated by heterogeneity: reaching more distant markets requires disproportionately greater blueprint quality because the boundary benefit falls while the boundary organizational cost rises.

4.2 The scope elasticity

The key sufficient statistic for the transformation from blueprint quality to scope is the scope elasticity.

Definition 4 (Scope elasticity)**.** *The scope elasticity at blueprint quality* z *is*

$$\varepsilon_\Phi(z) \equiv \frac{z\,\Phi'(z)}{\Phi(z)} = \frac{d \log \Phi(z)}{d \log z}. \tag{20}$$

In standard scale models with $\Phi(z) = z^\gamma$, the elasticity is constant: $\varepsilon_\Phi(z) = \gamma$ for all z. In the present environment, the scope elasticity is governed by how quickly the required blueprint function $H(d)$ rises with heterogeneity.

Define the generalized elasticity of required blueprint quality with respect to heterogeneity by

$$\mathcal{E}_H(d) \equiv \frac{d\,H'(d)}{H(d)}. \tag{21}$$

Assumption 4.2 (Increasing elasticity of required blueprint quality)**.** *The function* $\mathcal{E}_H(d)$ *is weakly increasing on* $(0, \bar{d})$.

Assumption 4.2 requires that the blueprint quality needed to absorb successively more distant markets rises at an increasingly elastic rate. Economically, this is exactly the organizational complexity mechanism: the further the firm expands from its core, the harder it becomes to convert additional blueprint quality into additional scope.

Proposition 4.2 (Declining scope elasticity). *Suppose Proposition 4.1 holds and Assumption 4.2 is satisfied. Then for every interior z,*

$$\varepsilon_\Phi(z) = \frac{1}{\mathcal{E}_H(\Phi(z))}. \tag{22}$$

In particular, $\varepsilon_\Phi(z)$ is weakly decreasing in z. If $\mathcal{E}_H$ is strictly increasing, then $\varepsilon_\Phi(z)$ is strictly decreasing.

Proposition 4.2 is the paper's main distributional mechanism. High-z firms have already absorbed the easy, near-core markets. Further expansion requires entering markets for which the required blueprint quality rises more than proportionally with distance. As a result, the same proportional improvement in blueprint quality translates into a smaller proportional expansion in scope for larger firms than for smaller firms.

4.3 Tail compression

The tail implications now follow from the shape of the scope mapping.

Corollary 4.1 (Bounded-support case). *Under Proposition 4.1, if $\bar{d} < \infty$ and $\lim_{d \uparrow \bar{d}} H(d) = \infty$, then the scope distribution has bounded support contained in $[0, \bar{d}]$, regardless of whether the blueprint distribution has unbounded support.*

The bounded-support case is the strongest possible form of compression: even arbitrarily large capability differences cannot push scope beyond the finite heterogeneity frontier.

The more interesting case is when the heterogeneity space is unbounded and scope is itself unbounded, but grows sublinearly in capability.

Proposition 4.3 (Tail compression). *Suppose $\bar{d} = \infty$, and blueprint quality z is Pareto with tail index $\alpha_z > 0$:*

$$\Pr(z > x) = \left(\frac{\underline{z}}{x}\right)^{\alpha_z} \qquad \text{for all } x \geq \underline{z}. \tag{23}$$

Assume moreover that the required-blueprint function H is regularly varying with index $\nu > 1$, i.e.

$$H(d) = d^\nu L_H(d), \tag{24}$$

where L_H is slowly varying. Then the induced scope distribution $d^* = \Phi(z)$ is regularly varying with tail index

$$\alpha_{d^*} = \nu \alpha_z > \alpha_z.$$

More precisely, there exists a slowly varying function $\widetilde{L}$ such that

$$\Pr(d^* > m) = \widetilde{L}(m)\, m^{-\nu\alpha_z}. \tag{25}$$

Proposition 4.3 delivers the core result: organizational complexity makes the right tail of firm scope thinner than the right tail of firm capability. A Pareto distribution of blueprint quality is transformed into a Pareto distribution of scope with a strictly larger tail index. The reason is exactly the declining scope elasticity in Proposition 4.2: better firms do expand more, but the mapping from capability to scope becomes flatter as firms move outward from their organizational core.

Corollary 4.2 (Super-polynomial compression)**.** *Suppose z is Pareto with tail index $\alpha_z > 0$ as in (23), and suppose that H grows faster than any power of d, i.e.*

$$\lim_{d\to\infty} \frac{H(d)}{d^p} = \infty \qquad \textit{for every } p > 0.$$

Then for every $\beta > 0$,

$$\Pr(d^* > m) = o(m^{-\beta}) \qquad \textit{as } m \to \infty.$$

In particular, the scope distribution is thinner than any Pareto tail.

Corollary 4.2 is relevant for environments in which portability decays approximately exponentially in heterogeneity. In those cases, the scope distribution is not merely thinner than the capability distribution; it is dramatically thinner.

Relation to standard models. In canonical models based on span-of-control (Lucas, 1978) or monopolistic competition (Melitz, 2003), the mapping from capability to size is typically log-linear or power-like. A firm with $10\times$ the capability is 10^γ times as large. There is no endogenous speed limit. In the present model, the speed limit comes from heterogeneity: the marginal market for a large firm is intrinsically more difficult than the marginal market for a small firm. The right tail is compressed not by exogenous decreasing returns, but by endogenous organizational difficulty.

4.4 Illustration

The mechanism is transparent in a simple benchmark specification. Suppose the required-blueprint function is

$$H(d) = \exp\left(\frac{d}{\lambda}\right)\left(1 + \kappa\frac{d^2}{2}\right)^{\phi}, \qquad \lambda, \kappa, \phi > 0. \tag{26}$$

This reduced form is consistent with two ingredients emphasized throughout the paper: portability decays with distance, and organizational congestion rises with the complexity of the footprint.

Under (26),

$$\mathcal{E}_H(d) = \frac{d}{\lambda} + \phi \frac{\kappa d^2}{1 + \kappa d^2/2},$$

which is strictly increasing in d. Therefore Proposition 4.2 applies, so the scope elasticity declines monotonically with blueprint quality. Moreover, H grows faster than any power of d, so Corollary 4.2 applies: the induced scope distribution is thinner than any Pareto tail. Section 6 calibrates this specification to firm-level scope data and quantifies the degree of tail compression.

The economic message is simple. The model does not require a thin-tailed capability distribution to generate a thin-tailed scope distribution. A fat right tail of blueprint quality can coexist with a much thinner right tail of firm scope because the organizational difficulty of the marginal market rises endogenously with expansion.

5 Organizational Complexity, Market Overlap, and Concentration

Section 4 characterized how organizational complexity compresses the mapping from blueprint quality to firm scope. This section asks how that compression shapes concentration—not merely through the firm size distribution, but through the geometry of firm presence across characteristic space.

The key conceptual point is that aggregate concentration depends on two distinct objects: how large firms are, and how widely the same firm spans distant parts of the market space. A firm that dominates a narrow niche contributes to local but not aggregate concentration. A firm that operates across many distant markets contributes to both. Organizational complexity governs the second margin: by limiting how far any firm can extend from its organizational core, it restricts the ability of a single firm to be present in many distant market slices simultaneously. Figure 8 previews this logic visually: under organizational complexity, footprints are narrow and localized, generating little distant overlap; under frictionless scalability, the same firms span wide arcs of characteristic space, creating regions where many footprints coincide.

To formalize this, we introduce a minimal multi-firm layer. Firm-level behavior is inherited entirely from Sections 2 and 4; the new ingredient is a cross-section of firms with heterogeneous cores in characteristic space. We define an overlap kernel that unifies local and aggregate concentration into a single object, and show that organizational complexity steepens its decay across characteristic distance.

5.1 Multi-firm environment

Characteristic space. Let the characteristic space be the one-dimensional torus $\mathbb{T}_L \equiv \mathbb{R}/L\mathbb{Z}$ of circumference $L > 0$, equipped with Lebesgue measure and the torus metric

$$\delta(s, s') \;\equiv\; \min\big\{|s - s'|,\; L - |s - s'|\big\}. \tag{27}$$

The torus avoids boundary effects and ensures stationarity. Empirically, $\mathbb{T}_L$ represents a stylized product space, technology space, or geographic ring along which markets are arranged.[1]

Firms and footprints. A finite collection of N firms populates the economy. Firm i is characterized by a pair (c_i, z_i), where $c_i \in \mathbb{T}_L$ is the firm's organizational core—the location in characteristic space at which its blueprint originates—and $z_i \geq 0$ is its blueprint quality. Heterogeneity of market s relative to firm i is measured by distance from the core:

$$d_i(s) \;\equiv\; \delta(s, c_i). \tag{28}$$

By the cutoff result of Proposition 2.4 and the scope function of Section 4, firm i's optimal footprint is the interval

$$M_i \;=\; \big\{s \in \mathbb{T}_L : \delta(s, c_i) \leq R_i\big\}, \qquad R_i \;\equiv\; \Phi(z_i), \tag{29}$$

where R_i is the firm's scope radius, determined by the scope function Φ from Section 4. We assume $R_i < L/2$ for all i, so that no firm wraps fully around the torus.

Core locations. Core locations $\{c_i\}_{i=1}^N$ are treated as exogenous. We do not model entry, exit, or endogenous core placement. This is deliberately minimal: the goal is to isolate the geometric consequences of scope compression for concentration, not to characterize equilibrium core configurations.

Figure 8 illustrates the multi-firm environment for a configuration of seven firms. Panel (A) shows footprints under organizational complexity: each firm's scope radius $R_i = \Phi(z_i)$ is compressed by the declining-elasticity mechanism of Section 4, producing narrow intervals centered on each firm's core. Footprints are largely non-overlapping, and no firm spans a large fraction of the characteristic space. Panel (B) shows the same firms under frictionless scalability, where $\Phi(z) = z$ and scope radii are proportional to blueprint quality. Several firms now cover wide arcs,

[1] The one-dimensional torus is a benchmark for tractability. The definitions and the overlap-kernel identity (Proposition 5.1) extend to any compact metric space with a transitive isometry group; the closed-form coverage kernel in Section 5.4 is specific to the interval-footprint structure on $\mathbb{T}_L$.

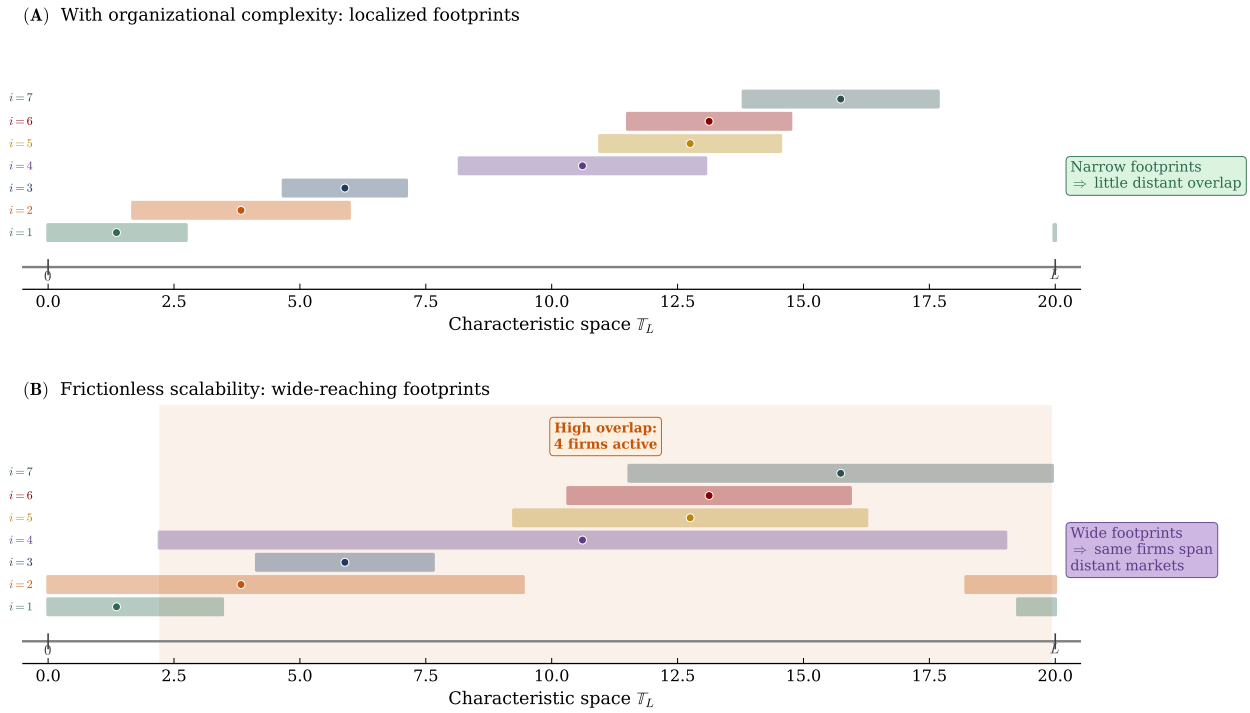


Figure 8: Footprint overlap in characteristic space under organizational complexity versus frictionless scalability.

Notes: Each horizontal bar represents a firm's footprint M_i on the characteristic space $\mathbb{T}_L$, centered at its organizational core (dot). Panel (A): Under organizational complexity, scope radii $R_i = \Phi(z_i)$ are compressed, producing narrow, localized footprints with little overlap between distant market locations. Panel (B): Under frictionless scalability ($\Phi(z) = z$), the same firms have much wider footprints, creating regions where multiple firms are simultaneously active across distant parts of the characteristic space.

and a region near the center of the space is served by four firms simultaneously. The contrast between the two panels is the geometric mechanism that the remainder of this section formalizes: organizational complexity limits the reach of individual firms, reducing same-firm overlap at distant market locations.

5.2 Local shares and concentration

Market-level activity. Let $y_i(s) \geq 0$ denote the gross activity of firm i at market s, with $y_i(s) = 0$ whenever $s \notin M_i$. Total activity at s is

$$Y(s) \equiv \sum_{i=1}^{N} y_i(s). \tag{30}$$

At any market s with $Y(s) > 0$, define firm i's local activity share by

$$\omega_i(s) \;\equiv\; \frac{y_i(s)}{Y(s)}, \qquad \sum_{i=1}^{N} \omega_i(s) = 1. \tag{31}$$

Definition 5 (Local concentration)**.** *The local Herfindahl–Hirschman index at market s is*

$$\mathrm{HHI}(s) \;\equiv\; \sum_{i=1}^{N} \omega_i(s)^2. \tag{32}$$

Average local concentration is

$$\overline{\mathrm{HHI}}_{\mathrm{loc}} \;\equiv\; \frac{1}{L}\int_{\mathbb{T}_L} \mathrm{HHI}(s)\, ds \;=\; \frac{1}{L}\int_{\mathbb{T}_L} \sum_{i=1}^{N} \omega_i(s)^2\, ds. \tag{33}$$

Definition 6 (Aggregate concentration)**.** *Define firm i's aggregate activity share by*

$$\bar{\omega}_i \;\equiv\; \frac{1}{L}\int_{\mathbb{T}_L} \omega_i(s)\, ds. \tag{34}$$

The aggregate Herfindahl–Hirschman index is

$$\mathrm{HHI}_{\mathrm{agg}} \;\equiv\; \sum_{i=1}^{N} \bar{\omega}_i^2. \tag{35}$$

Both objects are standard. The local HHI measures concentration within a market slice; the aggregate HHI measures concentration economy-wide. The overlap kernel defined below provides a single function that connects them.

5.3 The overlap kernel

Definition 7 (Overlap kernel)**.** *The same-firm overlap kernel is the function $\mathcal{O} : \mathbb{T}_L \to [0, 1]$ defined by*

$$\mathcal{O}(x) \;\equiv\; \frac{1}{L}\int_{\mathbb{T}_L} \sum_{i=1}^{N} \omega_i(s)\, \omega_i(s + x)\, ds. \tag{36}$$

The overlap kernel measures the average extent to which the same firm is present at two market locations separated by characteristic distance x. When $x = 0$, the integrand reduces to $\sum_i \omega_i(s)^2 =$

$\mathrm{HHI}(s)$. When x is large, markets s and $s+x$ are organizationally distant, and few firms span both, so $\mathcal{O}(x)$ is small.

Proposition 5.1 (Overlap kernel identity)**.** *The overlap kernel satisfies:*

(i) ***Local concentration.***

$$\mathcal{O}(0) \;=\; \overline{\mathrm{HHI}}_{\mathrm{loc}}. \tag{37}$$

(ii) ***Aggregate concentration.***

$$\frac{1}{L}\int_{\mathbb{T}_L} \mathcal{O}(x)\,dx \;=\; \mathrm{HHI}_{\mathrm{agg}}. \tag{38}$$

Proof. See Appendix A.7.

Proposition 5.1 shows that a single function—the overlap kernel—encodes both local and aggregate concentration. The value at zero is average local HHI; the average over characteristic distance is aggregate HHI. Aggregate concentration is therefore a summary of same-firm overlap across all characteristic distances, and any force that steepens the decay of $\mathcal{O}$ lowers aggregate concentration while potentially leaving local concentration unchanged.

Remark 5.1 (Interpretation as a correlation function)**.** *The overlap kernel is the autocorrelation function of firm presence. In spatial statistics, the rate at which an autocorrelation decays measures the range of dependence. Here, a slowly decaying $\mathcal{O}$ means that knowing a firm is present at s is informative about its presence at distant s'—i.e., the same firms span many distant markets. A rapidly decaying $\mathcal{O}$ means firm presence is localized.*

5.4 A tractable benchmark: interval footprints with exogenous cores

To connect the overlap kernel to organizational complexity, specialize to a benchmark in which the kernel admits a closed form.

Assumption 5.2 (Benchmark multi-firm economy)**.**

(i) Firms' organizational cores $\{c_i\}_{i=1}^N$ are drawn independently and uniformly on $\mathbb{T}_L$.

(ii) Blueprint qualities $\{z_i\}_{i=1}^N$ are drawn independently from a distribution F_z, independently of $\{c_i\}$.

(iii) The scope function Φ is the mapping from Section 4, so that scope radii $R_i = \Phi(z_i)$ inherit the declining-elasticity property of Proposition 4.2.

(iv) Market-level activity is symmetric within the footprint: $y_i(s) = y(z_i, \delta(s, c_i))$ for $s \in M_i$ and $y_i(s) = 0$ otherwise, where $y(z, d) > 0$ for $d < \Phi(z)$.

Under Assumption 5.2, the local share $\omega_i(s)$ at any market s depends only on which firms' footprints cover s and their respective activity levels there. To isolate the geometric mechanism cleanly, it is useful to work with the coverage overlap kernel, which tracks whether—not how intensively—the same firm is present at two locations.

Definition 8 (Coverage overlap kernel)**.** *The coverage overlap kernel is*

$$\mathcal{K}(x) \;\equiv\; \frac{1}{L}\int_{\mathbb{T}_L} \sum_{i=1}^{N} \mathbf{1}\{s \in M_i\}\,\mathbf{1}\{s+x \in M_i\}\,ds. \tag{39}$$

$\mathcal{K}(x)$ counts, on average across locations s, how many firms simultaneously serve both s and $s+x$. It is the extensive-margin analogue of the share-weighted overlap kernel $\mathcal{O}$.

Proposition 5.2 (Coverage overlap in the interval-footprint benchmark)**.** *Under Assumption 5.2, the coverage overlap kernel is*

$$\mathcal{K}(x) \;=\; \frac{1}{L}\sum_{i=1}^{N}(2R_i - x)_+, \tag{40}$$

where $(a)_+ \equiv \max\{a, 0\}$ *and* $R_i = \Phi(z_i)$*. Taking expectations over the random configuration of cores and blueprint qualities,*

$$\mathbb{E}\big[\mathcal{K}(x)\big] \;=\; \frac{N}{L}\,\mathbb{E}\big[(2R - x)_+\big] \;=\; \frac{N}{L}\,\mathbb{E}\big[(2\Phi(z) - x)_+\big], \tag{41}$$

where the expectation is over a single draw $z \sim F_z$*.*

Proof. See Appendix A.7.

Equation (41) is the section's key formula. It expresses the expected coverage overlap at characteristic distance x as a function of the distribution of scope radii $R = \Phi(z)$. Three properties are immediate:

(a) $\mathbb{E}[\mathcal{K}(x)]$ is weakly decreasing in x: overlap falls with characteristic distance.
(b) $\mathbb{E}[\mathcal{K}(x)]$ is increasing in the scope radius R: firms with larger footprints generate more overlap.
(c) $\mathbb{E}[\mathcal{K}(x)] = 0$ for all $x \geq 2\max_i R_i$: overlap vanishes beyond twice the largest firm's radius.

Organizational complexity enters through property (b): by compressing $\Phi(z)$, it reduces the scope radii and thereby steepens the decay of $\mathbb{E}[\mathcal{K}(x)]$ in characteristic distance.

5.5 Organizational complexity steepens overlap decay

The following result formalizes how scope compression translates into overlap and concentration.

Proposition 5.3 (Organizational complexity reduces overlap and concentration)**.** *Let Φ and $\widetilde{\Phi}$ be two scope functions satisfying $\widetilde{\Phi}(z) \geq \Phi(z)$ for all z, with strict inequality on a set of positive F_z-measure. Denote the associated scope radii by $R_i = \Phi(z_i)$ and $\widetilde{R}_i = \widetilde{\Phi}(z_i)$, and the corresponding coverage overlap kernels by $\mathcal{K}$ and $\widetilde{\mathcal{K}}$. Then:*

(i) ***Pointwise overlap comparison.*** *For all $x \geq 0$,*

$$\mathbb{E}\big[\widetilde{\mathcal{K}}(x)\big] \;\geq\; \mathbb{E}\big[\mathcal{K}(x)\big], \tag{42}$$

with strict inequality for all $x \in \big(0,\, 2\max \mathrm{supp}(\widetilde{R})\big)$.

(ii) ***Aggregate concentration comparison.***

$$\mathbb{E}\left[\frac{1}{L}\int_{\mathbb{T}_L} \widetilde{\mathcal{K}}(x)\,dx\right] \;>\; \mathbb{E}\left[\frac{1}{L}\int_{\mathbb{T}_L} \mathcal{K}(x)\,dx\right]. \tag{43}$$

In particular, let Φ be the scope function under organizational complexity (Section 4) and $\widetilde{\Phi}$ the frictionless benchmark with $\widetilde{\Phi}(z) = z$. Then organizational complexity reduces the coverage overlap kernel pointwise and lowers aggregate coverage concentration.

Proof. See Appendix A.7.

Proposition 5.3 delivers the section's main result. Organizational complexity limits concentration not merely by making firms smaller, but by limiting the spatial reach of individual firms across characteristic space. A firm with compressed scope is a local competitor: it dominates markets near its core but does not extend far enough to contribute to concentration at distant market locations. By contrast, a firm with frictionless scalability can span many distant market slices, generating same-firm overlap at large characteristic distances and thereby raising aggregate concentration.

The overlap kernel makes this distinction precise. Under organizational complexity, $\mathcal{O}(x)$ and $\mathcal{K}(x)$ decay rapidly in x: the same firm is rarely present at two markets that are organizationally far apart. Under frictionless scalability, the kernels decay more slowly, meaning the same firms appear across a wider swath of the economy. Integrating, aggregate concentration is lower under organizational complexity. Section 6 calibrates the model to firm-level scope data and quantifies the magnitude of this concentration gap.

Remark 5.3 (Scope compression versus core clustering)**.** *Proposition 5.3 compares economies that differ in the scope function Φ while holding the distribution of cores fixed. It does not address endogenous core clustering. If firms with similar organizational technologies cluster in characteristic space, local concentration could rise even as organizational complexity limits each firm's reach. The proposition isolates the scope channel: holding the geography of cores constant, scope compression steepens the overlap decay and lowers aggregate concentration.*

Section 4 showed that organizational complexity compresses the right tail of the scope distribution. In the present framework, this tail compression has a geometric interpretation: the largest scope radii—which generate overlap at the greatest characteristic distances—are precisely the ones most compressed by organizational complexity. Aggregate concentration therefore falls for two reasons: there are fewer very large firms (thinner tails), and those large firms reach less far across characteristic space (steeper overlap decay). The overlap kernel captures both channels in a single object.

Connection to aggregation. The overlap kernel provides a natural bridge between local and aggregate concentration. An economy in which each firm dominates a narrow niche can have high local HHI—many local monopolies—while aggregate HHI is low because no single firm accounts for a large share of the whole. Conversely, an economy of a few platform firms, each spanning most of characteristic space, may have moderate local HHI but high aggregate HHI. The overlap kernel encodes exactly which regime obtains: a slowly decaying $\mathcal{O}$ indicates platform-like economies with high aggregate concentration; a rapidly decaying $\mathcal{O}$ indicates localized economies with dispersed aggregate activity. Organizational complexity, by compressing scope, pushes the economy toward the second regime.

We emphasize that this section studies how organizational complexity shapes the geometry of multi-firm overlap in characteristic space. It takes firm cores and blueprint qualities as given and does not model strategic interaction, entry, or pricing. The concentration objects are Herfindahl indices of activity shares, not welfare-relevant markups or consumer surplus. A fuller treatment would endogenize core locations and embed local competition within each market slice. That extension is natural but distinct from the geometric aggregation question studied here.

6 Empirical Relevance

This section connects the organizational complexity framework to data. We calibrate the benchmark specification from Section 4 to match the distribution of firm scope documented by Hoberg

and Phillips (2025), then use the calibrated model to quantify tail compression and the local-versus-aggregate concentration gap. The exercise is an illustration, not a structural estimation: the goal is to show that the mechanism is quantitatively plausible and to discipline the key model objects with firm-level evidence.

6.1 Calibration

We use the benchmark required-blueprint function from Section 4:

$$H(d) = \exp\left(\frac{d}{\lambda}\right)\left(1+\kappa\frac{d^2}{2}\right)^{\phi}, \qquad \lambda, \kappa, \phi > 0,$$

and draw blueprint quality from a Pareto distribution with tail index α_z and lower bound $\underline{z} = 1$. The scope of a firm with blueprint quality z is $d^*(z) = \Phi(z) = H^{-1}(z)$, computed by numerical inversion.

We target three moments of the firm scope distribution from Hoberg and Phillips (2025, Table 1), who measure scope as the number of text-based product market segments (D2V-Scope) for U.S. public firms over 1989–2017: mean scope of 7.5, median scope of 6.0, and standard deviation of 5.8. Table 1 reports the calibrated parameter values obtained by minimizing the equally weighted sum of squared percentage deviations from these three targets.

The calibrated Pareto tail index $\alpha_z = 1.36$ is in the range of standard estimates for the upper tail of the firm productivity distribution. The portability decay parameter $\lambda = 18.1$ and the congestion parameters $\kappa = 0.005$ and $\varphi = 1.96$ imply that both channels—portability decay and organizational congestion—contribute meaningfully to scope resistance. At the median firm's scope, organizational complexity (the $(1+\kappa d^2/2)^{\varphi}$ term) accounts for 34% of the total log blueprint requirement $\log H(d)$; this share rises to 58% at the 99th percentile of scope, confirming that the complexity channel is quantitatively important precisely where it matters most—for the largest firms.

6.2 Scope distribution and related expansion

Figure 9 shows the model-implied scope distribution under the calibrated parameters. Panel (a) plots the density: the distribution is right-skewed, with mass concentrated between 1 and 15 industries and a tail extending to approximately 32—matching the range in Hoberg and Phillips (2025, Table 1). Panel (b) plots the complementary CDF on log-log scale. The scope distribution decays with an approximate slope of -1.7, substantially steeper than the capability distribution from

Table 1: Calibrated parameters and targeted moments

	Target (HP)	Model
Panel A: Targeted moments		
Mean scope	7.546	7.356
Median scope	6.000	6.090
Std. dev. of scope	5.800	5.767
Panel B: Calibrated parameters		
Pareto tail index (α_z)	—	1.357
Portability decay (λ)	—	18.108
Congestion weight (κ)	—	0.005
Congestion elasticity (φ)	—	1.958
Panel C: Non-targeted moments		
P99 scope	—	26.11
P90 scope	—	14.98
Skewness	—	1.47

Notes: Moments are from Hoberg and Phillips (2025, Table 1). D2V-Scope measures the number of text-based product market segments for U.S. public firms, 1989–2017 ($N = 101{,}535$ firm-year observations). Parameters are chosen to minimize equally weighted squared percentage deviations from the three targeted moments, using 100,000 simulated firms drawn from a Pareto distribution with lower bound $\underline{z} = 1$. The required-blueprint function is $H(d) = \exp(d/\lambda)\,(1 + \kappa\, d^2/2)^{\varphi}$. At the median scope, organizational complexity accounts for 34% of total blueprint resistance $\log H(d)$; this share rises to 58% at the 99th percentile.

which it is generated.

The model's cutoff structure (Proposition 2.4) implies that scope expansion is concentrated in related markets: firms absorb activities close to their organizational core before reaching into more heterogeneous environments. This matches the central finding of Hoberg and Phillips (2025), who document that average product market scope among U.S. public firms increased by roughly 60% between 1989 and 2017, with essentially all of the increase accounted for by expansion into closely related product markets. Crucially, this expansion occurred without an increase in the number of traditional operating segments reported in Compustat—consistent with the model's prediction that scope growth occurs along the intensive margin of relatedness, not through diversification into organizationally distant activities.

The model also predicts that scope expansion operates through intangible investment—improving the non-rival blueprint z or investing in standardization—rather than physical capital, which enters only through rival market-level inputs $x(s)$. Hoberg and Phillips (2025) test this directly using plausibly exogenous variation in scope-expansion incentives and find that firms with stronger incentives to expand scope subsequently increase R&D spending and pursue more ac-

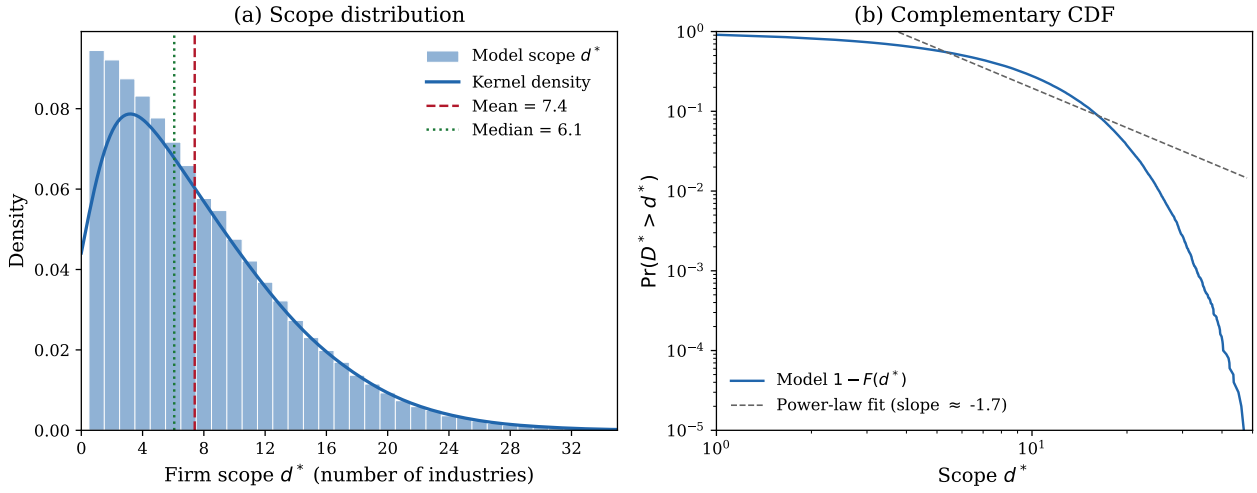


Figure 9: Calibrated scope distribution

Notes: Panel (a) plots the density of firm scope $d^* = \Phi(z)$ for 100,000 simulated firms under the calibrated parameters (Table 1). Dashed and dotted lines mark the mean (7.4) and median (6.1). Panel (b) plots the complementary CDF on log-log scale; the dashed line shows the best-fit power law with slope ≈ -1.7.

quisitions, but show no increase in capital expenditures. Separately, Argente et al. (2025) measure product-level standardization in the consumer packaged goods sector and find that firms reusing more product characteristics across their product lines achieve both greater scope and flatter marginal cost schedules.

6.3 Tail compression: quantifying the missing giants

Figure 10 plots the complementary CDFs of capability and scope jointly, making the tail compression mechanism from Proposition 4.3 quantitatively explicit. The capability distribution (Pareto with $\alpha_z = 1.36$) decays with slope -1.36 on the log-log scale. The scope distribution, generated by passing capability through the concave scope function $\Phi(z) = H^{-1}(z)$, decays much more steeply—with an approximate slope of -1.7 in the body of the distribution, curving away from any power law in the far right tail as Corollary 4.2 predicts.

The concrete implication is summarized in the annotation: a firm with $10\times$ the blueprint quality of another achieves only $2.1\times$ the scope. The "missing giants" in the right tail of the firm scope distribution are not absent because highly capable firms do not exist, but because organizational complexity makes the last units of scope disproportionately expensive. This provides a structural explanation for why the empirical firm size distribution exhibits a thinner right tail than standard models with log-linear capability-to-size mappings would predict.

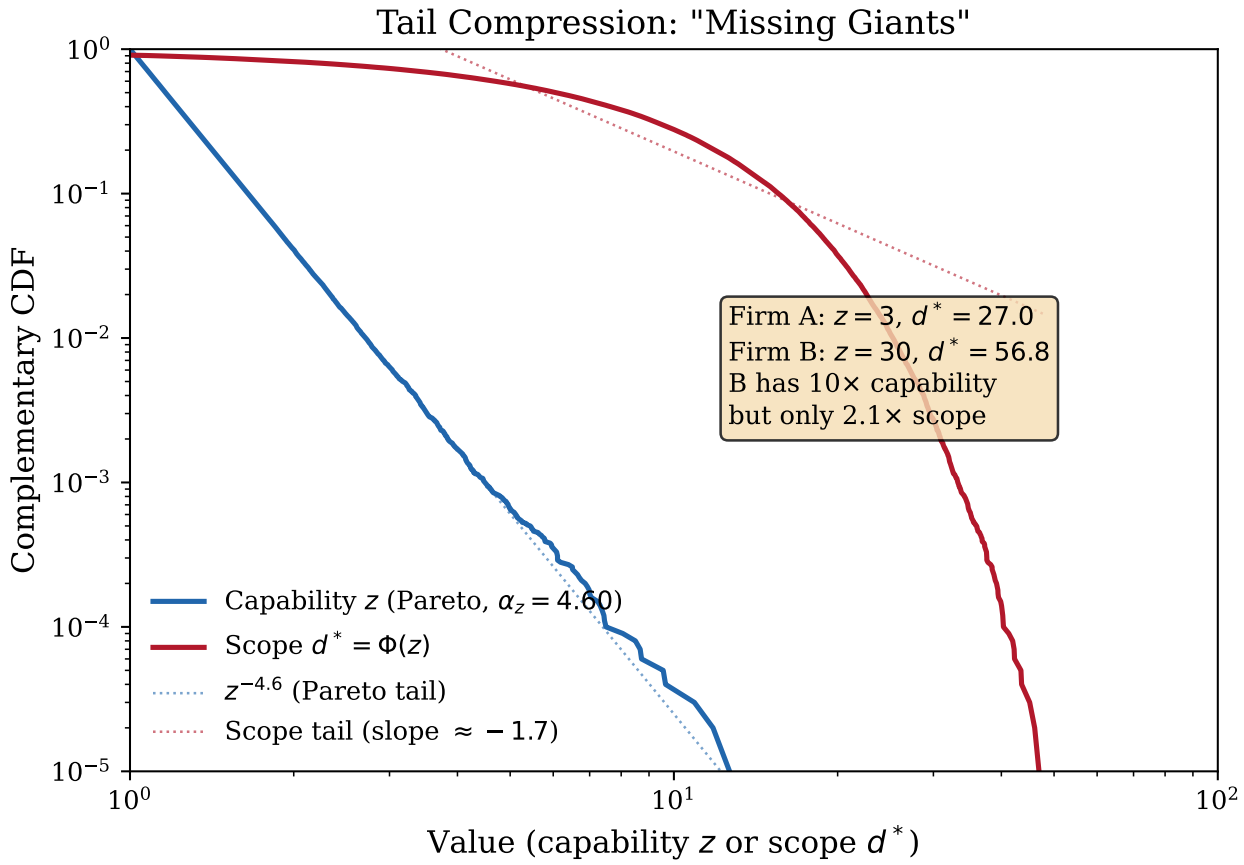


Figure 10: Tail compression: "missing giants"

Notes: Complementary CDFs of blueprint quality z (Pareto, $\alpha_z = 1.36$) and implied scope $d^* = \Phi(z)$ under calibrated parameters (Table 1). Dotted lines show power-law reference slopes. The annotation illustrates the compression: a firm with 10× the blueprint quality achieves only 2.1× the scope.

6.4 Local dominance versus aggregate concentration

To quantify the concentration implications of organizational complexity, we embed the calibrated scope function in the multi-firm torus model of Section 5. We place N firms with organizational cores uniformly distributed on a torus of circumference $L = 150$, assign each firm a scope radius $R_i = \Phi(z_i)$ drawn from the calibrated distribution, and compute the expected coverage overlap kernel $\mathbb{E}[\mathcal{K}(x)]$ from Proposition 5.2. As a benchmark, we compare against a frictionless economy in which $\Phi(z) = z$ (scope equals capability).

Figure 11 presents the results. Panel (a) plots the expected coverage overlap kernel against characteristic distance. Both models generate similar local overlap at $x = 0$: the same number of firms is active in any given market. But the calibrated model's kernel decays to zero by $x \approx 35$,

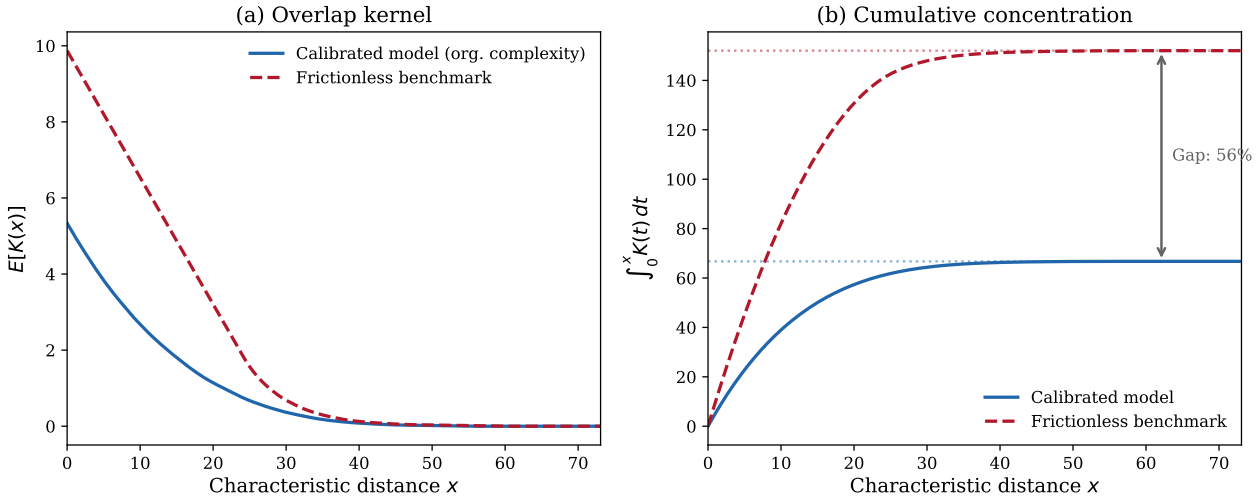


Figure 11: Coverage overlap and aggregate concentration: calibrated model versus frictionless benchmark

Notes: Panel (a) plots the expected coverage overlap kernel $\mathbb{E}[\mathcal{K}(x)]$ against characteristic distance x for the calibrated model (solid) and a frictionless benchmark with $\Phi(z) = z$ (dashed). Panel (b) plots the cumulative integral $\int_0^x \mathcal{K}(t)\,dt$, whose plateau is proportional to aggregate coverage concentration. The gap between the two curves at the plateau is 56%.

while the frictionless kernel persists to $x \approx 70$. Organizational complexity compresses the reach of individual firms, eliminating same-firm overlap at moderate and large characteristic distances.

Panel (b) makes the aggregation precise. The cumulative integral $\int_0^x \mathcal{K}(t)\,dt$—whose plateau is proportional to aggregate coverage concentration—is 56% lower under organizational complexity than under frictionless scalability. Both economies feature similar local concentration (firms dominate their core markets), but the calibrated model generates substantially lower aggregate concentration because the same firms do not appear everywhere.

This result connects directly to the empirical finding of Hoberg and Phillips (2025), who construct scope-adjusted HHI measures and find that, once multi-market scope is properly accounted for, horizontal concentration has not increased over their sample period—despite the well-documented rise in traditional industry-level HHI. In the language of the model, the increase in traditional HHI reflects rising local concentration (firms dominating their core markets), while scope-adjusted HHI captures aggregate concentration across the full characteristic space, which organizational complexity holds in check.

6.5 Organizational technology and scalability

Two additional predictions of the model are qualitatively supported by existing evidence but have not yet been tested directly with the calibrated specification.

First, Proposition 2.5 predicts that headquarters attention is concentrated in low-heterogeneity core activities, generating a within-firm effectiveness gradient. Direct evidence on this prediction is scarce because it requires observing how a single firm allocates managerial resources across activities of varying organizational distance. The prediction is, however, consistent with the finding of Bloom et al. (2019) that management practices vary substantially across establishments within the same firm, and with the Hoberg and Phillips (2025) finding that scope-expanding firms do not increase their reported operating segments—suggesting peripheral activities are absorbed into existing organizational structures rather than managed as distinct divisions.

Second, Proposition 3.2 predicts that a firm's ability to expand scope cheaply depends on the quality of its organizational technology, not merely on its current size. Argente et al. (2025) provide evidence consistent with this prediction: using plausibly exogenous demand variation (Chinese import competition and regional housing shocks), they find that more standardized firms—conditional on size—expand scope more aggressively in response to demand shocks.

A sharper test of the core–periphery prediction would require firm-level data linking the organizational distance of each activity from the firm's core to the intensity of managerial resources allocated to it. Estimating the concavity of the scope function directly—for example by regressing text-based scope measures on proxies for organizational capability—would provide a direct test of the tail compression mechanism. And the overlap kernel itself can in principle be estimated from the same text-based data used to measure scope, providing a new empirical object for studying the geometry of multi-firm competition.

7 Conclusion

This paper has proposed organizational complexity as the binding constraint on firm scope. The central insight is that the cost of expansion depends not on how many markets a firm serves, but on the cumulative heterogeneity of those markets relative to the firm's organizational core. This shift—from count to complexity—resolves a puzzle that count-based theories cannot address: why firms can scale across vast numbers of similar activities at negligible marginal cost, yet face sharp limits when expanding into materially distinct environments. The framework delivers a unified account of bounded scope, related diversification, core–periphery internal structure, and the discontinuous emergence of platform-type organizations.

The organizational complexity mechanism also reshapes how we think about the firm size distribution. In standard models, a fat-tailed distribution of productivity translates directly into a fat-tailed distribution of firm size. Here, that mapping is mediated by organizational constraints:

larger firms face more heterogeneous marginal markets, so capability differences translate into progressively smaller scope differences as firms grow. The result is tail compression—fewer very large firms than the productivity distribution alone would predict. The "missing giants" in the right tail are not absent because highly capable firms do not exist, but because organizational complexity prevents capability from fully expressing itself as scope.

The same mechanism that compresses individual firm scope also limits aggregate concentration through a geometric channel. Because organizational complexity restricts the reach of individual firms across characteristic space, same-firm overlap at distant market locations decays steeply. Aggregate concentration—which integrates this overlap across all characteristic distances—is therefore lower than it would be under frictionless scalability, even when local concentration is high. This distinction between localized dominance and economy-wide reach provides a structural lens through which to interpret the coexistence of rising local market power with moderate aggregate concentration.

The quantitative illustration disciplines these predictions with data. Calibrated to match the distribution of firm scope in Hoberg and Phillips (2025), the model generates a scope distribution whose mean, median, and dispersion match the data, with organizational complexity accounting for over half of the blueprint resistance faced by the largest firms. A firm with ten times the organizational capability of another achieves only twice the scope, and organizational complexity reduces aggregate concentration by 56% relative to a frictionless benchmark. These magnitudes suggest that the mechanism is not merely a theoretical possibility but a quantitatively important force shaping the cross-section of firm scope and the geometry of market competition.

References

Antràs, P. and E. Rossi-Hansberg (2009): "Organizations and Trade," Annual Review of Economics, 1, 43–64.

Argente, D., S. Moreira, E. Oberfield, and V. Venkateswaran (2025): "Scalable Expertise: How Standardization Drives Scale and Scope," Working Paper 34160, National Bureau of Economic Research.

Bernard, A. B., S. J. Redding, and P. K. Schott (2010): "Multiple-Product Firms and Product Switching," American Economic Review, 100, 70–97.

——— (2011): "Multiproduct Firms and Trade Liberalization," The Quarterly Journal of Economics, 126, 1271–1318.

Bloom, N., R. Sadun, and J. Van Reenen (2019): "What Drives Differences in Management Practices?" American Economic Review, 109, 1648–1683.

Bloom, N. and J. Van Reenen (2007): "Measuring and Explaining Management Practices across Firms and Countries," Quarterly Journal of Economics, 122, 1351–1408.

Bresnahan, T. F. and S. Greenstein (1999): "Technological Competition and the Structure of the Computer Industry," Journal of Industrial Economics, 47, 1–40.

Crouzet, N., J. Eberly, A. Eisfeldt, and D. Papanikolaou (2024): "Intangible Capital, Firm Scope, and Growth," Working paper.

Dyrda, S., G. Hong, M. A. Sajid, and J. B. Steinberg (2025): "Intangible Capital, Multinational Production, and Productivity Growth," Working paper.

Garicano, L. (2000): "Hierarchies and the Organization of Knowledge in Production," Journal of Political Economy, 108, 874–904.

Garicano, L. and E. Rossi-Hansberg (2006): "Organization and Inequality in a Knowledge Economy," The Quarterly Journal of Economics, 121, 1383–1435.

Gawer, A. and M. A. Cusumano (2014): "Industry Platforms and Ecosystem Innovation," Journal of Product Innovation Management, 31, 417–433.

Hoberg, G. and G. M. Phillips (2025): "Scope, Scale, and Concentration: The 21st-Century Firm," The Journal of Finance, 80, 415–466.

Lucas, R. E. (1978): "On the Size Distribution of Business Firms," The Bell Journal of Economics, 9, 508–523.

McGrattan, E. R. and E. C. Prescott (2010): "Unmeasured Investment and the Puzzling U.S. Boom in the 1990s," American Economic Journal: Macroeconomics, 2, 88–123.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," Econometrica, 71, 1695–1725.

Radner, R. (1993): "The Organization of Decentralized Information Processing," Econometrica, 61, 1109–1146.

Rochet, J.-C. and J. Tirole (2003): "Platform Competition in Two-Sided Markets," Journal of the European Economic Association, 1, 990–1029.

Zandt, T. V. (1999): "Real-Time Decentralized Information Processing as a Model of Organizations with Boundedly Rational Agents," Review of Economic Studies, 66, 633–658.

Appendix

A Proofs

A.1 Proofs for Section 2.1.

Proof of Proposition 2.1. We proceed by contradiction. Suppose there exists an optimal solution $(M^\star, z^\star, \{x^\star, \ell^\star, h^\star\}_{s\in M^\star})$ with

$$C(M^\star) = \int_{M^\star} d(s)\,\mu(ds) = \infty.$$

Step 1: Total gross payoff is finite at the candidate optimum. Let $(M^\star, z^\star, \{x^\star, \ell^\star, h^\star\})$ be the candidate optimal solution. By Assumption 1.7 (Conditional Integrability), for the fixed blueprint quality $z^\star$, there exists an integrable envelope $\overline{R}_{z^\star} \in L^1(S,\mu)$. Therefore, the total gross payoff over the served set $M^\star$ is finite at the candidate optimum:

$$\int_{M^\star} R(z_s^\star, x^\star(s))\,\mu(ds) \;\le\; \int_{M^\star} \overline{R}_{z^\star}(s)\,\mu(ds) \;\equiv\; B(z^\star) \;<\; \infty.$$

Note that while $B(z^\star)$ is finite for the fixed $z^\star$, we do not assume it is uniformly bounded for all z.

Step 2: Positive payoff in any $d(s) > 0$ market requires $h(s) > 0$. Fix a market s with $d(s) > 0$. By Assumption 2.2,

$$\theta(A(0,\ell), d(s)) = 0 \qquad \forall\, \ell \ge 0.$$

Therefore, if $h^\star(s) = 0$, then $a^\star(s) = A(0, \ell^\star(s))$ and the deployment constraint implies

$$z_s^\star = \theta(a^\star(s), d(s))\, z^\star = \theta(A(0,\ell^\star(s)), d(s))\, z^\star = 0.$$

Since $x^\star(s)$ is feasible (by (11)), Assumption 2.3 yields

$$R(z_s^\star, x^\star(s)) = R(0, x^\star(s)) \le 0.$$

Thus, for any market with $d(s) > 0$,

$$R(z_s^\star, x^\star(s)) > 0 \quad\Longrightarrow\quad h^\star(s) > 0. \tag{44}$$

Step 3: Any nontrivial optimum uses strictly positive total HQ attention. Let total HQ attention at $(M^\star, z^\star, h^\star)$ be

$$H^\star \;\equiv\; \chi(z^\star) + \int_{M^\star} h^\star(s)\,\mu(ds) \;\geq\; 0.$$

If $H^\star = 0$, then $\chi(z^\star) = 0$ implies $z^\star = 0$ (since χ is increasing with $\chi(0) = 0$), and $h^\star(s) = 0$ for μ-a.e. $s \in M^\star$, implying $z_s^\star = 0$ for μ-a.e. $s \in M^\star$. Hence, by Assumption 2.3,

$$\int_{M^\star} R(z_s^\star, x^\star(s))\,\mu(ds) = \int_{M^\star} R(0, x^\star(s))\,\mu(ds) \leq 0.$$

But the firm can always choose $M = \emptyset$ and $z = 0$, achieving payoff 0 at zero cost. Therefore, any optimal solution that attains a strictly positive objective value must satisfy $H^\star > 0$.

Step 4: Infinite complexity implies arbitrarily large unit HQ cost. By assumption, $q'(C) > 0$ and $\lim_{C\to\infty} q(C) = \infty$. Since $C(M^\star) = \infty$, we have

$$q(C(M^\star)) = \infty,$$

in the sense that for every $B > 0$ there exists a (finite) complexity level $\tilde{C}$ large enough such that $q(\tilde{C}) > B$, and $C(M^\star)$ exceeds $\tilde{C}$.

Step 5: Contradiction. The objective value at the candidate optimum equals

$$\int_{M^\star} R(z_s^\star, x^\star(s))\,\mu(ds) - q\big(C(M^\star)\big)\Big(\chi(z^\star) + \int_{M^\star} h^\star(s)\,\mu(ds)\Big) \;=\; \int_{M^\star} R(z_s^\star, x^\star(s))\,\mu(ds) - q\big(C(M^\star)\big)\,H^\star.$$

By Step 1, the first term is bounded above by $B(z^\star) < \infty$. By Step 3, $H^\star > 0$ for any nontrivial optimum. Since $q(C) \to \infty$ as $C \to \infty$ (Assumption 1.6) and $C(M^\star) = \infty$, for any constant $\bar{B} > 0$ there exists a finite $\tilde{C}$ such that $q(\tilde{C}) > \bar{B}$, and $C(M^\star) > \tilde{C}$. In particular, choosing $\bar{B} > B(z^\star)/H^\star$ yields

$$q\big(C(M^\star)\big)\,H^\star > B(z^\star),$$

so that the objective value is strictly negative. This contradicts feasibility of $(M, z) = (\emptyset, 0)$, which yields value zero. Therefore $C(M^\star) < \infty$ at any optimum.

□

Proof of Corollary 2.1. By definition of $M_\tau^\star$, we have $d(s) \geq \tau$ for all $s \in M_\tau^\star$. Therefore

$$\int_{M^\star} d(s)\,\mu(ds) \;\geq\; \int_{M_\tau^\star} d(s)\,\mu(ds) \;\geq\; \int_{M_\tau^\star} \tau\,\mu(ds) \;=\; \tau\,\mu(M_\tau^\star).$$

Rearranging gives

$$\mu(M_\tau^\star) \;\leq\; \frac{1}{\tau}\int_{M^\star} d(s)\,\mu(ds).$$

By Proposition 2.1, $C(M^\star) = \int_{M^\star} d(s)\,\mu(ds) < \infty$, hence $\mu(M_\tau^\star) < \infty$. □

Remark A.1 (Complexity versus market count)**.** *A common way to ensure finite scope in continuum environments is to impose a uniform lower bound on market heterogeneity, $d(s) \geq \underline{d} > 0$ μ-a.e. Under such an assumption, total organizational complexity and market measure coincide up to scale:*

$$C(M) \;=\; \int_M d(s)\,\mu(ds) \;\geq\; \underline{d}\,\mu(M),$$

so that $C(M) < \infty$ if and only if $\mu(M) < \infty$. In this case, bounding organizational complexity is equivalent to bounding the number (measure) of markets served, and the firm faces a hard span-of-control constraint whereby each additional market entails a strictly positive irreducible burden.

By contrast, when heterogeneity is allowed to vary across markets and may become arbitrarily small—while remaining integrable—finite complexity does not restrict the measure of markets served. Specifically, if $d \in L^1(S,\mu)$ and $d(s) > 0$ μ-a.e., then $C(M) < \infty$ is compatible with $\mu(M) = \infty$. Nevertheless, finite complexity still imposes sharp discipline on the number of materially distinct markets. For any threshold $\tau > 0$,

$$C(M) \;\geq\; \int_{\{s\in M:\, d(s)\geq\tau\}} d(s)\,\mu(ds) \;\geq\; \tau\cdot\mu\big(\{s\in M:\; d(s)\geq\tau\}\big),$$

which implies

$$\mu\big(\{s\in M:\; d(s)\geq\tau\}\big) \;\leq\; \frac{C(M)}{\tau}.$$

Thus, while the firm may serve a large (even unbounded) collection of markets with very small heterogeneity, it can serve only a finite mass of markets whose heterogeneity exceeds any fixed threshold τ. This distinction allows the model to accommodate scalable, platform-type organizations that operate across many highly similar markets, while preserving a meaningful notion of bounded organizational scope.

A.2 Proofs for Section 2.2

In proving Proposition 2.2 we will make use of the following standard result.

Lemma A.1 (Monotonicity under parametric maximization)**.** *Let $\mathcal{X}$ be a nonempty set and let $f : \mathcal{X} \times \mathbb{R}_+ \to \mathbb{R}$. Define the value function*

$$V(\lambda) \equiv \sup_{x \in \mathcal{X}} f(x, \lambda).$$

If $f(x, \lambda)$ is weakly decreasing in λ for each fixed $x \in \mathcal{X}$, then $V(\lambda)$ is weakly decreasing in λ.

Proof. Let $0 \leq \lambda_1 < \lambda_2$. For any $x \in \mathcal{X}$, monotonicity of f in λ implies $f(x, \lambda_1) \geq f(x, \lambda_2)$. Taking the supremum over $x \in \mathcal{X}$ on both sides yields

$$\sup_{x \in \mathcal{X}} f(x, \lambda_1) \geq \sup_{x \in \mathcal{X}} f(x, \lambda_2),$$

i.e. $V(\lambda_1) \geq V(\lambda_2)$. Since $\lambda_1 < \lambda_2$ was arbitrary, V is weakly decreasing in λ. □

Proof of Proposition 2.2. We conduct the proof in several steps:

Step 1: Formal definition of the market-level value function. For any market $s \in S$, blueprint quality $z \geq 0$, and organizational complexity level $C \geq 0$, define the market-level inner problem:

$$\Pi_s(z; C) \equiv \sup_{\substack{x \geq 0,\, \ell \geq 0,\, h \geq 0 \\ x + \ell \leq \bar{X}(s)}} \{R(z_s, x) - q(C) \cdot h\}, \tag{45}$$

subject to the deployment constraint

$$z_s = \theta(A(h, \ell), d(s)) \cdot z. \tag{46}$$

Here, $R : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}$ is the market payoff function, $q : \mathbb{R}_+ \to \mathbb{R}_+$ is the unit cost of headquarters attention, $\theta : \mathbb{R}_+ \times \mathbb{R}_+ \to [0, 1]$ is the portability function, and $A : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_+$ is the adaptation production function. The blueprint creation cost $q(C)\chi(z)$ is common across markets and is handled at the outer level.

Step 2: Parametric dependence on organizational complexity. For fixed (z, s), the value function $\Pi_s(z; C)$ depends on the complexity parameter C only through the unit cost $q(C)$, which multiplies the headquarters input h in the objective. All other primitives and constraints of the inner

problem—R, θ, A, $d(s)$, and $\bar{X}(s)$—as well as the feasible set, are independent of C. More precisely, we can rewrite (45) as

$$\Pi_s(z;C) = \sup_{(x,\ell,h)\in\mathcal{F}_s} \{R(\theta(A(h,\ell),d(s))\cdot z, x) - q(C)\cdot h\},$$

where $\mathcal{F}_s \equiv \{(x,\ell,h)\in\mathbb{R}^3_+ : x+\ell \le \bar{X}(s)\}$.

Step 3: Monotonicity of organizational complexity. By Assumption 2.1, $d(s) > 0$ for μ-almost every $s \in S$. Since $\mu(\Delta) > 0$ and $\Delta \subseteq S \setminus M$ is measurable, we have

$$C(M\cup\Delta) = \int_{M\cup\Delta} d(s)\,\mu(ds) = \int_M d(s)\,\mu(ds) + \int_\Delta d(s)\,\mu(ds) = C(M) + \int_\Delta d(s)\,\mu(ds).$$

Moreover, $d(s) > 0$ for μ-a.e. $s \in \Delta$ and $\mu(\Delta) > 0$, hence

$$\int_\Delta d(s)\,\mu(ds) > 0.$$

Therefore,

$$C(M\cup\Delta) > C(M). \tag{47}$$

Step 4: Monotonicity of the unit cost function. By Assumption 1.6 we have that $q(\cdot)$ is strictly increasing. Combined with (47), we obtain

$$q(C(M\cup\Delta)) > q(C(M)). \tag{48}$$

Step 5: Application of monotonicity under parametric maximization. Fix an incumbent market $s \in M$ and blueprint quality $z \ge 0$. For any feasible allocation $(x,\ell,h)\in\mathcal{F}_s$, define

$$f_{s,z}\big((x,\ell,h),\lambda\big) \;\equiv\; R(\theta(A(h,\ell),d(s))\cdot z,\; x) - \lambda\, h, \qquad \lambda \ge 0.$$

Define the parametric value function

$$V_{s,z}(\lambda) \;\equiv\; \sup_{(x,\ell,h)\in\mathcal{F}_s} f_{s,z}\big((x,\ell,h),\lambda\big).$$

By construction, the market-level value function satisfies

$$\Pi_s(z;C) = V_{s,z}\big(q(C)\big).$$

For each fixed $(x, \ell, h) \in \mathcal{F}_s$, the mapping $\lambda \mapsto f_{s,z}((x, \ell, h), \lambda)$ is weakly decreasing (strictly decreasing if $h > 0$, and constant if $h = 0$). Therefore, by Lemma A.1, the function $\lambda \mapsto V_{s,z}(\lambda)$ is weakly decreasing, and hence $\Pi_s(z; C)$ is weakly decreasing in $q(C)$.

Step 6: Conclusion. By Assumption 1.6 and $C(M \cup \Delta) > C(M)$ by (47), we have $q(C(M \cup \Delta)) > q(C(M))$ by (48).

By Step 5, the map $\lambda \mapsto \Pi_s(z; \cdot)$ evaluated at C such that $q(C) = \lambda$ is weakly decreasing in λ. Therefore,

$$\Pi_s(z; C(M \cup \Delta)) \leq \Pi_s(z; C(M)),$$

as claimed.

Step 7: Strictness of the inequality. Let $\lambda_1 \equiv q(C(M))$ and $\lambda_2 \equiv q(C(M \cup \Delta))$. Since $q(\cdot)$ is strictly increasing and $C(M \cup \Delta) > C(M)$, we have $\lambda_2 > \lambda_1$. The inequality is strict if and only if every optimal solution to the inner problem at λ_1 uses strictly positive headquarters attention.

Sufficiency. Suppose every optimal solution at λ_1 uses strictly positive headquarters attention. Let (x^*, l^*, h^*) be such an optimal allocation at λ_1. Its value at the new higher cost λ_2 is strictly lower:

$$f_{s,z}((x^*, l^*, h^*), \lambda_2) \;<\; f_{s,z}((x^*, l^*, h^*), \lambda_1) \;=\; \Pi_s(z; C(M)).$$

Now consider any candidate allocation $(x, l, h) \in \mathcal{F}_s$ for the problem at λ_2.

1. If $h > 0$, then $f_{s,z}((x, l, h), \lambda_2) < f_{s,z}((x, l, h), \lambda_1) \leq \Pi_s(z; C(M))$.
2. If $h = 0$, then $f_{s,z}((x, l, h), \lambda_2) = f_{s,z}((x, l, h), \lambda_1)$. However, because we assumed every optimal solution at λ_1 uses $h > 0$, any allocation with $h = 0$ must be strictly suboptimal at λ_1. Therefore, $f_{s,z}((x, l, 0), \lambda_1) < \Pi_s(z; C(M))$.

In both cases, the value is strictly less than the original maximum. Since this holds for all feasible allocations, the new supremum must be strictly lower:

$$\Pi_s(z; C(M \cup \Delta)) \;<\; \Pi_s(z; C(M)).$$

Necessity. Conversely, if there exists an optimal solution $(\hat{x}, \hat{l}, \hat{h})$ at λ_1 with $\hat{h} = 0$, then

$$f_{s,z}((\hat{x}, \hat{l}, 0), \lambda_2) \;=\; f_{s,z}((\hat{x}, \hat{l}, 0), \lambda_1) \;=\; \Pi_s(z; C(M)).$$

Therefore $\Pi_s(z; C(M \cup \Delta)) \geq \Pi_s(z; C(M))$. Since Step 7 established the weak inequality $\leq$, it must be an equality, and strictness fails. □

Proof of Proposition 2.3. We provide a rigorous proof establishing that the marginal value of blueprint quality aggregates across markets and is monotone in scope.

Step 1: Formalization of the firm's value function. Fix organizational complexity $C \geq 0$ and a served set $M \subseteq S$. The firm's total value from operating markets in M, as a function of blueprint quality $z \geq 0$, is given by

$$V(z; M, C) \equiv \int_{s \in M} \Pi_s(z; C)\, \mu(ds), \tag{49}$$

where $\Pi_s(z; C)$ is the market-level value function defined in Definition 2.

Step 2: Differentiability and the envelope theorem. Under regularity conditions (continuity and differentiability of R, θ, A), the market-level value function $\Pi_s(z; C)$ is differentiable in z for μ-almost every $s \in M$.

By the envelope theorem for parametric optimization, the derivative of $\Pi_s(z; C)$ with respect to z is given by

$$\frac{\partial \Pi_s(z; C)}{\partial z} = \frac{\partial R(z_s^*, x^*)}{\partial z_s} \cdot \frac{\partial z_s^*}{\partial z}, \tag{50}$$

where $(x^*, \ell^*, h^*) \in \arg\max$ solves the inner problem at (z, C), and $z_s^* = \theta(A(h^*, \ell^*), d(s)) \cdot z$.

The envelope theorem states that we can ignore the direct effects of z on the choice variables (x, ℓ, h) when computing $\frac{\partial \Pi_s}{\partial z}$, because these variables are chosen optimally and the first-order conditions hold.

Step 3: Computing the marginal benefit at the market level. From the deployment constraint $z_s = \theta(a(s), d(s)) \cdot z$, we have

$$\frac{\partial z_s^*}{\partial z} = \theta(A(h^*, \ell^*), d(s)) \equiv \theta^*(s). \tag{51}$$

By Assumption 2.4 (blueprint effectiveness), $R(z_s, x)$ is weakly increasing in z_s, so

$$\frac{\partial R(z_s^*, x^*)}{\partial z_s} = R_1(z_s^*, x^*) \geq 0. \tag{52}$$

Combining (50), (51), and (52), we obtain

$$\frac{\partial \Pi_s(z; C)}{\partial z} = R_1(z_s^*, x^*) \cdot \theta^*(s) \geq 0. \tag{53}$$

This is the market-level marginal benefit of increasing blueprint quality.

Step 4: Aggregation across markets. The total marginal value of blueprint quality under served set

M is obtained by differentiating (49):

$$\frac{\partial V(z; M, C)}{\partial z} = \frac{\partial}{\partial z} \int_{s \in M} \Pi_s(z; C)\, \mu(ds). \tag{54}$$

We now establish conditions under which differentiation can be passed through the integral.

Lemma A.2 (Dominated convergence for derivatives)**.** *Suppose:*

(i) For each $s \in M$, $\Pi_s(z; C)$ is differentiable in z.

(ii) There exists an integrable function $g \in L^1(M, \mu)$ such that for all z in a neighborhood of the evaluation point,

$$\left| \frac{\partial \Pi_s(z; C)}{\partial z} \right| \leq g(s) \quad \text{for } \mu\text{-a.e. } s \in M.$$

Then

$$\frac{\partial}{\partial z} \int_{s \in M} \Pi_s(z; C)\, \mu(ds) = \int_{s \in M} \frac{\partial \Pi_s(z; C)}{\partial z}\, \mu(ds).$$

Proof of Lemma A.2. This is a standard application of Lebesgue's dominated convergence theorem for differentiation under the integral sign. For any $\Delta z \to 0$,

$$\frac{V(z + \Delta z; M, C) - V(z; M, C)}{\Delta z} = \int_{s \in M} \frac{\Pi_s(z + \Delta z; C) - \Pi_s(z; C)}{\Delta z}\, \mu(ds).$$

By the mean value theorem, for each s,

$$\frac{\Pi_s(z + \Delta z; C) - \Pi_s(z; C)}{\Delta z} = \frac{\partial \Pi_s(\xi_s; C)}{\partial z}$$

for some ξ_s between z and $z + \Delta z$. By condition (ii), this is bounded by $g(s) \in L^1(M, \mu)$. By dominated convergence, taking $\Delta z \to 0$ yields

$$\lim_{\Delta z \to 0} \frac{V(z + \Delta z; M, C) - V(z; M, C)}{\Delta z} = \int_{s \in M} \lim_{\Delta z \to 0} \frac{\Pi_s(z + \Delta z; C) - \Pi_s(z; C)}{\Delta z}\, \mu(ds) = \int_{s \in M} \frac{\partial \Pi_s(z; C)}{\partial z}\, \mu(ds)$$

□

Step 5: Verifying the domination condition. We now show that condition (ii) of Lemma A.2 holds under our assumptions. From (53), the marginal benefit at market s is

$$\frac{\partial \Pi_s(z; C)}{\partial z} = R_1(z_s^*, x^*) \cdot \theta(A(h^*, \ell^*), d(s)).$$

By assumption, $\theta(a, d) \in (0, 1]$ for all (a, d), so $\theta(A(h^*, \ell^*), d(s)) \leq 1$.

We need to bound $R_1(z_s^*, x^*)$. Since $z_s^* = \theta^*(s) \cdot z \leq z$ and $x^* \leq \bar{X}(s)$, if R has bounded derivatives on compact sets (which holds under regularity), there exists a function $\bar{g}(s)$ depending only on $\bar{X}(s)$ and z such that

$$|R_1(z_s^*, x^*)| \leq \bar{g}(s).$$

If $\bar{g}$ is integrable (which follows from Assumption 1.7: the integrability envelope $\bar{R} \in L^1(S, \mu)$ implies regularity of R sufficient for this), then condition (ii) holds with $g(s) = \bar{g}(s)$.

Step 6: Establishing monotonicity in scope. By Lemma A.2 and (53), we have

$$\frac{\partial V(z; M, C)}{\partial z} = \int_{s \in M} \frac{\partial \Pi_s(z; C)}{\partial z} \, \mu(ds) = \int_{s \in M} R_1(z_s^*, x^*) \cdot \theta^*(s) \, \mu(ds). \tag{55}$$

Now consider two served sets $M_1 \subseteq M_2 \subseteq S$. We have

$$\begin{aligned} \frac{\partial V(z; M_2, C)}{\partial z} &= \int_{s \in M_2} R_1(z_s^*, x^*) \cdot \theta^*(s) \, \mu(ds) \\ &= \int_{s \in M_1} R_1(z_s^*, x^*) \cdot \theta^*(s) \, \mu(ds) + \int_{s \in M_2 \setminus M_1} R_1(z_s^*, x^*) \cdot \theta^*(s) \, \mu(ds) \end{aligned} \tag{56}$$

By Assumption 2.4 (blueprint effectiveness), $R_1(z_s^*, x^*) \geq 0$ for all s, and $\theta^*(s) \geq 0$ by definition. Therefore,

$$\int_{s \in M_2 \setminus M_1} R_1(z_s^*, x^*) \cdot \theta^*(s) \, \mu(ds) \geq 0.$$

From (56), this implies

$$\frac{\partial V(z; M_2, C)}{\partial z} \geq \frac{\partial V(z; M_1, C)}{\partial z}. \tag{57}$$

Step 7: Economic interpretation and final statement. Equation (57) establishes the main result: the marginal return to blueprint quality is weakly increasing in the scope of operations. Because the blueprint is non-rival, an increase in z raises productivity in every market where it is deployed. Serving a larger set of markets $M_2 \supseteq M_1$ increases the number of markets that benefit from the marginal improvement, and since each market's marginal benefit is nonnegative (by Assumption 2.4 (blueprint effectiveness)), the aggregate marginal return increases.

Formally, the marginal return to blueprint quality is proportional to

$$\frac{\partial V(z; M, C)}{\partial z} = \int_{s \in M} R_1(z_s^*(s), x^*(s)) \cdot \theta(A(h^*(s), \ell^*(s)), d(s)) \, \mu(ds),$$

which is a sum (integral) of market-level marginal benefits, each weighted by the portability factor $\theta^*(s) \in (0, 1]$. This completes the proof. □

Remark A.2 (Integrability and regularity)**.** *The proof requires technical conditions ensuring that differentiation can be passed through the integral. These conditions are satisfied under standard regularity assumptions:*

- *$R(z_s, x)$ is continuously differentiable with $R_1(z_s, x) \geq 0$ (Assumption 2.4),*
- *The derivative R_1 is bounded on compact sets,*
- *The integrability envelope $\bar{R} \in L^1(S, \mu)$ (Assumption 1.7) ensures sufficient regularity.*

These assumptions are economically mild and standard in models with continua of markets.

Remark A.3 (Strict inequality)**.** *The inequality in (57) is strict if and only if $M_2 \setminus M_1$ has positive measure and*

$$R_1(z_s^*, x^*) \cdot \theta^*(s) > 0$$

for a positive-measure subset of $M_2 \setminus M_1$. This holds generically when the additional markets contribute positive value and the blueprint has positive marginal product.

Proof of Proposition 2.4. Fix blueprint quality $z \geq 0$. With z fixed, the firm's choice of footprint $M \subseteq S$ solves

$$W(M) \equiv \int_{s \in M} \Pi_s\big(z; C(M)\big)\, \mu(ds) - q\big(C(M)\big)\, \chi(z), \tag{58}$$

where $C(M) \equiv \int_{s \in M} d(s)\, \mu(ds)$. By Proposition 2.2 (Organizational congestion), for each market s the function $\Pi_s(z; C)$ is weakly decreasing in C, and by Assumption 2.6 it is weakly decreasing in $d(s)$ for any fixed C.

Step 1 (No–crossing). Fix any $t \geq 0$ and define

$$A_t \equiv \{s \in M^\star : d(s) > t\}, \qquad B_t \equiv \{s \notin M^\star : d(s) < t\}.$$

We claim that it is impossible that both $\mu(A_t) > 0$ and $\mu(B_t) > 0$.

Suppose instead that $\mu(A_t) > 0$ and $\mu(B_t) > 0$. Let $\varepsilon \equiv \min\{\mu(A_t), \mu(B_t)\} > 0$. By Assumption 1.1, the measure space (S, Σ, μ) is non-atomic, so there exist measurable subsets $A \subseteq A_t$ and $B \subseteq B_t$ with $\mu(A) = \mu(B) = \varepsilon$. Define the swapped footprint

$$M' \equiv (M^\star \setminus A) \cup B.$$

Let $R \equiv M^\star \setminus A$ denote the set of retained markets.

Complexity comparison. Since $d(s) > t$ on A and $d(s) < t$ on B, we have

$$C(M') = C(M^\star) - \int_A d(s)\,\mu(ds) + \int_B d(s)\,\mu(ds).$$

Because $d(s) > t$ for all $s \in A$ and $\mu(A) = \varepsilon > 0$,

$$\int_A d(s)\,\mu(ds) > t\,\mu(A) = t\varepsilon.$$

Similarly, $d(s) < t$ for all $s \in B$ and $\mu(B) = \varepsilon$ imply

$$\int_B d(s)\,\mu(ds) < t\,\mu(B) = t\varepsilon.$$

Therefore $C(M') < C(M^\star)$, and since $q(\cdot)$ is strictly increasing (Assumption 1.6),

$$q(C(M')) < q(C(M^\star)). \tag{59}$$

Value comparison. Recall the firm's objective from (9):

$$W(M) = \int_{s\in M} \Pi_s\big(z; C(M)\big)\,\mu(ds) - q\big(C(M)\big)\chi(z).$$

We compare $W(M')$ and $W(M^\star)$ by decomposing each into contributions from retained markets, swapped markets, and the blueprint cost.

Write

$$W(M') = \int_R \Pi_s\big(z; C(M')\big)\,\mu(ds) + \int_B \Pi_s\big(z; C(M')\big)\,\mu(ds) - q(C(M'))\chi(z), \tag{60}$$

$$W(M^\star) = \int_R \Pi_s\big(z; C(M^\star)\big)\,\mu(ds) + \int_A \Pi_s\big(z; C(M^\star)\big)\,\mu(ds) - q(C(M^\star))\chi(z). \tag{61}$$

(i) Retained markets. By Proposition 2.2, $\Pi_s(z;\cdot)$ is weakly decreasing in organizational complexity. Since $C(M') < C(M^\star)$, for every $s \in R$,

$$\Pi_s\big(z; C(M')\big) \geq \Pi_s\big(z; C(M^\star)\big),$$

and hence

$$\int_R \Pi_s\big(z; C(M')\big)\,\mu(ds) \geq \int_R \Pi_s\big(z; C(M^\star)\big)\,\mu(ds). \tag{62}$$

(ii) Swapped markets. Define

$$\underline{\pi}_B \equiv \operatorname*{ess\,inf}_{s\in B} \Pi_s\big(z; C(M^\star)\big), \qquad \overline{\pi}_A \equiv \operatorname*{ess\,sup}_{s\in A} \Pi_s\big(z; C(M^\star)\big).$$

By construction, $d(s) < t$ for all $s \in B$ and $d(s) > t$ for all $s \in A$. Assumption 2.6 (heterogeneity ordering) implies that $\Pi_s(z; C)$ is weakly decreasing in $d(s)$ for any fixed C. Therefore $\underline{\pi}_B \geq \overline{\pi}_A$, and

$$\int_B \Pi_s\big(z; C(M^\star)\big)\,\mu(ds) \geq \varepsilon\,\underline{\pi}_B \geq \varepsilon\,\overline{\pi}_A \geq \int_A \Pi_s\big(z; C(M^\star)\big)\,\mu(ds). \tag{63}$$

Since $C(M') < C(M^\star)$, Proposition 2.2 gives $\Pi_s(z; C(M')) \geq \Pi_s(z; C(M^\star))$ for all $s \in B$. Thus

$$\int_B \Pi_s\big(z; C(M')\big)\,\mu(ds) \geq \int_B \Pi_s\big(z; C(M^\star)\big)\,\mu(ds). \tag{64}$$

(iii) Blueprint cost. From (59),

$$-q(C(M'))\chi(z) \geq -q(C(M^\star))\chi(z), \tag{65}$$

with strict inequality whenever $\chi(z) > 0$.

Conclusion of Step 1. Subtracting (61) from (60) and applying (62), (63), (64), and (65):

$$\begin{aligned}
W(M') - W(M^\star) &= \left[\int_R \Pi_s(z; C(M'))\,\mu(ds) - \int_R \Pi_s(z; C(M^\star))\,\mu(ds)\right] \\
&\quad + \left[\int_B \Pi_s(z; C(M'))\,\mu(ds) - \int_A \Pi_s(z; C(M^\star))\,\mu(ds)\right] \\
&\quad + \left[q(C(M^\star)) - q(C(M'))\right]\chi(z) \\
&\geq 0 + \left[\int_B \Pi_s(z; C(M^\star))\,\mu(ds) - \int_A \Pi_s(z; C(M^\star))\,\mu(ds)\right] + \left[q(C(M^\star)) - q(C(M'))\right]\chi(z) \\
&\geq 0 + 0 + 0 = 0.
\end{aligned}$$

Moreover, the final term is strictly positive whenever $\chi(z) > 0$ (i.e., whenever $z > 0$), contradicting the optimality of $M^\star$. Therefore, it is impossible that both $\mu(A_t) > 0$ and $\mu(B_t) > 0$.

Step 2 (Cutoff characterization). Define

$$d^\star \equiv \operatorname*{ess\,sup}_{s\in M^\star} d(s) \in [0,\infty].$$

We show that $M^\star = \{s : d(s) \le d^\star\}$ up to μ-null sets.

If $d(s) < d^\star$, choose t such that $d(s) < t < d^\star$. By definition of essential supremum, $\mu(\{u \in M^\star : d(u) > t\}) > 0$. Step 1 then implies $\mu(\{u \notin M^\star : d(u) < t\}) = 0$, and since $d(s) < t$, it follows that $s \in M^\star$ μ-almost everywhere.

If $d(s) > d^\star$, choose t with $d^\star < t < d(s)$. Then $\mu(\{u \in M^\star : d(u) > t\}) = 0$ by definition of $d^\star$, which implies $M^\star \subseteq \{d \le t\}$ μ-almost everywhere. Hence $s \notin M^\star$ μ-almost everywhere.

Combining these two cases yields

$$s \in M^\star \quad \Longleftrightarrow \quad d(s) \le d^\star \qquad \mu\text{-a.e.},$$

establishing the cutoff structure of the optimal footprint. □

A.3 Proofs for Section 2.3

Proof of Proposition 2.5. Fix blueprint quality $z \ge 0$, organizational complexity C, and a served set M. For each $s \in M$, the market-level inner problem (Definition 2) is

$$\Pi_s(z;C) \;\equiv\; \sup_{x\ge 0,\ \ell\ge 0,\ h\ge 0} \Big\{ R\big(z_s, x\big) - q(C)\, h \Big\} \quad \text{s.t.} \quad x+\ell \le \overline{X}(s), \;\; z_s = \theta\big(A(h,\ell), d(s)\big)\, z. \tag{66}$$

Fix two markets $s_1, s_2 \in M$ and write $d_i \equiv d(s_i)$ with $d_1 \le d_2$. Let $h^\star(s)$ denote an optimal (possibly set-valued) choice of headquarters attention in market s.

Step 1 (Reduce to an outer problem in h). For each market s and each $h \ge 0$, define the conditional value (optimizing out (x,ℓ) given h):

$$V(h;s) \;\equiv\; \sup_{x\ge 0,\ \ell\ge 0} \Big\{ R\big(\theta(A(h,\ell), d(s))\, z,\; x\big) \Big\} \quad \text{s.t.} \quad x+\ell \le \overline{X}(s). \tag{67}$$

Then the inner problem (66) is equivalent to

$$\Pi_s(z;C) \;=\; \sup_{h\ge 0} \Big\{ V(h;s) - q(C)\, h \Big\}. \tag{68}$$

Define the outer objective

$$\Psi(h; s) \equiv V(h; s) - q(C)\, h.$$

An optimizer $h^\star(s)$ is any element of $\arg\max_{h \geq 0} \Psi(h; s)$.

Step 2 (Decreasing differences in (h, d)). We show that $\Psi(h; s)$ has decreasing differences in $(h, d(s))$, equivalently increasing differences in $(h, -d(s))$.

By Assumption 2.7 (adaptation regularity), the functions R, θ, and A are continuously differentiable with $R_1 \geq 0$, $\partial_a \theta \geq 0$, and $\partial_h A > 0$. By the envelope theorem applied to (67), at any h where a conditional optimizer $(x^\star(s; h), \ell^\star(s; h))$ is well-defined we have

$$\frac{\partial}{\partial h} V(h; s) = R_1\big(z_s(h), x^\star(s; h)\big)\, \partial_a \theta\big(A(h, \ell^\star(s; h)), d(s)\big)\, \partial_h A\big(h, \ell^\star(s; h)\big)\, z, \tag{69}$$

where $z_s(h) = \theta(A(h, \ell^\star(s; h)), d(s))\, z$.

By Assumption 2.7, $R_1 \geq 0$ and $\partial_h A > 0$. Assumption 2.8 (single-crossing in heterogeneity) states that for any fixed $a \geq 0$, $\partial_a \theta(a, d)$ is weakly decreasing in d. This is the key channel generating the comparative static: holding the adaptation level a fixed, the marginal effectiveness of adaptation is lower in more heterogeneous markets.

Since $\bar{X}(s)$ is constant across served markets by assumption, the conditional optimizer $(x^\star(s; h), \ell^\star(s; h))$ depends on $d(s)$ only through the portability function θ. It follows that for fixed h, the composite $R_1(z_s(h), x^\star(s; h)) \cdot \partial_a \theta(A(h, \ell^\star(s; h)), d(s)) \cdot \partial_h A(h, \ell^\star(s; h))$ is weakly decreasing in $d(s)$, because the single-crossing property of $\partial_a \theta$ (Assumption 2.8) is the only channel through which $d(s)$ affects the marginal benefit of h.

Therefore, holding h fixed, the marginal benefit of h is weakly lower in more heterogeneous markets:

$$d_1 \leq d_2 \quad \Longrightarrow \quad \frac{\partial}{\partial h} V(h; s_1) \;\geq\; \frac{\partial}{\partial h} V(h; s_2) \qquad \text{for all } h \geq 0. \tag{70}$$

Since $\Psi(h; s) = V(h; s) - q(C)h$ and $q(C)h$ does not depend on $d(s)$, the same property holds for Ψ:

$$d_1 \leq d_2 \quad \Longrightarrow \quad \frac{\partial}{\partial h} \Psi(h; s_1) \;\geq\; \frac{\partial}{\partial h} \Psi(h; s_2). \tag{71}$$

Condition (71) is exactly decreasing differences in (h, d).

Step 3 (Monotone comparative statics). By Step 2, $\Psi(h; s)$ has increasing differences in $(h, -d(s))$. Under standard regularity conditions ensuring existence of a maximizer of (68) (e.g. compactness

of the feasible set for h or coercivity of Ψ), the monotone comparative statics theorem (Milgrom–Shannon or Topkis) implies that the set of maximizers $\arg\max_{h\geq 0}\Psi(h;s)$ is nonincreasing in $d(s)$ with respect to the strong set order.

In particular, for any selections $h^\star(s_i)\in\arg\max_{h\geq 0}\Psi(h;s_i)$,

$$d_1 \leq d_2 \quad\Longrightarrow\quad h^\star(s_1)\ \geq\ h^\star(s_2).$$

This establishes the core–periphery allocation: headquarters attention is weakly decreasing in heterogeneity. □

A.4 Proofs for Section 4

Proof of Proposition 4.2. By Proposition 4.1, $\Phi(z)=H^{-1}(z)$ for every interior z. Since H is C^1 and strictly increasing with $H'(d)>0$, the inverse function theorem gives

$$\Phi'(z)=\frac{1}{H'(\Phi(z))}.$$

Using the identity $z=H(\Phi(z))$, the scope elasticity is

$$\varepsilon_\Phi(z)=\frac{z\,\Phi'(z)}{\Phi(z)}=\frac{H(\Phi(z))}{\Phi(z)\,H'(\Phi(z))}=\frac{1}{\mathcal{E}_H(\Phi(z))}.$$

This proves (20). Since Φ is strictly increasing in z (Proposition 4.1) and $\mathcal{E}_H$ is weakly increasing in d (Assumption 4.2), the composition $z\mapsto\mathcal{E}_H(\Phi(z))$ is weakly increasing. Taking reciprocals, $\varepsilon_\Phi(z)$ is weakly decreasing. If $\mathcal{E}_H$ is strictly increasing, then ε_Φ is strictly decreasing. □

Proof of Proposition 4.3. See Appendix A.6. □

A.5 Proofs for Section 3

Proof of Lemma 3.1. Fix $\bar{C}<\infty$ and define

$$\bar{q}\ \equiv\ q(\bar{C})\ <\ \infty.$$

Let $(\bar{x},\bar{h},\bar{\ell})$ be given by Assumption 3.3(iii), and define

$$\bar{a}\ \equiv\ A(\bar{h},\bar{\ell})\ >\ 0.$$

Choose any $\alpha \in (0,1)$, for example $\alpha = \frac{1}{2}$. By Assumption 3.3(i), $\theta(\bar{a}, 0) = 1$. By Assumption 3.3(ii), the map $d \mapsto \theta(\bar{a}, d)$ is continuous. Hence there exists $\varepsilon(\bar{C}) > 0$ such that

$$\theta(\bar{a}, d) \;\geq\; \alpha \qquad \text{for all } d \in [0, \varepsilon(\bar{C})].$$

Now fix any market s with $d_0(s) \leq \varepsilon(\bar{C})$, any complexity level $C \in [0, \bar{C}]$, and any blueprint quality $z \geq 0$. Evaluate the extended market-level problem (15) at the feasible policy

$$x = \bar{x}, \qquad h = \bar{h}, \qquad \ell = \bar{\ell}, \qquad \sigma = 0.$$

This policy is feasible because $\bar{x} + \bar{\ell} \leq \inf_{u \in S} \overline{X}(u) \leq \overline{X}(s)$ by Assumption 3.3(iii), and $\sigma = 0$ is admissible. Since $g(0) = 1$, the induced effective deployment is

$$z_s = \theta\big(A(\bar{h}, \bar{\ell}), d_0(s) g(0)\big) z = \theta(\bar{a}, d_0(s)) z \;\geq\; \alpha z.$$

Therefore

$$\begin{aligned} \hat{\Pi}_s(z; C) &\geq R(\theta(\bar{a}, d_0(s)) z, \bar{x}) - q(C)\bar{h} - \psi(z) c(0) \\ &\geq R(\alpha z, \bar{x}) - \bar{q}\,\bar{h}, \end{aligned}$$

where we used $q(C) \leq q(\bar{C}) = \bar{q}$ and $c(0) = 0$.

Define

$$\underline{\pi}(z; \bar{C}) \;\equiv\; R(\alpha z, \bar{x}) - \bar{q}\,\bar{h}.$$

Then for all $z \geq 0$, all $C \in [0, \bar{C}]$, and all s with $d_0(s) \leq \varepsilon(\bar{C})$,

$$\hat{\Pi}_s(z; C) \;\geq\; \underline{\pi}(z; \bar{C}).$$

By Assumption 3.4(i),

$$\underline{\pi}(z; \bar{C}) \xrightarrow[z \to \infty]{} \infty.$$

By Assumption 3.4(ii),

$$\frac{R(\alpha z, \bar{x})}{\chi(z)} \xrightarrow[z \to \infty]{} \infty,$$

hence also

$$\frac{\underline{\pi}(z; \bar{C})}{\chi(z)} = \frac{R(\alpha z, \bar{x})}{\chi(z)} - \frac{\bar{q}\,\bar{h}}{\chi(z)} \xrightarrow[z \to \infty]{} \infty.$$

Since $\underline{\pi}(z;\bar{C}) \to \infty$, there exists $z_1(\bar{C}) < \infty$ such that $\underline{\pi}(z;\bar{C}) > 0$ for all $z \geq z_1(\bar{C})$. This proves both parts of the lemma. □

Proof of Proposition 3.1. Fix $\rho \geq 1$, $z \geq 0$, and $\bar{C} < \infty$. By definition, every candidate footprint entering $\mathcal{K}^\rho(z;\bar{C})$ satisfies $M \subseteq S^\rho$. Therefore

$$\mu(M) \leq \mu(S^\rho).$$

Taking the supremum over all such M yields

$$\mathcal{K}^\rho(z;\bar{C}) \leq \mu(S^\rho).$$

By Assumption 3.2(ii), $\mu(S^\rho) < \infty$. Hence

$$\mathcal{K}^\rho(z;\bar{C}) \leq \mu(S^\rho) < \infty.$$

The final statement follows immediately. □

Proof of Proposition 3.2. Fix $\bar{C} < \infty$. Let $\varepsilon(\bar{C})$, $z_1(\bar{C})$, and $\underline{\pi}(\cdot;\bar{C})$ be given by Lemma 3.1.

Now fix any target scale $K < \infty$ and define

$$\delta_K \equiv \min\left\{\varepsilon(\bar{C}), \frac{\bar{C}}{K}\right\} > 0.$$

By Assumption 3.2(iv),

$$\mu\big(S^\rho_{\delta_K}\big) = \mu(S^\rho \cap \{s \in S : d_0(s) \leq \delta_K\}) \xrightarrow[\rho\to\infty]{} \infty.$$

Hence there exists $\rho_K < \infty$ such that for all $\rho \geq \rho_K$,

$$\mu\big(S^\rho_{\delta_K}\big) \geq K.$$

Fix any $\rho \geq \rho_K$. Because the market space is non-atomic (Assumption 1.1), there exists a measurable subset

$$M^\rho_K \subseteq S^\rho_{\delta_K}$$

such that

$$\mu(M_K^\rho) = K.$$

Consider the feasible common recipe

$$x(s) = \bar{x}, \qquad h(s) = \bar{h}, \qquad \ell(s) = \bar{\ell}, \qquad \sigma(s) = 0 \qquad \text{for all } s \in M_K^\rho,$$

where $(\bar{x}, \bar{h}, \bar{\ell})$ is the recipe from Assumption 3.3(iii). Because $d_0(s) \le \delta_K$ on M_K^ρ and $g(0) = 1$,

$$C(M_K^\rho, \sigma) = \int_{M_K^\rho} d_0(s)\, \mu(ds) \le \delta_K\, \mu(M_K^\rho) = \delta_K K \le \bar{C}.$$

Hence the complexity cap is satisfied.

Moreover, since $\delta_K \le \varepsilon(\bar{C})$, every market in M_K^ρ lies in the near-core region covered by Lemma 3.1. Therefore, for every $z \ge z_1(\bar{C})$ and every $s \in M_K^\rho$,

$$\hat{\Pi}_s(z; C(M_K^\rho, \sigma)) \;\ge\; \underline{\pi}(z; \bar{C}).$$

Integrating over M_K^ρ and subtracting the common blueprint creation cost gives

$$\begin{aligned}\hat{V}^\rho(M_K^\rho, z, x, \ell, h, \sigma) &= \int_{M_K^\rho} \hat{\Pi}_s(z; C(M_K^\rho, \sigma))\ \mu(ds) - q(C(M_K^\rho, \sigma))\, \chi(z) \\ &\ge K\, \underline{\pi}(z; \bar{C}) - q(\bar{C})\chi(z).\end{aligned}$$

By Lemma 3.1,

$$\frac{\underline{\pi}(z; \bar{C})}{\chi(z)} \xrightarrow[z\to\infty]{} \infty.$$

Hence there exists $z_K \ge z_1(\bar{C})$ such that for all $z \ge z_K$,

$$K\, \underline{\pi}(z; \bar{C}) > q(\bar{C})\chi(z).$$

For all such z,

$$\hat{V}^\rho(M_K^\rho, z, x, \ell, h, \sigma) > 0.$$

Thus M_K^ρ is admissible in the definition of $\mathcal{K}^\rho(z; \bar{C})$, so

$$\mathcal{K}^\rho(z; \bar{C}) \ge \mu(M_K^\rho) = K \qquad \text{for all } \rho \ge \rho_K,\ z \ge z_K.$$

This proves the first claim.

The equivalent cofinal statement is merely a restatement of the same quantifiers: for every target $K < \infty$, one can find thresholds (ρ_K, z_K) beyond which $\mathcal{K}^\rho(z; \bar{C}) \geq K$ uniformly. Hence the bounded-complexity scale frontier diverges in the joint limit. □

A.6 Proofs for Section 4

Throughout this appendix, define

$$H(d) \equiv \frac{K(d)}{B(d)}, \qquad d \in [0, \bar{d}).$$

Proof of Proposition 4.1. By Assumption 4.1, the functions B and K are positive and C^1 on $[0, \bar{d})$, with $B'(d) < 0$ and $K'(d) > 0$ for all interior d. Hence H is C^1 and

$$H'(d) = \frac{K'(d)B(d) - K(d)B'(d)}{B(d)^2}. \tag{72}$$

Because $B(d) > 0$, $K(d) > 0$, $K'(d) > 0$, and $-B'(d) > 0$, the numerator in (72) is strictly positive. Therefore

$$H'(d) > 0 \qquad \text{for all } d \in (0, \bar{d}),$$

so H is strictly increasing on $[0, \bar{d})$.

Now fix any blueprint quality z in the range of H. By the maintained characterization in Assumption 4.1, an interior optimal cutoff $d^*(z)$ must satisfy

$$z\, B(d^*(z)) = K(d^*(z)),$$

which is equivalent to

$$H(d^*(z)) = z.$$

Since H is strictly increasing, this equation has at most one solution. Because z lies in the range of H, it has exactly one solution, namely

$$d^*(z) = H^{-1}(z).$$

This proves part (i).

Part (ii) follows immediately, because in the benchmark ordered environment scope is the cut-off:

$$\Phi(z) = \mu(M^*(z)) = d^*(z) = H^{-1}(z).$$

Since H^{-1} is strictly increasing, Φ is strictly increasing in z.

For part (iii), suppose $\bar{d} < \infty$ and

$$\lim_{d \uparrow \bar{d}} H(d) = \infty.$$

Because H is strictly increasing and unbounded as $d \uparrow \bar{d}$, its inverse is well defined on $[H(0), \infty)$ and satisfies

$$\lim_{z \to \infty} H^{-1}(z) = \bar{d}.$$

Hence

$$\lim_{z \to \infty} \Phi(z) = \lim_{z \to \infty} H^{-1}(z) = \bar{d}.$$

This proves the proposition. □

Proof of Proposition 4.2. By Proposition 4.1,

$$\Phi(z) = H^{-1}(z)$$

for every interior z. Since H is C^1 and strictly increasing with $H'(d) > 0$, the inverse function theorem implies that Φ is differentiable and

$$\Phi'(z) = \frac{1}{H'(\Phi(z))}.$$

Using also the identity

$$z = H(\Phi(z)),$$

we obtain

$$\begin{aligned}\varepsilon_{\Phi}(z) &= \frac{z\,\Phi'(z)}{\Phi(z)} \\ &= \frac{H(\Phi(z))}{\Phi(z)\,H'(\Phi(z))} \\ &= \frac{1}{\dfrac{\Phi(z)\,H'(\Phi(z))}{H(\Phi(z))}} \\ &= \frac{1}{\mathcal{E}_H(\Phi(z))}.\end{aligned}$$

This proves (22).

Now, by Proposition 4.1, Φ is increasing in z. By Assumption 4.2, $\mathcal{E}_H(d)$ is weakly increasing in d. Therefore the composition

$$z \mapsto \mathcal{E}_H(\Phi(z))$$

is weakly increasing in z. Taking reciprocals preserves weak monotonicity in the opposite direction because the composition is strictly positive. Hence

$$\varepsilon_{\Phi}(z) = \frac{1}{\mathcal{E}_H(\Phi(z))}$$

is weakly decreasing in z.

If $\mathcal{E}_H$ is strictly increasing, then

$$z_1 < z_2 \quad \Rightarrow \quad \Phi(z_1) < \Phi(z_2) \quad \Rightarrow \quad \mathcal{E}_H(\Phi(z_1)) < \mathcal{E}_H(\Phi(z_2)),$$

so $\varepsilon_{\Phi}(z)$ is strictly decreasing. This proves the proposition. □

Proof of Corollary 4.1. Under the assumptions of Proposition 4.1,

$$\Phi(z) = H^{-1}(z)$$

and, by part (iii) of that proposition,

$$\lim_{z\to\infty} \Phi(z) = \bar{d} < \infty.$$

Therefore

$$0 \leq \Phi(z) < \bar{d} \qquad \text{for every finite } z,$$

and the support of the induced scope distribution is contained in $[0, \bar{d}]$ (or in $[0, \bar{d})$ if one keeps the endpoint excluded). Hence the scope distribution has bounded support. □

Proof of Proposition 4.3. Let $d^* \equiv \Phi(z)$ denote the random variable describing firm scope. Because H is strictly increasing and $\Phi = H^{-1}$, for every sufficiently large m,

$$\{d^* > m\} = \{\Phi(z) > m\} = \{z > H(m)\}.$$

Therefore

$$\Pr(d^* > m) = \Pr(z > H(m)). \tag{73}$$

By assumption,

$$\Pr(z > x) = \left(\frac{\underline{z}}{x}\right)^{\alpha_z} \qquad \text{for all } x \geq \underline{z}.$$

Since $H(m) \to \infty$ as $m \to \infty$, for all sufficiently large m we may substitute $x = H(m)$ into the Pareto tail formula to obtain

$$\begin{aligned}\Pr(d^* > m) &= \left(\frac{\underline{z}}{H(m)}\right)^{\alpha_z} \\ &= \underline{z}^{\alpha_z}\, H(m)^{-\alpha_z}.\end{aligned}$$

Now use the regular variation assumption

$$H(m) = m^{\nu} L_H(m), \qquad \nu > 1,$$

where L_H is slowly varying. Then

$$\begin{aligned}\Pr(d^* > m) &= \underline{z}^{\alpha_z}\, (m^{\nu} L_H(m))^{-\alpha_z} \\ &= \underline{z}^{\alpha_z}\, m^{-\nu\alpha_z}\, L_H(m)^{-\alpha_z}.\end{aligned}$$

Define

$$\widetilde{L}(m) \equiv \underline{z}^{\alpha_z} L_H(m)^{-\alpha_z}.$$

Because positive powers and reciprocals of slowly varying functions are slowly varying, $\widetilde{L}$ is slowly

varying. Hence

$$\Pr(d^* > m) = \widetilde{L}(m)\, m^{-\nu\alpha_z}.$$

Therefore the scope distribution is regularly varying with tail index

$$\alpha_{d^*} = \nu\alpha_z.$$

Since $\nu > 1$, we have

$$\alpha_{d^*} = \nu\alpha_z > \alpha_z.$$

Thus the right tail of scope is strictly thinner than the right tail of blueprint quality. This proves the proposition. □

Proof of Corollary 4.2. Fix any $\beta > 0$. Choose $p > \beta/\alpha_z$. By hypothesis,

$$\frac{H(m)}{m^p} \xrightarrow[m\to\infty]{} \infty.$$

Equivalently, there exists a positive function $a(m) \to \infty$ such that

$$H(m) = m^p a(m)$$

for all sufficiently large m.

Using again the identity

$$\Pr(\Phi(z) > m) = \Pr(z > H(m)),$$

and the Pareto tail of z, we obtain for all sufficiently large m,

$$\begin{aligned}\Pr(\Phi(z) > m) &= \left(\frac{\underline{z}}{H(m)}\right)^{\alpha_z} \\ &= \underline{z}^{\alpha_z}\, m^{-p\alpha_z}\, a(m)^{-\alpha_z}.\end{aligned}$$

Because $a(m) \to \infty$, we have $a(m)^{-\alpha_z} \to 0$. Moreover, since $p\alpha_z > \beta$,

$$m^{-p\alpha_z} = o(m^{-\beta}).$$

Therefore

$$\Pr(\Phi(z) > m) = o(m^{-\beta}) \qquad \text{as } m \to \infty.$$

Since $\beta > 0$ was arbitrary, the scope distribution is thinner than any Pareto tail. □

A.7 Proofs for Section 5

Proof of Proposition 5.1. We establish each part in turn.

Part (i): Local concentration. Evaluating the overlap kernel at $x = 0$:

$$\mathcal{O}(0) \;=\; \frac{1}{L}\int_{\mathbb{T}_L} \sum_{i=1}^{N} \omega_i(s)\,\omega_i(s+0)\,ds \;=\; \frac{1}{L}\int_{\mathbb{T}_L} \sum_{i=1}^{N} \omega_i(s)^2\,ds \;=\; \overline{\mathrm{HHI}}_{\mathrm{loc}},$$

where the last equality is the definition of average local concentration (33).

Part (ii): Aggregate concentration. Integrate $\mathcal{O}(x)$ over $x \in \mathbb{T}_L$ and exchange the order of integration:

$$\begin{aligned}\frac{1}{L}\int_{\mathbb{T}_L} \mathcal{O}(x)\,dx &= \frac{1}{L}\int_{\mathbb{T}_L} \left[\frac{1}{L}\int_{\mathbb{T}_L} \sum_{i=1}^{N} \omega_i(s)\,\omega_i(s+x)\,ds\right] dx \\ &= \frac{1}{L^2}\sum_{i=1}^{N} \int_{\mathbb{T}_L}\int_{\mathbb{T}_L} \omega_i(s)\,\omega_i(s+x)\,ds\,dx. \end{aligned} \tag{74}$$

The exchange is justified by Fubini's theorem: for each i, the integrand $\omega_i(s)\omega_i(s+x)$ is nonnegative and jointly measurable, and the double integral is bounded above by L^2 since $\omega_i \le 1$.

For each i, apply the substitution $s' = s + x$ (with $ds' = dx$ on the torus, since Lebesgue measure is translation-invariant on $\mathbb{T}_L$):

$$\int_{\mathbb{T}_L}\int_{\mathbb{T}_L} \omega_i(s)\,\omega_i(s+x)\,ds\,dx = \int_{\mathbb{T}_L} \omega_i(s)\,ds \cdot \int_{\mathbb{T}_L} \omega_i(s')\,ds'. \tag{75}$$

To verify (75), note that for fixed s, the inner integral over x is

$$\int_{\mathbb{T}_L} \omega_i(s+x)\,dx \;=\; \int_{\mathbb{T}_L} \omega_i(s')\,ds',$$

which is independent of s. Therefore

$$\int_{\mathbb{T}_L}\int_{\mathbb{T}_L} \omega_i(s)\,\omega_i(s+x)\,ds\,dx \;=\; \int_{\mathbb{T}_L} \omega_i(s)\,ds \cdot \int_{\mathbb{T}_L} \omega_i(s')\,ds' \;=\; \big(L\bar{\omega}_i\big)^2,$$

where $\bar{\omega}_i \equiv \frac{1}{L}\int_{\mathbb{T}_L} \omega_i(s)\,ds$ is firm i's aggregate share.

Substituting back into (74):

$$\frac{1}{L}\int_{\mathbb{T}_L} \mathcal{O}(x)\,dx \;=\; \frac{1}{L^2}\sum_{i=1}^{N}(L\bar{\omega}_i)^2 \;=\; \sum_{i=1}^{N}\bar{\omega}_i^2 \;=\; \mathrm{HHI}_{\mathrm{agg}}. \qquad \square$$

Proof of Proposition 5.2. We prove the closed-form expression (40), then derive the expectation (41).

Step 1: Single-firm overlap measure. Fix firm i with core $c_i \in \mathbb{T}_L$ and scope radius $R_i < L/2$. Firm i's footprint is the arc

$$M_i \;=\; \{s \in \mathbb{T}_L : \delta(s, c_i) \le R_i\},$$

which has Lebesgue measure $|M_i| = 2R_i$.

For a given displacement $x \ge 0$, define the overlap set

$$O_i(x) \;\equiv\; \{s \in \mathbb{T}_L : s \in M_i \text{ and } s + x \in M_i\}.$$

We claim that

$$|O_i(x)| \;=\; (2R_i - x)_+. \tag{76}$$

To verify, note that $s \in M_i$ requires $\delta(s, c_i) \le R_i$, i.e., $s \in [c_i - R_i,\, c_i + R_i]$ (modulo L). Similarly, $s + x \in M_i$ requires $\delta(s + x, c_i) \le R_i$, i.e., $s \in [c_i - R_i - x,\, c_i + R_i - x]$. The overlap set $O_i(x)$ is therefore the intersection of two arcs of length $2R_i$ whose centers are displaced by x:

$$O_i(x) \;=\; [c_i - R_i,\, c_i + R_i] \;\cap\; [c_i - R_i - x,\, c_i + R_i - x].$$

Since both arcs have length $2R_i$ and their centers are separated by x, the intersection has length

$$|O_i(x)| \;=\; \max\{2R_i - x,\, 0\} \;=\; (2R_i - x)_+,$$

provided $x \le 2R_i < L$, which holds since $R_i < L/2$. When $x > 2R_i$, the arcs are disjoint and $|O_i(x)| = 0$. This establishes (76).

Step 2: Summing over firms. By definition of the coverage overlap kernel,

$$\mathcal{K}(x) \;=\; \frac{1}{L}\int_{\mathbb{T}_L}\sum_{i=1}^{N}\mathbf{1}\{s\in M_i\}\,\mathbf{1}\{s+x\in M_i\}\,ds \;=\; \frac{1}{L}\sum_{i=1}^{N}|O_i(x)| \;=\; \frac{1}{L}\sum_{i=1}^{N}(2R_i-x)_+.$$

This is (40).

Step 3: Expectation. Under Assumption 5.2, cores $\{c_i\}$ and blueprint qualities $\{z_i\}$ are drawn independently. Note that (40) does not depend on the core locations $\{c_i\}$—only on the scope radii $\{R_i\}$—because the overlap of two concentric arcs displaced by a fixed x depends only on the arc length, not on the center.

Taking expectations over $\{z_i\}$:

$$\mathbb{E}[\mathcal{K}(x)] \;=\; \frac{1}{L}\sum_{i=1}^{N}\mathbb{E}\big[(2R_i-x)_+\big] \;=\; \frac{N}{L}\,\mathbb{E}\big[(2R-x)_+\big],$$

where the second equality uses the iid assumption on $\{z_i\}$, so that $R_i=\Phi(z_i)$ are iid with common distribution. Substituting $R=\Phi(z)$ gives (41). □

Proof of Proposition 5.3. Part (i): Pointwise comparison. Fix $x\geq 0$. By Proposition 5.2,

$$\mathbb{E}[\mathcal{K}(x)] \;=\; \frac{N}{L}\,\mathbb{E}\big[(2\Phi(z)-x)_+\big], \qquad \mathbb{E}[\widetilde{\mathcal{K}}(x)] \;=\; \frac{N}{L}\,\mathbb{E}\big[(2\widetilde{\Phi}(z)-x)_+\big].$$

Since $\widetilde{\Phi}(z)\geq\Phi(z)$ for all z, we have $2\widetilde{\Phi}(z)-x\geq 2\Phi(z)-x$ for all z, and therefore

$$(2\widetilde{\Phi}(z)-x)_+ \;\geq\; (2\Phi(z)-x)_+ \qquad \text{for all } z.$$

Taking expectations preserves the inequality:

$$\mathbb{E}\big[(2\widetilde{\Phi}(z)-x)_+\big] \;\geq\; \mathbb{E}\big[(2\Phi(z)-x)_+\big],$$

establishing the weak inequality in (42).

For strict inequality, fix $x\in(0,\,2\max\operatorname{supp}(\widetilde{R}))$. By assumption, there exists a set $A\subseteq\operatorname{supp}(F_z)$ with $F_z(A)>0$ such that $\widetilde{\Phi}(z)>\Phi(z)$ for all $z\in A$. For any $z\in A$ with $2\widetilde{\Phi}(z)>x$ (which holds for a positive-measure subset of A since $x<2\max\operatorname{supp}(\widetilde{R})$), the function

$r \mapsto (2r - x)_+$ is strictly increasing at $r = \widetilde{\Phi}(z)$, so

$$(2\widetilde{\Phi}(z) - x)_+ \;>\; (2\Phi(z) - x)_+.$$

Since this strict inequality holds on a set of positive F_z-measure, the expected value is strictly larger:

$$\mathbb{E}\big[(2\widetilde{\Phi}(z) - x)_+\big] \;>\; \mathbb{E}\big[(2\Phi(z) - x)_+\big].$$

Part (ii): Aggregate concentration comparison. Integrate the pointwise inequality from part (i) over $x \in \mathbb{T}_L$:

$$\mathbb{E}\left[\frac{1}{L}\int_{\mathbb{T}_L} \widetilde{\mathcal{K}}(x)\,dx\right] \;=\; \frac{1}{L}\int_{\mathbb{T}_L} \mathbb{E}[\widetilde{\mathcal{K}}(x)]\,dx \;\geq\; \frac{1}{L}\int_{\mathbb{T}_L} \mathbb{E}[\mathcal{K}(x)]\,dx \;=\; \mathbb{E}\left[\frac{1}{L}\int_{\mathbb{T}_L} \mathcal{K}(x)\,dx\right],$$

where the exchange of expectation and integration is justified by nonnegativity and Fubini's theorem.

Strict inequality follows because the pointwise strict inequality from part (i) holds for all x in the open interval $(0,\ 2\max \mathrm{supp}(\widetilde{R}))$, which has positive Lebesgue measure. Therefore the integral is strictly larger under $\widetilde{\Phi}$. □

Additional properties of the coverage overlap kernel. For completeness, we record several properties of $\mathcal{K}$ that are used in the discussion following Proposition 5.2.

Lemma A.3 (Monotonicity and support of the coverage kernel)**.** *Under Assumption 5.2:*

(i) $\mathcal{K}(x)$ *is weakly decreasing and convex on* $[0, L/2]$.
(ii) $\mathcal{K}(x) = 0$ *for all* $x \geq 2\max_i R_i$.
(iii) $\mathcal{K}(0) = \frac{1}{L}\sum_{i=1}^N 2R_i = \frac{2}{L}\sum_{i=1}^N \Phi(z_i)$.
(iv) The integral of the coverage kernel satisfies

$$\frac{1}{L}\int_0^{L/2} \mathcal{K}(x)\,dx \;=\; \frac{1}{L^2}\sum_{i=1}^N 2R_i^2. \tag{77}$$

Proof. All properties follow from the closed form $\mathcal{K}(x) = \frac{1}{L}\sum_i (2R_i - x)_+$.

Part (i). Each summand $(2R_i - x)_+$ is piecewise linear, weakly decreasing, and convex in x. A nonnegative linear combination of convex decreasing functions is convex and decreasing.

Part (ii). If $x \geq 2\max_i R_i$, then $2R_i - x \leq 0$ for all i, so every summand is zero.

Part (iii). At $x = 0$: $(2R_i - 0)_+ = 2R_i$, so $\mathcal{K}(0) = \frac{1}{L}\sum_i 2R_i$.

Part (iv).

$$\int_0^{L/2} \mathcal{K}(x)\,dx = \frac{1}{L}\sum_{i=1}^{N}\int_0^{L/2}(2R_i - x)_+\,dx \;=\; \frac{1}{L}\sum_{i=1}^{N}\int_0^{2R_i}(2R_i - x)\,dx$$
$$= \frac{1}{L}\sum_{i=1}^{N}\left[2R_i x - \frac{x^2}{2}\right]_0^{2R_i} \;=\; \frac{1}{L}\sum_{i=1}^{N}\left(4R_i^2 - 2R_i^2\right) \;=\; \frac{1}{L}\sum_{i=1}^{N} 2R_i^2.$$

Dividing by L gives (77). □

Lemma A.4 (Aggregate coverage concentration under Pareto blueprint quality). *Suppose* $z \sim \mathrm{Pareto}(\alpha_z, \underline{z})$ *and* $\Phi(z) = z^{1/\nu}$ *for some* $\nu > 1$. *Then* $R = \Phi(z)$ *has tail index* $\nu\alpha_z$ *and*

$$\mathbb{E}[R^2] \;=\; \frac{\alpha_z\,\underline{z}^{2/\nu}}{\alpha_z - 2/\nu} \qquad \textit{provided } \alpha_z > 2/\nu.$$

The expected aggregate coverage concentration $\frac{1}{L}\int_0^{L/2}\mathbb{E}[\mathcal{K}(x)]\,dx = \frac{2N}{L^2}\mathbb{E}[R^2]$ *is therefore finite and decreasing in* ν*: stronger organizational complexity (larger* ν*) lowers aggregate coverage concentration.*

Proof. Since $z \sim \mathrm{Pareto}(\alpha_z, \underline{z})$, the random variable $R = z^{1/\nu}$ satisfies $\Pr(R > r) = \Pr(z > r^\nu) = (\underline{z}/r^\nu)^{\alpha_z} = \underline{z}^{\alpha_z}\, r^{-\nu\alpha_z}$ for $r \geq \underline{z}^{1/\nu}$, confirming the tail index $\nu\alpha_z$.

The second moment is

$$\mathbb{E}[R^2] \;=\; \mathbb{E}[z^{2/\nu}] \;=\; \int_{\underline{z}}^{\infty} x^{2/\nu}\,\alpha_z \underline{z}^{\alpha_z} x^{-\alpha_z - 1}\,dx \;=\; \alpha_z \underline{z}^{\alpha_z}\int_{\underline{z}}^{\infty} x^{2/\nu - \alpha_z - 1}\,dx.$$

The integral converges if and only if $\alpha_z > 2/\nu$, in which case

$$\mathbb{E}[R^2] \;=\; \alpha_z \underline{z}^{\alpha_z}\cdot\frac{\underline{z}^{2/\nu - \alpha_z}}{\alpha_z - 2/\nu} \;=\; \frac{\alpha_z\,\underline{z}^{2/\nu}}{\alpha_z - 2/\nu}.$$

By Lemma A.3(iv), expected aggregate coverage concentration is $\frac{2N}{L^2}\mathbb{E}[R^2]$, which is increasing in $\mathbb{E}[R^2]$. Since $\mathbb{E}[R^2] = \frac{\alpha_z \underline{z}^{2/\nu}}{\alpha_z - 2/\nu}$ is decreasing in ν (both the numerator $\underline{z}^{2/\nu}$ decreases and the denominator $\alpha_z - 2/\nu$ increases as ν rises), aggregate coverage concentration is decreasing in ν. □